

Sensors and Electronic Instrumentation Advances

Proceedings

**of the 12th International Conference
on Sensors and Electronic
Instrumentation Advances
(SEIA' 2026)**

**30 September - 2 October 2026
Granada, Spain**



Sensors and Electronic Instrumentation Advances:

**Proceedings of the XII International Conference
on Sensors and Electronic Instrumentation Advances**

30 September - 2 October 2026

Granada, Spain

Edited by Sergey Y. Yurish



Sergey Y. Yurish, *Editor*
Sensors and Electronic Instrumentation Advances
SEIA' 2026 Conference Proceedings

Copyright © 2026

by International Frequency Sensor Association (IFSA) Publishing, S. L.

E-mail (for orders and customer service enquires): ifsa.books@sensorsportal.com

Visit our Home Page on <http://www.sensorsportal.com>

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (IFSA Publishing, S. L., Barcelona, Spain).

Neither the authors nor International Frequency Sensor Association Publishing accept any responsibility or liability for loss or damage occasioned to any person or property through using the material, instructions, methods or ideas contained herein, or acting or refraining from acting as a result of such use.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identifying as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

ISBN: 978-84-09-90031-2

BN-20260922-XX

BIC: TJFC

Contents

Foreword	5
Compact UVN + LWIR LIBS Instrument with Machine-Learning Analysis for Real-Time In-Situ Resource Identification on Lunar and Martian Surfaces.....	6
<i>C. S.-C. Yang F. Jin, S. B. Trivedi, Amber R. Benecchi and Thirimachos Bourlai</i>	
Dependence of Segment Anything Model (SAM) Performance on Part Geometry in Bin Picking Applications.....	8
<i>M. Franaszek, P. Pilipchak and K. S. Saidi</i>	
Low-Complexity Trajectory Estimation for Table-Slide Rehabilitation Using an IMU Sensor	14
<i>Y. Uchida, E. Ohkubo, M. Nishi and T. Funayama</i>	
Future Transformation Prediction of Land Use/Land Cover Using Machine Learning and Neural-Network-Based Transition Potential Modeling in the Middle Rio Grande Region 2025-2040	18
<i>Omar. S. Belhaj, Stanley. T. Mubako, Raed Aldouri, Craig Tweedie, Elhadi. A. Hadia, Rabia. E. Al Mamlook, William Hargeove, Elizabeth Walsh and Alex Mayer</i>	
Current Status and Future Development of National Standard Time UTC(NTSC)	25
<i>Shuhong Zhao, Yanrong Xue, Shanshan Bai, Yu Zhang, and Sufang Liu</i>	
Spatially-Resolved Conductivity Mapping of Ultra-Thin Glass Using Guard Ring Electrodes and Silver Conductive Ink Contact Layers	30
<i>D. Otto, M. Eiserbeck, D. Käßler and G. Bausch</i>	
A Fundamental Evaluation of a Ventilator Operation Education System Applying Eye-Tracking Technology	35
<i>N. Katsura, S. Sakakida, R. Kurihara, Y. Uchida and T. Wabe, A. Fujii</i>	
Estimating Outdoor Eye-Level Daylight Exposure from Satellite Data: Protocol for a Proof-of-Concept Experiment.....	40
<i>S. W. de Vries, J. Mardaljevic, H. S. M. Kort and J. van Duijnhoven</i>	
Investigation of the Flow Rate Dependence of Blood Coagulation Measurement in an Extracorporeal Circulation Circuit Using LED Photoacoustic Imaging.....	44
<i>Takahiro Wabe, Naoki Katsura, Yohsuke Uchida, Ryo Suzuki and Akimitsu Fujii</i>	
Long-Term Objective Evaluation of Posture during Self-Study Using Image Analysis Technology and Impact on Learning Effectiveness	50
<i>Takahiro Wabe and Yasutaka Uchida</i>	
A Deterministic On-Demand Simulation Framework for Parametric Analysis of Coincidence-Based Frequency Measurement	56
<i>Avalos-Gonzalez Daniel, Misael Reyes I., Islas-Pereda S., José M. Ramírez, Pérez-Pérez Blanca Dalila, Zamora-Machado Marlene and Martínez-Quintero Luis</i>	
A Cuffless Palm-Site On-Demand Continuous Blood Pressure Monitoring Prototype with Deep Learning-Based Estimation	63
<i>Iyappan Gunasekaran, Lawson Michael, Kannie Wy Chan</i>	
Fast Electrical Impedance Spectroscopy for Transient Dynamic Processes.....	69
<i>Lukas Lauritsch and Marco Da Silva</i>	
Predictive Anemometric Surveillance: Forecasting the IEC 61400-12-1 Annex K Wind-Speed Discrepancy for Proactive Sensor Fault Detection.....	73
<i>Máximo Paz Iaconis, Lucas Carmusciano and Juan Manuel Gonzalez Kapnik</i>	
Deep Latent Clustering of Kinematic Waveforms for Unbiased Functional Profiling.....	78
<i>Shmuel Amour and Sara Rosenblum</i>	
Simulation-Based Study of a GMR Sensor Device for Event-Driven Tracking of Magnetic Microparticles	84
<i>A. Shahryari, F. Pardo, J.A. Boluda, C. Reig, R. Afonso and S. Cardoso</i>	

A Low-Cost SDR-Based Telecommunications Laboratory Using Repurposed TV Boxes	89
<i>Arquimedes Lopes da Silva, Diego F. C. Siqueira, Leandro L. Alves, Maria Clara R. N. and Eliézer Alves Teixeira</i>	
Computer Vision Analysis of Sanding Marks in Wood.....	94
<i>M. Colina, J. Ramos, L. Marín and C. Reig</i>	
Deterministic Modeling and Advanced Micro-Doppler Analysis of Quadrotor Drones in the X-Band	99
<i>A. Le Dréo, A. Khenchaf, J.-C. Cexus and P. Pouliguen</i>	
Impact of Sparsity and Algorithm Parameters on SFT-based Wideband Spectrum Sensing	108
<i>N. Luberne, A. Khenchaf, F. Comblet, B. Gorisse and M. Bauche</i>	
Design and Evaluation of an ESP32-Based Wireless Sensor Network for Subsurface Temperature Monitoring in Road Energy-Harvesting Applications	113
<i>T. Bača, P. Šarařin, O. Karpiř and M. Formanek</i>	
AI-Based Virtual Sensors as Cybersecurity Mechanisms in Industrial IoT Systems: A Systematic Literature Review, Taxonomy, and Cyber-Resilience Framework.....	118
<i>I. Gabriřov, G. Koman and M. Kubina</i>	
Utilisation of Ultrasonic Anemometers for Acquiring Airflow Data in Tunnel Fire Spread Models....	124
<i>M. Hodoř, L. Formanek and P. Daniřovič</i>	
Distributed VOC and Environmental Monitoring System for Road Tunnel Applications.....	130
<i>M. Hodoř, L. Formanek and P. Daniřovič</i>	
Wind Feature Engineering for Tree-Sway Energy Harvesting Assessment in Forest IoT Nodes	136
<i>P. Kolok, J. Kapitulik, M. Chochul and J. řumský</i>	
Structural Optimization of a Circular Restoring Spring for a Low-Frequency Electromagnetic Standard Vibrator	141
<i>X. F. Zhang, T. Xu, T. T. Du and M. C. Xu</i>	
Hydrostat – Real-Time Monitoring of Water Quality Parameters	144
<i>M. Kubařčík, A. Tupý, M. Tebeřk and P. řevčík</i>	
Measurement Error Analysis of Eddy-Current Sensors under Tilted Probe Installation	149
<i>X. F. Zhang, T. Xu and L. K. Wu</i>	
Ultra-Low-Power Odor Classification on the ESP32-C6 Using Integer-Based Models	152
<i>M. Kubařčík, P. Kolok, L. řechovič and M. Chochul</i>	
Design and Preliminary Validation of a Modular ESP32-H2 Sensor Node.....	158
<i>Andrej Brnuřk, Andrej řimkovič, Peter řarařin, Matuř Formanek and Tomř Bača</i>	
A Comparative Study of Variational Autoencoder Architectures for Anomaly-Based Fault Detection in an Industrial Control System	161
<i>Ietezaz ul Hassan, Krishna Panduruv, Daniel Riordan and Prof. Joseph Walsh</i>	
Resistance-Based Thermal Sensing and Heat Propagation Modeling of a PTC-Heated Shape Memory Alloy Actuator	166
<i>D. Arriaga-Ventura, Rogelio Bustamante-Bello and Javier Izquierdo Reyes</i>	
Cost-Effective Hybrid Raman-Infrared Non-Contact Thermometry for Industrial Process Monitoring.....	173
<i>V. Multian, C. W. Rischau and J. Teyssier</i>	
Ultra-Thin Inkjet-Printed Metal/Polymer Composite Strain Sensors for Wearable Health Monitoring.....	177
<i>A. Khan, A. S. Nittala, S. Ali, A. Hassan and A. Bermak</i>	
Atmospheric Hydrometeor Monitoring Using 24 GHz Microwave Attenuation	180
<i>B. H. Kwon and Z. Seo</i>	
Wearable Heartbeat Sensors from Laser-Induced Graphene on a Poly(Dimethylsiloxane) Substrate with Poly(Ethylene Glycol) and Triton X Additives.....	183
<i>M. Spasenovič, A. Gavran, M. Tomič, I. Peřić, T. Vičentič and M. V. Pergal</i>	
Design and Validation of a Time-Frequency Data Analysis and Data Processing Platform.....	188
<i>Y. R. Xue, S. H. Zhao, L. D. Meng, Y. F. Sun, Y. W. Li, Y. Liu, and X. H. Li</i>	

Foreword

Over more than a decade, the SEIA Conference Series has developed into an international forum for researchers, engineers, developers, and practitioners working across the broad and rapidly evolving fields of sensors, measurement science, electronic instrumentation, and related technologies. The conference continues its principal mission of bringing together fundamental research, advanced engineering solutions, and practical applications, while encouraging the exchange of ideas between academia, research institutions, and industry.

The contributions collected in this volume reflect the remarkable diversity and interdisciplinary character of contemporary sensing and instrumentation research. They address developments ranging from novel sensing principles and measurement methods to intelligent sensor systems, embedded electronics, signal processing, artificial intelligence, machine learning, IoT technologies, industrial monitoring, environmental sensing, biomedical applications, robotics, telecommunications, and advanced metrology.

A particularly visible trend in this year's contributions is the increasing integration of sensing technologies with AI and data-driven methods. Machine learning is no longer used only as a post-processing tool; it is becoming an integral component of sensor systems themselves, enabling autonomous interpretation, anomaly detection, predictive monitoring, intelligent classification, and real-time decision support. At the same time, advances in low-power embedded platforms, wireless sensor networks, edge computing, virtual sensors, and distributed measurement systems are creating new opportunities for deploying intelligent sensing solutions in environments where reliability, energy efficiency, security, and scalability are essential.

The papers presented at SEIA' 2026 also demonstrate the continuing importance of measurement science and instrumentation fundamentals. Accurate calibration, traceability, uncertainty evaluation, signal acquisition and processing, sensor characterization, and reliable measurement remain essential foundations upon which new intelligent and connected systems must be built.

The applications represented in these Proceedings include industrial and infrastructure monitoring, environmental measurements, healthcare and rehabilitation, wearable sensing, intelligent manufacturing, transportation, energy systems, remote and planetary sensing, time and frequency measurements, and advanced scientific instrumentation. Such diversity clearly illustrates how sensor and measurement technologies have become enabling elements of modern technological development across virtually every area of science and engineering.

SEIA' 2026 also places particular emphasis on transforming research results into practical and scalable solutions. Progress in sensors and instrumentation increasingly depends not only on achieving improved sensitivity or accuracy, but also on addressing integration, robustness, interoperability, power consumption, manufacturability, cybersecurity, and long-term reliability. These challenges require close cooperation between specialists from different disciplines and reinforce the importance of international scientific meetings such as SEIA.

The preparation of these Proceedings would not have been possible without the contributions of many people. I would like to express my sincere gratitude to all authors for submitting their research and sharing their results with the international community. My special appreciation goes to the members of the IPC and the reviewers for their time, expertise, and careful evaluation of the submitted papers. Each paper has undergone careful peer review to ensure scientific quality and relevance.

I hope that these Proceedings will serve not only as a record of the scientific results presented in Granada, but also as a useful reference for researchers and engineers and as a source of new ideas, collaborations, and future developments in sensors, electronic instrumentation, and measurement technologies.

Prof., Dr. Sergey Y. Yurish,
SEIA' 2026 Conference Chairman

Compact UVN + LWIR LIBS Instrument with Machine-Learning Analysis for Real-Time In-Situ Resource Identification on Lunar and Martian Surfaces

C. S.-C. Yang¹, F. Jin¹, S. B. Trivedi¹, Amber R. Benecchi² and Thirimachos Bourlai²

¹ Brimrose Technology Corporation, Sparks, MD 21152, USA

² Multispectral Imagery Lab, College of Engineering, University of Georgia, Athens, GA 30602, USA

Tel.: 14104727070

E-mail: claytony@brimrose.com

Summary: Real-time identification of water, volatiles, minerals, and rare earth elements is essential for future lunar and Martian in-situ resource utilization. This paper presents the Fast Elemental and Molecular Interrogation Sensor (FEMIS), a compact instrument that combines ultraviolet-visible-near-infrared (UVN) and long-wave infrared (LWIR) laser-induced breakdown spectroscopy with machine-learning analysis. By collecting atomic/ionic UVN emission and molecular LWIR emission from the same laser-induced plasma, FEMIS provides complementary elemental and molecular information in a single measurement. Representative tests on regolith simulants, water-bearing samples, nitrate hydrates, organic analogs, and rare-earth-containing compounds showed diagnostic signatures under both ambient and near-vacuum conditions. Combined UVN + LWIR machine-learning models improved material discrimination and concentration prediction, supporting trace-level detection down to about 100 ppm. These results demonstrate FEMIS as a compact AI-enabled chemical sensor for autonomous planetary resource assessment.

Keywords: LIBS, Infrared sensors, Chemical sensors, Machine learning, Lunar ISRU, planetary instrumentation, Autonomous sensing.

1. Introduction

Rapid identification of local resources is a key requirement for sustained exploration of the Moon and Mars. Water ice, hydrated minerals, volatiles, rare earth elements, and construction-relevant regolith components must be characterized in real time with minimal sample handling. FEMIS addresses this need by extending conventional LIBS into a dual-band sensing approach that acquires elemental information in the ultraviolet-visible-near-infrared (UVN) and molecular information in the long-wave infrared (LWIR) from the same laser-induced plasma event [1,2].

2. Instrument Concept and Representative Results

FEMIS was evaluated on lunar and Martian regolith simulants, major mineral components, water-bearing samples, hydroxyl-containing compounds, nitrate hydrates, organic analog materials, and rare-earth-containing compounds. The UVN channel covers approximately 200-850 nm and provides atomic and ionic signatures, while the LWIR channel covers approximately 5-12 μm and captures molecular vibrational features in the fingerprint region. Representative tests showed UVN emission signatures from elements relevant to lunar regolith, including Ca, Si, Mg, Na, Ti, Fe, Al, Mn, H, and O. The LWIR spectra revealed broad molecular features associated

with silicate structures near 9.2 μm , water bending near 6.3 μm , and nitrate-related emission near 7.6 μm . These dual-band measurements demonstrate the value of combining elemental and molecular contrast in a single laser shot, as illustrated in Fig. 1.

3. Low-pressure validation and intelligent analysis

A vacuum/gas chamber was integrated with the FEMIS measurement system to evaluate performance under lunar- and Martian-relevant pressure conditions. Spectra were acquired from ambient pressure down to near vacuum (~ 0.001 atm). Although both UVN and LWIR signal intensities decreased by roughly a factor of six under vacuum, the key diagnostic spectral features remained identifiable, confirming that the dual-band LIBS approach remains viable under low-pressure planetary environments.

For automated interpretation, convolutional and residual neural-network models [3] were trained using UVN spectra, LWIR spectra, and combined UVN + LWIR inputs. Importantly, the AI analysis showed that the LWIR channel retained stronger prediction capability even under vacuum conditions. The LWIR-based model achieved a higher Kendall's tau value [4] than the UVN-based model, 0.6604 versus 0.4121, indicating that LWIR molecular features remain more predictive and robust after pressure-induced changes in plasma emission.




Fig. 1. Simultaneous (a–d) UV-Vis-NIR and (e) LWIR LIBS spectra of Cerium(III) nitrate hexahydrate, Lunar Highland Regolith Simulant Exolith LHS-1, and their 50/50 mixture. The spectra show Ce III emission lines, H₂O and NO₃ molecular features, and elemental/silicate signatures from the lunar regolith simulant, demonstrating LIBS-based multi-component analysis for lunar exploration.

Overall, the combined UVN + LWIR models provided the best concentration prediction and material discrimination, accurately identifying target materials across a wide concentration range from pure samples down to approximately 100 ppm. This result shows that FEMIS can use LWIR molecular information to maintain reliable AI-assisted material identification even in low-pressure ISRU-relevant environments.

4. Compact Instrument Implementation

The built compact FEMIS instrument is shown in Fig. 2, with approximate dimensions of 90 × 60 × 20 cm³, a mass of approximately 16 kg, and steady-state power consumption near 20 W. It incorporates a compact Q-switched laser, beam expansion and focusing optics, a UVN spectrometer, an LWIR grating spectrometer, and a cooled MCT detector. Embedded control electronics coordinate laser firing, detector timing, spectral acquisition, preprocessing, and real-time material identification. Although the current configuration was built as a laboratory-compatible prototype, its optical and electronic subsystems are modular, and the overall dimensions and mass can be readily reduced through tighter optomechanical packaging and further component miniaturization.




Fig. 2. Schematic of the compact FEMIS instrument combining a UVN spectrometer and an LWIR grating spectrometer for dual-band elemental and molecular LIBS sensing.

5. Conclusions

FEMIS provides a practical route toward real-time, non-contact elemental and molecular sensing for lunar and Martian resource utilization. The instrument detects regolith-related elements, water, hydroxyl-bearing compounds, nitrates, rare-earth-related signatures, and organic analog materials under both ambient and near-vacuum conditions. When combined with machine-learning analysis, the dual-band UVN + LWIR approach improves interpretation of complex spectra and supports sensitive material identification down to the 100 ppm level. These results position FEMIS as a compact AI-enabled planetary field instrument for autonomous resource mapping and future low-SWaP exploration payloads.

Acknowledgements

This material is based upon work supported by the National Aeronautics and Space Administration under STTR Contract No. 80NSSC23PB380.

References

- [1]. C. S.-C. Yang, F. Jin, S. Trivedi, U. Hommerich, et al., Long-wave infrared laser-induced breakdown spectroscopy of complex gas molecules in the vicinity of a laser-induced plasma, *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, Vol. 294, 2023, Article 122536.
- [2]. C. S.-C. Yang, F. Jin, S. Trivedi, U. Hommerich, et al., Comparative long-wave infrared laser-induced breakdown spectroscopy employing 1-D and 2-D focal plane array detectors, *Sensors*, Vol. 23, Issue 3, 2023, Article 1366.
- [3]. K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2016, pp. 770–778.
- [4]. M. G. Kendall, A new measure of rank correlation, *Biometrika*, Vol. 30, Issues 1–2, 1938, pp. 81–93.

Dependence of Segment Anything Model (SAM) Performance on Part Geometry in Bin Picking Applications

M. Franaszek, P. Piliptchak and K. S. Saidi

National Institute of Standards and Technology, Gaithersburg, MD 20899, USA
E-mail: marek@nist.gov

Summary: The Segment Anything Model (SAM) can segment novel, previously unseen objects from 2D images. This capability is especially important in automated bin picking applications, where low-volume and high-variability scenarios are common. When prompted by a point, SAM outputs three masks with varying confidence scores. Prior research indicates that SAM's performance depends on the type of parts in a bin: for parts that are visually decomposable into subcomponents, correctly segmented masks are more often associated with a medium confidence score. Conversely, for parts that are not readily visually decomposable into subcomponents, correct masks are typically associated with the highest confidence score. We investigated the transition from visually decomposable to non-decomposable part categories in simulation by using computer-aided design (CAD) models with gradually altered geometric features. Differences in the dimensions of these geometric features cause corresponding differences in grayscale images. We found that this transition occurs within a relatively narrow range of feature modifications and image intensities. This demonstrates the strong sensitivity of SAM to the specific types of parts within a bin. Consequently, prior knowledge of a part's complexity can be leveraged to better tune the processing of SAM's output.

Keywords: Foundation models, Segment anything model, Robotic bin picking.

1. Introduction

The successful deployment of bin picking systems to automate manufacturing tasks, such as assembly or kitting, depends strongly on vision system performance. The vision system must segment the acquired depth images to determine the six degree-of-freedom (6DoF) pose of individual parts, enabling the integrated robot to execute tasks successfully.

Accurate instance segmentation of unorganized piles of manufacturing parts remains especially challenging. While early segmentation algorithms required model retraining for different parts [1], the emergence of foundation models such as the Segment Anything Model (SAM) has made retraining for novel parts almost obsolete [2]. When prompted by a point in a query image or a bounding box, SAM outputs three binary segmentation masks with distinct confidence scores.

Continued research on prompts has led to the adoption of Large Language Models (LLMs), which allow the use of text-based prompts to obtain segmentation masks for specified objects [3-5]. These new Promptable Concept Segmentation (PCS) techniques have yielded promising results for semantic segmentation when tested on Common Objects in Context (COCO)-style databases [6]. However, extensive testing on images of industrial parts randomly filling a bin still needs to be conducted.

Another kind of prompt in the form of a computer-aided design (CAD) model was recently proposed [7]. It is especially useful for bin picking applications where CAD models of the parts filling the bins are usually available. A modified version of SAM (SAM3) [5] was trained on synthetic images created with the NVIDIA Isaac Sim simulator [8], which was

used to build multi-object scenes utilizing CAD models from the ABC dataset of industrial parts [9]. It was tested on two standard industrial benchmark datasets, T-LESS and ITODD, which do not contain many heavy clutter scenes typical of bin picking applications [10, 11].

Currently, using point prompts in depth images remains the easiest way to obtain instance segmentation masks. To maximize the efficiency of robotic path planning, the vision system should provide as many correct masks as possible, prompted from points scattered across the entire image area. Previous work demonstrated that SAM's performance depends heavily on the type of parts in the bin [12]. For parts that are visually decomposable into subcomponents (as in Fig. 1(a)), accurate masks from SAM's multi-mask output were most often associated with the medium (second-highest) confidence score. For parts lacking obvious subcomponents (as in Fig. 1(b)), accurate masks typically correlated with the highest confidence score.




Fig. 1. Examples of two types of parts: a) visually decomposable into subcomponents (drawn with different colors); b) not readily decomposable into subcomponents.

Previously, the distinction between the two types of parts was based on subjective, unquantifiable

evaluation. In this study, a more objective metric is formulated. We systematically investigate how SAM's performance shifts during the transition from one type of part to another by gradually altering the geometric features of the part within a simulated bin. Our results suggest that this transition occurs over a relatively narrow range of geometric changes. Those changes result in corresponding changes in the 8-bit grayscale intensity of depth images. We found that two compact, adjacent regions on a depth image whose intensities differ by less than 15 will most likely be segmented as one region that stretches over both adjacent areas. Thus, a part represented in a depth image as a patch where no compact subregions can be identified with intensity differences greater than $\Delta I = 15$, should be classified as non-decomposable part.

Patches consisting of subregions with intensity differences much larger ($\Delta I \gg 15$) correspond to decomposable parts for which binary masks associated with the medium score will be more accurate. Our results reveal a specific range of differences in geometric features, along with a corresponding range of intensities in a grayscale image, where no single strategy (strongest vs. medium score mask) is preferred. In such mixed regions, to maximize the number of correctly segmented parts in a bin, masks associated with both the strongest and medium confidence scores may need to be analyzed.

Thus, prior knowledge of a part's dimensions and shape can be utilized to optimize the processing of SAM's output, thereby improving overall bin picking system performance.

2. Simulation

We generated two types of grayscale images: one showing a bin filled with screws, and the other showing a mosaic of tiles with different intensities. The first type yielded a realistic-looking pile of screws, but to check SAM's performance, many images had to be visually evaluated. The second type shows artificial scenes composed of geometric primitives for which a metric gauging SAM's performance can be calculated by running autonomous image processing algorithms. An 8-bit grayscale was used for all images, and binary masks output by SAM associated with the largest confidence score were analyzed.

2.1. Generation of Screws Images

We use a physics-based simulator to simulate dropping parts into a bin and generate depth images derived from the known distances between the simulated parts and bin bottom. Using the CAD models of an M12 and M8 hex screw, we generated a series of $N = 20$ models for each with a gradually decreasing head diameter, $D_h(n)$, while keeping the stem diameter D_s , stem length L_s , and head length L_h fixed (see Fig. 2). The deformation rate, $\alpha(n)$, is defined as a ratio:

$$\alpha(n) = \frac{D_h(n) - D_s}{D_h - D_s}, \quad (1)$$

where $n = 0, \dots, N$, and $D_h(0) = D_h$ corresponds to the head diameter in the original CAD model, while $D_h(N) = D_s$, as illustrated in Fig. 2, and $D_h(n)$ are equally distributed between D_s and D_h . In Table 1, the relevant dimensions of the unmodified CAD models for both screws are provided, where $L_{tot} = L_h + L_s$ is the total length of the screw.

Table 1. Screw dimensions in [mm].

	L_{tot}	L_s	L_h	D_s	D_h
M8	26	18	8	8	13
M12	32	20	12	12	18

For each n -th CAD model, we dropped a fixed number of parts into a bin and saved the resulting depth image in orthonormal projection. The initial pose of each part was randomized, and this step was repeated for $j = 1, \dots, 10$ to obtain ten distinct instances of unorganized piles for each part.




Fig. 2. Two limiting cases of the CAD models: the original with head diameter $D_h(0)$ and the most reduced head diameter where $D_h(N) = D_s$.

For each (n, j) saved depth image, we selected prompt points from a regular grid. We used the same bin for dropping M12 and M8 screws; therefore, the grid size and the number of dropped parts were set differently to ensure a similar complexity in the resulting unorganized pile. For M8 screws, the grid size was 12×14 and the number of dropped parts was 54; for M12 screws, the corresponding parameters were 11×12 and 48. Points falling on the background (the bin bottom) were ignored, as shown in Fig. 3. The accepted $M_j(n)$ points were used as prompts for SAM segmentation.

The resulting $M_j(n)$ binary masks were visually evaluated, and only those containing both the stem and the head of the screw were counted. This procedure enabled tracking the transition between the two categories of parts (visually decomposable into

subcomponents and with no obvious subcomponents) as the deformation rate $\alpha(n)$ was swept across its range. Accepted masks fell into two categories: the first, where the entire screw was unobstructed and the mask covered the full stem and head facing the camera; and the second, where only a portion of the screw was visible due to partial occlusion by neighboring parts. Example of masks from the first category are shown in Fig. 4 while masks belonging to the second category are shown in Fig. 5. While some masks in the second category may not always be good candidates for picking, they were included in the total count of correct masks since they represent a challenging test for segmentation. Not all masks containing both stem and head were accepted: in some cases, the segmentation successfully crossed the boundary between stem and head but ultimately failed to produce correct masks (e.g., due to under-segmentation or leaking). Examples of rejected cases are shown in Fig. 6. The masks in Fig. 6a) and d) are under-segmented while masks in Fig. 6b) and c) leaked.




Fig. 3. Depth image displaying the grid of points: yellow circles denote ignored background points, while $M_j(n)$ blue squares indicate active prompt points for SAM.

In this way, the number of accepted masks, $K_j(n)$, was determined. For each n -th deformed CAD model and each j -th instance of a pile, a fraction $\rho_j(n) = K_j(n)/M_j(n)$ was calculated from which the corresponding means $\bar{\rho}(n)$ and standard deviations $\delta(n)$ were calculated for M8 and M12 screws.

2.2. Generation of Images with a Mosaic of Tiles

We start with an image completely filled with the background (intensity 0) and all perimeter pixels set to maximum (intensity 255). Such a scene mimics the image of an empty bin. The entire background area is then divided into 11×9 squares, and the intensity in each square is set to a random number from (0,255). Then, each square tile is further divided into two rectangles of unequal areas, and the intensity inside the smaller rectangle is increased by $\Delta I(n) = \Delta I_{max} \alpha(n)$ where $0 \leq \alpha \leq 1$ plays the same role as deformation factor in (1), and ΔI_{max} corresponds to $D_h - D_s$ for

screws. Finally, to simulate noise, a random integer from $\{-1,0,1\}$ is added to the intensities of all pixels in all rectangles while keeping all final values clipped to the (0,255) interval. Examples of such images are shown in Fig. 7.




Fig. 4. Example of correctly segmented masks from the first category: unobstructed masks of M8 screws, deformation factor $\alpha = 0.098$.




Fig. 5. Example of correctly segmented masks from the second category: obstructed masks of M12 screws, deformation factor $\alpha = 0.065$.




Fig. 6. Examples of discarded masks which contain both stem and head: a) under segmented mask from the first category; b) leaked mask from the first category; c) leaked mask from the second category; d) under segmented mask from the second category.




Fig. 7. Example of images with mosaic of tiles created for two different instances j and $\alpha = 1$, $\Delta I_{max} = 38$.

Each square tile may be viewed as a representation of a screw, with the smaller rectangle corresponding to the screw's head and the larger rectangle acting as its stem. At the center of each rectangle, we set a prompt point and perform segmentation using SAM. Since the boundaries of each square tile are known, the binary masks output by SAM can be directly compared with the corresponding tile. For each mask and square tile, the Intersection over Union (IoU) can be calculated. A mask output by SAM is accepted as correct if $\text{IoU} > 0.95$. For each given $\Delta I(n)$, we repeat the entire procedure for $j = 1, \dots, 10$, assigning different randomly selected intensities to each square tile for each j , as shown in Fig. 7. The mean fraction of accepted masks $\bar{\rho}(n)$ and standard deviations $\delta(n)$ were calculated similarly to calculations for screws.

3. Results

Fig. 8 presents examples of two depth images of M12 screws corresponding to deformation factors $\alpha(0)$ and $\alpha(N)$, overlaid with their complete SAM masks. Examples of types of masks overlaid on the depth image of screws are shown in Fig. 9. Four types of masks are shown: masks labeled as a) are under-segmented and they overlap with only the heads of the screws; masks labeled as b) are also under-segmented and they overlap with the stems of the screws; the mask labeled as c) shows leaked segmentation; and the masks labeled as d) are correct and they overlap with the entire visible portions of the screws, including the heads and stems. Locations of prompt points resulting in the displayed masks are marked by dots. In Fig. 10, a similar plot is shown for

the image of mosaic of tiles, labels a) – d) and color coding have the same meaning for square tiles and their subcomponents (rectangles) as for screws, heads and stems in Fig. 9.

Finally, the dependence of the mean ratio $\bar{\rho}$ of correctly segmented masks vs. deformation factor α for two sizes of screws is shown in Fig. 11. Similar plots for images of mosaics of tiles, generated for two different values of ΔI_{max} , are shown in Fig. 12. The mid-point α_{mid} for deformation factor is defined as such for which $\bar{\rho}(\alpha_{mid}) = (\bar{\rho}_{max} + \bar{\rho}_{min})/2$.




Fig. 8. Depth images of a pile of M12 screws for two limiting values of the deformation factor: a) $\alpha = 1$; b) $\alpha \approx 0$. Complete SAM output masks are displayed in color.

In Table 2, characteristics of screws and tiles at deformation factor equal to mid-point α_{mid} are displayed. $\Delta D_{max} = D_h - D_s$ is the largest difference between geometric features of the subcomponents of the decomposable part (diameters of head and stem for screws, see Fig. 2) and $\Delta D = D_h(\alpha) - D_s$ is calculated for a given deformation factor α . The corresponding difference in the image intensity (assuming eight-bit gray scale and the highest point in a scene determined by a bin height $h = 48$ mm) is in the rightmost column. The smaller ΔD_{max} is for M8 screws and the larger is for M12 screws. For mosaic of tiles, only the differences in image intensities ΔI are available since the intensities in the square tiles were already randomly selected from the $\langle 0, 255 \rangle$ interval. Under these conditions, two ΔI_{max} values match closely the two ΔD_{max} values.




Fig. 9. Examples of four types of masks overlaid on depth image of screws: a) incomplete, only head segmented; b) incomplete, only stem segmented; c) leaked segmentation; d) complete, correct masks of entire screw. Dots represent locations of prompt points for SAM.




Fig. 10. Examples of four types of masks overlaid on image of mosaic with tiles, labels a) – d) and color coding have the same meaning as in Fig. 9. Dots represent locations of prompt points for SAM.

Table 2. Screws and tiles characteristics at mid-point.

	α_{mid}	$\Delta D(\alpha_{mid})$	$\Delta I(\alpha_{mid})$
$\Delta D_{max} = 5 \text{ mm}$	0.32	1.6 mm	10
$\Delta D_{max} = 6 \text{ mm}$	0.21	1.3 mm	8
$\Delta I_{max} = 32$	0.42	--	13
$\Delta I_{max} = 38$	0.39	--	15




Fig. 11. Dependence of $\bar{\rho}$ on α for images of screws, error bars given by δ . Dotted lines show location of mid-point α_{mid} .




Fig. 12. Dependence of $\bar{\rho}$ on α for images of mosaic of tiles, error bars given by δ . Dotted lines show location of mid-point α_{mid} .

4. Discussion

The distinction between the two types of parts – visually decomposable into subcomponents versus not readily decomposable, as shown in Fig. 1 – was originally based on a subjective, unquantifiable impression. By gradually altering the geometric features of the parts, we were able to quantitatively study the transition between both part types in a more systematic manner. The dependence of $\bar{\rho}$ on α plotted in Fig. 11 and Fig. 12 indicates that this transition is described by a smooth function, with no abrupt shift from maximum to minimum values of $\bar{\rho}$ at a single deformation factor α . Thus, within a certain range $[\alpha_{mid} - \beta, \alpha_{mid} + \beta]$, there is no single preferred selection strategy for the masks output by SAM: whether to select those associated with the highest confidence score (as shown in Fig. 8b) or those associated with the medium score when very few correct masks receive the highest score (as shown in Fig. 8a). To maximize the total number of correctly segmented masks and afford optimal robotic path planning options, processing masks associated with both the highest and medium confidence scores may be necessary, though this could lead to increased total cycle time in automated bin picking.

The width 2β of α interval inside which both strategies may be needed can be interpreted as a degree of separation between the two approaches: for $\alpha < \alpha_{mid} - \beta$, utilizing masks with the highest confidence is advantageous, whereas for $\alpha > \alpha_{mid} + \beta$, the opposite holds true, and a better strategy relies on masks with medium confidence scores. Larger values of β widen the α interval where masks from both confidence levels must be analyzed. Conversely, smaller β values expand the regions where a single strategy suffices, albeit with diminishing benefit because when α approaches α_{mid} , the contribution from the opposite strategy is increasing.

Dimensional differences among subcomponents of decomposable parts, ΔD (such as the difference between head and stem diameters in screws) induce corresponding intensity variations across subregions in depth images. Simplified grayscale tile-mosaic images, such as the one shown in Fig. 7, reproduced

SAM's performance on realistic bin-picking screw images remarkably well. The masks shown in Fig. 9 and Fig. 10 exhibit identical characteristics: for example, correctly segmented masks (type d) can be generated by a prompt point placed on either the head or the stem. The primary advantage of utilizing these synthetic images is that SAM's segmentation output can be evaluated both automatically and precisely.

The rightmost column in Table 2 lists ΔI values calculated at $\alpha = \alpha_{mid}$. All four values, derived from distinct grayscale images and scene configurations, fall within a relatively narrow interval of $\langle 8, 15 \rangle$. The physical interpretation of ΔI is that when an image contains a compact region composed of two subregions with an intensity difference larger than ΔI , the mask with the highest confidence score will likely cover only one subregion rather than the composite entity. In such cases, a medium-confidence mask prompted from a single point is more likely to encompass the entire region. While both scenarios occur with high frequency, exceptions can occur, as is typical with point-prompted segmentation. In industrial bin picking, CAD models of the target parts are usually readily available, and recent developments – such as the CAD-prompted segmentation framework based on SAM3 – may resolve the ambiguity inherent to point prompts [7]. However, this emerging technique still warrants extensive evaluation on complex, heavily cluttered bin-picking scenes containing significant object occlusions.

5. Conclusions

Visual inspection of parts to be picked from a bin by a robotic arm integrated with a vision system may provide a first hint of whether they belong to the class of visually decomposable objects for which, paradoxically, not the largest score masks but the medium score masks provide more accurate segmentation. However, this subjective evaluation may lead to false conclusions. The knowledge of geometric dimensions of possible subcomponents, such as ΔD in Table 2, may not be sufficient to support the visual evaluation since what really matters is the resulting difference in the intensity ΔI in depth image. This difference depends not only on ΔD but also on the bit depth of the grayscale and the entire context of the scene, e.g., bin height. Other factors, like camera distance, lighting conditions, or part color and finish

also affect formation of real depth image, but they were not accounted for in our simplified simulation. More testing is required, but the initial results obtained in this paper suggest that an intensity difference of $\Delta I \gtrsim 15$ may act as a barrier prohibiting the further expansion of a binary mask initiated from a point.

References

- [1]. F. Sultana, A. Sufian, P. Dutta, Evolution of image segmentation using deep convolutional neural network: A survey, *Knowledge-Based Systems*, Vol. 201-202, 2020, 106062.
- [2]. A. Kirillov, et al., Segment anything, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 4015-4026.
- [3]. X. Lai, et al., LISA: Reasoning segmentation via large language model, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 9579-9589.
- [4]. J. Li, et al., SAM3-I: Segment anything with instructions, in *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026, pp. 27232-27249.
- [5]. N. Carion, et al., SAM 3: Segment anything with concepts, in *Proceedings of the 14th International Conference on Learning Representations (ICLR)*, 2026.
- [6]. COCO dataset, <https://cocodataset.org>
- [7]. Z. Tang, et al., CAD-prompted SAM3: Geometry-conditioned instance segmentation for industrial objects, *arXiv*, 2026, arXiv:2602.20551.
- [8]. NVIDIA, Isaac Sim documentation, <https://docs.isaacsim.omniverse.nvidia.com>
- [9]. S. Koch, et al., ABC: A big CAD model dataset for geometric deep learning, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9601-9611.
- [10]. T. Hodan, et al., T-LESS: An RGB-D dataset for 6D pose estimation of texture-less objects, in *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2017, pp. 880-888.
- [11]. B. Drost, et al., Introducing MVTec ITODD – A dataset for 3D object recognition in industry, in *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017, pp. 2200-2208.
- [12]. M. Franaszek, et al., Investigating the ambiguity of SAM when applied to depth and RGB images in bin picking applications, in *Proceedings of the 11th International Conference on Sensors and Electronic Instrumentation Advances (SEIA' 2025)*, 24-26 September 2025, Ponta Delgada, São Miguel (Azores Islands), Portugal, 2025, pp. 64-69.

Low-Complexity Trajectory Estimation for Table-Slide Rehabilitation Using an IMU Sensor

Y. Uchida¹, E. Ohkubo¹, M. Nishi¹ and T. Funayama²

¹ Teikyo University of Science, Dept. of Life & Health Sciences, Adachi-ku,
120-0045 Tokyo, Japan

² Teikyo University of Science, Dept. of Occupational therapy, Uenohara-shi,
409-0193, Yamanashi, Japan
Tel.: + 81 369101010
E-mail: uchida@ntu.ac.jp

Summary: Accurate trajectory estimation using inertial measurement units (IMUs) is challenging because small errors in acceleration and orientation estimation accumulate during double integration, resulting in substantial trajectory drift. This study proposes a low-complexity IMU-based trajectory estimation method for repetitive table-slide rehabilitation exercises. The proposed method combines QUEST-based initial attitude estimation, quaternion-based extended Kalman filtering, gravity compensation, motion-frequency-based high-pass filtering, and double integration. The method was evaluated using IMU data acquired during repetitive table-slide movements and compared with a conventional direct integration approach. Experimental results showed that the proposed method substantially reduced trajectory drift and reconstructed stable repetitive trajectories while preserving the characteristics of the rehabilitation movement. These results demonstrate the effectiveness of the proposed method for trajectory estimation during repetitive table-slide rehabilitation exercises

Keywords: IMU, Trajectory estimation, Rehabilitation, Table-slide exercise, Quaternion, Extended Kalman filter.

1. Introduction

Inertial measurement units (IMUs) are widely used for human motion analysis and rehabilitation because they are compact, lightweight, and inexpensive. Unlike optical motion-capture systems, IMUs do not require external cameras or dedicated measurement environments, making them suitable for clinical and home-based rehabilitation applications [1, 2].

However, accurate trajectory estimation using IMUs remains challenging. Since position is obtained by double integration of acceleration, small errors caused by sensor bias, orientation estimation, and gravity compensation accumulate over time, resulting in significant trajectory drift [3, 4]. Although various methods have been proposed to improve orientation estimation and reduce drift, accurate trajectory reconstruction remains difficult, particularly when using only IMU measurements [5-8].

This study focuses on the table-slide exercise, a repetitive upper-limb rehabilitation task in which the forearm moves back and forth on a tabletop. Because the movement is predominantly planar and repetitive, it provides a suitable target for a simple trajectory estimation method based on IMU measurements.

This paper proposes a low-complexity IMU-based trajectory estimation method for repetitive table-slide rehabilitation exercises. The proposed method combines QUEST-based initial attitude estimation, quaternion-based extended Kalman filtering, gravity compensation, motion-frequency-based high-pass filtering, and double integration. The effectiveness of the proposed method was evaluated by comparing the

reconstructed trajectories with those obtained using a conventional direct integration method.

2. Proposed Trajectory Estimation Method

The overall procedure of the proposed trajectory estimation method is illustrated in Fig. 1. The method consists of five processing steps: (1) QUEST-based initial attitude estimation, (2) quaternion-based extended Kalman filtering, (3) gravity compensation, (4) motion-frequency-based high-pass filtering, and (5) double integration for trajectory reconstruction.




Fig. 1. Overview of the proposed trajectory estimation method.

2.1. Orientation Estimation and Gravity Compensation

The IMU measured tri-axial acceleration and angular velocity. The raw sensor outputs were converted into physical units, and a constant acceleration bias was estimated using an initial stationary period.

The initial sensor orientation was estimated using the QUEST algorithm. The orientation was subsequently updated using a quaternion-based extended Kalman filter, which fused the acceleration and angular velocity measurements. The estimated quaternion was used to transform the measured acceleration from the sensor coordinate system to the global coordinate system. The gravitational acceleration component was then removed to obtain the motion acceleration used for trajectory estimation.

2.2. High-Pass Filtering

After gravity compensation, a high-pass filter was applied to the forward-direction acceleration to reduce slowly varying and offset components before double integration. The cutoff frequency was set to 0.3 Hz according to the dominant frequency of the repetitive table-slide motion. The filtered acceleration was subsequently used for trajectory reconstruction.

2.3. Double Integration

The filtered acceleration was integrated to estimate velocity and position. The reconstructed trajectory was obtained by double integration of the filtered acceleration signal.

3. Experimental Conditions

3.1. Participants

One healthy adult volunteer participated in this preliminary study. The participant had no history of upper-limb disorders and provided informed consent before the experiment. The research was approved by the Ethics Committee of Teikyo University of Science.

3.2. Measurement System

A wearable IMU (SensorTile.box PRO, STMicroelectronics) was attached to the forearm support board used during the table-slide exercise (Fig. 2). The IMU measured tri-axial acceleration and angular velocity. Sensor data were recorded at the original sampling frequency of 66.7 Hz and processed offline using Python.

The exercise was considered a constrained planar motion with negligible vertical displacement and only small changes in orientation. Therefore, the IMU

motion was assumed to be predominantly confined to the horizontal plane.

The movement was repeated several times at a comfortable, approximately constant rhythm. The experiment was simultaneously recorded using a video camera for visual verification.




Fig. 2. IMU sensor and board used during the experiment.

3.3. Table-Slide Task

The participant performed repeated table-slide movements in five directions while the forearm rested on the tabletop. The exercise was regarded as a constrained planar motion with negligible vertical displacement and only small changes in orientation. Therefore, the IMU motion was assumed to be predominantly confined to the horizontal plane.

The movement was repeated several times at a comfortable, approximately constant rhythm. The experiment was simultaneously recorded using a video camera for visual verification.

Fig. 3 shows the frequency spectrum of the forward-direction acceleration during the repetitive table-slide exercise. The dominant motion frequency was approximately 0.366 Hz; therefore, the cutoff frequency of the high-pass filter was set to 0.3 Hz.

The recorded IMU data were processed using the proposed trajectory estimation method. The reconstructed trajectories obtained by the proposed method were compared with those estimated using the conventional direct integration method. In addition, the forward-direction acceleration signals before and after high-pass filtering were compared to evaluating the effect of the preprocessing step.




Fig. 3. Frequency spectrum of the forward-direction acceleration during the repetitive table-slide exercise. The dominant motion frequency was approximately 0.366 Hz, and the high-pass filter cutoff frequency was set to 0.3 Hz.

4. Results

4.1. Comparison of Reconstructed Trajectories

Fig. 4 compares the trajectories reconstructed using the conventional direct integration method and the proposed method. As shown in Fig. 4(a), the conventional method exhibited a substantial drift, and the forward displacement increased continuously as the integration error accumulated. The final forward displacement reached approximately 19 m despite the actual movement being confined to a tabletop.

In contrast, the proposed method reconstructed a stable repetitive trajectory, as shown in Fig. 4(b). The estimated trajectory remained within approximately 0.30 m in the forward direction, which is consistent with the expected range of the table-slide exercise. These results demonstrate that the proposed method effectively suppressed drift while preserving the repetitive characteristics of the rehabilitation movement.




Fig. 4. Comparison of reconstructed trajectories obtained using (a) the conventional direct integration method and (b) the proposed trajectory estimation method.

4.2. Effect of High-Pass Filtering

Fig. 5 shows the forward-direction acceleration before and after applying the high-pass filter with a cutoff frequency of 0.3 Hz. The repetitive motion pattern was preserved after filtering, whereas slowly varying and offset components were noticeably reduced.

These results indicate that the high-pass filtering process effectively suppressed low-frequency components while maintaining the characteristics of the rehabilitation movement.




Fig. 5. Forward-direction acceleration before and after high-pass filtering with a cutoff frequency of 0.3 Hz.

5. Discussion

The results demonstrated that the proposed trajectory estimation method substantially reduced trajectory drift during repetitive table-slide movement compared with the conventional direct integration approach. As shown in Fig. 4, the conventional method exhibited a large accumulated drift, whereas the proposed method reconstructed a stable planar trajectory consistent with the repetitive forward and backward motion. This improvement can be attributed to the combination of orientation estimation, gravity compensation, and motion-frequency-based high-pass filtering.

The table-slide exercise is characterized by repetitive forward and backward movement with relatively small changes in orientation. Fig. 5 showed that the overall acceleration waveform remained similar after high-pass filtering, indicating that the acceleration associated with the intended movement was largely preserved. At the same time, the mean acceleration decreased from 0.0320 to 0.0015 m/s², and both the standard deviation and the root mean square value were slightly reduced. These results suggest that the high-pass filter reduced the slowly varying and offset components of the acceleration signal while preserving the motion-related fluctuations. Although the difference in the acceleration waveform was relatively small, even a small residual acceleration offset accumulates during double integration, leading to substantial trajectory drift. Therefore, reducing this offset contributed to the improved trajectory reconstruction observed in Fig. 4.

The proposed method exploits the repetitive characteristics of table-slide motion by selecting the high-pass filter cutoff frequency according to the dominant movement frequency. This task-specific design reduced trajectory drift while preserving the

motion characteristics required for estimating relative trajectories during rehabilitation exercises.

Several limitations should be considered. First, this preliminary study included only one healthy participant. Additional experiments involving a larger number of participants, including patients undergoing upper-limb rehabilitation, are necessary to evaluate the robustness and clinical applicability of the proposed method. Second, the trajectory evaluation in this study was primarily based on qualitative comparison with the recorded video. Future work should include quantitative validation using an external motion capture system or another high-accuracy reference measurement.

Although the cutoff frequency of the high-pass filter was selected according to the dominant frequency of the table-slide movement in this study, the optimal cutoff frequency may vary depending on the characteristics and speed of different rehabilitation exercises. Future studies will investigate adaptive parameter selection and evaluate the applicability of the proposed method to a wider range of upper-limb rehabilitation tasks.

6. Conclusions

In this study, a low-complexity trajectory estimation method for repetitive table-slide rehabilitation using a wearable IMU was proposed. By combining QUEST-based orientation estimation, gravity compensation, motion-frequency-based high-pass filtering, and double integration, the proposed method substantially reduced trajectory drift compared with conventional direct integration. The results suggest that the proposed method is suitable for estimating relative trajectories during repetitive rehabilitation exercises without external positioning systems. Future work includes quantitative validation, evaluation with patients, and adaptive filter parameter selection.

Acknowledgements

This study was partly supported by a JSPS Grant-in-Aid for Scientific Research (Issue no. JP23K11207).

References

- [1]. A. M. Sabatini, Quaternion-based extended Kalman filter for determining orientation by inertial and magnetic sensing, *IEEE Transactions on Biomedical Engineering*, Vol. 53, Issue 7, 2006, pp. 1346-1356.
- [2]. S. O. H. Madgwick, A. J. L. Harrison, R. Vaidyanathan, Estimation of IMU and MARG orientation using a gradient descent algorithm, in *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR)*, 2011, pp. 1-7.
- [3]. H. A. Hashim, Attitude determination and estimation using vector observations: Review, challenges and comparative results, *arXiv*, 2020, arXiv:2001.03787.
- [4]. A. H. A. Hafez, U. Demircioglu, H. Bakir, R. Alzubi, A comprehensive review of reinforcement learning algorithms for control systems: Applications, challenges, and future directions, *IEEE Access*, Vol. 14, 2026, pp. 42809-42837.
- [5]. O. J. Woodman, An introduction to inertial navigation, Technical Report UCAM-CL-TR-696, *University of Cambridge Computer Laboratory*, Cambridge, 2007.
- [6]. SensorTile.box PRO, STMicroelectronics, <https://www.st.com/en/evaluation-tools/steval-mkboxpro.html>
- [7]. Q. Wang, et al., Interactive wearable systems for upper body rehabilitation: A systematic review, *Journal of NeuroEngineering and Rehabilitation*, Vol. 14, 2017, 20.
- [8]. Y. Uchida, E. Ohokubo, M. Nishi, T. Funayama, et al., Machine-learning-based evaluation of upper limb motion using 8×8 time-of-flight and LiDAR sensors: SVM classification and deep neural network analysis, *Sensors & Transducers*, Vol. 270, Issue 3, 2025, pp. 30-39.

Future Transformation Prediction of Land Use/Land Cover Using Machine Learning and Neural-Network-Based Transition Potential Modeling in the Middle Rio Grande Region 2025-2040

Omar. S. Belhaj^{1,4}, Stanley. T. Mubako², Raed Aldouri¹, Craig Tweedie¹, Elhadi. A. Hadia³, Rabia. E. Al Mamlook^{5,6}, William Hargeove¹, Elizabeth Walsh¹ and Alex Mayer¹

¹ University of Texas at El Paso, 500 West University Avenue, 79968, El Paso, Texas, U.S.A.

² California Department of Water Resources 715 P St., 6th Floor, Sacramento, CA 95814 C: 916-873-4784, U.S.A.

³ Libyan Center for Sustainable Development Research, Khoms, Libya

⁴ Elmergeb University, College of Science, Khoms, Libya

⁵ Trine University, Department of Business Administration, Indiana, P.O. Box 49008, U.S.A.

⁶ Zawia University, Department of Industrial Engineering and Engineering Management, Al Zawia, P.O. Box 16418, Libya
Tel.: + 19152749938
E-mail: obelhaj@miners.utep.edu

Summary: Predicting future ecosystem changes, including the location, timing, extent, patterns, and trajectories of resource transformation, is essential for effective environmental protection and sustainable development. In recent years, geospatial projection tools and predictive modeling techniques have become increasingly important for forecasting land use and land cover (LULC) dynamics. This study employs the TerrSet to analyze and predict future LULC changes in the Middle Rio Grande region (MRGR) along the U.S.–Mexico border. Using LULC classification datasets from 2010, 2015, and 2020, the model projects LULC changes over 15 years (2025–2040) to identify future spatial patterns and trends. The results indicate that agricultural land is projected to decline by approximately 6.5 %, particularly in areas surrounding major metropolitan centers where continued urban expansion is expected. Developed open space and developed land are projected to increase by approximately 5 % and 16 %, respectively, reflecting ongoing urbanization. Surface water bodies are projected to decrease by approximately 12 %, while evergreen forests and shrublands are expected to decline modestly by about 0.4 % each. These projected LULC changes may have significant implications for regional water resources, including increased competition for limited water allocations among agricultural, industrial, municipal, and ecological sectors, as well as the potential reallocation of water resources to support expanding urban development. The findings provide valuable information for planners, policymakers, land developers, ranchers, resource managers, funding agencies, and other stakeholders involved in land-use planning, water resource management, environmental conservation, and the development of resilient strategies for future water supply, treatment, and monitoring in this rapidly changing transboundary region.

Keywords: Future prediction, Machine learning, Land change modeler, Change patterns, Change tendencies, Ecosystems.

1. Introduction

Anthropogenic activities have led to significant changes in resource availability across various locations worldwide over time [1]. These activities have significantly transformed land use and/ land cover (LULC), especially in regions with rapid population growth, urbanization, and development [2-5]. Measuring and monitoring LULC transformations provide clear insights for scientists and guidance for planning authorities, decision-makers, and resource managers in sustainable environmental management [3]. With free access to the USGS Landsat archive and advances in remote sensing techniques, LULC change dynamics using high-frequency temporal datasets have become possible [6]. Abundant LULC change research projects on environmental changes are designed using multi-temporal remote sensing data frameworks [7]. These research projects consistently revealed that

human activities, in combination with natural resources, are crucial drivers of LULC dynamics and change across all spatial and temporal scales [8].

In addition, future LULC transformations are among the attributes considered to play a significant role in decision-making processes and future planning. Future projections indicate ongoing urbanization, underscoring the need for sustainable land management practices [9]. Using remote sensing data and several methods and models [10-12]. These methods and models can involve time interpolation or extrapolation and provide a basis for addressing and segregating the complex cluster of biophysical and socioeconomic factors that affect the rate, quantity, extent, and location of LULC conversions [13].

Beyond baseline mapping, predictive models quantify the impacts of LULC on global climate, carbon cycling, biodiversity, water budgets, and ecosystem services [14]. Understanding specific model components and data requirements is essential

to maximizing their utility in environmental research and policymaking [15]. LULC models are broadly categorized into spatial description and spatial transition frameworks [16]. Spatial description models capture structural landscape dynamics using statistical regression analyses [15]. Conversely, spatial transition models integrate explicit geospatial metadata, including geographic coordinates and pixel-level state arrangements, to project landscape trajectories over time [17].

To simulate these spatial transitions, researchers deploy Cellular Automata (CA), Artificial Neural Networks (ANN), and Weights of Evidence (WoE) architectures [18]. CA models excel at neighborhood-driven pixel simulations, whereas ANN frameworks map complex, non-linear driving variables without assuming normal data distributions [15]. Concurrently, WoE models apply Bayesian probabilities to isolate the influence of individual empirical conditioning factors on historical changes [19]. The TerrSet Land Change Modeler (LCM) bridges the gap between these statistical descriptions and spatial transition models [20]. The LCM first uses empirical drivers, such as multilayer perceptron neural networks, to estimate baseline LULC transition potentials [18]. It then couples these potentials with a spatial transition architecture, leveraging Markov Chain matrices to project precise pixel allocations and multi-temporal geographic arrangements [20].

This study extends previous Middle Rio Grande Region (MRGR) LULC classification work [13, 24, 26] by applying LCM-based transition-potential modeling to generate future LULC scenarios for the 2025–2040 period for regional planning. The study area expands from the southwest U.S.A. and the southeast Mexico borderland to find future projections and trajectories of LULC using random forest in ArcGIS and LCM in TerrSet liberaGIS Geospatial Monitoring and Modeling Software in Edrisi that combines Markov chain analysis with machine learning-based transition potential modeling, typically using multilayer perceptron neural networks to estimate spatial transition likelihoods [21–23].

1.1. Study Area

This study was conducted in MRGR, the dryland covers 36988 km² from southern New Mexico to far west Texas in the US and northern Chihuahua in Mexico, as shown in Fig. 1. This area consists of various terrain features and includes three metropolitan cities: El Paso (Texas, USA), Las Cruces (New Mexico, USA), and Ciudad Juárez (Chihuahua, Mexico). In addition, this area experiences intensive agricultural activity concentrated along the Rio Grande valley. The shrublands and forests occupy the mountains and the Chihuahua desert ecosystem. The region faces resource challenges, including water scarcity and growing water needs across sectors such as agriculture, urban areas, and industry [24, 25].




Fig. 1. The study area.

2. Data and Methodology

2.1. Data Preparation and Processing

The preparation and processing steps included downloading Landsat 5 Thematic Mapper (TM) images for 2010. Also, Landsat 8 images, Operational Land Imager (OLI) for 2015 and 2020. The satellite scenes were captured between the second half of May and the first week of July, which is considered the “leaf-on” season in this area. Atmospheric data correction is applied. The images are clipped to the study area extent. Minimum noise fraction transformation is performed to reduce the inherent spectral dimensionality and noise in multispectral data.

2.2. Images LULC Classification and Assessment

Image LULC classification for 2010 and 2015 was previously conducted as part of the LULC classification and measuring in the MRGR study using a combined classification in ArcGIS 10.8.2 [13, 24, 26]. However, the 2020 LULC classification was performed using Random Forest, followed by object-based classification in ArcGIS Pro 3.7. The maps were created for the targeted years. These maps were finalized in ArcGIS Pro 3.7 by using object appearance in the imagery to eliminate misclassifications and improve classification accuracy. Classification accuracy assessment was performed for individual LULC categories and the overall classification, focusing on analyses of 2010, 2015, and 2020 using high-resolution imagery from the National Agriculture Imagery Program (NAIP) and Google Pro Maps.

2.3. Modeling Software Validation

TerrSet software validation was applied to the LULC classification results of the 2010, 2015, and 2020 maps [24, 26] to assess the software’s predictive performance and ensure that the projection results are

accurate and usable. The two-step processes validate the LULC categories produced by the software, support, and strengthen the results of future prediction analysis.

2.4. LULC Future Prediction Analysis

Following the generation of historical LULC classification maps for the baseline years 2010, 2015, and 2020, rigorous accuracy assessment and validation procedures were applied to evaluate the reliability of the classified maps and establish a reliable basis for future LULC prediction. Following validation, the Land Change Modeler (LCM) module within the TerrSet Geospatial Monitoring and Modeling System was used to develop a predictive framework to simulate LULC dynamics over a 15-year period (2025–2040), with projections for 2025, 2030, 2035, and 2040.

The LCM utilized historical LULC transitions and transition potentials derived from the 2015 and 2020 classification maps to estimate the spatial and temporal patterns of future LULC change [27, 28]. The model identified areas with potential for transition among LULC categories and generated projected LULC maps for each future milestone. In addition, the model calculated the projected areal extent of each LULC category, enabling quantitative assessment of future changes in landscape composition.

A post-classification comparative analysis was subsequently conducted to quantify changes among the projected LULC classes across the four forecast periods. This analysis evaluated both net changes in the total area of individual LULC classes and class-to-class transitions, thereby identifying the magnitude and direction of projected LULC change. The resulting maps, transition statistics, and area estimates were then analyzed to characterize the spatial and temporal trajectories of LULC change throughout the 2025–2040 prediction period.

3. Results and Discussion

3.1. LULC Class Measurement

The baseline historical analysis between 2010 and 2020 revealed dynamic LULC transformations across the study area (Table 1), which were used as the basis for predicting future LULC change. The results show that the shrubs remained the dominant landscape class, experiencing a net increase of 566 Km² (~1.74 %) despite a temporary peak in 2015. Conversely, Evergreen Forest experienced the largest absolute reduction, contracting by 587 Km² (~27 %) over the ten years. Urban expansion was evident, with the Developed Area increasing from 1022 Km² to 1186 Km², representing a notable ~16 % growth. This urbanization, alongside environmental factors, likely contributed to steady declines in Agriculture (~8 %)

and Water bodies (~41 %), while Developed Open Space remained at around 40 Km².

Table 1. LULC change 2010-2020.

Year	LULC category areas (km ²)					
	Agriculture	Developed open space	Developed area	Water	Evergreen forest	Shrubs
2010	1111	40	1022	135	2155	32525
2015	1091	40	1078	106	1148	33825
2020	1023	40	1186	80	1568	33091

3.2. Accuracy Assessment of LULC Classification

Following completion of ground-truth data collection, reference data processing, and accuracy assessment, confusion matrices were generated to evaluate the performance of the LULC classifications for 2010, 2015, and 2020. The assessment compared the classified LULC maps with independent reference samples to quantify the agreement between the mapped and observed LULC classes. The results demonstrated consistently high classification performance across all three study years, with an overall classification accuracy of approximately 99 % and a Kappa coefficient of approximately 0.96. These results indicate a high level of agreement between the classified maps and the reference data, demonstrating the reliability of the LULC classifications for subsequent change detection and predictive modeling. The consistently high accuracy across the three temporal datasets also provides confidence that the observed changes in LULC patterns are primarily representative of actual landscape changes rather than classification errors. Consequently, the validated 2010, 2015, and 2020 LULC datasets were considered sufficiently reliable for subsequent temporal change analysis and future LULC prediction using the Land Change Modeler (LCM).

3.3. Modeling Software Validation

TerrSet validation was applied to the LULC classification results of the 2010, 2015, and 2020 maps [24, 26] to assess the software's predictive performance. This process is completed in two steps:

1. Comparing the 2020 real LULC classification map with the 2020 LULC prediction map created by the model. The step was completed using the 2010 and 2015 LULC-classified maps, as shown in Table 1, to produce the 2020 LULC prediction map and compare it with the 2020 real LULC-classified map. Results showed 90 % overall accuracy for the 2020 LULC predicted map, as detailed in Table 2.

- Run the validation panel in the LCM to assess the quality of the 2020 LULC predicted map compared to the 2020 real LULC classified map. The validation process applies a three-way cross-tabulation between the 2020 real LULC-classified map and the predicted 2020 LULC map. The results showed category matching between the two maps. Therefore, the validation results indicate that the model can make confident predictions about future LULC transformation in the region.

Table 2. 2020 future LULC prediction change accuracy.

LULC category	2020 actual classes km ²	2020 predicted classes km ²	difference km ²	% difference	% Accuracy
Agriculture	1052	1023	29	3	97
Open space	40	40	0	0	100
Developed area	1247	1186	61	5	95
Water	98	80	18	23	78
Evergreen forest	1137	1568	431	27	73
Shrubs	33414	33091	323	1	99
Total area	36988	36988			
Total accuracy					90

3.3. Future LULC Prediction Analysis and Discussion

The predictive modeling framework simulated a 15-year future trajectory of LULC transformation from 2025 to 2040, as shown in Tables 3 and 4. Also, illustrated in Figs. 2, 3, 4, and 5 are projected continued LULC alterations across the study domain.

Table 3. LULC future prediction changes 2025-2040.

Year	Future prediction LULC categories areas (km ²)					
	Agriculture	Developed open space	Developed area	Water	Evergreen forest	Shrubs
2025	1038	40	1392	90	1135	33293
2030	1020	40	1441	88	1267	33132
2035	992	41	1541	82	1122	33210
2040	971	42	1617	79	1130	33149

The most prominent projected trend is the continued expansion of developed areas, which increases from 1392 km² in 2025 to 1617 km² in 2040, representing a net increase of 225 km² (~ 16.2 %). The sustained increase indicates continued urban and infrastructure development across the region, especially around the metropolitan areas of Las

Cruces, El Paso, and Ciudad Juárez, where continued urban growth and outward expansion are expected to transform surrounding agricultural and shrubland areas [26, 29]. The resulting conversion and fragmentation of these lands may increase anthropogenic pressure on land and water resources and contribute to localized hydrological stress. Continued expansion of the built environment may also reduce landscape connectivity and increase the fragmentation of remaining natural and agricultural areas, highlighting the need for integrated regional land-use planning and sustainable growth strategies. This expansion corresponds to a marginal rise in developed open space, which increases from 40 km² to 42 km² after 2030, representing about a 5 % increase by 2040 [30].

Table 4. LULC future category prediction changes 2025-2040

LULC Category	2025 km ²	2040 km ²	Net Change km ²	Horizon % Change (Relative to 2025)	Average Annual Rate of Change
Agriculture	1038	971	-67	-6 %	-0.43 % / year
Developed Open Space	40	42	(+2)	(+5 %)	+0.33 % / year
Developed Area	1392	1617	(+225)	(+16 %)	+1 % / year
Water	90	79	-11	-12 %	-0.81 % / year
Evergreen Forest	1135	1130	-5	-0.4 %	-0.03 % / year
Shrubs	33293	33149	-144	-0.4 %	-0.03 % / year




Fig. 2. MRGR LULC classes prediction 2025.

In contrast, agricultural land is projected to decline steadily, from 1038 km² in 2025 to 971 km² in 2040, a net reduction of 67 km² (~6.5 %). This decline suggests that agricultural areas may increasingly be converted to developed land or other LULC categories as regional development continues. Also, the surface area of water bodies decreases slightly from 90 km² to 79 km² (~12 %) due to continued reductions in snowpack and streamflow [24, 26, 29].




Fig. 3. MRGR LULC classes prediction 2030.




Fig. 4. MRGR LULC classes prediction 2035.




Fig. 5. MRGR LULC classes prediction 2040.

Evergreen Forest exhibits structural fluctuations, peaking at 1267 km² in 2030 before contracting to 1130 km² in 2040, a decline of about 0.4 %, due to expected droughts and wildfires [24, 26, 29].

The dominant Shrubs class exhibits a downward trend, decreasing by a net total of 144 km² from

33293 km² in 2025 to 33149 km² by about 0.4 % in 2040 over the projection horizon, thereby providing significant areas for development [24, 26, 29].

The results support the areas, trends, and patterns of the LULC changes that occurred in the MRGR in the last few decades through 1994-2015, as stated in O. Belhaj et al.'s (2024 study for classifying and measuring the LULC in the MRGR and cited in this study [24-26, 29]. Ultimately, these LCM metrics quantify the long-term patterns of human-induced land conversion across the regional landscape.

5. Conclusion

This study established a comprehensive spatiotemporal framework to quantify historical land-use/land-cover (LULC) dynamics from 2010 to 2020 and project future landscape change through 2040 across a 36988 km² region. Historical analyses revealed a highly dynamic landscape characterized by sustained urban expansion, with Developed Areas increasing by approximately 16 % (164 km²). This growth was accompanied by notable declines in key natural and agricultural resources, including an approximately 8 % reduction in Agricultural land and a 41 % loss of surface Water bodies.

These observed trends were incorporated into the TerrSet liberaGIS Land Change Modeler (LCM) to develop and validate a predictive framework for forecasting future LULC dynamics. Model projections indicate that urbanization will continue to intensify between 2025 and 2040, with Developed Areas expected to expand by an additional 16 % (225 km²), reaching approximately 1617 km² by 2040. This expansion is projected to occur primarily at the expense of agricultural land and surface Water bodies, which are forecast to decline by approximately 6 % and 12 %, respectively, relative to the 2025 baseline. Natural vegetation, dominated by shrublands, is also projected to undergo spatial reorganization, resulting in a net loss of approximately 144 km² (0.4 %).

Overall, the findings indicate that continued urban expansion will place increasing pressure on agricultural resources, surface water, and natural ecosystems, with important implications for regional environmental sustainability. The historical and projected LULC patterns presented in this study provide a robust scientific foundation to support land-use planning, guide sustainable urban development, protect riparian corridors and other ecologically sensitive areas, and inform long-term environmental management and policy decisions across the Middle Rio Grande region.

Abbreviations and Acronyms

LULC: Land Use/ Land Cover, MRGR: Middle Rio Grande Region, TM: Thematic Mapper, OLI: Operational Land Imager, CA: Cellular Automata, ANN: Artificial Neural Networks, WoE: Weights of

Evidence, LCM: Land Change Modeler, NAIP: National Agriculture Imagery Program.

Acknowledgements

This study was made possible through support from the National Institute of Food and Agriculture, U.S. Department of Agriculture, under award number 2015-68007-23130. The opinions in this paper are those of the authors and do not necessarily reflect the USDA's views.

References

- [1]. Y. H. Mir, S. Mir, M. A. Ganie, J. A. Bhat, et al., Overview of land use and land cover change and its impacts on natural resources, in Ecologically Mediated Development: Promoting Biodiversity Conservation and Food Security, *Springer Nature Singapore*, 2025, pp. 101-130.
- [2]. S. Naeem, S. A. Mahmood, Exploring the nexus between LULC transformation, land surface temperature, and drought in a rapidly urbanizing landscape: The case of Multan, *DYSONA – Applied Science*, Vol. 7, Issue 1, 2026, pp. 98-109.
- [3]. B. Bairwa, R. Sharma, A. Kundu, S. S. Sammen, et al., Predicting changes in land use and land cover using remote sensing and land change modeler, *Frontiers in Environmental Science*, Vol. 13, 2025, 1540140.
- [4]. N. Bozali, Spatiotemporal simulation of land use and land cover changes in Türkiye through a CA-Markov framework, *Scientific Reports*, Vol. 16, Issue 1, 2026, 5320.
- [5]. M. K. Mosleh, Integrating the CA-Markov model and geospatial techniques for spatiotemporal prediction of land use/land cover dynamics in Qus District, Egypt, *Modeling Earth Systems and Environment*, Vol. 11, Issue 5, 2025, pp. 1-15.
- [6]. P. A. D. Santos, M. Adami, M. C. A. Picoli, V. H. R. Prudente, et al., Land use and land cover products for agricultural mapping applications in Brazil: Challenges and limitations, *Remote Sensing*, Vol. 17, Issue 13, 2025, 2324.
- [7]. L. Hu, X. Duan, J. Liu, Application and assessment of a CA-ANN model for land use change simulation and multi-temporal prediction in Guiyang city, China, *Sustainability*, Vol. 18, Issue 3, 2026, 1518.
- [8]. S. Thapa, D. J. Gibson, R. Li, Forest transition under climate pressure: Land use land cover change in the Greater Shawnee National Forest, *Remote Sensing*, Vol. 18, Issue 7, 2026, 1079.
- [9]. Z. Tahir, M. Haseeb, S. A. Mahmood, S. Batool, et al., Predicting land use and land cover changes for sustainable land management using CA-Markov modelling and GIS techniques, *Scientific Reports*, Vol. 15, Issue 1, 2025, 3271.
- [10]. H. I. Gündüz, Land-use land-cover dynamics and future projections using GEE, ML, and QGIS-MOLUSCE: A case study in Manisa, *Sustainability*, Vol. 17, Issue 4, 2025, 1363.
- [11]. G. Amin, I. Imtiaz, E. Haroon, N. U. Saqib, et al., Assessment of machine learning algorithms for land cover classification in a complex mountainous landscape, *Journal of Geovisualization and Spatial Analysis*, Vol. 8, Issue 2, 2024, 34.
- [12]. A. Amgoth, H. P. Rani, K. V. Jayakumar, Exploring LULC changes in Pakhal Lake area, Telangana, India using QGIS MOLUSCE plugin, *Spatial Information Research*, Vol. 31, Issue 4, 2023, pp. 429-438.
- [13]. O. Belhaj, Land use/land cover change patterns and trends in two dryland regions, PhD Thesis, *The University of Texas at El Paso*, El Paso, 2022.
- [14]. Y. Ren, J. Wang, L. Zhang, Future land use/land cover and its impacts on ecosystem services: A framework coupled with structure, function, and dynamics, *Environmental Monitoring and Assessment*, Vol. 196, Issue 8, 2024, 712.
- [15]. M. M. Rahman, G. Szabó, A comprehensive review on land use/land cover (LULC) change modeling for urban development: Current status and future prospects, *Sustainability*, Vol. 15, Issue 2, 2023, 903.
- [16]. P. Choudhary, C. P. Devatha, A. Azhoni, Integration of multi-layer perceptron neural network and cellular Automata-Markov chain approach for the prediction of land use land cover in land change modeler, *Ecological Modelling*, Vol. 506, 2025, 111162.
- [17]. N. Aquilué, M. De Cáceres, M. J. Fortin, A. Fall, et al., Machine learning in modelling land-use and land cover-change (LULCC): A review of foundational elements and future frontiers, *Science of The Total Environment*, Vol. 820, 2022, 153123.
- [18]. M. I. J. Putra, S. Phinn, Integrating artificial neural network-cellular automata and Bayesian weight of evidence for spatiotemporal analysis of urban change in Brisbane City, *Computational Urban Science*, Vol. 6, Issue 1, 2026, 13.
- [19]. J. M. Silva, M. Paegelow, Integrating weights of evidence and spatial transition probabilities in environmental modeling, *Environmental Monitoring and Assessment*, Vol. 197, Issue 3, 2025, 241.
- [20]. J. R. Eastman, J. V. Solórzano, TerrSet 2024: Geospatial Monitoring and Modeling System User Manual, Clark Labs, *Clark University*, Worcester, 2024.
- [21]. A. B. S. Güner, F. Sunbul, Comparative modeling of land-use change using LCM and GeoFLUS: Implications for urban expansion and regional-scale geotechnical risk screening, *Applied Sciences*, Vol. 16, Issue 4, 2026, 2082.
- [22]. P. P. Jaishi, Land use/cover change detection and modeling using a multi-layer perception neural network deep learning model for Orlando, Florida metropolitan area, Master's Thesis, *The University of West Florida*, Pensacola, 2025.
- [23]. M. Arifuzzaman, M. R. Hasan, T. J. Toma, S. B. Hassan, et al., An advanced decision tree-based deep neural network in nonlinear data classification, *Technologies*, Vol. 11, Issue 1, 2023, 24.
- [24]. O. Belhaj, S. Mubako, C. Tweedie, R. Aldouri, et al., Land use/land cover change classification and measurement in the Middle Rio Grande region, USA–Mexico using remote sensing and geographic information systems, *Sensors & Transducers*, Vol. 267, Issue 4, 2024, pp. 20-33.
- [25]. O. S. Belhaj, S. T. Mubako, R. Aldouri, C. Tweedie, et al., Change detection of land use land cover in the Middle Rio Grande region, US-Mexico border land 1994-2015, in *Proceedings of the 11th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2025, pp. 21-27.

- [26]. O. Belhaj, S. Mubako, R. Aldouri, C. Tweedie, et al., Measuring land use/land cover change in the Middle Rio Grande region, USA–Mexico using remote sensing and geographic information systems, in *Proceedings of the 10th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2024, pp. 162-168.
- [27]. J. R. Eastman, TerrSet: Geospatial Monitoring and Modeling System, Clark Labs, *Clark University*, Worcester, 2020.
- [28]. R. G. Pontius, N. R. Malizia, Effect of category aggregation on optical validation of land-change models, *International Journal of Geographical Information Science*, Vol. 18, Issue 3, 2004, pp. 229-254.
- [29]. O. Belhaj, S. Mubako, R. Aldouri, C. Tweedie, et al., Determination of change in surface waterbodies in the Middle Rio Grande Basin by Modified Normalized Difference Water Index (MNDWI) 1994-2020, Libya, *Libyan Journal of Ecological & Environmental Sciences and Technology*, Vol. 4, Issue 2, 2022, pp. 67-80.
- [30]. J. Dou, R. Xia, K. Zhang, C. Xu, et al., Landscape fragmentation of built-up land significantly impact on water quality in the Yellow River Basin, *Journal of Environmental Management*, Vol. 371, 2024, 123232.

Current Status and Future Development of National Standard Time UTC(NTSC)

Shuhong Zhao¹, Yanrong Xue¹, Shanshan Bai¹, Yu Zhang¹ and Sufang Liu²

¹ National Time Service Center, Chinese Academy of Sciences, Xi'an, China

² University of Chinese Academy of Sciences, Beijing, China

Tel.: + 086 83890304

E-mail: zsh@ntsc.ac.cn

Summary: This paper proposes a high-precision time generation algorithm for UTC(k) based on the collaborative optimization of two frequency control strategies. By introducing a weighted switching function, the method governs the switching between the two strategies. Taking into account both the accuracy of the steered signal relative to the reference and the stability of the steered signal, this method enables smooth switching between different control regions, thereby ensuring an optimal balance between accuracy and stability for the system. This approach aims to overcome the limitations of existing Bang-Bang control and LQG control strategies, offering a solution that ensures the accuracy and stability of the system time UTC(k) while achieving optimal control. Experimental results indicate that the time differences between the Test-steered time and the national standard time UTC(NTSC) are maintained within ± 0.5 ns over a two-month period. Additionally, Frequency stabilities of 1.3×10^{-15} at 1 day and 6.78×10^{-16} at 5 days are achieved.

Keywords: UTC(k), Frequency control, Bang-bang control, LQG control, Weighted switching function.

1. Introduction

Coordinated Universal Time (UTC) is produced by the International Bureau of Weights and Measures (BIPM), which collects clock data from over 400 atomic clocks maintained by nearly 90 timing laboratories worldwide, generates the free atomic time scale (EAL) through a weighted average, and then calibrates its scale interval using the SI second as realized by more than twenty cesium fountain primary frequency standards and optical clocks to form the strictly uniform International Atomic Time (TAI). On this basis, leap seconds are inserted according to the Universal Time (UT1) determined by the International Earth Rotation and Reference Systems Service (IERS) so that the difference between UTC and UT1 is maintained within ± 0.9 s, and the final results are published in the monthly Circular T bulletin with a latency of approximately 40–50 days.

As the legally established global unified time reference, UTC is essentially a post-processed “paper time”; its fundamental significance lies in providing the traceability source for the local physical realizations UTC(k) independently maintained by national timing laboratories. It is through the real-time UTC(k) signals essential to critical fields such as satellite navigation, deep-space tracking, telemetry and command, communication synchronization, and precision metrology that this paper reference is transformed into continuous, high-precision time and frequency support in practice.

The generation of UTC(k), the local physical realization independently maintained by each timing laboratory, can be summarized as follows. A real-time time scale is first produced by applying a weighted average algorithm to the internal time-difference

measurements of a local clock ensemble comprising commercial cesium clocks and active hydrogen masers. High-precision remote comparison links, such as GNSS common-view and Two-Way Satellite Time and Frequency Transfer (TWSTFT), are then used to continuously monitor the offset of this scale with respect to TAI or UTC. On this basis, frequency control techniques are employed to finely adjust the local scale by combining UTC – UTC(k) offsets published in the BIPM Circular T (highly accurate, monthly latency) with near-real-time link data (low latency, noisier). In this way, UTC(k) closely approximates UTC while ensuring the output of continuous real-time physical signals.

Frequency control algorithms directly affect the performance of UTC(k), including its phase noise level, long-term frequency stability, and accuracy relative to international UTC. Commonly used frequency control algorithms primarily employ bang-bang and LQG algorithms. Bang-bang control is an on-off strategy that uses extreme control inputs to quickly return a system to its target state, leading to rapid responses but potentially causing oscillations in UTC(k) timekeeping. In contrast, LQG (Linear Quadratic Gaussian) control is an optimal strategy based on state estimation, which is effective in noisy environments but requires high model accuracy and is sensitive to noise and uncertainties.

This paper introduces a high-precision time generation method for UTC(k) that optimizes two frequency control strategies using a weighted switching function. The weighted switching function dynamically evaluates the system state based on the control error and system stability to select the most appropriate frequency control strategy. Advanced control algorithms, particularly in addressing

nonlinear factors such as atomic clock frequency drift, can significantly enhance the reliability of UTC(k) and maintain time differences within sub-nanosecond levels.

2. An Optimal Control Algorithm

To achieve optimal system control under varying conditions, the present invention innovatively constructs a multi-mode hybrid control architecture that realizes seamless switching between bang-bang control and LQG control strategies via a dynamic weighted switching function. The characteristics and switching logic of each control strategy are as follows:

2.1. Bang-Bang control

Bang-Bang control is a switching control strategy whose basic idea is to control the system by switching between two extreme control inputs (usually between the maximum and minimum values) to achieve the desired system response. The mathematical model of Bang-Bang control is generally relatively concise; assume its state equation is:

$$\dot{x}(t) = f(x(t), u(t)), \quad (1)$$

where $x(t)$ is the state variable of the system, and $u(t)$ is the control input.

In Bang-Bang control, the control input $u(t)$ can only take either the maximum value u_{max} or the minimum value u_{min} , and its choice is based on the error between the system state $x(t)$ and the desired state $x_{desired}$.

$$u(t) = \begin{cases} u_{max} & x(t) < x_{desired} \\ u_{min} & x(t) > x_{desired} \\ 0 & x(t) = x_{desired} \end{cases} \quad (2)$$

From this it can be seen that when the current system state $x(t)$ is less than the target state $x_{desired}$, the control input u_{max} is set to the maximum value u_{max} , that is, the control system rapidly drives the state toward the target. When the current system state $x(t)$ is greater than the target state $x_{desired}$, the control input $u(t)$ is set to the minimum value u_{min} , that is, it drives the system in the opposite direction to reduce the state error, eventually making the system approach the target state. When the system state $x(t)$ is close to the target state $x_{desired}$, the control input $u(t)$ can be set to zero (or a dead-zone value) to stop further adjustment. The Bang-Bang control strategy is simple and efficient, and is suitable for systems that require a fast response; however, it may also cause problems such as oscillations. In practical applications, the performance of Bang-Bang control is influenced by system dynamics, the magnitude of the error, and the required target accuracy.

2.2. LQG Control

LQG (Linear Quadratic Gaussian) control is a control method that combines optimal control with state estimation. It is typically applied to linear systems, especially in the presence of noise and uncertainties. To compute the optimal control value using the LQG control method, the quadratic cost function must be minimized:

$$J = E[\sum_{t=0}^{\infty} [x(t)^T W_Q x(t) + u(t)^T W_R u(t)]], \quad (3)$$

where $x(t)$ is the state vector of the system, $u(t)$ is the control input, W_Q is the state weighting matrix (usually used to penalise deviations of the system state from the target), and W_R is the control input weighting matrix (usually used to penalise control energy). E denotes the expectation operator, which accounts for the randomness of noise. The user can adjust the parameters W_Q and W_R in the above equation to achieve either fast response ($W_R \gg W_Q$) or slow response ($W_R \ll W_Q$). The steady-state Riccati equation is used to obtain the linear quadratic regulator (LQR) gain $\hat{K}_0$:

$$\hat{K}_0 = \phi(\tau)^T \hat{K}_0 \phi(\tau) + W_Q - \phi(\tau)^T \hat{K}_0 B (B^T \hat{K}_0 B + W_R)^{-1} B^T \hat{K}_0 \phi(\tau), \quad (4)$$

where B is the control matrix. From this, the control value is computed. Then a Kalman filter is used to predict the state estimate for the next time step $\hat{x}(t)$ and the filter gain $\hat{G}_0$, which is applied to steer the atomic clock.

$$u(t) = -(B^T \hat{K}_0 B + W_R)^{-1} B^T \hat{K}_0 \phi(\tau) \hat{x}(t) = -\hat{G}_0 \hat{x}(t) \quad (5)$$

LQG control combines optimal control (LQR) and state estimation (Kalman filtering). It can effectively control the system while minimising the cost function, ensuring that the optimal control input is obtained even under noisy and uncertain conditions. It handles process noise (such as uncertainties and external disturbances) and measurement noise very well, which makes it highly suitable for practical applications involving noisy and partially observable systems. Through the Kalman filter, LQG can effectively estimate the system state, providing accurate estimates even when some states are not directly measurable or are corrupted by noise. However, LQG control is sensitive to the accuracy of the system model; if there is significant uncertainty or model mismatch, the performance of LQG may not meet expectations.

2.3. Adaptive Selection Based on a Weighted Switching Function

Based on the calculation result of the weighted switching function, the system can select the control

strategy that best meets the current error and stability requirements. A threshold μ_{thresh} is set, and the appropriate control strategy is selected according to the value of the weighted switching function. The detailed procedures and flowchart of this optimal control algorithm are shown in Fig. 1.




Fig. 1. Flowchart of an Optimal Control Algorithm.

Step 1: Calculation of atomic timescale

Based on the observational data from the laboratory atomic clocks, an atomic time scale TA is established using an atomic time algorithm, which serves as the reference for subsequent control actions.

Step 2: Calculation of the error term $e[k]$

The error between the current system output $y(t)$ and the reference atomic time TA is computed as:

$$e[k] = TA - x[k] \quad (6)$$

Step 3: Calculation of the signal variation rate $J_s[k]$

The signal variation rate, which reflects the rate of change of the system output, is used to characterise system stability. It is calculated using the discrete-time expression:

$$J_s[k] = \left(\frac{e[k+1] - 2e[k] + e[k-1]}{T^2} \right)^2, \quad (7)$$

where TT is the sampling period, $e[k+1]$ and $e[k]$ denote the error values at the current and next time instants, respectively.

Step 4: Calculation of the weighted switching function $\mu[k]$

Based on the error term and the signal variation rate, the weighted switching function $\mu[k]$ is computed as:

$$\mu[k] = \frac{w_e \cdot |e[k]| + w_s \cdot J_s[k]}{w_e \cdot |e[k]| + w_s \cdot J_s[k] + w_d}, \quad (8)$$

where w_e, w_s, w_d are weighting coefficients that adjust the respective contributions of the error term and the variation rate term to the control strategy selection.

Step 5: Selection of the control strategy μ_{thresh}

A threshold μ_{thresh} is predetermined. The control mode is selected according to the value of $\mu[k]$:

If $\mu[k] > \mu_{\text{thresh}}$, Bang-Bang control (limit control) is activated, applying the maximum or minimum control effort to achieve rapid frequency pulling;

If $\mu[k] \leq \mu_{\text{thresh}}$, LQG optimal control (Linear Quadratic Gaussian control) is employed to achieve high-precision steady-state regulation.

Step 6: Control execution and iterative update

The computed control quantity is fed to the phase fine-tuning instrument to adjust the frequency of the atomic clock. The steered signal is then used as the reference for internal comparison acquisition, and the clock comparison data of all atomic clocks relative to this reference are re-obtained. The procedure returns to Step 1, forming a closed-loop iteration.

3. Experimental Results

To verify the effectiveness of the adaptive weighted switching control strategy combining Bang-Bang control and LQG (Linear Quadratic Gaussian) optimal control proposed in this paper, an atomic clock steering experimental system was constructed. The system uses the 10 MHz signal output by the national standard time UTC(NTSC) as the external frequency reference. The disciplined clock is an active hydrogen maser, whose output frequency is adjusted in real time by a phase fine-tuning instrument. The phase difference between the disciplined hydrogen maser and the UTC(NTSC) reference signal is measured in real time by an SR620 time-interval counter and fed back to the controller to form a closed loop. The overall structure of the experimental system is presented in Fig. 2.




Fig. 2. Structure diagram of the experimental system.

As a reference, UTC(NTSC) itself exhibits exceptionally high performance. Since 2024, its maximum absolute deviation from UTC has remained below 1.5 ns, with a 5-day frequency stability of approximately 4.1×10^{-16} and a 30-day stability of approximately 2.5×10^{-16} . Therefore, UTC(NTSC) serves as a reliable benchmark for evaluating the performance of UTC(Test).

The experiment was carried out over a continuous two-month period under stable laboratory conditions (temperature controlled within ± 1 °C). The sampling interval was set to one hour, and the control law was updated every one hour to balance computational load and real-time performance. The experimental results, shown in Fig. 3, indicate that the time differences between the control time UTC(Test) and the national standard time UTC(NTSC) are maintained within ± 0.5 ns over a two-month period.




Fig. 3. Time differences between the control time UTC(Test) and UTC(NTSC).

Additionally, Fig. 4 demonstrates a stability of $1.3\text{E-}15@1\text{day}$, and $6.78\text{E-}16@5\text{day}$. This frequency stability value is comparable to the best reported results for laboratory-based UTC(k) steering, confirming the effectiveness of our optimal approach.




Fig. 4. Frequency instability of control time UTC(Test) compared to UTC(NTSC).

Collectively, the results confirm that the proposed optimal control strategy achieves state-of-the-art time-keeping performance, validating its effectiveness and practical viability for future deployment in time distribution systems.

4. Conclusion

We have proposed an adaptive control strategy based on a weighted switching function that explicitly balances two critical performance metrics: control accuracy (measured by the deviation from the desired reference) and system stability (assessed by the rate of change of the system output). The switching function integrates these two measures and, through predefined threshold zones, enables smooth and hysteresis-free transitions among bang-bang control, progressive control, and optimal control. This design effectively mitigates abrupt control actions and the associated oscillations that often plague conventional switching approaches.

The main advantages of the proposed method are its flexibility and robustness: by dynamically adjusting the control law according to real-time error and stability conditions, the system maintains satisfactory performance across a wide range of operating scenarios. Experimental results confirm that the approach significantly improves control accuracy and adaptive capability while preserving overall system stability. However, the method's performance relies on appropriate selection of the threshold parameters, which may require prior knowledge of system dynamics. Future work will focus on developing automatic threshold tuning mechanisms and extending the framework to nonlinear and time-varying systems, further enhancing its practical applicability.

Acknowledgements

This work was supported by the Natural Science Basic Research Program of Shaanxi (Program No.2026JC-YBQN-0051), and National Young Metrology Talent Program.

References

- [1]. G. Panfilo, A. Harmegnies, L. Tisserand, A new weighting procedure for UTC, *Metrologia*, Vol. 51, Issue 3, 2014, pp. 285-292.
- [2]. C. Grebing, A. Al-Masoudi, S. Dörscher, S. Häfner, et al., Realization of a timescale with an accurate optical lattice clock, *Optica*, Vol. 3, Issue 6, 2016, pp. 563-569.
- [3]. G. Rovera, M. Abgrall, M. Siccaldi, Characterization of an auxiliary offset generator for steering of H masers, in *Proceedings of the 26th European Frequency and Time Forum (EFTF)*, 2012, pp. 14-17.
- [4]. H. Hachisu, F. Nakagawa, Y. Hanado, T. Ido, Months-long real-time generation of a time scale based on an optical clock, *Scientific Reports*, Vol. 8, Issue 1, 2018, 4243.
- [5]. S. Zhao, S. Dong, S. Bai, W. Wu, et al., A UTC(NTSC) steering algorithm based on an atomic clock ensemble scale, in *Proceedings of the Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, 2023, pp. 1-3.

- [6]. S. Zhao, S. Dong, S. Bai, Z. Gao, Research on an optimized frequency steering algorithm, *Journal of Electronics & Information Technology*, Vol. 43, Issue 5, 2021, pp. 1457-1464 (in Chinese).
- [7]. S. Bai, S. Zhao, S. Dong, Z. Gao, Atomic clock frequency steering method based on optimal control theory, in *Proceedings of the IEEE Ultrasonics, Ferroelectrics, and Frequency Control Joint Symposium (UFFC-JS)*, 2024, pp. 1-4.
- [8]. P. A. Koppang, State space control of frequency standards, *Metrologia*, Vol. 53, Issue 3, 2016, R60.

(013)

Spatially-Resolved Conductivity Mapping of Ultra-Thin Glass Using Guard Ring Electrodes and Silver Conductive Ink Contact Layers

D. Otto¹, M. Eiserbeck², D. Käbler² and G. Bausch¹

¹ Leipzig University of Applied Sciences, Faculty of Engineering, Electronic Engineering Laboratory,
Eilenburger Straße 13, 04317 Leipzig, Germany

² Leipzig University of Applied Sciences, Research and Transfer Centre Leipzig, Electronic Engineering
Laboratory, Eilenburger Straße 13, 04317 Leipzig, Germany
Tel.: + 0049 341 3076 3145
E-mail: daniel.otto@htwk-leipzig.de

Summary: This paper presents a spatially-resolved conductivity mapping method for ultra-thin, brittle, and surface-rough dielectric samples. The method combines a three-terminal guard ring electrode configuration with silver conductive ink contact layers applied by tampon printing, and spring-loaded pogo-pin probes integrated into an automated XY gantry system. The key novelty is not spatially-resolved conductivity measurement per se, but a specific contact architecture that simultaneously addresses surface roughness, mechanical fragility, and high-field operation, a combination not covered by any existing single method. Validation measurements on SCHOTT AS87 eco aluminosilicate glass (100 μm , $\sigma = 5.6 \times 10^{-12}$ S/cm) using three guard ring configurations at defined spatial positions on two reference samples yield conductivity values of 5.5 to 5.8 pS/cm at 30 kV/cm for the 1.0 mm guard gap configuration, in excellent agreement with the manufacturer specification. The spatial variation across measurement positions confirms sample homogeneity and demonstrates the method's capability for position-resolved dielectric characterisation.

Keywords: Conductivity mapping, Guard ring electrodes, Dielectric materials, High electric field, Poole-Frenkel emission, Silver conductive ink, Tampon printing, Pogo-pin probe.

1. Introduction

The characterisation of volume conductivity in thin dielectric materials is essential for quality control and material development in applications ranging from display glass to insulating layers in power electronics and electrostatic chuck ceramics. Standard methods such as ASTM D257 use guarded electrode configurations to suppress surface leakage, but assume flat, polished, macroscopically uniform samples and require direct metal contact. These conditions are rarely met in practice for ultra-thin, rough, or fragile substrates.

Three specific challenges arise simultaneously in industrial thin-glass characterisation. Surface roughness prevents reliable direct electrode contact. Mechanical fragility of substrates below 200 μm thickness limits the applicable contact force. High-field operation above 10 kV/cm is required to resolve conductivity values below 10^{-11} S/cm within practical measurement times.

The key novelty of this work is not spatially-resolved conductivity mapping per se, which exists in numerous forms [9], but a specific contact and scanning architecture: tampon-printed compliant silver ink as a soft conformal electrode, combined with guard ring geometry, spring-loaded pogo-pin probing, and automated XY-grid scanning, enabling high-field operation on samples that are simultaneously rough, brittle, and large-area. The architecture supports large-area coverage exceeding 100 mm and tolerates surface roughness through the compliant ink contact layer. High-field operation up to 50 kV/cm is achieved

without mechanical damage, as spring-loaded pogo-pin probes limit contact force to 1.13 ± 0.23 N [10]. Automated XY scanning provides conductivity maps at mm-scale spatial resolution.

2. Measurement Method

2.1. Guard Ring Electrode Configuration

The measurement employs a three-terminal guard ring configuration following ASTM D257 to suppress edge field effects and surface leakage currents [1, 8]. The guard ring surrounds the central measurement electrode, diverting leakage to a separate terminal so that only the volume current through the sample contributes to the conductivity measurement. Without guarding, surface leakage and fringe-field contributions at the electrode edge would cause systematic overestimation of the bulk conductivity, particularly at field strengths above 10 kV/cm where surface currents can become comparable in magnitude to the bulk current in high-resistivity samples [5]. Fig. 1 shows a cross-sectional finite-element simulation (FEMM) of the electric field distribution, illustrating that the guard ring effectively confines the measurement field to the volume beneath the main electrode.

Three configurations with guard gaps of 0.1, 0.5 and 1.0 mm, with the effective electrode diameters 10.1, 10.5 and 11.0 mm, were investigated.




Fig. 1. Cross-sectional FEMM simulation of the electric field distribution in the guard ring configuration. The guard electrodes (left and right, outer) divert leakage current away from the main measurement electrode (centre), confining the measured current to the sample volume directly beneath the main electrode.

The effective measurement area A_{eff} is corrected for each guard gap following ASTM D257 [4] with

$$A_{eff} = \pi * \left(r_{main} + \frac{g}{2} \right)^2, \quad (1)$$

where $r_{main} = 5.0$ mm and g is the radial guard gap width.

Conductivity σ is then calculated from measured current I_{meas} , applied voltage V_{appl} , sample thickness l_{DUT} , and effective electrode area A_{eff} , shown in equation (2).

$$\sigma = \frac{(I_{meas} * l_{DUT})}{(V_{appl} * A_{eff})} \quad (2)$$

2.2. Silver Conductive Ink Contact Layer

Direct metal contact with thin glass risks mechanical damage and produces unreliable contact resistance on surfaces with $R_a > 0.1$ μm . A silver conductive ink (Bectron CP 6662 [6], 72–74 % Ag, dry thickness < 10 μm) is applied to both sample faces by tampon printing, conforming to the surface topography and providing a low-resistance, mechanically compliant interface. The ink conductivity $\sigma_{paint} \approx 10^5$ S/cm renders the contact resistance negligible compared to the glass bulk resistance ($\sigma_{glass} \approx 10^{-12}$ to 10^{-8} S/cm). The sheet resistance of the ink layer (< 0.1 Ω/sq) remains orders of magnitude below any glass sample resistance at all applied field strengths, confirming that the ink does not dominate or systematically bias the conductivity measurement [6].

Fig. 2 shows a SCHOTT AS87 eco sample with three tampon-printed guard ring electrode sets, each corresponding to a different guard gap configuration. The three functional patterns, each consisting of a central main electrode disk and a surrounding guard ring, are clearly visible. The ink edges are sharply defined and conform to the glass surface. Fig. 3 shows a representative cross-sectional height profile of a dried ink layer measured by contact profilometry. The ink layer reaches approximately 3 μm dry thickness over a width of 1.0 ± 0.025 mm, consistent with the manufacturer specification (< 10 μm dry thickness [6]), and confirms that the ink layer is thin relative to the 100 μm glass thickness and does not mechanically or electrically interfere with the sample.



Fig. 2. SCHOTT AS87 eco aluminosilicate glass sample (100 μm) with tampon-printed silver conductive ink guard ring electrode sets. The three functional patterns correspond to guard gap configurations of 0.1, 0.5 and 1.0 mm and are visible as concentric disk-and-ring structures. Additional ink shapes visible in the image are artefacts of the tampon printing pattern and carry no electrical function.




Fig. 3. Cross-sectional height profile of the tampon-printed silver ink contact layer (Bectron CP 6662) on the glass surface, measured by contact profilometry. The ink layer reaches approximately 3 μm dry thickness over a width of 1.0 ± 0.025 mm, consistent with the manufacturer specification and confirming conformal deposition on the glass surface.

2.3. Automated Measurement System

Spring-loaded pogo-pin probes (HARWIN P25-0822 [10], 1.13 ± 0.23 N contact force, < 50 m Ω contact resistance) contact the printed ink pads without damaging the glass substrate. The probes are mounted on a three-axis XY gantry with ± 1 mm positional repeatability and force-controlled Z-approach. The electrometer (Keysight B2985A [7], 0.01 fA resolution, ± 1000 V) is connected in three-terminal guarded configuration. A 1000:1 voltage divider (1 M Ω , 0.1 % precision) monitors the applied voltage independently of the source output. The entire assembly is housed in an aluminum shielding enclosure (800 $\times$ 500 $\times$ 700 mm) to suppress electromagnetic interference. Eight BME280 sensors monitor temperature and relative humidity throughout each measurement. Fig. 4 shows the probe head positioned over the glass sample during a measurement.

2.4. Measurement Protocol

Applied voltage is stepped from 100 V to 500 V (10–50 kV/cm for 100 μm samples) in steps of 25 V,

yielding 17 voltage levels. At each level, the current is recorded at 5 s intervals for 600 s (120 data points per level). The Keysight B2985A integration time is set to 100 NPLC (2000 ms at 50 Hz). All measurements are performed at 21–22 °C and 45–47 % relative humidity (monitored by 8 BME280 sensors). The steady-state conductivity at each voltage level is computed from the mean of the final 30 data points ($t = 450\text{--}595$ s), minimising displacement-current contributions. Each measurement position is characterised in triplicate; the standard deviation over the three repeats is used as the primary repeatability indicator. Joule heating is negligible: at 500 V and the measured currents (< 500 nA), dissipated power is below 0.25 mW.



Fig. 4. Interior view of the automated measurement system. The spring-loaded pogo-pin probe head (centre) is positioned over the silver ink electrode pads on the glass sample. The three-terminal guarded probe configuration is connected to the Keysight B2985A electrometer. The XY gantry allows automated positioning across the sample surface.

3. Validation Results

Two SCHOTT AS87 eco aluminosilicate glass samples (B and C, 79×69 mm, 100 μm thickness, $\sigma_{\text{spec}} = 5.6 \times 10^{-12}$ S/cm [11]) were characterised using the three guard ring configurations (gaps: 0.1 mm, 0.5 mm, 1.0 mm). Each configuration is printed at a defined spatial position on the sample ($x = 23$ mm, 43 mm, 63 mm; $y = 30$ mm, measured from the lower-left corner). Each position was measured in triplicate.

The field-dependent conductivity is described by the combined Arrhenius–Poole-Frenkel model:

$$\sigma(T, E) = \sigma_0 * \exp\left(\frac{-Ea}{k_B T}\right) * \exp(\beta * \sqrt{E}), \quad (3)$$

where σ_0 is a pre-exponential factor, Ea the thermal activation energy, k_B the Boltzmann constant, T the absolute temperature, and β the Poole-Frenkel coefficient [2, 3].

Fig. 5 shows the raw measured current I as a function of applied field strength E for all

measurement runs on samples B and C, colour-coded by the instantaneous sample temperature recorded by the BME280 sensors (19–25 °C). The point cloud shows the prediction of the combined Arrhenius–Poole-Frenkel model (3) evaluated over the measured temperature range. All measured data from both samples lies within this prediction band across the full field range from 10 to 50 kV/cm, confirming that the combined temperature- and field-dependent model describes the measured conduction behaviour.




Fig. 5. Measured current I vs. applied field strength E for all measurement runs on samples B and C, colour-coded by measurement temperature (19–25 °C). The point cloud shows the prediction band of the combined Arrhenius–Poole-Frenkel model (3) evaluated over the measured temperature range. All measured data from both samples lies within the model band throughout the full field range, confirming consistency with Poole-Frenkel charge transport.

Fig. 6 shows the measured conductivity σ as a function of applied field strength E for all six measurement configurations. All curves exhibit the characteristic field-enhanced increase consistent with Eq. (3). The close agreement between samples B and C across all field strengths confirms inter-sample reproducibility. A systematic separation is visible between the three guard ring gap configurations: the 0.1 mm gap yields the highest apparent σ , the 1.0 mm gap the lowest. This trend is consistent with the guard-gap-dependent A_{eff} correction (1): larger guard gaps result in a larger A_{eff} and correspondingly lower computed σ for the same measured current. The shaded uncertainty bands for each curve reflect the standard deviation across three consecutive measurements at each position and confirm that the separation between configurations is statistically significant and not attributable to measurement scatter. The 1.0 mm configuration, whose effective area is best defined by the A_{eff} correction, yields results closest to the manufacturer specification across the entire field range.

At the reference field of 30 kV/cm, the 1.0 mm guard gap configuration yields $\sigma = 5.5$ pS/cm (sample B) and $\sigma = 5.8$ pS/cm (sample C), in excellent agreement with the manufacturer datasheet value of 5.6×10^{-12} S/cm [11]. The deviation from the datasheet value is below 4 % for both samples, well within the

combined measurement uncertainty. The combined measurement uncertainty (Type A + B) is estimated at approximately 7 %, with the dominant Type B contributions from the voltage divider calibration accuracy (0.1 %) and the glass substrate thickness tolerance; the Type A repeatability component is quantified from the triplicate measurements at each position. The Type A component, derived from the standard deviation across three consecutive measurements, is below 2 % for the 1.0 mm configuration on sample C and below 32 % for sample B, where a larger spread between repeats was observed at this field strength. This difference in repeatability between the two samples is attributed to slight variations in pogo-pin contact quality between measurement runs and not to a material inhomogeneity, as the mean values remain consistent.




Fig. 6. Conductivity σ vs. applied field strength E for all three guard ring configurations on two SCHOTT AS87 eco reference samples (B and C, 100 μm). Shaded bands indicate $\pm 1\sigma$ repeatability over three consecutive measurements. The dash-dotted horizontal line marks the manufacturer specification of 5.6 pS/cm [11]; the dashed vertical line marks the 30 kV/cm reference field. The field-dependent increase is consistent with Poole-Frenkel charge transport.

Fig. 7 shows the spatially-resolved conductivity map at 30 kV/cm across the three measurement positions on each sample. The spatial variation across the three positions per sample is below $\pm 25\%$ of the mean value and is consistent between samples B and C, confirming both the homogeneity of the reference glass and the absence of systematic position-dependent bias introduced by the tampon-printed electrode geometry. The observed position-to-position variation is dominated by the guard-gap-dependent A_{eff} correction rather than genuine spatial conductivity differences in the material: when comparing configurations with the same gap width across samples B and C, the inter-sample spread is below 3 %, confirming excellent material homogeneity within the measured area. These results validate the method's capability to distinguish electrode-geometry-induced offsets from genuine spatial conductivity variations, which is a prerequisite for reliable spatially-resolved conductivity mapping.




Fig. 7. Spatially-resolved conductivity map of two SCHOTT AS87 eco aluminosilicate glass samples (B and C, 79 $\times$ 69 mm, 100 μm) at 30 kV/cm. Each measurement position ($x = 23, 43, 63$ mm; $y = 30$ mm) carries a different guard ring configuration (gaps: 0.1 mm, 0.5 mm, 1.0 mm). Values in pS/cm, averaged over three repeats; $\pm$ standard deviation shown below each position. The systematic decrease with guard gap width is consistent with the A_{eff} correction (1); the 1.0 mm gap results (5.5–5.8 pS/cm) are in excellent agreement with the manufacturer specification of 5.6 pS/cm [11].

4. Conclusion

A spatially-resolved conductivity mapping method for ultra-thin dielectric samples has been presented and validated. The combination of tampon-printed silver ink contact layers, three-terminal guard ring electrode configurations, and spring-loaded pogo-pin probing enables high-field (up to 50 kV/cm) conductivity characterisation on brittle, rough, and large-area samples without mechanical damage. Validation on SCHOTT AS87 eco aluminosilicate glass (100 μm) using three guard ring geometries at six spatial positions on two reference samples confirms excellent agreement with the manufacturer specification ($\sigma = 5.5\text{--}5.8$ pS/cm at 30 kV/cm vs. 5.6 pS/cm datasheet), with inter-sample and inter-position repeatability within $\pm 25\%$. The method is directly applicable to ultra-thin glass characterisation in display, photovoltaic, and power electronics applications. A systematic raster scan using a uniform guard ring geometry across a $\geq 3 \times 3$ position grid (enabling fully comparable spatial conductivity maps) is planned for a forthcoming journal article.

Acknowledgements

This work was carried out within the project EffElek+ and supported by the German Federal Ministry of Education and Research (BMBF) through Projektträger Jülich (PtJ) under Grant No. 03XP0620B. The authors thank enaDyne GmbH for providing reference glass samples and technical support as project partner.

References

- [1]. Standard test methods for DC resistance or conductance of insulating materials, ASTM D257-14, *ASTM International*, West Conshohocken, 2021.

- [2]. J. Frenkel, On pre-breakdown phenomena in insulators and electronic semi-conductors, *Physical Review*, Vol. 54, Issue 9, 1938, pp. 647-648.
- [3]. S. M. Sze, K. K. Ng, Physics of Semiconductor Devices (3rd Ed.), *Wiley-Interscience*, Hoboken, 2006.
- [4]. C. Freye, F. Jenau, Model-based accuracy enhancements for guarded conductivity measurements: Determination of effective electrode areas utilising numerical field simulation, *High Voltage*, Vol. 3, Issue 4, 2018, pp. 217-225.
- [5]. Dielectric and resistive properties of solid insulating materials - Part 3-1: Determination of resistive properties (DC methods), IEC 62631-3-1:2023, *International Electrotechnical Commission*, Geneva, 2023.
- [6]. Bectron CP 6662 — Technical data sheet, TDS N°1047, Rev. 2, *ELANTAS Europe*, Hamburg, 2021.
- [7]. Measuring insulating material resistivity using the B2985A/87A — Technical overview, Doc. No. 5992-0124EN, *Keysight Technologies*, Santa Rosa, 2014.
- [8]. H.-G. Lee, J.-G. Kim, Volume and surface resistivity measurement of insulating materials using guard-ring terminal electrodes, *Energies*, Vol. 13, Issue 11, 2020, 2811.
- [9]. J. Krupka, D. Nguyen, J. Mazierska, Microwave and RF methods of contactless mapping of the sheet resistance and the complex permittivity of conductive materials and semiconductors, *Measurement Science and Technology*, Vol. 22, Issue 8, 2011, 085703.
- [10]. P25-0822 spring-loaded contact — Product data sheet, *HARWIN plc*, Havant, 2023.
- [11]. AS 87 eco — Aluminosilicate glass, product data sheet (preliminary), *SCHOTT AG*, Mainz, 2025.

A Fundamental Evaluation of a Ventilator Operation Education System Applying Eye-Tracking Technology

N. Katsura¹, S. Sakakida¹, R. Kurihara¹, Y. Uchida² and T. Wabe³, A. Fujii⁴

¹ Teikyo Heisei University, 2-51-4 Higashi-Ikebukuro, Toshima-ku, 170-8445, Tokyo, Japan

² Tokyo University of Technology, 5-23-22 Nishikamata, Ota-ku, 144-8535, Tokyo, Japan

³ Junshin Gakuen University, 1-1-1 Chikushigaoka, Minami-ku, Fukuoka-shi, 815-8510, Fukuoka, Japan

⁴ Tokyo Medical & Health College, 5-13-4 Nakakasai, Edogawa-ku, 134-0083, Tokyo, Japan

Tel.: +81 90 40979626

E-mail: n.katsura@thu.ac.jp

Summary: Clinical engineers are required to possess advanced skills, but the quality of equipment operation depends on experience, making it difficult to transfer skills to new recruits. This study focused on ventilators to objectively analyze the eye movement behavior of novice operators as an initial step toward establishing educational indicators. Compared to blood purification devices and other equipment, ventilators have a simpler circuit configuration and do not involve complex processes such as priming. Therefore, we believe they are suitable as a model for more purely extracting and analyzing the empirical elements of operation techniques and contributing to the construction of effective educational indicators.

Keywords: Educational indicators, Ventilator, Gaze measurement, Clinical engineer, Proficiency, Experience.

1. Introduction

In the operation of medical devices, the ability to selectively filter essential information from a patient's clinical status and the surrounding environment to make rapid and appropriate decisions is of paramount importance. To establish objective educational indicators, it is necessary to clarify how operator proficiency and clinical experience influence operational behavior [4-6]. For this purpose, we selected a mechanical ventilator as the target device, as its circuit configuration and operational framework are relatively straightforward and do not involve a priming process [1-3].

Regarding the evaluation of proficiency in medical device operation, previous literature has reported quantitative assessments using motion-tracking systems and the extraction of characteristics related to the situational awareness of expert nurses through gaze analysis [7, 8]. However, these prior studies have primarily focused on observations within specific Areas of Interest (AOIs) or the classification of proficiency itself; they have not yet comprehensively captured the shifts in gaze behavior throughout the entirety of an operation.

Therefore, as a fundamental study aimed at developing an educational system that can objectively facilitate the learning of expert experiential knowledge, we conducted an analysis using gaze sensors. As an initial step toward clarifying the effect of clinical experience on gaze behavior during ventilator circuit assembly, this study first analyzed the eye movements of novice clinical engineers with no prior clinical experience in ventilator management.

2. Equipment

In this study, a Servo-i ventilator (SIEMENS, Model-No. 6449701E407E, Serial-No. 106821S11) was utilized (Fig. 1). The ventilator was driven by a compressor (MODEL-PE, N-tech Co., Ltd., Model 71R637-P170). To provide the gas supply, air from the compressor was distributed and connected to both the oxygen and air supply ports of the ventilator. A ventilator circuit manufactured by Fisher & Paykel Healthcare was employed (Fig. 2). The complete experimental setup, comprising the ventilator, compressor, and breathing circuit, is illustrated in Fig. 1. The gas flow path follows the arrows shown in the figure: inspiratory gas flows through the blue tube (inspiratory limb), while expiratory gas flows through the white tube (expiratory limb).



Fig. 1. Gas flow in a ventilator circuit.

For gaze measurement, the Gazohound gaze analysis system was used (Fig. 3). This system consists of two camera units: an eye camera (Fig. 3-①) to capture eye movements and a front camera (Fig. 3-②) to record the subject's field of view. The eye camera is equipped with a near-infrared CMOS sensor, and the front camera utilizes a color CMOS sensor. Gaze analysis is performed by tracking eye movements captured via near-infrared irradiation.



Fig. 2. Ventilator circuit.



Fig. 3. Line of sight analysis system.

3. Experimental Methods

3.1. Experimental Design and Subjects

To evaluate the effect of clinical engineering proficiency on eye movement behavior, this study enrolled two novice clinical engineers with no prior experience in ventilator management as participants (hereafter referred to as the **novice group**). The experiment was conducted in a quiet environment designed to simulate a real clinical setting [11, 12].

3.2. Gaze Measurement System Configuration

The Gazohound gaze analysis system (manufactured by GAZO Co., Ltd.) was utilized for gaze measurement (Fig. 3). This system is composed of two primary components: an eye-tracking camera using a near-infrared CMOS sensor (Fig. 3-①) and a color CMOS front camera to record the subject's field of view (Fig. 3-②). During measurement, input devices were assigned via software, and the following sampling conditions were established to optimize both data accuracy and processing load:

- **Eye Tracking:** To capture rapid eye movements, the resolution was set to 320×200 pixels with a frame rate of 120 fps;
- **Field-of-View Recording:** To ensure visibility of the surrounding environment, the resolution was set to 1280×720 pixels with a frame rate of 30 fps.

Input devices (USB 2.0 Camera and FHD Camera) were assigned within the software, and microphone input was disabled.

3.3. Optimization and Calibration of the Gaze Measurement System

To ensure measurement accuracy, an individual optimization process was conducted for each subject prior to task execution. First, the pupil detection range was defined (Fig. 4); this was adjusted to exclude unnecessary background, such as facial contours, while ensuring the pupil remained within the detection range even when the subject moved their eyes in any direction. Subsequently, calibration was performed to align specific points within the field-of-view video with the fixation points, thereby ensuring tracking stability and precision (Fig. 5).




Fig. 4. Pupil detection range setting.



Fig. 5. Completion confirmation screen.

4. Experimental Tasks and Procedures

4.1. Experimental Tasks

The experimental task involved the assembly of a ventilator circuit (Fisher & Paykel Healthcare) to a mechanical ventilator (Servo-i, SIEMENS) [13-15]. Participants were required to complete a series of

assembly steps, connecting the inspiratory limb (blue) and expiratory limb (white) tubes to the designated ports on the ventilator unit and compressor (Fig. 6). The specific procedures were as follows:

1. Pre-calibration: Prior to the start of measurements, gaze calibration and optimization of the pupil detection range were performed according to the individual ocular characteristics of each participant. Adjustments were specifically made to ensure the pupil remained within the detection range even during vertical and horizontal eye movements, thereby ensuring the stability of the eye-tracking data.

2. Task Execution: Participants performed the assembly process by connecting the blue inspiratory and white expiratory tubes to the specified ports on the ventilator and compressor units.

3. Data Recording: Gaze data synchronized with field-of-view video were continuously recorded throughout the entire process, from the start of the task to its completion. The acquired time-series data were subsequently utilized for spatial distribution analysis via heatmaps, as well as for the qualitative and quantitative analysis of fixation patterns.



Fig. 6. Ventilator circuit assembly.

4.2. Data Collection and Analysis Methodology

Gaze data synchronized with field-of-view (FOV) video were continuously recorded throughout the entire process, from the commencement to the completion of the task. The acquired time-series data were subjected to the following analyses:

- **Spatial Distribution Analysis:** Visualization of fixation points using heatmaps;
- **Gaze Pattern Analysis:** Qualitative and quantitative analysis of specific fixation sites and gaze transition patterns relative to task proficiency.

5. Analysis of Eye Movement During Ventilator Circuit Assembly

The present study investigated the influence of clinical experience on eye movement behavior during

ventilator circuit assembly. Two novice clinical engineers (CEs) with no prior clinical experience in ventilator management (hereafter referred to as the novice group) participated in this study. Eye movements during the task were continuously recorded using an eye-tracking device [9, 10]. The collected time-series data were analyzed to compare gaze locations and fixation patterns associated with the participants' level of experience.

5.1. Spatial Distribution of Fixation Points

The heat maps of the novice group (Figs. 7 and 8) revealed that fixation points were widely distributed not only on the task-relevant objects but also throughout the entire workspace. In particular, Fig. 7 demonstrated frequent gaze shifts away from the ventilator toward the participant's own hands, the ventilator frame.




Fig. 7. Heatmap: Novice Subject 1.




Fig. 8. Heatmap: Novice Subject 2.

5.2. Visual Feedback during Manual Operation

A notable characteristic observed in the novice group (Fig. 8) was prolonged and direct fixation on their own hands during manipulation.

6. Discussion

This study evaluated the eye movement behavior of novice clinical engineers during ventilator circuit assembly to understand their visual attention characteristics. The primary finding was that novice

participants exhibited a widely distributed gaze pattern and prolonged fixation on their own hands, rather than focusing efficiently on task-relevant components [19-21].

Frequent gaze shifts toward task-irrelevant areas further suggest that the spatial relationships among circuit components and the assembly procedure had not yet been fully internalized [16-18]. As a result, participants repeatedly searched for information necessary for subsequent operations [22, 23]. Such inefficient gaze behavior may contribute to prolonged task completion time and increase the likelihood of operational errors, potentially compromising patient safety [24-27].

These findings indicate that novice training should focus not only on teaching procedural knowledge but also on developing efficient visual attention strategies. Educational approaches incorporating visualization of expert gaze behavior may facilitate the acquisition of efficient gaze allocation strategies during ventilator circuit assembly [28].

7. Future Work

This study investigated the eye movement behavior of novice clinical engineers during ventilator circuit assembly. Although prolonged fixation on the hands and broadly distributed gaze patterns were observed, the absence of an experienced comparison group limited our ability to determine whether these characteristics were attributable to differences in clinical experience.

Future studies should include experienced clinical engineers (expert group) to compare gaze behavior between novice and expert participants. Quantitative comparisons of fixation locations, fixation duration, gaze transition patterns, and the frequency of hand-directed fixations are expected to clarify experience-dependent visual strategies and identify gaze behaviors characteristic of expert performance.

Furthermore, incorporating expert gaze behavior into educational materials may provide an effective approach for supporting skill acquisition in novice clinical engineers. Such findings could contribute to the development of evidence-based educational programs aimed at improving the efficiency and safety of ventilator circuit assembly training.

References

- [1]. T. Murakami, et al., A study on the relationship between proficiency and motion information in medical device operation, in *Proceedings of the Multimedia, Distributed, Cooperative and Mobile Symposium (DICOMO)*, 2019, pp. 1590-1593 (in Japanese).
- [2]. R. Okane, et al., Characteristics of observation by skilled nurses in mechanical ventilator attachment settings: Comparison with novice nurses using eye tracking, *Japanese Journal of Ergonomics*, Vol. 58, Issue 4, 2022, pp. 186-194 (in Japanese).
- [3]. J. Reason, Human error: Models and management, *BMJ*, Vol. 320, Issue 7237, 2000, pp. 768-770.
- [4]. Institute of Medicine, To Err is Human: Building a Safer Health System, *National Academies Press*, Washington, 2000.
- [5]. L. T. Kohn, et al., Making health care safer: A critical analysis of patient safety practices, *Agency for Healthcare Research and Quality*, Rockville, 2001.
- [6]. P. Carayon, et al., Work system design for patient safety: The SEIPS model, *Quality and Safety in Health Care*, Vol. 15, Issue Supplement 1, 2006, pp. i50-i58.
- [7]. A. Gegenfurtner, et al., Expertise differences in the comprehension of visual representations: A meta-analysis of eye-tracking research, *Educational Research Review*, Vol. 6, Issue 3, 2011, pp. 236-252.
- [8]. T. Tien, et al., Eye-tracking in surgery: A systematic review, *Surgical Endoscopy*, Vol. 28, Issue 6, 2014, pp. 1755-1766.
- [9]. A. T. Duchowski, Eye Tracking Methodology: Theory and Practice (3rd Ed.), *Springer*, Cham, 2017.
- [10]. K. Holmqvist, et al., Eye Tracking: A Comprehensive Guide to Methods and Measures, *Oxford University Press*, Oxford, 2011.
- [11]. S. Brams, et al., To go with the flow or not: Eye movement modeling examples enhance simulator-based surgical training, *Journal of Surgical Education*, Vol. 78, Issue 4, 2021, pp. 1300-1310.
- [12]. T. Drew, M. L.-H. Võ, J. M. Wolfe, The invisible gorilla strikes again: Sustained inattention blindness in expert observers, *Psychological Science*, Vol. 24, Issue 9, 2013, pp. 1848-1853.
- [13]. K. Catchpole, et al., Patient safety and medical device technology, *British Journal of Anaesthesia*, Vol. 111, Issue 1, 2013, pp. 11-17.
- [14]. J. Sweller, Cognitive load theory, learning difficulty, and instructional design, *Learning and Instruction*, Vol. 4, Issue 4, 1994, pp. 295-312.
- [15]. F. Paas, et al., Cognitive load theory and instructional design: Recent developments, *Educational Psychologist*, Vol. 38, Issue 1, 2003, pp. 1-4.
- [16]. D. M. Gaba, The future of simulation in healthcare, *Quality and Safety in Health Care*, Vol. 13, Issue Supplement 1, 2004, pp. i2-i10.
- [17]. C. Pruden, G. B. Beecham, M. Waseem, Human factors in medical simulation, in StatPearls, StatPearls Publishing, Treasure Island, 2025, <https://www.ncbi.nlm.nih.gov/books/NBK559226/>
- [18]. S. Shappell, et al., The Human Factors Analysis and Classification System - HFACS, *Ashgate Publishing*, Farnham, 2000.
- [19]. J. N. Vickers, Perception, Cognition, and Decision Training: The Quiet Eye in Action, *Human Kinetics*, Champaign, 2007.
- [20]. E. A. Henneman, J. L. Marquard, D. L. Fisher, A. Gawlinski, Eye tracking: A novel approach for evaluating and improving the safety of healthcare processes in the simulated setting, *Simulation in Healthcare*, Vol. 12, Issue 1, 2017, pp. 51-56.
- [21]. E. A. Henneman, A. Gawlinski, Eye-tracking technology: An innovative approach for teaching and evaluating patient safety behaviors, *Nursing Education Perspectives*, Vol. 37, Issue 6, 2016, pp. 356-357.
- [22]. C. M. Schulz, et al., Situation awareness in anesthesia: Concept and research, *Anesthesiology*, Vol. 118, Issue 3, 2013, pp. 729-742.
- [23]. F. Wang, et al., An algorithmic approach to determine expertise development using object-related gaze

- pattern sequences, *Behavior Research Methods*, Vol. 54, Issue 1, 2022, pp. 493-507.
- [24]. B. Law, M. S. Atkins, A. E. Kirkpatrick, A. J. Lomax, Eye gaze patterns differentiate novice and experts in a virtual laparoscopic surgery training environment, in *Proceedings of the Symposium on Eye Tracking Research & Applications (ETRA)*, 2004, pp. 41-48.
- [25]. K. Rayner, Eye movements in reading and information processing: 20 years of research, *Psychological Bulletin*, Vol. 124, Issue 3, 1998, pp. 372-422.
- [26]. M. A. Khawaja, et al., Using eye tracking to measure cognitive load, in *Proceedings of the 27th International Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA)*, 2009, pp. 317-320.
- [27]. P. M. Fitts, M. I. Posner, *Human Performance*, Brooks/Cole, Belmont, 1967.
- [28]. K. A. Ericsson, et al., The role of deliberate practice in the acquisition of expert performance, *Psychological Review*, Vol. 100, Issue 3, 1993, pp. 363-406.

Estimating Outdoor Eye-Level Daylight Exposure from Satellite Data: Protocol for a Proof-of-Concept Experiment

S. W. de Vries¹, J. Mardaljevic², H. S. M. Kort¹ and J. van Duijnhoven¹

¹Eindhoven University of Technology, Building Lighting Group, Eindhoven, The Netherlands

²Loughborough University, School of Architecture, Building & Civil Engineering, Leicestershire, UK

E-mail: s.w.d.vries@tue.nl

Summary: Light impacts human health and well-being, making quantifications of personal light exposure (PLE) valuable for related research and practice. PLE is typically measured using wearable light dosimeters, which can be burdensome to deploy at scale. As such, for population-level research, a scalable estimate of sufficient accuracy may be preferable to individual measurements. This abstract presents the protocol for an ongoing proof-of-concept experiment investigating whether eye-level outdoor daytime PLE can be estimated remotely, using only a pocket-worn smartphone as subject-side input. Building on evidence that Copernicus Atmosphere Monitoring Service (CAMS) data can accurately determine unobstructed horizontal daylight levels, we combine CAMS data with Climate-Based Daylight Modelling (CBDM) simulations to recreate estimates of outdoor eye-level PLE. A single participant walks routes at an open and a dense urban site in Eindhoven, the Netherlands, under varying sky conditions, wearing helmet-mounted instruments that record approximate viewing direction, sky-dome obstructions, and eye-level light exposure as ground truth. Models of varying complexity, differing in how they represent viewing direction and obstructions within the field of view, are compared against these measurements. The outcomes serve to demonstrate feasibility and to inform a data-collection protocol for a larger, geographically distributed campaign in which the approach can be validated.

Keywords: Personal light exposure, Satellite data, Light measurements, Light simulation.

1. Introduction

The light an individual's eyes are exposed to over time, known as personal light exposure (PLE), affects their health and well-being. Light modulates various physiological and behavioural processes such as alertness, sleep, and mood primarily through activation of the intrinsically photosensitive retinal ganglion cells in the human eye [1, 2], making recordings of PLE valuable for studies on these outcomes (e.g. [3, 4]). Light exposure outdoors is also recognised to play a role in delaying the onset and progression of myopia, rendering PLE informative for myopia research too [5, 6].

PLE is typically quantified using small wearable light sensors, also known as (light) dosimeters [7]. These devices should ideally be positioned at the eye [8], yet even compact spectacle-mounted dosimeters remain bulky and rather obtrusive for the wearer [9-11]. Body-worn placements are better tolerated but introduce discrepancies relative to eye-level measurements [12-14]. Regardless of placement, any method that relies on individual measurements with a wearable device is burdensome for participants, as well as for researchers: devices must be purchased, distributed, maintained, charged, and periodically retrieved for data collection. Consequently, most existing PLE datasets are sparse and originate from regions with sufficient financial and research infrastructure to support such studies, leaving other regions underrepresented. For research questions posed at the level of populations rather than individuals, this suggests a different approach: a scalable estimate of sufficient accuracy that can be

applied anywhere may be more valuable than high-fidelity individual measurements that cannot easily be deployed at scale.

Accordingly, we propose a measurement-simulation approach to estimate eye-level PLE without a dedicated light-measuring device. As an initial step, we limit the scope to outdoor daytime conditions; extending the approach to indoor environments is a subsequent step. In line with the scalability requirement above, we restrict subject-side inputs to data already collected by a device people carry anyway: a pocket-worn smartphone. Using timestamped location data from this phone, we determine daylight availability from the Copernicus Atmosphere Monitoring Service (CAMS) [15], as previous research has demonstrated that such data can yield accurate measures of unobstructed horizontal daylight exposure [16]. We then use a Climate-Based Daylight Modelling (CBDM) approach [17, 18] to reconstruct sky dome luminance distributions from these data, model environmental obstructions from geospatial data sources, and simulate the measurements (photopic illuminance and melanopic equivalent daylight illuminance (m-EDI)) an eye-level dosimeter would record. Simulation accuracy is expected to depend strongly on two factors: how well an individual's viewing direction can be inferred, given that a pocket-worn phone cannot observe where a person is looking, and how well obstructions within the corresponding field of view are captured by the available geospatial data.

Validating this approach requires a large dataset spanning diverse geographical locations, since accuracy is expected to vary with the surrounding built

and natural environment, and with sky conditions. Designing such a campaign requires practical experience with the full measurement-simulation chain. We therefore first conduct a proof-of-concept experiment, the protocol of which is described in this abstract, in which CAMS-based estimates are compared against measurements from a dosimeter worn at eye level, alongside reference measurements of the wearer's viewing direction and the obstructions within their field of view. The dataset is deliberately small: the aim is to establish whether the approach produces estimates worth evaluating at scale, and to surface methodological issues. Its outcomes will guide the design of a larger, geographically distributed campaign through which the validity of the approach can be assessed.

2. Experiment Protocol

A single participant walks a route at two sites in Eindhoven, the Netherlands (51.44° N, 5.47° E) during Autumn: an open field surrounded by trees, and a dense urban environment. Measurements are conducted under clear, mixed, and overcast skies, with date and time selected from weather forecasts and actual sky conditions classified by a weather station within 2 km of both sites. The participant carries a pocket-worn smartphone (Motorola Moto G05) logging location coordinates over time via global navigation satellite systems (GNSS). Ground-truth data on approximate viewing direction, obstructions, and the resulting PLE are collected by devices placed on a helmet worn by the participant (Fig. 1a).

As a proxy for viewing direction, the head's azimuth angle is recorded by a digital compass (HP215 IMU motion data logger) mounted on top of the helmet. These recordings are compared against the bearing (movement over ground) derived from the logged coordinates to quantify the directional error.

The logged coordinates are pre-processed before bearing estimation (e.g. [19, 20]), and we additionally consider whether the smartphone's inertial and magnetometer sensors can improve the viewing direction estimate [21, 22]. Periods without movement require a separate treatment, since bearing is undefined; we compare retaining the last known bearing against replacing the viewing direction with an azimuthal average.

Obstructions are captured by a fisheye camera (12.3 MP Sony IMX477 sensor with an Arducam 180° fisheye M12 lens) mounted on top of the helmet, imaging the sky dome. In each image the obstructed fraction of the sky dome is identified by segmenting sky from non-sky, for which several established approaches are available [23-27]. These fractions are then compared to obstruction estimates based on geospatial data sources: GNSS signal quality [28, 29], Google Street View (GSV) imagery [30-32], or 3D geometric models [33, 34]. These yield implementations of varying spatial fidelity, from a 3D geometric model included directly in the simulations, to a sky-dome mask derived from GSV data (an approach also seen in sunlight exposure modelling [35]), to a single scalar attenuation factor from GNSS data (Fig. 1b).

Finally, we investigate how uncertainty in viewing direction and reconstruction of obstructions propagates into CAMS-based PLE estimates. Ground-truth PLE (photopic illuminance and m-EDI), the outcome variable to be simulated, is measured by a dosimeter (ActLumus, Condor Instruments) at the front of the helmet, just above the eyes. Combining the different options for approximating viewing direction and obstructions yields a set of models of varying complexity. Comparing each model's output against the reference measurements allows us to investigate which input combinations produce estimates that warrant evaluation at scale.




Fig. 1. Measurement setup: (a) Equipment worn by the participant, (b) Simulation of (eye-level) personal light exposure (PLE) using a Copernicus Atmosphere Monitoring Service (CAMS) based sun and sky model, combined with estimates of surrounding obstructions (through an attenuation factor, sky-dome mask, or 3D geometric model) and viewing direction.

3. Conclusions

Quantifying PLE with dosimeters is burdensome for participants and researchers, thereby constraining the scale and geographic spread of the resulting

datasets. The approach outlined here instead derives outdoor eye-level PLE from CAMS data using CBDM and geospatial data, with a pocket-worn smartphone as the only subject-side input. The experiment protocol described in this abstract does not establish the

accuracy of that approach: with a single participant, two sites, and a single season, the dataset is small by design and cannot support such a claim. What it does provide is a first end-to-end test of the full measurement-simulation chain and a direct comparison of models differing in how viewing direction and obstructions are represented. Its outcomes will indicate which input combinations warrant evaluation at scale, and will inform the design of the geographically distributed campaign required for validation.

References

- [1]. D.-J. Dijk, S. N. Archer, Light, sleep, and circadian rhythms: Together again, *PLOS Biology*, Vol. 7, Issue 6, 2009, e1000145.
- [2]. R. J. Lucas, et al., Measuring and using light in the melanopsin age, *Trends in Neurosciences*, Vol. 37, Issue 1, 2014, pp. 1-9.
- [3]. M. G. Figueiro, et al., Circadian-effective light and its impact on alertness in office workers, *Lighting Research & Technology*, Vol. 51, Issue 2, 2019, pp. 171-183.
- [4]. A. Didikoglu, et al., Associations between light exposure and sleep timing and sleepiness while awake in a sample of UK adults in everyday life, *Proceedings of the National Academy of Sciences*, Vol. 120, Issue 42, 2023, e2301608120.
- [5]. R. Ashby, et al., IMI: The role of light in refractive development and myopia: Evidence from animal and human studies, *Investigative Ophthalmology & Visual Science*, Vol. 66, Issue 15, 2025, 5.
- [6]. L. Wen, et al., Objectively measured near work, outdoor exposure and myopia in children, *British Journal of Ophthalmology*, Vol. 104, Issue 11, 2020, pp. 1542-1547.
- [7]. J. van Duijnhoven, et al., Measuring light exposure in daily life: A review of wearable light loggers, *Building and Environment*, Vol. 274, 2025, 112771.
- [8]. S. W. de Vries, M. Gkaintatzi-Masouti, J. van Duijnhoven, J. Mardaljevic, et al., Recommendations for light-dosimetry field studies based on a meta-analysis of personal light levels of office workers, *Lighting Research & Technology*, Vol. 57, Issue 1, 2025, pp. 47-70.
- [9]. E. Balajadia, et al., Usability and acceptability of a corneal-plane α -opic light logger in a 24-h field trial, *Digital Biomarkers*, Vol. 7, Issue 1, 2023, pp. 139-149.
- [10]. O. Stefani, et al., Wearable light loggers in field conditions: Corneal light characteristics, user compliance, and acceptance, *Clocks & Sleep*, Vol. 6, Issue 4, 2024, pp. 619-634.
- [11]. J. Zauner, A. Biller, M. Spitschan, Impact of light logger and dosimeter placement on wearability and appeal in real-life settings, *Open Research Europe*, Vol. 5, 2025, 293.
- [12]. S. W. de Vries, J. Mardaljevic, J. van Duijnhoven, Impact of wear position on dosimeter performance: Measurement validity under simulated indoor illumination, *npj Biological Timing and Sleep*, Vol. 3, 2026, 19.
- [13]. S. W. de Vries, J. Mardaljevic, J. van Duijnhoven, Impact of wear position on dosimeter performance: A hybrid measurement-simulation approach to quantify in-situ factors, *npj Biological Timing and Sleep*, Vol. 3, 2026, 20.
- [14]. J. Zauner, S. W. de Vries, A. Didikoglu, J. van Duijnhoven, et al., Sensor placement causes outcome-dependent bias in ambulatory light-exposure estimates, *bioRxiv*, 2026, 10.1101/2026.01.21.634125.
- [15]. V.-H. Peuch, et al., The Copernicus atmosphere monitoring service: From research to operations, *Bulletin of the American Meteorological Society*, Vol. 103, Issue 12, 2022, pp. E2650-E2668.
- [16]. J. Mardaljevic, E. Brembilla, M. Eames, Daylight solar radiation AMY data derived from satellite remote sensing: Validation against ground measurements and comparison with TMYs, *Building Services Engineering Research & Technology*, Vol. 46, Issue 5, 2025, pp. 653-691.
- [17]. J. Mardaljevic, Simulation of annual daylighting profiles for internal illuminance, *Lighting Research & Technology*, Vol. 32, Issue 3, 2000, pp. 111-118.
- [18]. E. Brembilla, J. Mardaljevic, Climate-based daylight modelling for compliance verification: Benchmarking multiple state-of-the-art methods, *Building and Environment*, Vol. 158, 2019, pp. 151-164.
- [19]. F. Chazal, D. Chen, L. Guibas, X. Jiang, et al., Data-driven trajectory smoothing, in *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*, 2011, pp. 251-260.
- [20]. J. J. Early, A. M. Sykulski, Smoothing and interpolating noisy GPS data with smoothing splines, *Journal of Atmospheric and Oceanic Technology*, Vol. 37, Issue 3, 2020, pp. 449-465.
- [21]. Z.-A. Deng, G. Wang, Y. Hu, D. Wu, Heading estimation for indoor pedestrian navigation using a smartphone in the pocket, *Sensors*, Vol. 15, Issue 9, 2015, pp. 21518-21536.
- [22]. J. Perul, V. Renaudin, Head: Smooth estimation of walking direction with a handheld device embedding inertial, GNSS, and magnetometer sensors, *NAVIGATION: Journal of the Institute of Navigation*, Vol. 67, Issue 4, 2020, pp. 713-726.
- [23]. C. Grimmond, S. Potter, H. Zutter, C. Souch, Rapid methods to estimate sky-view factors applied to urban areas, *International Journal of Climatology*, Vol. 21, Issue 8, 2001, pp. 903-913.
- [24]. L. Chapman, J. Thornes, A. Bradley, Rapid determination of canyon geometry parameters for use in surface radiation budgets, *Theoretical and Applied Climatology*, Vol. 69, Issue 1-2, 2001, pp. 81-89.
- [25]. J. Luo, S. Etz, A physical model-based approach to detecting sky in photographic images, *IEEE Transactions on Image Processing*, Vol. 11, Issue 3, 2002, pp. 201-212.
- [26]. U. M. Moin, J.-I. Tsutsumi, Rapid estimation of sky view factor and its application to human environment, *Journal of the Human-Environment System*, Vol. 7, Issue 2, 2004, pp. 83-87.
- [27]. K. A. Nice, et al., Sky pixel detection in outdoor imagery using an adaptive algorithm and machine learning, *Urban Climate*, Vol. 31, 2020, 100572.
- [28]. L. Chapman, J. E. Thornes, A. V. Bradley, Sky-view factor approximation using GPS receivers, *International Journal of Climatology*, Vol. 22, Issue 5, 2002, pp. 615-621.
- [29]. L. Chapman, J. E. Thornes, Real-time sky-view factor calculation and approximation, *Journal of Atmospheric*

- and *Oceanic Technology*, Vol. 21, Issue 5, 2004, pp. 730-741.
- [30]. R. Carrasco-Hernandez, A. R. D. Smedley, A. R. Webb, Using urban canyon geometries obtained from Google Street View for atmospheric studies: Potential applications in the calculation of street level total shortwave irradiances, *Energy and Buildings*, Vol. 86, 2015, pp. 340-348.
- [31]. L. Zeng, J. Lu, W. Li, Y. Li, A fast approach for large-scale sky view factor estimation using street view images, *Building and Environment*, Vol. 135, 2018, pp. 74-84.
- [32]. F.-Y. Gong, Z.-C. Zeng, E. Ng, L. K. Norford, Spatiotemporal patterns of street-level solar radiation estimated using Google Street View in a high-density urban environment, *Building and Environment*, Vol. 148, 2019, pp. 547-566.
- [33]. W. Sirko, et al., Continental-scale building detection from high resolution satellite imagery, *arXiv*, 2021, arXiv:2107.12283.
- [34]. R. Peters, et al., Automated 3D reconstruction of LoD2 and LoD1 models for all 10 million buildings of the Netherlands, *Photogrammetric Engineering & Remote Sensing*, Vol. 88, Issue 3, 2022, pp. 165-170.
- [35]. X. Li, Y. Yoshimura, W. Tu, C. Ratti, A pedestrian-level strategy to minimize outdoor sunlight exposure, in *Artificial Intelligence, Machine Learning, and Optimization Tools for Smart Cities* (P. M. Pardalos, et al., Eds.), *Springer*, Cham, 2022, pp. 123-134.

Investigation of the Flow Rate Dependence of Blood Coagulation Measurement in an Extracorporeal Circulation Circuit Using LED Photoacoustic Imaging

Takahiro Wabe¹, Naoki Katsura², Yohsuke Uchida³, Ryo Suzuki⁴ and Akimitsu Fujii⁵

¹ Department of Medical Engineering, Faculty of Health, Junshin Gakuen University,
1-1-1 Chikushigaoka, Minami-ku, Fukuoka-shi, Fukuoka, Japan

² Department of Medical Science, Faculty of Health and Medical Science, Teikyo Heisei University, Japan

³ Department of Clinical Engineering, School of Health Sciences, Tokyo University of Technology, Japan

⁴ Faculty of Pharmaceutical Sciences, Teikyo University, Tokyo, Japan

⁵ Tokyo University of Technology Graduate School of Medical Technology, Tokyo, Japan

Tel.: + 81925541255, fax: + 81925522707

E-mail: wabe.t@junshin-u.ac.jp

Summary: Blood coagulation during extracorporeal circulation, including hemodialysis and ECMO, presents critical clinical challenges. Traditional pressure sensor monitoring remains reactive and fails to identify early thrombus formation. Building upon our previously reported LED-based optoacoustic imaging system, this paper advances the technology toward clinical implementation. We defined and met the operational parameters necessary for bedside deployment, evaluating system robustness under various simulated clinical conditions. The proposed system non-invasively images micro-clots and localized structural changes by capturing photoacoustic waves resulting from optical absorption in blood. Validation tests in simulated extracorporeal loops confirmed the reliable detection of incipient coagulation before complete circuit blockage. This compact, cost-effective, and safe LED-based platform enables continuous bedside monitoring. Consequently, optoacoustic imaging provides a viable paradigm shift from reactive detection to proactive prevention, substantially mitigating clinical risks in extracorporeal therapies.

Keywords: Blood clotting, Photoacoustic imaging, LED, Circuit in an extracorporeal device, Temporal observation, Predictive maintenance, Flow rate.

1. Introduction

1.1. Blood Coagulation and Clinical Challenges in Extracorporeal Circulation Therapy

Extracorporeal circulation therapies – including extracorporeal membrane oxygenation (ECMO) and continuous blood purification – are indispensable life-support modalities for severe respiratory failure, circulatory collapse, acute kidney injury, and refractory hyperinflammatory conditions such as severe COVID-19. In ECMO applications, these systems temporarily replace pulmonary and cardiac functions, providing a critical bridge to organ recovery. However, these life-saving interventions fundamentally rely on diverting patient blood through artificial extracorporeal circuits.

A primary clinical hurdle in managing these circuits is controlling and detecting intra-circuit blood coagulation. Contact between blood components and artificial inner surfaces rapidly triggers platelet activation and the intrinsic coagulation cascade. In severe infections, systemic inflammatory responses markedly exacerbate thrombogenicity. Consequently, intra-circuit clot formation occurs frequently; these thrombi risk complete circuit occlusion and threaten fatal systemic thromboembolic complications if released into patient circulation.

To mitigate these risks, clinical protocols combine bio-functional surface modifications, such as heparin coatings, with systemic continuous anticoagulation. Despite these measures, completely preventing thrombosis remains challenging. Anticoagulants themselves can elicit foreign-body reactions, and dosing must be meticulously balanced against the patient's elevated bleeding risk. Currently, clinical monitoring for intra-circuit coagulation relies heavily on hydraulic pressure sensors. However, a detectable pressure spike usually indicates advanced thrombus accumulation or near-total occlusion. Such retrospective detection mandates emergency circuit replacement, discarding the blood volume within the circuit, and temporary cessation of therapy. This process imposes a severe operational burden on medical staff and heightens patient risks, particularly regarding allogeneic blood transfusions. Therefore, establishing a highly sensitive, real-time "predictive maintenance" platform capable of detecting pre-thrombotic states before irreversible occlusion occurs is an urgent clinical priority [1-5].

1.2. Feasibility and Clinical Significance of Photoacoustic Monitoring under Dynamic Flow Conditions

To overcome these diagnostic limitations, photoacoustic (PA) imaging has emerged as a

promising modality for non-invasively monitoring hematological dynamics with high optical contrast. PA imaging leverages the photoacoustic effect, where pulsed optical irradiation induces rapid thermoelastic expansion, generating transient acoustic waves detectable by ultrasonic transducers. For clinical translation into bedside platforms, our group developed a portable, cost-effective PA monitoring system utilizing light-emitting diodes (LEDs) as illumination sources, replacing bulky, high-cost pulsed lasers. Preliminary micro-channel experiments demonstrated the feasibility of this system for non-destructively characterizing progressive blood coagulation. Within clinical circuits, the drip chamber represents the primary measurement window due to its propensity for localized flow stagnation.

Crucially, bridging the vast divide between controlled laboratory experiments and actual bedside implementation requires proving that the system functions robustly within the volatile environment of real clinical practice. In hospital intensive care units, systemic blood flow velocities are rarely static; they fluctuate dynamically due to patient hemodynamic instability, therapeutic flow adjustments, and changing circuit resistance. Flow velocity variations exert complex physical impacts on PA signals. Mechanistically, altered shear rates dynamically shift red blood cell orientation, aggregation kinetics, and local hematocrit distributions. If these flow-induced acoustic artifacts obscure signals driven by actual clot formation, the technology remains a mere laboratory concept, strictly unusable at the bedside.

Consequently, rigorously validating system stability and signal robustness against severe flow fluctuations is not merely an incremental technical test, but the absolute, paramount prerequisite for establishing true clinical utility. Demonstrating that our LED-based PA system can isolate early coagulation signatures amidst complex flow dynamics proves that this technology transcends experimental feasibility and is directly ready for practical deployment in hospital environments. This study systematically evaluates the system in a clinical mock loop, establishing a vital academic and practical foundation for real-time, proactive safety monitoring in actual hospital care [6-13].

2. Method

2.1. Construction of Experimental Apparatus and Extracorporeal Circulation Circuit

To replicate clinical environments, a simulated extracorporeal circulation loop was constructed using medical-grade tubing, an air trap chamber (drip chamber), and a clinical roller pump. The circulating fluid temperature was precisely maintained at 37.0 ± 0.5 °C using a thermostatic water bath to mimic physiological conditions during therapies such as continuous renal replacement therapy (CRRT) and

hemodialysis. The photoacoustic (PA) imaging system comprised a commercial ultrasound/photoacoustic unit (AcousticX, CYBERDYNE Inc.) integrated with dual high-output LED arrays as the optical excitation source. The LED arrays emitted near-infrared light at a peak wavelength of 850 nm with an energy density of approximately 200 μ J/pulse. This wavelength was selected due to the high optical absorption coefficients of both oxygenated (HbO_2) and deoxygenated hemoglobin (Hb), which serve as the dominant primary endogenous chromophores in whole blood. A linear array ultrasound transducer was utilized to detect thermoelastically generated acoustic waves. Based on prior frequency optimization studies comparing 10 MHz and 7 MHz center frequencies, a 7 MHz transducer was selected. Although 10 MHz transducers offer superior axial resolution, the 7 MHz configuration yields approximately a 4- to 5-fold increase in acoustic sensitivity at deeper tissue depths, ensuring adequate signal-to-noise ratio (SNR) across the wall thickness and lumen of the air trap chamber. The target measurement site was situated at the drip chamber – a critical circuit component structurally susceptible to micro-flow stagnation and thrombus initiation. To guarantee measurement reproducibility and eliminate motion artifacts induced by manual operation, the LED arrays and the 7 MHz transducer were rigidly mounted on a custom 3D-printed fixture. The illumination-to-reception path angle was set to 40°, maintaining a constant acoustic coupling distance of 1.0 cm between the transducer face and the outer wall of the drip chamber.

2.2. Blood Sample Preparation and Coagulation Induction Protocol

To minimize biological variability, adhere to biohazard safety standards, and facilitate high-throughput testing, sterile sheep whole blood (Kohjin Bio Co., Ltd.; Model No. 12070210) was selected as the substitute biological matrix. The stored blood was preserved in Alsever's solution, an anticoagulant that chelates free calcium ions to prevent spontaneous thrombosis during transportation. To reactivate the intrinsic and extrinsic coagulation cascades without altering native coagulation factors or platelet membrane phospholipids, calcium ions were replenished. Calcium gluconate (solubility limit: 3.3 g/100 mL in aqueous solution) was employed as the recalcification agent. Specifically, 1.1 g of calcium gluconate was completely dissolved in 33 mL of 0.9 % physiological saline. For each experimental trial, 33 mL of this reagent was rapidly bolus-injected into 50 mL of circulating sheep blood within the reservoir, initiating a uniform, reproducible coagulation cascade. PA signal acquisition commenced immediately upon injection and continued continuously for 10 minutes, or until the roller pump automatically tripped due to safety threshold limits triggered by elevated venous circuit pressure resulting from downstream occlusion.

2.3. Selection of Operating Conditions Based on Clinical Settings

To quantitatively systematically evaluate the impact of hydrodynamic forces on PA signal evolution, operational flow rates were set based on clinical hemodialysis protocols: Low flow condition: 100 mL/min; Moderate flow condition: 200 mL/min; High flow condition: 300 mL/min. These dynamic conditions were systematically cross-evaluated against historical baseline data acquired at 60 mL/min (the minimal operational boundary in continuous therapies) to establish a comprehensive flow-dependent acoustic response matrix.

2.4. Criteria for Circuit Coagulation Assessment

To evaluate the predictive sensitivity of PA imaging relative to conventional diagnostics, three traditional physical indicators were continuously recorded as reference benchmarks: Drip chamber fluid level drop: Discontinuous liquid level shifts caused by regional flow resistance changes or fibrin mesh accumulation. Venous pressure variations: Deviations recorded by in-line piezoresistive pressure transducers. Microbubble formation: Cavitation or microbubble generation adjacent to the roller pump head caused by localized negative pressure during flow impairment. The temporal latency between these macro-physical manifestations and the onset of micro-structural PA signal alterations was rigorously quantified.

2.5. Acquisition and Multidimensional

Signal Processing of PA B-mode Images Rather than relying on single-element A-mode signals, this study evaluated two-dimensional (2D) PA B-mode grayscale images to retain spatial acoustic density distributions. Raw radiofrequency (RF) acoustic signals generated by AcousticX were reconstructed into 8-bit grayscale pixel matrices representing local acoustic energy densities. Image sequences were pre-analyzed using ImageJ (National Institutes of Health) to assess spatial brightness profiles over time. Temporal differential matrix analysis and spatial variance filtering were applied across all pixels from 10 min. This processing pipeline automatically delineated Regions of Interest (ROIs) within the drip chamber corresponding to initial fibrin cross-linking and erythrocyte aggregation, removing subjective operator bias. The mean pixel intensity, rate of intensity change, and spatial variance within the identified ROIs were tracked across all flow velocity regimes (100, 200, and 300 mL/min) and merged with 60 mL/min baseline dataset to characterize the acoustic scattering dynamics of blood under varying shear conditions.

3. Results

3.1. A-Mode Acoustic Signal Characteristics and Hydrodynamic Circuit Responses

To quantitatively characterize photoacoustic (PA) waveform evolution during progressive blood coagulation, acoustic radiofrequency (RF) signals were captured in real-time by placing the ultrasound transducer against the outer lumen of the drip chamber. Time zero (0 min) was defined as the moment when 1.1 g of calcium gluconate in 33 mL of physiological saline was rapidly bolus-injected to reactivate the coagulation cascade.

Across all evaluated flow rate regimes (100, 200, and 300 mL/min), the overall hemodynamic progression exhibited strong qualitative consistency with baseline observations previously acquired at 60 mL/min. During the initial phase (0 min), circulating blood retained fluid mobility with stable baseline acoustic profiles. However, at approximately 10 min post-injection, a distinct coagulation transition manifested, marked by exponential increases in blood viscosity and localized micro-thrombus aggregation. As fibrin cross-linking intensified, circuit hydraulic resistance escalated sharply, ultimately causing total cessation of blood flow across all operational flow rates by 15 min.

Throughout the terminal phase leading to flow arrest, the circuit experienced severe hydraulic destabilization. Occlusion downstream induced a profound negative pressure state upstream, particularly around the roller pump head. This condition precipitated three distinct macro-physical criteria confirming irreversible thrombosis and total loss of perfusion function:

- 1) Drip chamber fluid level drop: A vertical drop of approximately 1.8 cm in the blood-air interface relative to baseline;
- 2) Circuit pressure collapse: A drastic shift in venous line pressure from an initial operating baseline of +60 mmHg down to extreme negative pressures ranging between -1 and -10 mmHg;
- 3) Microbubble nucleation: Widespread cavitation and microbubble accumulation visual within both the drip chamber and flexible PVC tubing due to localized negative pressure gradients.

Fig. 1 illustrates representative PA A-mode RF waveforms recorded at 10 min post-injection – prior to complete flow arrest – for flow rates of (a) 60 mL/min, (b) 100 mL/min, (c) 200 mL/min, and (d) 300 mL/min. The vertical axis denotes elastic wave signal amplitude (PA intensity), whereas the horizontal axis corresponds to acoustic propagation distance from the transducer face. High-amplitude local peaks indicate regions of heightened optical absorption corresponding to dense fibrin-erythrocyte aggregates or structural boundary reflections. Spatially, the left region of each profile represents the transducer-chamber interface, while the right region corresponds to the internal luminal blood phase.




Fig. 1. Measurement of photoacoustic phenomenon using A mode ultrasonic imaging.

3.2. B-Mode Tomographic Evaluation and Multidimensional Brightness Mapping

(1) Interpretation of B-Mode Tomographic Images

Cross-sectional B-mode tomographic reconstructions of the drip chamber interior are presented in Fig. 2. Spatially, the acoustic boundary on the left corresponds to the transducer contact wall, whereas the right region reflects the blood-filled chamber cavity. High-brightness (hyperechoic/high PA amplitude, rendered in white) regions delineate localized thrombi where high optical absorption by accumulated red blood cells triggered robust thermoelastic expansion. Conversely, dark or

low-intensity regions (black) indicate non-coagulated, homogeneous blood or background acoustic noise below the detection threshold.




Fig. 2. B mode image by photoacoustic imaging.

(2) 3D Brightness Mapping Analysis via ImageJ

To eliminate subjective visual bias and objectively quantify spatiotemporal variations in PA intensity, two-dimensional B-mode grayscale matrices were transformed into three-dimensional surface intensity maps using ImageJ (National Institutes of Health) (Fig. 3). The dashed black line marks the physical boundary of the drip chamber wall, dividing the transducer position (left) from the blood phase (right). The spatial axes were defined as follows: the x-axis represents horizontal depth into the blood lumen (mm), the y-axis represents vertical spatial distribution (mm), and the z-axis represents 8-bit signal brightness values ranging from 0 to 255 gradations.

Signal intensity was color-coded using a false-color thermal scale: a z-value of 0 (minimum PA intensity) is depicted in black, transitioning continuously through orange and yellow as signal amplitude increases, and peaking at white for maximum intensity ($z = 255$). High z-axis peaks in the color map physically signify localized zones of high-density erythrocyte entrapped within fibrin networks, confirming the progressive spatial growth of thrombi.

3.3. Synthesis of Flow-Rate Independence and Diagnostic Feasibility

These empirical findings confirm that the proposed LED-based PA imaging system effectively tracks blood coagulation dynamics even under dynamic flow fluctuations typical of clinical extracorporeal therapies. Across all tested flow velocity regimes

(60–300 mL/min), PA wave amplitudes and spatial profiles exhibited consistent, reproducible shifts starting at $t = 10$ min, propagating radially from the chamber wall inward. By decoupling flow-induced acoustic artifacts from true coagulation signatures, this system achieves highly sensitive, early-stage thrombus detection prior to irreversible pressure drops. Consequently, this technology establishes a robust diagnostic foundation for transition from reactive, pressure-based alarm systems to proactive, real-time predictive monitoring in clinical perfusion.




Fig. 3. 3D diagram of luminance part by ImageJ.

4. Conclusions

This study conclusively demonstrates that the portable, light-emitting diode (LED)-based photoacoustic (PA) imaging system provides a highly sensitive, non-invasive, and robust real-time predictive monitoring platform for intra-circuit blood coagulation in extracorporeal circulation therapies. By replacing bulky, expensive pulsed lasers with dual high-output near-infrared LED arrays emitting at a peak wavelength of 850 nm integrated with a 7 MHz linear ultrasound transducer, the developed hardware architecture overcomes traditional bedside implementation barriers. This setup achieves an optimal balance between optical absorption sensitivity and acoustic penetration depth, establishing a clinically viable and cost-effective solution for continuous perfusion monitoring at the drip chamber.

The experimental evaluation within a clinical mock circulation loop under standardized thermal (37.0 ± 0.5 °C) and recalcified coagulation conditions establishes three fundamental scientific and clinical breakthroughs:

(1) Unquestionable Diagnostic Latency Superiority over Conventional Physical Indicators: Across all operational conditions, the LED-PA

platform reliably detects micro-structural coagulation transitions at $t=10$ min post-recalcification – a full 5 minutes prior to total hydraulic occlusion and flow arrest ($t = 15$ min). Crucially, this acoustic manifestation precedes all macro-physical failure metrics currently used in clinical intensive care units. Specifically, local PA amplitude escalation occurs well before the onset of a 1.8 cm drip chamber fluid level drop, severe venous line pressure collapse (dropping from an operational baseline of +60 mmHg down to extreme negative states between -1 and -10 mmHg), or pump-head microbubble nucleation caused by localized negative pressure gradients. This 5-minute diagnostic window proves that photoacoustic monitoring successfully eliminates the dangerous time lag inherent in reactive, pressure-sensor-based alarm systems.

(2) Complete Flow-Rate Independence and Hydrodynamic Artifact Immunity: A major technical hurdle preventing the bedside deployment of optical and acoustic sensors has been signal interference from dynamic blood flow fluctuations. Varying shear rates alter red blood cell orientation, aggregation kinetics, and local hematocrit distributions, creating potential acoustic scattering artifacts. Our cross-evaluation across clinical flow velocity regimes (100, 200, and 300 mL/min) alongside low-flow baseline boundaries (60 mL/min) establishes that the temporal evolution ($t = 10$ min onset), spatial radial propagation patterns, and amplitude growth of PA signals remain completely consistent and flow-independent. The system successfully decouples hydrodynamic shear noise from true biochemical clot formation, proving that volatile patient hemodynamics or therapeutic flow adjustments will not compromise diagnostic reliability in actual intensive care environments.

(3) Objective Spatial Characterization of Early Thrombus Aggregation: By advancing from single-element A-mode radiofrequency waveforms to two-dimensional (2D) B-mode tomographic reconstructions and three-dimensional (3D) surface brightness intensity mapping via ImageJ, this research provides a fully quantitative framework for tracking spatial thrombus growth. Automated temporal differential matrix analysis and spatial variance filtering successfully isolate hyperechoic, high-brightness Regions of Interest (ROIs) within localized micro-flow stagnation zones inside the drip chamber. Signal intensity escalation to peak 8-bit saturation values ($z = 255$ gradations) directly tracks the physical accumulation of dense erythrocyte-fibrin meshworks, eliminating operator subjectivity and establishing a repeatable metric for evaluating pre-thrombotic states.

The physical mechanism underlying this high diagnostic precision relies on the strong optical absorption coefficients of both oxygenated (HbO_2) and deoxygenated (Hb) hemoglobin at 850 nm. As the intrinsic and extrinsic coagulation cascades progress, the phase transition of whole blood into a cross-linked fibrin matrix densely traps circulating red blood cells

into localized clusters. Under pulsed LED irradiation, these concentrated hemoglobin aggregates undergo rapid thermoelastic expansion, generating high-amplitude acoustic waves that stand out against homogeneous background blood.

From a clinical perspective, these results mark a critical shift in extracorporeal life support management. Conventional clinical practice operates reactively: intervention occurs only after pressure spikes signify advanced thrombus accumulation or total circuit occlusion. Emergency circuit replacements forced by retrospective detection require discarding the blood volume within the synthetic tubing, interrupting vital organ support, increasing allogeneic blood transfusion risks, and imposing severe operational burdens on medical personnel. In contrast, the LED-PA platform enables a proactive "predictive maintenance" paradigm. By identifying pre-thrombotic states 5 minutes before physical flow impairment, clinicians can adjust continuous anticoagulation dosages, execute fluid flushes, or perform elective circuit swaps long before catastrophic failure occurs.

While this study establishes baseline diagnostic performance using a standardized sheep whole blood model, future work will advance this platform toward direct clinical deployment. Integrating deep-learning computer vision pipelines directly into the AcousticX processing framework will enable automated clot volume estimation and automated risk scoring. Furthermore, validating signal stability in clinical trials involving patient blood subjected to complex inflammatory responses, circulating cytokines, and pharmacological anticoagulants will complete the transition to bedside care.

In summary, this research proves that LED-based photoacoustic imaging is a mature, robust, and clinically deployable diagnostic technology. By establishing complete signal stability under severe flow fluctuations and offering unmatched predictive sensitivity over conventional hydraulic pressure sensors, this technology establishes the foundation for a new generation of proactive extracorporeal life-support systems that maximize therapeutic safety and elevate patient care worldwide.

Acknowledgments

T. Wabe and Y. Uchida of Teikyo Univ. Of Sci. would like to thank Mr. N. Sato of CYBERDYNE, INC Research and Development Dept, for his cooperation in this experiment.

References

- [1]. N. Tang, D. Li, X. Wang, Z. Sun, Abnormal coagulation parameters are associated with poor prognosis in patients with novel coronavirus pneumonia, *Journal of Thrombosis and Haemostasis*, Vol. 18, Issue 4, 2020, pp. 844-847.
- [2]. P. Mehta, D. F. McAuley, M. Brown, E. Sanchez, et al., COVID-19: Consider cytokine storm syndromes and immunosuppression, *The Lancet*, Vol. 395, Issue 10229, 2020, pp. 1033-1034.
- [3]. J. Utley, Pathophysiology of cardiopulmonary bypass: Current issues, *Journal of Cardiac Surgery*, Vol. 5, Issue 3, 1990, pp. 177-189.
- [4]. Y. Mori, Blood compatible material, *Kobunshi Ronbunshu*, Vol. 42, Issue 10, 1985, pp. 601-615 (in Japanese).
- [5]. L. Gott, J. D. Whiffen, R. C. Dutton, Heparin bonding on colloidal graphite surfaces, *Science*, Vol. 142, Issue 3597, 1963, pp. 1297-1298.
- [6]. T. Wabe, R. Suzuki, K. Maruyama, Y. Uchida, Possibility for temporal observation of thrombus generated in extracorporeal circulator circuit by photoacoustic imaging using LED, in *Proceedings of the 5th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2019, pp. 157-160.
- [7]. T. Wabe, R. Suzuki, K. Maruyama, Y. Uchida, The measurement of blood coagulation process in extracorporeal circuit using LED photoacoustic imaging, *Sensors & Transducers*, Vol. 237, Issue 9-10, 2019, pp. 88-94.
- [8]. T. Wabe, R. Suzuki, K. Maruyama, Y. Uchida, Possibility for a feature value extraction method using photoacoustic imaging to evaluate predictive maintenance on blood clotting in an extracorporeal circuit, in *Proceedings of the 6th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2020, pp. 70-74.
- [9]. T. Wabe, R. Suzuki, K. Maruyama, Y. Uchida, Possibility of extracting feature value from the changes in brightness over time of blood coagulation in the extracorporeal circuit, *Sensors & Transducers*, Vol. 246, Issue 7, 2020, pp. 64-70.
- [10]. T. Wabe, R. Suzuki, A. Fujii, Y. Uchida, et al., Basic study on blood coagulation measurement in an extracorporeal circulation circuit by LED photoacoustic method using an extracorporeal circulation device, in *Proceedings of the 7th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2021, pp. 175-178.
- [11]. T. Wabe, R. Suzuki, A. Fujii, Y. Uchida, et al., Basic research on blood coagulation measurement of extracorporeal circulation circuit using photoacoustic imaging of LED light source, *Sensors & Transducers*, Vol. 253, Issue 6, 2021, pp. 1-8.
- [12]. T. Wabe, R. Suzuki, A. Fujii, Y. Uchida, et al., Selection of measurement site for predictive maintenance of blood coagulation in an extracorporeal circulation circuit using LED photoacoustic imaging and an extracorporeal circulation device, in *Proceedings of the 8th International Conference on Sensors and Electronic Instrumentation Advances (SEIA)*, 2022, pp. 55-59.
- [13]. T. Wabe, R. Suzuki, A. Fujii, Y. Uchida, et al., Measurement site selection and validation in extracorporeal circuit for predictive maintenance of blood coagulation using photoacoustic imaging of LED light source, *Sensors & Transducers*, Vol. 259, Issue 5, 2022, pp. 45-51.

Long-Term Objective Evaluation of Posture during Self-Study Using Image Analysis Technology and Impact on Learning Effectiveness

Takahiro Wabe¹ and Yasutaka Uchida²

¹ Department of Medical Engineering, Faculty of Health, Junshin Gakuen University, 1-1-1 Chikushigaoka, Minami-ku, Fukuoka-shi, Fukuoka, Japan

² Teikyo University of Science, Tokyo, Japan

Tel.: + 81925541255, fax: + 81925522707

E-mail: wabe.t@junshin-u.ac.jp

Summary: This study aimed to quantitatively demonstrate the importance of study posture. We measured and analyzed the self-study posture of medical national examination candidates over a year using motion capture and image analysis. Based on three-dimensional coordinates of the head, trunk, and lower limbs, we determined that the group that consistently maintained good posture consistently achieved high levels of performance on mock examinations. Conversely, the group with a high frequency of poor posture remained at the middle to lower levels, and this trend was consistent with the final pass/fail results. Therefore, this study suggests a close correlation between study posture and national examination performance, possibly through maintaining concentration and reducing physical fatigue.

Keywords: Learning posture, National examination mock exam scores, Motion capture, Image analysis, Longitudinal behavioral analysis.

1. Introduction

In educational settings, posture instruction has historically been framed around etiquette and physical health. However, recent advances in neuroscience and behavioral science demonstrate that posture substantially influences learning motivation, focus, and higher-order cognitive functions. These insights suggest that maintaining proper posture directly enhances learning efficiency, sparking renewed interest across educational philosophy, psychology, and physiology.

Physiologically, postural shifts modulate cerebral cortical arousal via the brainstem reticular formation. Sustaining an upright posture relies on deep trunk musculature – such as the transversus abdominis and multifidus muscles – which stabilizes the pelvis and spine. This muscular engagement mitigates physical fatigue and prevents cognitive decline during prolonged study. Indeed, empirical evidence indicates that students with superior trunk muscle strength sustain concentration longer during extended academic tasks [1-5].

Conversely, insufficient trunk muscle strength induces postural collapse, causing localized musculoskeletal discomfort (e.g., neck and low-back pain) and mental fatigue. This decline in physical integrity directly compromises attentional capacity. Consequently, optimizing learning outcomes requires an integrated assessment that incorporates musculoskeletal mechanics alongside cognitive and psychological dimensions. Recent technological developments further support this approach through information and communication technology (ICT) and artificial intelligence (AI) [6]. Advanced systems now estimate real-time engagement by tracking posture

dynamics and behavioral logs, demonstrating strong correlations between posture variations and task accuracy. Furthermore, computer vision algorithms analyzing webcam footage can automatically detect head tilt and upper-body slouching to infer inattention or drowsiness. Such objective, automated metrics transcend subjective evaluation, unlocking new avenues for personalized learning interventions.

Despite these advancements, empirical studies directly linking objective biomechanical tracking to high-stakes academic performance remain limited. To address this gap, this study quantifies kinematic postural changes during prolonged study sessions using motion capture technology. The target cohort comprises students preparing for the Clinical Engineer National Examination – a rigorous medical credentialing exam with a recent pass rate of 65.7 %. We comprehensively analyze the interplay among dynamic posture metrics, trunk muscle strength, self-reported and physiological concentration levels, and actual academic performance on standardized mock examinations. By modeling postural stability and change patterns, this research elucidates the structural relationships between physical and psychological determinants of learning, establishing a robust scientific foundation for evidence-based pedagogical interventions.

2. Method

2.1. Motion Capture Setup and Image Processing Infrastructure

To precisely quantify dynamic postural variations during prolonged cognitive tasks, we established a

non-contact video analysis pipeline using computer vision. Recording was conducted using a high-resolution, fixed-position camera positioned perpendicular to the subject's sagittal plane under controlled, standardized illumination to minimize ambient noise and flickering.

To ensure computational reproducibility and cross-platform stability, software dependencies were managed within a unified Anaconda virtual environment [7, 8]. Image processing and feature extraction were built using Python and OpenCV. Input video frames underwent background subtraction, edge detection, and Gaussian blurring for noise reduction. Key anatomical feature points were tracked dynamically across frame sequences using the Lucas-Kanade optical flow algorithm. Preprocessed time-series coordinate data were subsequently transferred to MATLAB (MathWorks) for mathematical modeling, feature reduction, and advanced statistical evaluation.

Fig. 1 illustrates the anatomical feature points and the corresponding spatial kinematic model.




Fig. 1. Example configuration of a posture measurement system for explaining the method.

To establish objective postural baselines, each participant's static and dynamic alignments were recorded prior to testing to define reference states for "good" (optimal) and "poor" (suboptimal) posture:

(1) Good Posture Baseline

Defined as an anatomically neutral alignment minimizing musculoskeletal load. This state was indexed across 16 primary spatial landmarks (yielding 33 bilateral keypoints in 3D space: coordinates,

resulting in 99 kinetic variables). These landmarks spanned three functional regions:

- 1) Head: Nose, eyes, ears, and mouth;
- 2) Torso and Upper Extremities: Shoulders, elbows, wrists, and digits;
- 3) Lower Extremities: Iliopsoas/hip joints, knees, ankles, heels, and toes.

(2) Poor Posture Baseline

Defined by retroversion of the pelvis, posterior torso inclination, and marked forward head translation (slouching/kyphotic posture).

Because postural coordinates covary during prolonged task engagement, feature selection was performed to reduce collinearity. From the full variable set, geometric metrics – such as the cervical angle (ear–shoulder–hip) and forward head displacement magnitude – were extracted to optimize classification sensitivity.

2.2. Dimensionality Reduction and Classification

To model temporal posture degradation while mitigating redundant data, machine learning workflows were implemented in MATLAB. High-dimensional time-series landmark trajectories were processed using Principal Component Analysis (PCA) to derive uncorrelated composite factors while preserving global variance. The first two principal components, capturing the primary vectors of spine collapse and fatigue-induced slouching, were extracted as input features.

Using these principal component scores and key geometric metrics as feature vectors, a Support Vector Machine (SVM) classifier was constructed. The SVM optimizes a maximum-margin separating hyperplane in high-dimensional space, providing robust generalization performance even with moderate sample sizes. The classifier was trained using labeled reference datasets of optimal and suboptimal postures.

To capture complex, non-linear trajectories of postural collapse that linear boundaries fail to segregate, a Radial Basis Function (RBF) kernel was applied. Hyperparameters were tuned via 5-fold cross-validation. The distance of temporal data points from the SVM decision boundary served as a continuous quantitative index of fatigue accumulation and postural deterioration over time. Time-series trajectories were rendered in 2D and 3D feature spaces to evaluate inter- and intra-individual biomechanical trends across extended study sessions.

3. Results

3.1. Longitudinal Postural Control Measurement and Baseline Evaluation

Dynamic posture measurements were conducted across six testing periods over the study duration (May, July, September, November, and January/February of

the subsequent year). To assess postural stability, the Euclidean distances were calculated between each participant's time-series spatial coordinates and the predefined reference coordinates corresponding to optimal ("good") and suboptimal ("poor") postures.

Descriptive statistics of the dynamic posture metrics over time are summarized in Table 1. A one-way analysis of variance (ANOVA) was performed to evaluate time-series variations in mean Euclidean distances across the measurement months. The analysis revealed no statistically significant differences across periods, indicating that global postural control characteristics remained macroscopically stable without marked deterioration or seasonal fluctuations throughout the academic year.

To establish quantitative ground truth labels for classification, normative reference video clips were compiled for optimal and suboptimal postures. Individual discrimination thresholds (cutoff values) separating good and poor postures were calculated across all 99 spatial variables. Table 2 presents the key evaluation metrics and their corresponding calculated cutoff values.

Table 1. An example of the obtained data.

	nose		
	x	y	z
Ave.	0.51	0.35	-1.11
	left eye(inner)		
	x	y	z
Ave.	0.53	0.33	-1.12
	left eye		
	x	y	z
Ave.	0.54	0.33	-1.12

Table 2. Selected average values for "good posture" and "poor posture".

	nose		
	x	y	z
Optimal posture	0.49	0.33	-0.61
Suboptimal posture	0.44	0.50	-1.51
Cut-off point for optimal posture	0.47	0.42	-1.06
	left eye(inner)		
	x	y	z
Optimal posture	0.52	0.30	-0.59
Suboptimal posture	0.47	0.49	-1.56
Cut-off point for optimal posture	0.49	0.40	-1.07
	left eye		
	x	y	z
Optimal posture	0.54	0.30	-0.59
Suboptimal posture	0.49	0.48	-1.55
Cut-off point for optimal posture	0.51	0.39	-1.07

3.2. Identification of Key Postural Indicators via Principal Component Analysis

Principal Component Analysis (PCA) was applied to the 3D coordinate trajectories of the 16 anatomical landmarks (99 total spatial variables) to isolate the

principal features capturing postural breakdown during extended study sessions.

Evaluation of eigenvector loadings and contribution ratios demonstrated that a two-variable subset – specifically, the depth-axis (Z-axis) displacement of the left shoulder and the right ear – exhibited the highest variance concentration, accounting for a cumulative contribution ratio of 0.72. These two metrics reflect anteroposterior displacement and head-neck alignment. These results statistically confirm that the dynamic coupling between anteroposterior shoulder displacement and ear translation serves as a dominant marker of postural collapse, capturing trunk inclination and forward head translation during task engagement.

3.3. Spatial Classification and Academic Performance Association via Support Vector Machines

To quantitatively evaluate and classify individual alignment states, a Support Vector Machine (SVM) classifier was constructed using the two primary PCA-derived indicators and as input feature vectors.

Fig. 2 illustrates the decision space and optimal decision boundary derived from the SVM model using a non-linear Radial Basis Function (RBF) kernel, with on the vertical axis and on the horizontal axis. Participants were color-coded according to their national examination outcomes (orange markers: "Pass" group; blue markers: "Fail" group).

When national exam performance attributes were mapped onto the two-dimensional feature space, the "Pass" and "Fail" cohorts segregated into distinct, localized cluster regions defined by the SVM hyperplanes. This spatial separation demonstrates a robust quantitative association between dynamic biomechanical stability during prolonged study and high-stakes cognitive performance.




Fig. 2. Results from Support Vector Machines.

4. Consideration

4.1. Anatomical, Biomechanical, and Ergonomic Validity of Derived Postural Indicators

The Principal Component Analysis (PCA) conducted in this study established that a specific subset of kinematic spatial variables – namely, the depth-axis displacement of the left shoulder left shoulder Z and the right ear right ear Z – yielded a cumulative variance contribution ratio of 0.72. This robust mathematical concentration of variance reflects fundamental biomechanical and ergonomic principles governing upper-body structural collapse during prolonged seated cognitive tasks.

The Z-axis measures translational movement along the sagittal (anteroposterior) plane. In biomechanical terms, the relative spatial displacement between the auditory meatus (head landmark) and the acromion process (shoulder girdle landmark) directly captures the kinematic progression of forward head posture (FHP) and thoracic hyperkyphosis ("slouching" or "rounded back"). During extended desk-bound activities, such as writing, reading, or interacting with digital displays, progressive mental fatigue and visual strain induce a gradual anterior translation of the center of mass of the head. In an adult, the human head weighs approximately 4.5-5.5 kg. As the head shifts forward relative to the cervical spine, the gravitational moment arm increases exponentially, imposing disproportionate mechanical loads on the posterior cervical musculature and upper trapezius.

When muscular fatigue impairs the endurance of these posterior stabilization muscle groups, compensatory strategies emerge. To offset the growing mechanical torque, the upper trunk sways posteriorly or the scapulae undergo protraction, abduction, and internal rotation, yielding the classic "rounded shoulder" phenotype. Furthermore, asymmetrical postural habits – such as unilateral trunk rotation during writing or asymmetrical desk configurations – cause the sagittal displacement of the shoulder joint left shoulder Z and the combined axial rotation or lateral flexion of the cranium right ear Z to covary dynamically.

Consequently, the identification of these two coupled Z-axis parameters provides a mathematically minimal yet biomechanically optimal feature space. This feature pair sensitively detects the early kinematic hallmarks of fatigue-induced compensatory movements, isolating functional upper-body degradation from background movement noise in complex, three-dimensional task environments.

4.2. Homeostatic Constancy, Longitudinal Stability, and Micro-Kinematic Sway

Analysis of the longitudinal measurements across six distinct testing periods (spanning May to February of the following academic year) revealed no statistically significant overall degradation or

improvement in mean Euclidean postural metrics. This apparent invariance over time can be interpreted through the biological concept of dynamic homeostasis within the central nervous system's postural control loops.

From a physiological perspective, human postural control relies on continuous sensorimotor feedback integration involving visual, vestibular, and somatosensory inputs. Unless an individual develops acute musculoskeletal pathology or structural spinal deformities, the neuromuscular system actively compensates for transient imbalances to maintain overall spatial equilibrium. Furthermore, the observational protocol itself must be evaluated considering potential behavioral artifacts. Because postural data were recorded in discrete sessions lasting several minutes each, participants may have exerted conscious motor control to rectify their alignment – a classic manifestation of the Hawthorne effect driven by camera awareness. Under short-term observation, conscious cortical override temporarily bypasses fatigue-induced subcortical relaxation, causing temporal average coordinates to converge within a normative physiological baseline.

However, the absence of macro-level longitudinal shifts in *mean values* does not signify an absence of biomechanical variance. Rather, it underscores the necessity of differentiating macro-level static posture from micro-level dynamic stability. While annual mean baselines remain globally stable due to homeostatic feedback, the *underlying dynamics* – such as high-frequency postural micro-sway, the latency to postural collapse during continuous study, and the frequency of sudden compensatory shifts – exhibit pronounced individual variance. As demonstrated in our classification analysis, these subtle, dynamic fluctuations under continuous cognitive load directly differentiate high-performing students from their peers.

4.3. Neuromuscular, Respiratory, and Neurocognitive Mechanisms Underlying Exam Performance

The primary finding of this study lies in the high-dimensional spatial separation achieved by the Support Vector Machine (SVM) model. Using the PCA-derived feature space left shoulder Z and right ear Z, the SVM cleanly segregated participants into distinct clusters corresponding to their final outcome ("Pass" vs. "Fail") on the Clinical Engineer National Examination. This clear margin of separation provides empirical evidence that physical posture maintenance capacity – rooted in deep trunk neuromuscular endurance – exerts a profound influence on long-term cognitive endurance and academic achievement.

Several interconnected physiological mechanisms elucidate this relationship:

1) Thoracic Biomechanics and Respiratory Function: Insufficient strength or endurance in key core stabilizing musculature (e.g., *transversus*

abdominis, *multifidus*, diaphragm, and pelvic floor) causes posterior pelvic tilt and thoracic kyphosis. This slouched alignment restricts the movement of the rib cage and constrains diaphragmatic excursion, reducing chest wall compliance. The resulting decrement in tidal volume leads to sub-clinical, chronic mild hypoxia and altered arterial blood gas dynamics. Reduced oxygen delivery to the cerebral cortex during intense cognitive processing accelerates mental exhaustion.

2) Somatosensory Integration and Cortical Arousal: The Ascending Reticular Activating System (ARAS) within the brainstem regulates global cortical arousal based on tonic somatosensory afferent signals originating from anti-gravity muscle spindles and articular mechanoreceptors. Slouching collapses structural anti-gravity tension, diminishing tonic proprioceptive input to the brainstem. This drop in sensory bombardment reduces ARAS firing rates, causing cortical arousal to wane and inducing drowsiness, attentional lapses, and diminished working memory capacity.

3) Vascular Compression and Nuchal Discomfort: Forward head posture forces hyper-contraction of the suboccipital and upper trapezius muscles to prevent forward cranial rotation. Sustained isometric tension impedes local microvascular perfusion, inducing nuchal rigidity, tension headaches, and localized discomfort. This noxious nociceptive signaling competes directly for central executive resources in the prefrontal cortex (PFC), exacerbating cognitive fatigue.

Students in the "Pass" group possessed superior neuromuscular endurance and postural stability. Consequently, during prolonged study sessions, they resisted fatigue-induced biomechanical collapse, maintaining optimal cortical arousal, unimpaired pulmonary ventilation, and sustained executive function throughout the year. Conversely, students in the "Fail" group exhibited compromised musculoskeletal support as early as the baseline measurements in May. Under continuous task demands, these students suffered frequent postural degradation below critical threshold values. The continuous physical effort required to compensate for poor mechanics generated a persistent drain on cognitive bandwidth. Accumulated over months of preparation, this continuous "leakage" of executive processing power manifested as reduced learning retention, lower mock examination scores, and ultimate failure on the high-stakes national credentialing exam.

5. Conclusions

In this study, we conducted a rigorous, longitudinal investigation into the complex interplay between objective upper-body kinematic postural dynamics, neuromuscular endurance, and long-term, high-stakes cognitive performance among students preparing for the Clinical Engineer National Examination. Historically, physical posture within educational

environments has been addressed primarily as a matter of etiquette or basic musculoskeletal ergonomics. However, by establishing a non-contact, highly reproducible computer-vision and machine-learning tracking framework, this research provides empirical evidence that dynamic postural stability serves as an essential, physiologically grounded foundation for sustained executive function, cognitive endurance, and academic success. Methodologically, our experimental design leveraged a combination of optical flow algorithms and multi-dimensional feature extraction within a unified Anaconda and MATLAB computational ecosystem. By tracking 33 key bilateral anatomical landmarks (yielding 99 kinetic variables across 3D space) under standardized conditions, we systematically mapped temporal posture degradation during prolonged study sessions. Principal Component Analysis (PCA) demonstrated that a remarkably minimal spatial feature subset – specifically, the depth-axis (Z-axis) displacement of the left shoulder_Z and the right ear_Z – captured 72% of the total variance associated with fatigue-induced structural collapse. Biomechanically, this coupling measures the dynamic interaction between forward head posture (FHP) and thoracic hyperkyphosis, effectively tracking the progressive anterior translation of the head's center of mass and the compensatory protraction of the shoulder girdle under continuous task engagement. Using these two primary PCA-derived indicators as input vectors, a Support Vector Machine (SVM) classifier employing a non-linear Radial Basis Function (RBF) kernel achieved a clear high-dimensional spatial separation between students who ultimately passed the national examination and those who failed. This distinct segregation proves that physical posture maintenance under continuous cognitive load is not merely a passive byproduct of sitting, but an active physiological determinant of learning retention and intellectual stamina. While macro-level longitudinal measurements across six testing periods over the academic year revealed global spatial stability due to central nervous system homeostatic feedback loops and short-term conscious motor control (e.g., the Hawthorne effect), micro-kinematic instabilities and latency to postural collapse clearly differentiated high-performing students from their lower-performing peers. The physiological and neurocognitive mechanisms underlying these findings are multifaceted: Respiratory and Thoracic Biomechanics: Insufficient core and trunk muscle endurance (e.g., transversus abdominis, multifidus) leads to pelvic retroversion and thoracic slouching. This mechanical collapse restricts chest wall expansion and constrains diaphragmatic excursion, reducing tidal volume and inducing sub-clinical, chronic mild hypoxia. The resultant decrease in cerebral oxygen delivery during intensive cognitive tasks accelerates mental fatigue. Somatosensory Integration and Central Arousal: The Ascending Reticular Activating System (ARAS) within the brainstem modulates global cortical arousal based on tonic somatosensory afferent signals from

anti-gravity muscle spindles and articular mechanoreceptors. Slouching diminishes this anti-gravity tension, leading to reduced ARAS firing rates, waning cortical arousal, attentional lapses, and diminished working memory capacity. Vascular Compression and Nociceptive Interference: Forward head posture demands sustained isometric hyper-contraction of the suboccipital and upper trapezius muscles to prevent forward cranial drop. This local microvascular compression causes nuchal rigidity and localized pain, generating noxious nociceptive signaling that competes directly for central executive processing resources within the prefrontal cortex. Students in the "Pass" cohort possessed superior neuromuscular endurance, enabling them to resist biomechanical collapse, maintain optimal cortical arousal, and preserve cognitive bandwidth across months of rigorous exam preparation. Conversely, students in the "Fail" cohort experienced persistent physical degradation, creating an ongoing "leakage" of executive capacity that ultimately impaired learning retention and exam outcomes. Despite these significant contributions, certain limitations must be recognized. The discrete laboratory recording setup could not entirely eliminate conscious motor corrections by participants. Furthermore, our spatial kinematic model focused primarily on upper-body Z-axis translation, omitting detailed lower-extremity kinetics and pelvic tilt dynamics.

While this investigation offers a novel, non-contact, objective framework linking posture kinematics to high-stakes academic outcomes, several limitations warrant consideration.

First, as noted, the laboratory measurement environment could not entirely eliminate the Hawthorne effect. While discrete camera recordings established baseline physical capacities, they may not fully capture the unconstrained, sub-conscious postural habits adopted during unmonitored, late-night study sessions. Future investigations should deploy ambient, non-intrusive computer vision pipelines or lightweight, wearable inertial measurement units (IMUs) capable of continuous, 24/7 longitudinal tracking in naturalistic home and classroom environments.

Second, while focusing on upper-body Z-axis translation optimized the signal-to-noise ratio and classification performance in this study, the current kinematic model omits lower-extremity alignment and full pelvic kinematics. Integrating full-body biomechanical models that combine 3D pelvic tilt, lumbar lordosis angles, and ground reaction force

distributions will yield a more comprehensive understanding of the kinetic chain during cognitive fatigue.

Addressing these limitations will advance educational technology beyond passive, verbal posture reminders. By integrating real-time computer vision with predictive machine learning algorithms, future systems can provide automated, non-invasive biofeedback. Such infrastructure will dynamically detect cognitive fatigue and postural decay, enabling personalized ergonomical and pedagogical interventions that optimize learner health and academic performance.

References

- [1]. E. Peper, R. Harvey, L. Mason, I.-M. Lin, Do better in math: How your body posture may change stereotype threat response, *NeuroRegulation*, Vol. 5, Issue 2, 2018, pp. 67-74.
- [2]. K. Inagaki, T. Shimizu, Y. Sakairi, Effects of posture regulation on mood states, heart rate and test performance in children, *International Journal of Sport and Health Science*, Vol. 16, 2018, pp. 195-207.
- [3]. T. Koseki, F. Kakizaki, S. Hayashi, N. Nishida, et al., Effect of forward head posture on thoracic shape and respiratory function, *Journal of Physical Therapy Science*, Vol. 31, Issue 1, 2019, pp. 63-68.
- [4]. H. Kano, Y. Matsuzaki, H. Takaya, T. Yagi, et al., An electroencephalographic study of brain activity produced by postural change from supine to sitting position, *Journal of Physiological Anthropology and Applied Human Science*, Vol. 22, Issue 6, 2003, pp. 273-281.
- [5]. M. Muehlhan, H. Budde, T. Rauh, A. Strobel, The effect of body posture on cognitive performance: A question of sleep quality, *Frontiers in Human Neuroscience*, Vol. 8, 2014, 171.
- [6]. I. Betto, R. Hatano, H. Nishiyama, Distraction detection of lectures in e-learning using machine learning based on human facial features and postural information, *Artificial Life and Robotics*, Vol. 28, Issue 1, 2023, pp. 166-174.
- [7]. J. A. O. Trejo, Basic uses of Python for research in our laboratory: Examples of data analysis and processing of histological samples based on the special graduate lecture for postgraduate students, *Juntendo Medical Journal*, Vol. 71, Issue 4, 2025, pp. 231-237.
- [8]. N. Shirahama, A quasi-experimental study of parallel Python-R learning in data science education: Implementation and assessment in a BYOD environment, *Journal of Intelligence Science in Local Research*, Vol. 1, Issue 2, 2024, pp. 1-19.

A Deterministic On-Demand Simulation Framework for Parametric Analysis of Coincidence-Based Frequency Measurement

Avalos-Gonzalez Daniel, Misael Reyes I., Islas-Pereda S., José M. Ramírez, Pérez-Pérez Blanca Dalila, Zamora-Machado Marlene and Martínez-Quintero Luis
Universidad Autónoma de Baja California, Blvd. Benito Juárez, 21280, Mexicali, Baja California, México
Tel.: +52 (686) 841-8235
E-mail: davalos@uabc.edu.mx

Summary: Coincidence-based frequency measurement methods generate temporal structures whose behavior depends on signal periods, pulse width, temporal offset, and observation interval. Systematic investigation of these factors is challenging because the associated parameter space is broad and difficult to explore experimentally. This work presents a deterministic on-demand simulation framework based on a unified time-domain parameterization, enabling direct and reproducible generation of measurement scenarios from their defining parameters. The framework generates coincidence-related temporal vectors, period estimates, and error distributions consistent with the ideal deterministic coincidence-based measurement model. Representative scenarios illustrate its capability to analyze the influence of the main temporal parameters on the resulting coincidence structures. Comparison with experimental measurements under equivalent nominal conditions showed consistent dominant error behavior, with 95.62 % of the experimental estimates remaining within the theoretical error interval for the evaluated condition. The proposed framework provides a reproducible computational infrastructure for systematic parametric exploration and future investigations involving error-reduction and coincidence-selection strategies.

Keywords: Coincidence-based frequency measurement, Time and frequency measurements, Deterministic simulation framework, Time-domain parameterization, Parametric exploration.

1. Introduction

Frequency measurement is widely used in electronic instrumentation, sensing systems, and frequency-domain transducers. In many applications, the variable of interest is represented as a frequency variation generated by resonant or oscillatory structures, requiring measurement techniques capable of detecting small frequency changes with adequate resolution and speed [1].

Among the available approaches, coincidence-based frequency measurement methods estimate frequency through the temporal comparison of periodic pulse trains. Early developments demonstrated that coincidence-based and rational-approximation techniques can provide fast frequency estimation while exploiting temporal relationships between periodic signals [2]. Instead of relying exclusively on long counting intervals, these methods exploit coincidence events between a reference signal and a measured signal, generating temporal structures that contain information about the relationship between both signals.

As the method evolved, it became evident that its behavior depends not only on the nominal signal periods but also on additional temporal parameters. Previous studies have shown that pulse width can modify coincidence patterns and the formation of coincidence packets [3], while geometric constraints associated with the temporal relationship between signals influence the structure of the generated measurement data [4]. Likewise, phase-related effects have been shown to alter the resulting coincidence distributions and measurement behavior [5].

Consequently, understanding the behavior of the method often requires evaluating large numbers of operating conditions.

Although previous studies have investigated specific aspects of coincidence-based measurement systems, including computational analysis of the data generated during coincidence processes [6], systematic exploration of the associated parametric space remains challenging. This difficulty is particularly relevant when variables such as pulse width and temporal offset must be analyzed over broad ranges of operating conditions. Furthermore, exhaustive repositories of precomputed simulations become increasingly difficult to maintain as parameter resolution and dimensionality increase.

To address these challenges, this work presents a deterministic on-demand simulation framework for coincidence-based frequency measurement. The framework adopts a unified time-domain parameterization and generates measurement scenarios directly from their defining parameters, enabling reproducible and systematic exploration of coincidence-based measurement conditions without relying on large precomputed simulation repositories.

2. Background and Research Gap

Coincidence-based frequency measurement methods have evolved through advances in frequency instrumentation and the analysis of temporal coincidence structures. This section reviews these developments and identifies the research gap addressed by the proposed framework.

2.1. Frequency-Based Instrumentation

Frequency-output sensing systems are widely used in instrumentation because frequency signals provide high noise immunity, digital compatibility, and suitability for long-distance transmission [1]. Consequently, frequency-to-digital conversion techniques have become fundamental components of modern sensing and measurement platforms. The development of smart sensors, frequency-output transducers, and virtual instrumentation environments has further increased the importance of frequency-based measurement techniques in modern instrumentation systems [7-10].

As coincidence-based measurement methods evolved, computational resources became increasingly important for analyzing the temporal structures generated during the measurement process. Previous work on virtual instrumentation and intelligent sensor interfaces demonstrated the advantages of software-assisted measurement environments capable of providing flexibility, automation, and advanced data analysis [7, 8]. Similarly, universal frequency-to-digital conversion architectures and standardized interfaces for smart sensors have highlighted the importance of computational support in frequency-based instrumentation [9, 10].

Under these conditions, computational frameworks can be viewed as part of the measurement infrastructure itself, providing controlled and reproducible environments for studying coincidence-based measurement behavior across a broad range of operating conditions.

2.2. Evolution of Coincidence-Based Measurement Methods

Coincidence-based frequency measurement methods estimate frequency through the temporal comparison of periodic signals. Early approaches treated coincidence events primarily as counting markers used to construct frequency estimates [11, 12]. Subsequent developments demonstrated that rational-approximation techniques could improve measurement performance while exploiting additional information contained in the temporal relationship between signals [2].

Further studies revealed that the generated data may include coincidence packets, geometric constraints, pulse-width effects, and temporal structures that influence the measurement process [3, 4, 13]. In particular, pulse width was shown to affect the formation of partial and complete coincidences [3, 14], while geometric analysis demonstrated that period relationships, pulse width, and initial temporal conditions modify the structure of the generated coincidence data [4]. Additional studies related to frequency-shift measurements and high-resolution applications confirmed the importance of understanding these temporal structures under realistic operating conditions [15, 16].

As a result, the method evolved from a simple counting procedure into a temporal phenomenon governed by the interaction of signal periods, pulse width, temporal offset, and observation interval. Understanding the behavior of the method therefore requires analysis of the internal structures generated by the coincidence process rather than consideration of nominal frequency values alone [4, 6].

Fig. 1 illustrates a representative oscilloscope capture obtained from a coincidence-based measurement system. The temporal relationship between the reference and measured signals, together with the resulting coincidence regions, highlights the physical phenomena that motivate the computational analysis performed by the proposed framework.




Fig. 1. Oscilloscope capture of an experimental coincidence-based measurement system.

2.3. Research Gap

Although previous studies have investigated specific aspects of coincidence-based measurement systems, most analyses remain focused on particular operating conditions or restricted parameter configurations. Studies have addressed jitter influence [2], pulse-width effects [3, 14], geometric constraints [4], phase or temporal-offset effects [5], and computational analysis of generated data [6]. More recent works have also explored frequency-estimation behavior in specific application domains and under particular measurement conditions [17, 18]. However, these contributions remain distributed across isolated regions of a much broader parametric space.

This limitation becomes increasingly significant because several influential variables, including signal periods, pulse width, temporal offset, and observation interval, are naturally continuous. Systematic evaluation of their combined influence rapidly increases the number of possible operating scenarios and may be difficult to reproduce experimentally, particularly when arbitrary temporal offsets must be considered [5, 6].

As parameter resolution and dimensionality increase, repositories of precomputed simulations become difficult to maintain and extend. This motivates the use of on-demand generation of coincidence-based measurement scenarios.

2.4. Relation to Previous Work

Previous studies in this research line addressed specific characteristics of coincidence-based frequency measurement, including pulse-width effects and coincidence-packet formation [3], geometric constraints and measurement optimization [4], phase-dependent measurement behavior [5], computational analysis of experimentally generated coincidence data [6], pulse-width optimization for a required resolution [14], and the occurrence of median fractions in experimental measurements [18]. In contrast, the present work does not focus on a single operating parameter or a particular estimation criterion. It provides a deterministic parameter-driven framework that generates coincidence structures, period estimates, and error distributions on demand from arbitrary combinations of signal periods, pulse width, temporal offset, and observation interval. Thus, simulation is used here as a reusable and reproducible infrastructure for systematic parametric exploration, rather than being limited to case-specific analysis of particular measurement conditions.

3. Coincidence-Based Measurement Model

The proposed framework is based on a coincidence-based measurement model in which two periodic pulse trains are compared over a finite observation interval. Previous studies have shown that the resulting coincidence structures depend on pulse width, initial temporal conditions, and geometric restrictions associated with the temporal relationship between signals [4, 5, 14].

The model is defined using a unified time-domain parameterization composed of the reference period T_r , measured period T_m , pulse width W , initial temporal offset Δt_ϕ , and observation interval T_{obs} . The corresponding frequencies are treated as derived quantities and are obtained from:

$$f = 1/T, \quad (1)$$

where f is the frequency and T is the period.

Under this formulation, coincidence events occur whenever pulses from both signals overlap in time. Coincidence-based frequency measurement relies on the existence of temporal relationships between integer multiples of the reference and measured periods. These relationships can be expressed as:

$$A \cdot T_r = B \cdot T_m \pm \varepsilon, \quad (2)$$

where A and B are positive integers, T_r is the reference period, T_m is the measured period, and ε represents the temporal mismatch associated with the coincidence process.

Coincidence events occur when this mismatch remains within the limits imposed by the pulse width.

The resulting coincidence sequence depends on the interaction among signal periods, pulse width, temporal offset, and observation interval, generating temporal structures that may include isolated coincidences, partial coincidences, or coincidence packets [3, 4, 13, 14].

The measurement process can therefore be represented as a mapping between the input parameter set and the corresponding coincidence structures generated during the observation interval.

The model adopts a unified time-domain representation in which all operating conditions are expressed as temporal quantities. This formulation provides the conceptual foundation for the deterministic simulation framework described in the following section. Such a representation is consistent with previous computational analyses showing that the behavior of coincidence-based measurement systems can be investigated through the temporal structures generated during the measurement process [6, 18].

4. Proposed Deterministic Simulation Framework

The proposed framework provides a computational implementation of the coincidence-based measurement model described in the previous section. For a given parameter set, the framework generates the temporal structures associated with the coincidence process and extracts the information required for subsequent analysis.

The framework first generates the periodic pulse sequences corresponding to the reference and measured signals. These signals preserve the temporal characteristics defined by the selected periods, pulse width, temporal offset, and observation interval. The generated pulse trains are then evaluated to identify coincidence-related events and the temporal structures produced by their interaction.

The resulting coincidence information is organized into paired temporal vectors containing the retained coincidence-related events. These vectors constitute the primary framework outputs and provide the basis for period estimation, error analysis, coincidence-structure evaluation, and graphical representation of the generated scenarios.

In addition to numerical outputs, the framework generates visual representations of the coincidence process, including oscilloscope-like signal plots, estimated-period distributions, and error histograms. These outputs facilitate interpretation of parameter-dependent behavior and support qualitative and quantitative analysis of the generated coincidence structures.

Because the framework is deterministic, identical parameter sets always produce identical outputs. This characteristic ensures reproducibility and enables systematic comparison among different operating conditions while preserving full control over the explored parameter space.

5. On-Demand Generation Versus Precomputed Repositories

The proposed framework was conceived as a generative environment rather than as a repository of previously computed measurement scenarios. This design choice is motivated by the characteristics of the coincidence-based measurement parameter space and by the scalability limitations associated with exhaustive simulation repositories.

Consider a parameter space defined by reference period, measured period, pulse width, temporal offset, and observation interval. If each variable is discretized into a finite number of values, the total number of required scenarios can be expressed as:

$$N_{total} = N_{Tr} \cdot N_{Tm} \cdot N_W \cdot N_{\Delta t\phi} \cdot N_{Tobs}, \quad (3)$$

where each term represents the number of discrete values assigned to the corresponding temporal parameter.

Because these variables are naturally continuous, increasing parameter resolution rapidly increases the number of possible operating conditions. Consequently, exhaustive repositories require progressively larger storage capacity and become more difficult to organize, maintain, transport, and extend.

The proposed framework adopts an alternative strategy in which measurement scenarios are generated directly from their defining parameters. Under this approach, the information required to reproduce a scenario is reduced to a compact parameter set, while the corresponding temporal structures are generated only when needed.

This on-demand generation strategy preserves reproducibility, parameter flexibility, and arbitrary temporal resolution without requiring storage of all possible operating conditions. As a result, the framework remains scalable even when broad regions of the parameter space must be explored systematically.

6. Simulation Methodology

Once a parameter set has been defined, the proposed framework generates the corresponding coincidence-based measurement scenario directly from its temporal parameters. This section describes the deterministic procedure used to construct the coincidence-related temporal structures and the outputs derived from them.

6.1. Scenario Generation

Each measurement scenario is defined by a parameter set, and for a given parameter set, the framework generates a temporal sequence associated with the measured signal. This sequence is represented by the vector vT_m , whose first element corresponds to T_m . Subsequent elements are obtained by successive

additions of T_m until the observation interval is reached.

The coincidence-based method compares elapsed time intervals through integer numbers of complete periods. Therefore, the temporal position of the coincidences selected at the beginning and end of a measurement introduces a bounded discrepancy between the elapsed times represented by the measured and reference signals. In this work, W denotes the pulse width, and each boundary coincidence may introduce a temporal deviation of up to approximately one pulse width. The deviations associated with the initial and final coincidences may partially compensate each other or, in the limiting cases, accumulate in the same direction, producing a total discrepancy approaching $2W$.

Fig. 2 illustrates both limiting cases using the positive pulse edges as temporal references. If the measurement interval is defined by the first and last coincidences shown in the figure, 10 complete periods of the measured signal and 11 complete periods of the reference signal are counted. However, the corresponding intervals $10T_m$ and $11T_r$ differ by approximately $2W$, with the measured-signal interval being longer. Conversely, when the two central coincidences are used as the measurement boundaries, the intervals correspond to $2T_m$ and $3T_r$, and the measured-signal interval is approximately $2W$ shorter. Intermediate coincidence positions produce discrepancies between these limiting cases.




Fig. 2. Representative limiting cases of accumulated temporal discrepancy between measured and reference intervals.

Accordingly, each element of vT_m is evaluated against the nearest reference-signal event. Elements whose temporal difference exceeds the established $2W$ bound are discarded, while the remaining elements are retained as coincidence-related events, satisfying

$$|vT_m - k \cdot T_r| < 2 \cdot W, \quad (4)$$

where k is a positive integer chosen such that $k \cdot T_r$ corresponds to the reference event nearest to the evaluated element of vT_m .

The original indices of the retained elements are subsequently used to construct the corresponding reference vector vT_r . The resulting pair of vectors vT_m, vT_r contains the coincidence-related temporal information associated with the ideal coincidence-based measurement model.

When multiple parameter values are specified, the framework automatically evaluates all parameter combinations using the same deterministic generation and filtering procedure.

6.2. Generated Outputs

The retained vectors vT_m and vT_r constitute the primary outputs of the framework. These vectors are subsequently used to generate estimated period values, estimation errors, coincidence distributions, and coincidence-packet structures.

In addition, the framework produces graphical outputs including oscilloscope-like signal representations, estimated-period plots, and error histograms. These outputs provide the basis for the representative analyses presented in the following section.

Because every scenario is defined exclusively by its parameter set, all evaluations can be reproduced directly without requiring precomputed repositories or external datasets.




Fig. 2. Representative limiting cases of accumulated temporal discrepancy between measured and reference intervals.

7. Representative Results

The representative results illustrate the information generated by the framework from the retained coincidence-related vectors vT_m and vT_r . These results are intended to show how the proposed framework

supports analysis of the temporal behavior of coincidence-based measurement conditions.

7.1. Temporal Representation of Generated Signals

The first output is an oscilloscope-like representation of the generated reference and measured pulse trains, Fig. 2 shows an example. This visualization allows direct inspection of the temporal relationship between signals and shows how pulse width and temporal offset affect the occurrence of coincidence regions.

7.2. Estimated Period Values

Each pair of corresponding coincidence-related events retained in vT_m and vT_r defines the integer numbers of complete measured and reference periods elapsed between the considered coincidence boundaries. These counts, denoted by B and A , respectively, are used to obtain an individual estimate of the measured period by rearranging (2):

$$\hat{T}_m = \frac{A}{B} T_r \pm \frac{\varepsilon}{B}, \quad (5)$$

where $\hat{T}_m$ is an individual estimate of the measured period, A is the number of complete reference periods T_r , B is the corresponding number of complete measured periods, and ε represents the temporal discrepancy associated with the coincidence boundaries, as defined in (2). Therefore, a single scenario produces a set of estimated period values rather than a unique estimate, as illustrated in Fig. 3.

The graphical representation of these estimates shows their dispersion with respect to the nominal measured period T_m . This output allows direct observation of the variability introduced by the coincidence process under the selected operating conditions.




Fig. 3. Dispersion of individual T_m estimates around the real measured period.

7.3. Error Distribution

The estimation error is obtained by comparing each estimated period value $\hat{T}_m$ with the nominal value T_m . The resulting error distribution is represented as a histogram, providing a compact view of the dispersion and concentration of estimation errors, Fig. 4 shows an example.




Fig. 4. Normalized error distribution obtained from simulated measurements.

This representation enables comparison between scenarios involving different pulse widths, temporal offsets, period relationships, or observation intervals.

7.4. Parametric Exploration

When more than one value is assigned to any input parameter, the framework evaluates all parameter combinations using the same deterministic generation and filtering procedure. This allows the influence of pulse width, temporal offset, period relationship, and observation interval to be analyzed systematically without experimental reconfiguration or precomputed repositories.

8. Experimental Comparison of Error Distributions

To assess the framework against a physical coincidence-based measurement system, simulated and experimental error distributions were compared under equivalent nominal conditions, with $W = 200$ ns and an ideal initial coincidence. Consequently, the initial-boundary error is eliminated, reducing the theoretical interval from the general $\pm 2W$ bound to $\pm W$ (± 200 ns). The experimental distribution in Fig. 5 shows an overall error range and pattern consistent with the simulated distribution in Fig. 4. Quantitatively, 95.62 % of the experimental estimates remain within $\pm W$, while 4.38 % exceed this interval.

Both distributions exhibit a continuous error pattern within the expected range. The small experimental excursions beyond $\pm W$ are consistent with non-ideal effects such as timing jitter, relative-phase fluctuations, oscillator drift, pulse-width variations, and acquisition uncertainty, which are not included in the ideal deterministic model. Overall, the experimental comparison supports the framework's capability to reproduce the dominant error behavior under the evaluated conditions.




Fig. 5. Normalized error distribution obtained from experimental measurements.

9. Conclusions

This work presented a deterministic on-demand simulation framework for coincidence-based frequency measurement systems. The proposed approach combines a unified time-domain parameterization with direct generation of measurement scenarios, enabling reproducible exploration of coincidence-based measurement conditions without relying on large repositories of precomputed simulations.

The framework generates coincidence-related temporal structures, estimated period values, and error distributions directly from a compact set of operating parameters. Representative scenarios demonstrated its capability to analyze the influence of pulse width, temporal offset, period relationship, and observation interval on the resulting coincidence structures. Comparison with experimental measurements under equivalent nominal conditions showed a similar overall error range and distribution behavior, with 95.62 % of the experimental estimates remaining within the theoretical error interval for the evaluated condition. The small excursions beyond this interval are consistent with non-ideal temporal effects not included in the deterministic model.

These results support the use of the proposed framework as a practical computational infrastructure for systematic investigation of coincidence-based measurement behavior and for future studies involving coincidence-selection criteria, uncertainty analysis, error-reduction strategies, and other processing approaches based on coincidence-generated data.

Acknowledgements

The authors acknowledge the support of the Facultad de Ingeniería, Universidad Autónoma de Baja California.

References

- [1]. F. N. Murrieta-Rico, V. Petranovskii, O. Y. Sergiyenko, D. Hernández-Balbuena, et al., Frequency domain sensors and frequency measurement techniques, *Applied Mechanics and Materials*, Vol. 756, 2015, pp. 575-584.
- [2]. O. Sergiyenko, D. Hernández-Balbuena, V. Tyrsa, P. L. A. Rosas-Méndez, et al., Analysis of jitter influence in fast frequency measurements, *Measurement*, Vol. 44, Issue 7, 2011, pp. 1229-1242.
- [3]. F. N. Murrieta-Rico, O. Y. Sergiyenko, V. Petranovskii, D. Hernández-Balbuena, et al., Pulse width influence in fast frequency measurements using rational approximations, *Measurement*, Vol. 86, 2016, pp. 67-78.
- [4]. D. Avalos-Gonzalez, O. Sergiyenko, D. Hernández-Balbuena, V. Tyrsa, et al., Constraints definition and application optimization based on geometric analysis of the frequency measurement method by pulse coincidence, *Measurement*, Vol. 126, 2018, pp. 184-193.
- [5]. J. de D. Sánchez-López, F. N. Murrieta-Rico, V. Petranovskii, J. Antúnez-García, et al., Effect of phase in fast frequency measurements for sensors embedded in robotic systems, *International Journal of Advanced Robotic Systems*, Vol. 16, Issue 4, 2019, 1729881419869727.
- [6]. F. N. Murrieta-Rico, E. I. Cota-Rivera, M. E. Raygoza-Limón, R. Jiménez-Sánchez, et al., Computational study of data generated during time-domain overlapping processes, in *Proceedings of the Mexican International Conference on Computer Science (ENC)*, 2023, pp. 1-5.
- [7]. S. Y. Yurish, N. V. Kirianaki, N. O. Shpak, Virtual instrumentation for frequency output smart sensors, in *Proceedings of the 11th IMEKO TC-4 Symposium on Trends in Electrical Measurements and Instrumentation*, 2001, pp. 191-194.
- [8]. S. Y. Yurish, F. Reverter, R. Pallàs-Areny, Measurement error analysis and uncertainty reduction for period- and time-interval-to-digital converters based on microcontrollers, *Measurement Science and Technology*, Vol. 16, Issue 8, 2005, pp. 1160-1166.
- [9]. S. Y. Yurish, High-speed universal frequency-to-digital converter for quasi-digital sensors and transducers, *Sensors & Transducers*, Vol. 80, Issue 6, 2007, pp. 1225-1229.
- [10]. S. Y. Yurish, IEEE 1451 standard and frequency output sensors: How to obtain a broad-based industry adoption?, *Sensors & Transducers*, Vol. 59, 2005, pp. 412-418.
- [11]. N. V. Kirianaki, S. Y. Yurish, N. O. Shpak, Methods of dependent count for frequency measurements, *Measurement*, Vol. 29, Issue 1, 2001, pp. 31-50.
- [12]. S. Y. Yurish, Novel modified method of the dependent count for high precision and fast measurements of frequency-time parameters of electric signals, in *Proceedings of the IEEE Instrumentation and Measurement Technology Conference (IMTC)*, 2008, pp. 876-881.
- [13]. D. Hernández-Balbuena, O. Sergiyenko, V. Tyrsa, L. Burtseva, et al., Signal frequency measurement by rational approximations, *Measurement*, Vol. 42, Issue 1, 2009, pp. 136-144.
- [14]. F. N. Murrieta-Rico, O. Y. Sergiyenko, V. Petranovskii, D. Hernández-Balbuena, et al., Optimization of pulse width for frequency measurement by the method of rational approximations principle, *Measurement*, Vol. 125, 2018, pp. 463-470.
- [15]. D. Avalos-Gonzalez, D. Hernández-Balbuena, V. Tyrsa, V. M. Kartashov, et al., Application of fast frequency shift measurement method for INS in navigation of drones, in *Proceedings of the 44th Annual Conference of the IEEE Industrial Electronics Society (IECON)*, 2018, pp. 3159-3164.
- [16]. F. N. Murrieta-Rico, D. Hernández-Balbuena, J. C. Rodríguez-Quinonez, O. Raymond-Herrera, et al., High resolution measurement of physical variables change for INS, in *Proceedings of the 25th IEEE International Symposium on Industrial Electronics (ISIE)*, 2016, pp. 912-917.
- [17]. F. N. Murrieta-Rico, V. Petranovskii, D. H. Galván, J. Antúnez-García, et al., Frequency shifts estimation for sensors based on optoelectronic oscillators, *IEEE Sensors Journal*, Vol. 21, Issue 10, 2021, pp. 11283-11290.
- [18]. F. N. Murrieta-Rico, O. Sergiyenko, P. Mercorelli, V. Petranovskii, Evaluating the existence of fractions during the frequency measurement process, in *Proceedings of the Joint Conference of the European Frequency and Time Forum and IEEE International Frequency Control Symposium (EFTF/IFCS)*, 2025, 11194426.

A Cuffless Palm-Site On-Demand Continuous Blood Pressure Monitoring Prototype with Deep Learning-Based Estimation

Iyappan Gunasekaran, Lawson Michael and Kannie Wy Chan

Hong Kong Centre of Cerebro-Cardiovascular Health Engineering (COCHE), Shatin, Hong Kong

Tel.: + 852 51057469

E-mail: igunasekaran@hkcoche.org

Summary: Cardiovascular diseases, particularly hypertension, are leading causes of death and disability worldwide, driving the need for accurate, continuous blood pressure (BP) monitoring. Conventional cuff-based devices are unsuitable for beat-to-beat tracking in daily life, while current PPG-based wearables often provide only snapshot measurements and are limited by calibration requirements, motion artefacts, skin-tone dependence, and subclinical accuracy. This work introduces an AI-powered, cuffless wearable for continuous, beat-to-beat BP monitoring on demand using palm-acquired PPG waveforms. A high-fidelity PPG sensor is combined with adaptive signal-quality assessment and deep neural networks trained on large physiological datasets to estimate arterial pressure in real time. By dynamically selecting the most reliable signal segments, the system delivers robust, motion-resilient, and skin-tone-independent BP readings. The CNN-LSTM model was first pretrained on the publicly available MIMIC-III physiological database, which contains synchronised PPG signals and invasively measured arterial blood pressure recordings from a broad clinical population. In the pretraining phase, the model achieved mean absolute errors of 3.75 mmHg for systolic BP and 3.21 mmHg for diastolic BP, indicating strong baseline performance under the database conditions. The pretrained model was then fine-tuned on proprietary datasets to better match the characteristics of the sensor, the measurement setting, signal-quality variability, and the target user group of the intended platform. The fine-tuned model achieved MAEs of 4.58 mmHg for SBP and 3.11 mmHg for DBP. The proposed technology can serve as a supportive tool for enabling non-invasive, real-time BP monitoring across clinical and everyday environments, enabling earlier detection, improved risk stratification, and better long-term management of hypertension and related cardiovascular conditions, alongside validation of signal quality, ethical approval, and prospective validation protocol studies that will be carried out in continuation of current research progress.

Keywords: Beat-to-beat prediction, Blood pressure, Cardiovascular disease, Continuous monitoring, Hypertension, Non-invasive.

1. Introduction

Cardiovascular diseases remain the leading cause of death and disability worldwide, with hypertension constituting a major, modifiable risk factor for myocardial infarction, stroke, heart failure, and chronic kidney disease [1-3]. Accurate blood pressure (BP) with beat-to-beat assessment is central to the diagnosis, risk stratification, and long-term management of hypertensive patients. While conventional cuff-based sphygmomanometers are well-established clinical reference standards, their size, intermittent measurement paradigm, and reliance on occlusive cuffs inherently limit their suitability for continuous, beat-to-beat, or truly wearable BP monitoring in daily life [4, 5].

In recent years, wearable devices incorporating optical sensors, particularly photoplethysmography (PPG), have emerged as promising solutions for noninvasive, cuffless BP estimation [6-9]. These systems aim to extend BP monitoring beyond the clinic, enabling ambulatory and home-based surveillance for 24 hours or longer [10]. However, current PPG-based approaches still present important limitations. Many implementations provide only snapshot or short interval measurements rather than genuine beat-to-beat monitoring in clinical standards [11, 12]; they frequently require repeated calibration

against a cuff-based reference [13]; and they often fail to achieve clinical-grade accuracy across diverse user populations [14]. Moreover, optical BP estimation is highly susceptible to motion artefacts, environmental lighting conditions, and skin tone-dependent variations in signal quality, which collectively compromise measurement reliability and can systematically disadvantage individuals with darker skin tones [15, 16].

These shortcomings underscore an urgent need for a skin-tone-independent, motion-robust, and precisely positionable wearable technology capable of delivering accurate, continuous, and cuffless BP estimation under real-world conditions. Advances in deep neural networks and large, physiologically rich datasets now offer the opportunity to move beyond traditional, hand-engineered features and toward end-to-end models that can learn complex mappings between PPG waveforms and arterial pressure, while simultaneously incorporating signal quality metrics and contextual information [17, 18].

In this work, we present a novel, non-invasive, continuous, beat-to-beat BP monitoring and prediction technology based on PPG waveforms acquired from the palm. The proposed system detects peripheral pulse waves using a compact PPG sensor and integrates adaptive signal quality evaluation with deep learning models trained on comprehensive

physiological datasets. By dynamically selecting and weighting the most reliable signal segments, the system enables beat-to-beat, real-time BP estimation without an occlusive cuff. Our approach is designed to maintain clinical-grade accuracy across a wide range of physiological states and user characteristics, including cardiac, diabetic, hypertensive, and pregnant populations, while preserving performance under challenging environmental and motion conditions.

1.1. Proposed Design

The proposed invention introduces a next-generation wearable blood pressure monitoring system that integrates PPG sensor signals of the palm. This smart system is designed to overcome the limitations of traditional cuff-based devices by combining mechanical sensing modalities with optical BP estimation methods that are inherently more robust to motion artefacts, ambient light interference, sensor misalignment, and vibration. In addition, it reduces pressure-related discomfort, minimises the need for frequent cuff-based recalibration, and accounts for user-specific biometric variations. A schematic representation of the invention is shown in Fig. 1.




Fig. 1. Schematic representation of the proposed invention.

1.2. The Prototype of the Invention

The PPG sensor (KMAX 1.2.0 30101) is interfaced with a high-resolution analogue-to-digital converter (ADC) to digitise physiological signals with high fidelity. The digitised signals are then processed by an AI-based signal quality evaluation module, which continuously analyses multiple input streams in real time and automatically selects the most reliable signal for beat-to-beat blood pressure estimation.

The core intelligence of the system is a deep neural network model trained on large-scale datasets comprising synchronised PPG waveforms and reference blood pressure measurements acquired using established gold-standard devices. Through this training process, the model learns complex spatiotemporal patterns within the waveforms and can infer blood pressure with high accuracy, while remaining robust to user specific characteristics and varying environmental conditions.

A prototype implementation of the invention (Fig. 2) has been developed and is currently undergoing advanced validation. Its principal novel features include:

Multiwavelength Sensor Technology: Integrates advanced multiwavelength photoplethysmography sensors with dynamic, real-time signal selection, continuously assessing and optimising signal quality for superior measurement reliability.

Effortless Positioning: Ensures consistent, high-quality signal acquisition without the need for precise alignment with specific arteries, simplifying device placement and improving user experience.

Beat-to-Beat Monitoring: Delivers continuous, real-time blood pressure estimation on a beat-to-beat basis, enabling detailed tracking of cardiovascular dynamics with every cardiac cycle.

Deep Neural Network Inference: Employs state-of-the-art deep learning architectures to provide robust, high-accuracy blood pressure prediction, even under variable physiological and environmental conditions.

1.3. Study Protocol

This study was designed to evaluate the accuracy of the proposed blood pressure monitoring technology. Participants were enrolled at two sites: the Hong Kong Centre for Cerebrovascular Health Engineering (COCHE), Hong Kong Science and Technology Parks (HKSTP), Science Park, Hong Kong, and Prince of Wales Hospital (PWH), Sha Tin, Hong Kong. The study protocol was reviewed and approved by the Hong Kong Science and Technology Parks Corporation (Approval number: HKSTP CREC ref: 2023 005). All participants provided written informed consent prior to initiation of the validation procedures.

COCHE investigators had full access to all study data and were responsible for data integrity and statistical analysis. Data underlying the study findings are available from the corresponding authors upon reasonable request and from other designated study supporters, subject to applicable ethical and regulatory requirements.

1.4. Statistical Methods and Data Preparation

The neural network at the core of this device is trained on a large, heterogeneous dataset of physiological signals collected from a diverse subject population. The dataset includes photoplethysmography (PPG) waveforms paired with corresponding reference blood pressure measurements recorded across a broad spectrum of physiological states and operating conditions to promote robustness and generalisability. All PPG–blood pressure pairs are temporally aligned with clinically validated measurements obtained using conventional cuff-based sphygmomanometers, which serve as the ground-truth labels for supervised learning of beat-to-beat blood pressure. Before model training, the data undergo comprehensive preprocessing, including signal filtering, normalisation, and artefact suppression, to improve quality and uniformity [19]. Both proprietary

in-house collections and carefully selected public datasets are incorporated to capture wide physiological variability. A custom signal quality index is used to exclude low-quality segments, ensuring that only high-fidelity signals are used to optimise the deep

learning model. By combining stringent quality control with broad coverage of physiological and noise conditions, the resulting network is optimised for reliable, real-time blood pressure estimation across diverse users, environments, and use scenarios.




Fig. 2. The A prototype development of the invention.

2. Methodology

The proposed system architecture is organised as a modular processing pipeline comprising the following key stages:

- Signal Acquisition:** PPG waveforms, which capture pressure-related volume fluctuations at the skin surface, are acquired using wearable sensors in direct contact with the user's body. These analogue signals are then routed to a high-resolution analogue-to-digital converter (ADC), which digitises the data with high precision for subsequent processing.
- Signal Quality Assessment:** A real-time signal quality evaluation module continuously assesses each input channel using predefined metrics, including noise level, signal-to-noise ratio (SNR), periodicity, and waveform continuity. Based on these metrics, the system dynamically selects, at any given moment, the sensor providing the highest-fidelity data for blood pressure estimation.
- Neural Network Processing:** The proposed approach uses two separate neural network models, each built around a hybrid Convolutional Neural Network–Long Short-Term Memory (CNN-LSTM) architecture [20]. One model is dedicated to estimating systolic blood pressure (SBP), while the other is dedicated to estimating diastolic blood pressure (DBP) [21]. This separation is justified by the fact that SBP and DBP arise from different physiological mechanisms and are reflected differently in the photoplethysmography (PPG) waveform [22].

SBP represents the highest arterial pressure reached during cardiac contraction and is therefore more strongly linked to the sharp systolic peak and rapid rising edge of the PPG signal. In contrast, DBP is the lowest arterial pressure during cardiac relaxation

and is more closely associated with the diastolic decay and the dicrotic notch [23]. Using independent models enables each network to learn features that are most relevant to its specific target, rather than forcing a single model to balance gradients from two distinct prediction tasks [24]. Both models share the same overall architecture, as shown in Table 1. The convolutional part includes three 1D convolutional layers, each followed by max-pooling. The first layer uses a relatively large kernel to extract broad, low-frequency morphological information from the PPG waveform, while the later layers use smaller kernels to progressively refine the learned representations [25]. This staged feature extraction reduces the signal dimensionality while preserving the most informative characteristics of the waveform.

Table 1. TCNN-LSTM Model Architecture Summary.

1.	1D Convolutional	Large kernel, followed by max pooling
2.	1D Convolutional	Medium kernel, followed by max pooling
3.	1D Convolutional	Small kernel, followed by max pooling
4.	LSM	Returns the full hidden state sequence to the subsequent layer
5.	LSTM	Returns the final hidden state
6.	Fully Connected (Dense)	96 neurons, ReLU activation
7.	Output Layer	1 neuron producing scalar SBP or DBP prediction

The features produced by the convolutional blocks are then passed to two LSTM layers, which are designed to capture temporal dependencies in the

physiological time series. Because LSTM units can retain information over extended sequences, they are well suited for modelling the quasi-periodic structure of the PPG signal and its relationship to cardiovascular behaviour [26]. The first LSTM layer returns the full sequence of hidden states to the second LSTM layer, allowing the network to build a deeper temporal representation before the final state is passed to the dense layers [27].

After temporal encoding, the extracted features are fed into a fully connected network consisting of two dense layers with 96 neurons each and ReLU activation. A final single-neuron output layer provides the predicted value for either SBP or DBP, depending on the model.

The input to both models is the raw PPG waveform after bandpass filtering between 0.5 Hz and 30 Hz and normalisation. Training is performed using mean squared error (MSE) as the loss function, while root mean squared error (RMSE) is used to monitor performance during training. Hyperparameters are selected through grid search on a separate validation set.

The models were first pretrained on the MIMIC-III open-source physiological database, a large-scale publicly available dataset comprising synchronised PPG waveforms and invasively recorded arterial blood pressure measurements obtained from a diverse clinical population. Following pretraining, the models were fine-tuned using data gathered from internal proprietary studies, enabling adaptation to the specific sensor characteristics, signal acquisition conditions, and target user population of the proposed device. On the MIMIC-III dataset, the pretrained models achieved mean absolute errors (MAEs) of 3.75 mmHg for SBP

and 3.21 mmHg for DBP, and the scatter plot of SBP and DBP estimation error distributions in mmHg, as shown in Fig. 3, demonstrates strong baseline predictive accuracy prior to device-specific fine-tuning. Following pre-training of the proposed network on the MIMIC III dataset, it was fine-tuned on an independent in-house dataset. The fine-tuned models using the a held out subject separated test data of 20 % of the subjects achieved mean absolute errors (MAEs) of 4.58 mmHg for SBP and 3.11 mmHg for DBP and a strong positive linear relationship between the predicted and reference values, yielding correlation coefficients of $r = 0.85$ for SBP and $r = 0.76$ for DBP, as shown in Fig. 4, demonstrating strong baseline predictive accuracy prior to device tuning with the local data sets. After analysing the Bland-Altman plot, it was observed that the data consistently falls when deploying the model directly on the serial plotter data sets, which are considered a classic domain shift, while the data was collected internally.

The proposed model was developed using a two-stage training strategy. First, the CNN-LSTM hybrid neural network was pretrained using the MIMIC-III open-source physiological database, which contains synchronised photoplethysmography waveforms and invasively recorded arterial blood pressure measurements from diverse clinical populations. This pretraining stage enabled the model to learn general physiological relationships between PPG waveform morphology and reference blood pressure values. On the MIMIC-III dataset, the pretrained model achieved mean absolute errors of 3.75 mmHg for systolic blood pressure and 3.21 mmHg for diastolic blood pressure, indicating strong baseline predictive performance under controlled database conditions.




Fig. 3. Scatter plot of actual versus predicted values for the MIMIC-III data set.




Fig. 4. Scatter plot of actual versus predicted values for the validation dataset. The left plot shows SBP, and the right plot shows DBP.

Following pretraining, the model was fine-tuned using internal proprietary datasets collected with the proposed device. This fine-tuning stage was essential for adapting the model to the specific sensor characteristics, acquisition environment, signal quality, and intended user population of the Smart BP Tracker devices [28]. Using the in-house dataset, the fine-tuned model achieved MAEs of 4.58 mmHg for SBP and 3.11 mmHg for DBP. The predicted values also demonstrated a positive linear association with the reference measurements, with Pearson correlation coefficients of $r = 0.85$ for SBP and $r = 0.76$ for DBP. These results suggest that transfer learning from a large public physiological dataset, followed by device-specific fine-tuning, improved the model's ability to generalise to signals acquired from the target wearable platform.

3. Conclusion & Future Work

The findings of this study indicate that the proposed deep learning-based wearable system is well suited for continuous, non-invasive, beat-to-beat blood pressure monitoring, offering clear advantages over traditional intermittent cuff-based methods for applications that require frequent or long-term assessment. By relying solely on PPG signals, the method is simpler, more comfortable, and better aligned with wearable healthcare use cases. Nevertheless, the variability observed in prediction differences underscores the need for further optimisation and rigorous validation. In particular, the proposed network should be evaluated against gold-standard devices for systolic and diastolic blood pressure estimation, with performance assessed in accordance with the Association for the Advancement of Medical Instrumentation (AAMI) criterion of a mean error and standard deviation within 5 ± 8 mmHg [29]. In addition, a substantial proportion of simultaneous measurements should fall within clinically acceptable agreement limits. Future work will therefore focus on improving the efficiency and robustness of the estimation algorithm, especially by addressing domain shift between the MIMIC datasets and the internal data platform. Further studies should also include a larger and more diverse population and examine factors such as skin tone, subgroup performance, measurement position at the wrist or finger, and other external conditions. Overall, the prototype and in-house serial plotter show promise as reliable and clinically applicable blood pressure monitoring tools, with potential to support future smart healthcare systems for early detection, long-term monitoring, and personalised cardiovascular care.

Acknowledgements

This study was supported by the InnoHK initiative of the Innovation and Technology Commission of the Hong Kong Special Administrative Region Government.

References

- [1]. T. G. Pickering, K. Davidson, W. Gerin, J. E. Schwartz, Masked hypertension, *Hypertension*, Vol. 40, Issue 5, 2002, pp. 795-796.
- [2]. S. D. Pierdomenico, F. Cuccurullo, Prognostic value of white-coat and masked hypertension diagnosed by ambulatory monitoring in initially untreated subjects: An updated meta analysis, *American Journal of Hypertension*, Vol. 24, Issue 1, 2011, pp. 52-58.
- [3]. C. Cuspidi, C. Sala, M. Tadic, M. Rescaldani, et al., Untreated masked hypertension and subclinical cardiac damage: A systematic review and meta-analysis, *American Journal of Hypertension*, Vol. 28, Issue 6, 2015, pp. 806-813.
- [4]. T. Kauzlaricova, M. Augustynek, J. Kubicek, Advancements in noninvasive blood pressure measurement: A comprehensive review of cuffed, cuffless, and contactless methods, *IEEE Transactions on Instrumentation and Measurement*, Vol. 75, 2026, 4003031.
- [5]. N. C. Hurley, E. S. Spatz, H. M. Krumholz, R. Jafari, et al., A survey of challenges and opportunities in sensing and analytics for risk factors of cardiovascular disorders, *ACM Transactions on Computing for Healthcare*, Vol. 2, Issue 1, 2020, pp. 1-42.
- [6]. K. B. Kim, H. J. Baek, Photoplethysmography in wearable devices: A comprehensive review of technological advances, current challenges, and future directions, *Electronics*, Vol. 12, Issue 13, 2023, 2923.
- [7]. G. Wang, M. Atef, Y. Lian, Towards a continuous non-invasive cuffless blood pressure monitoring system using PPG: Systems and circuits review, *IEEE Circuits and Systems Magazine*, Vol. 18, Issue 3, 2018, pp. 6-26.
- [8]. Z. Sastimoglu, S. Subramaniam, A. I. Faisal, W. Jiang, et al., Wearable PPG based BP estimation methods: A systematic review and meta-analysis, *IEEE Journal of Biomedical and Health Informatics*, Vol. 29, Issue 4, 2024, pp. 2439-2452.
- [9]. C. El-Hajj, P. A. Kyriacou, A review of machine learning techniques in photoplethysmography for the non-invasive cuff-less measurement of blood pressure, *Biomedical Signal Processing and Control*, Vol. 58, 2020, 101870.
- [10]. R. Mulkamala, G. S. Stergiou, A. P. Avolio, Cuffless blood pressure measurement, *Annual Review of Biomedical Engineering*, Vol. 24, Issue 1, 2022, pp. 203-230.
- [11]. N. Pilz, D. S. Picone, A. Patzak, O. S. Opatz, et al., Cuff-based blood pressure measurement: Challenges and solutions, *Blood Pressure*, Vol. 33, Issue 1, 2024, 2402368.
- [12]. Y. D'Mello, J. Skoric, S. Xu, P. J. Roche, et al., Real-time cardiac beat detection and heart rate monitoring from combined seismocardiography and gyrocardiography, *Sensors*, Vol. 19, Issue 16, 2019, 3472.
- [13]. Y. Shafiti, S. Hughes, W. Taylor, M. A. Imran, et al., Beyond the cuff: State-of-the-art on cuffless blood pressure monitoring, *Sensors*, Vol. 26, Issue 4, 2026, 1243.
- [14]. S. Knight, J. Lipoth, M. Namvari, C. Gu, et al., The accuracy of wearable photoplethysmography sensors for telehealth monitoring: A scoping review, *Telemedicine and e-Health*, Vol. 29, Issue 6, 2023, pp. 813-828.

- [15]. S. Hu, Y. Elsahar, J. Hou, V. M. Dwyer, Multispectral optoelectronic sensor to detect peripheral blood pulsatile variations with equivalent performance in light, medium and dark skin tones, *Scientific Reports*, Vol. 15, Issue 1, 2025, 1546.
- [16]. S. Sakib, Z. Hasan, N. Roy, A state-of-the-art survey of remote photoplethysmography for contactless health parameters sensing, *WIREs Data Mining and Knowledge Discovery*, Vol. 15, Issue 3, 2025, e70039.
- [17]. G. Nie, J. Zhu, G. Tang, D. Zhang, et al., A review of deep learning methods for photoplethysmography data, *arXiv*, 2024, arXiv:2401.12783.
- [18]. F. Carlucci, Optimization of deep neural networks for ppg-based blood pressure estimation on edge devices, PhD Thesis, *Politecnico di Torino*, Turin, 2024.
- [19]. A. Temko, Accurate heart rate monitoring during physical exercises using PPG, *IEEE Transactions on Biomedical Engineering*, Vol. 64, Issue 9, 2017, pp. 2016-2024.
- [20]. L. Saikevičius, V. Raudonis, G. Dervinis, V. Baranauskas, Non-contact vision-based techniques of vital sign monitoring: Systematic review, *Sensors*, Vol. 24, Issue 12, 2024, 3963.
- [21]. L. Wang, W. Zhou, Y. Xing, X. Zhou, A novel neural network model for blood pressure estimation using photoplethysmography without electrocardiogram, *Journal of Healthcare Engineering*, Vol. 2018, Issue 1, 2018, 7804243.
- [22]. J. Ara, Machine-learning models for predicting blood pressure and cardiac function using wearable sensor data, *International Journal of Scientific Interdisciplinary Research*, Vol. 6, Issue 2, 2025, pp. 102-142.
- [23]. D. Lee, H. Kwon, D. Son, H. Eom, et al., Beat-to-beat continuous blood pressure estimation using bidirectional long short-term memory network, *Sensors*, Vol. 21, Issue 1, 2020, 96.
- [24]. S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Computation*, Vol. 9, Issue 8, 1997, pp. 1735-1780.
- [25]. Y. Chen, M. Orlandi, P. Rapa, S. Benatti, et al., PhysioWave: A multi-scale wavelet-transformer for physiological signal representation, in *Advances in Neural Information Processing Systems 38*, *Curran Associates, Inc.*, Red Hook, 2026, pp. 169357-169399.
- [26]. Y. Xu, Z. Wang, X. Liu, Photoplethysmogram (PPG)-based biometric identification using 2D signal transformation and multi-scale feature fusion, *Sensors*, Vol. 25, Issue 15, 2025, 4849.
- [27]. M. Elgendi, Optimal signal quality index for photoplethysmogram signals, *Bioengineering*, Vol. 3, Issue 4, 2016, 21.
- [28]. Z. Meng, X. Yang, X. Liu, D. Wang, X. Han, Non-invasive blood pressure estimation combining deep neural networks with pre-training and partial fine-tuning, *Physiological Measurement*, Vol. 43, Issue 11, 2022, 11NT01.
- [29]. G. S. Stergiou, B. Alpert, S. Mieke, R. Asmar, et al., A universal standard for the validation of blood pressure measuring devices: Association for the Advancement of Medical Instrumentation/European Society of Hypertension/International Organization for Standardization (AAMI/ESH/ISO) Collaboration Statement, *Hypertension*, Vol. 71, Issue 3, 2018, pp. 368-374.

(020)

Fast Electrical Impedance Spectroscopy for Transient Dynamic Processes

Lukas Lauritsch and Marco Da Silva

Johannes Kepler University Linz, Institute of Measurement Technology, Altenberger Straße 69,
 4040 Linz, Austria
 Tel.: +43 732 2468 5924
 E-mail: lukas.lauritsch@jku.at

Summary: This paper presents the development and optimization of a single-channel system for high speed impedance spectroscopy as a basis for future multichannel expansion. The transient freezing of distilled water was investigated using a custom-built needle probe over a frequency range from 1 kHz to 1 MHz. The measured voltage gain magnitude clearly revealed the transition from liquid water to ice. A Debye-based two-point calibration routine was applied to the measured voltage gain magnitude to describe the frequency-dependent electrical response of water and ice. Finally, the maximum acquisition speed was evaluated using a reference RC circuit. A complete sweep comprising 100 frequency points between 10 kHz and 1 MHz was acquired within 12.75 ms, demonstrating the high speed acquisition capability of the developed system.

Keywords: Electrical impedance spectroscopy, Direct digital synthesis, Transimpedance amplifier, Phase transition, Transient dynamic processes, Debye relaxation.

1. Introduction

Impedance spectroscopy offers a promising solution for analyzing transient dynamic processes, such as multiphase flows or chemical reactions. Multi-frequency impedance measurement techniques have already been investigated for frequency-dependent material characterization in the past [1, 2]. This work focuses on the development and experimental validation of a single-channel electronics architecture optimized for rapid magnitude-based acquisition as a basis for future multichannel implementation. To validate the performance of the developed electronics, the transient freezing process of distilled water into ice was investigated as an initial benchmark measurement using a needle probe. In addition, the measurement capability at the maximum acquisition speed was briefly demonstrated using a simple RC circuit.

2. Measurement Electronics

A block diagram of the measurement electronics for electrical impedance spectroscopy is shown in Fig. 1. For signal generation, a direct digital synthesis (DDS) chip (AD5932) executes a frequency sweep using a sine voltage with a constant amplitude of 600 mVpp and zero DC offset. This chip executes frequency sweeps up to 25 MHz with a minimum time interval for each frequency of 40 ns. For instance, capturing a frequency response from 1 kHz to 1 MHz with 100 frequency points requires a theoretical total sweep time of 4 μ s. However, the practical acquisition time is primarily determined by the settling behavior of the analog electronics and the data-acquisition timing. The input voltage V_i is applied to the excitation electrode, while the measurement electrode is connected to a transimpedance amplifier with the

feedback impedance Z_f , consisting of $R_f = 1 \text{ M}\Omega$ and $C_f = 10 \text{ pF}$. The OPA656 is selected as the operational amplifier due to its large gain-bandwidth product of 230 MHz, making it suitable for high-impedance and high-frequency measurements. The output gain of the transimpedance amplifier is determined by

$$\frac{V_o(\omega)}{V_i(\omega)} = -\frac{Z_f(\omega)}{Z_x(\omega)} = -\frac{Y_x(\omega)}{Y_f(\omega)}, \quad (1)$$

where Z_x and Y_x represent the frequency-dependent impedance and admittance of the material between the electrodes, respectively [3]. Consequently, frequency-dependent changes in the material properties result in corresponding changes in the measured voltage gain.

The output voltage of the operational amplifier is subsequently fed into the logarithmic detector (AD8307), which acts as a high dynamic range AC-to-DC converter that compresses the wide signal amplitude variations into a stable DC voltage suitable for precision data acquisition. The detector was calibrated experimentally, using known sinusoidal input amplitudes over the investigated frequency range. The measured output voltage V_m was converted back to the corresponding signal amplitude according to

$$V_o = V_{\text{int}} 10^{V_m/V_Y}, \quad (2)$$

where $V_Y = 0.52 \text{ V/decade}$ is the detector slope and $V_{\text{int}} = 96 \text{ }\mu\text{V}$ is the extrapolated intercept voltage, exemplary here given for 100 kHz.

Finally, this processed DC voltage is captured by a high speed USB data acquisition board (MCC USB-1608GX-2AO-OEM) featuring a 16-bit analog-to-digital converter (ADC) resolution and a sampling rate of 500 kS/s.




Fig. 1. Schematic block diagram of the single-channel impedance spectroscopy measurement electronics.

It should be noted that the proposed system measures only the magnitude of the voltage response, since no phase information is acquired. Therefore, the subsequent analysis is based on the frequency-dependent voltage gain magnitude.

3. Experimental Results

3.1. Freezing Process

As a primary experimental validation, the system was utilized to monitor the transient freezing process of distilled water into ice over a period of 35 minutes. Since this phase transition is a relatively slow process, a measurement interval of 5 s was chosen, and the measurement was conducted using a frequency sweep from 1 kHz to 1 MHz. For the measurement, a custom-built needle probe was employed, which features a double coaxial geometry. The probe consists of two

stainless steel electrodes with diameters of 1 mm and an electrode spacing of 3 mm. The probe was immersed in distilled water inside a glass container. The glass container was placed in a temperature controlled bath (LAUDA ECO RE 415 S) to carry out the freezing process under controlled conditions. During the experiment, the bath temperature was decreased from +10 °C to -12 °C.

The resulting time-dependent voltage gain spectra, calculated using equation (1) and (2), are shown in Fig. 2. In the plotted frequency response, the phase transition from liquid water to ice is clearly observable. At lower frequencies, the transition is particularly pronounced due to the strong decrease in the conductive contribution. Furthermore, a distinct shift is visible in the higher frequency range, reflecting the fundamental change in the dielectric response of the medium as liquid water transforms into the ice structure [2].




Fig. 2. Voltage gain as a function of frequency, plotted at selected time intervals to monitor the phase transition.

3.2. Debye-Based Two-Point Calibration

For a quantitative model-based evaluation of the frequency-dependent electrical response, a Debye-based two-point calibration using air and water was applied to the measured voltage gain magnitude. Since no phase information was measured, the calibration was based on the magnitude response. The frequency-dependent complex permittivity ϵ was

described using the single relaxation time Debye model:

$$\epsilon(\omega) = \epsilon_{\infty} + \frac{\Delta\epsilon}{1+j\omega\tau}, \quad (3)$$

where $\Delta\epsilon = \epsilon_s - \epsilon_{\infty}$, ϵ_s is the low-frequency or static permittivity, ϵ_{∞} is the high-frequency permittivity and τ is the characteristic relaxation time of the medium [4, 5]. The probe admittance was then calculated as

$$Y_{\text{probe}}(\omega) = C_0 \left(\frac{\sigma}{\epsilon_0} + j\omega\epsilon(\omega) \right), \quad (4)$$

where C_0 denotes the empty probe capacitance, and σ is the effective electrical conductivity [6]. Air can be treated as approximately non-dispersive in the investigated frequency range ($\epsilon_s = \epsilon_\infty$) [7]. Consequently, no Debye relaxation term is required for the air reference, and (3) and (4) reduce to:

$$Y_{\text{air}}(\omega) = j\omega C_0 \epsilon_{\text{air}} \quad (5)$$

The empty probe capacitance C_0 was estimated from the low-frequency region of the air measurement. The Debye relaxation parameters assumed for air, water, and ice are summarized in Table 1.

Table 1. Debye relaxation parameters for air, water and ice according to [5, 8, 9].

Parameter	Air	Water	Ice
ϵ_s	1.00059	84	95
ϵ_∞	1.00059	5	3.08
τ (s)	-	8.08×10^{-12}	6×10^{-5}
σ ($\mu\text{S}/\text{cm}$)	0	11.56	5×10^{-3}

Based on the voltage gain relation in (1), a two-point calibration using air and water as reference media was performed as

$$|Y_{\text{probe}}(\omega)| = A(\omega) + B(\omega) \cdot \left| \frac{V_o}{V_i} \right|_{\text{meas}}(\omega), \quad (6)$$

where the frequency-dependent calibration coefficients $A(\omega)$ and $B(\omega)$ were determined from the air and liquid water reference measurement. The predicted ice gain magnitude was then obtained by

inverting (6), using $|Y_{\text{ice}}(\omega)|$ from (3) and (4) with the ice parameters listed in Table 1:

$$\left| \frac{V_o}{V_i} \right|_{\text{model}}(\omega) = \frac{|Y_{\text{ice}}(\omega)| - A(\omega)}{B(\omega)} \quad (7)$$

Fig. 3 compares the measured voltage gain magnitude with the calibrated Debye-based model at two characteristic stages of the freezing process: the initial liquid state ($t = 0$ min) and the fully frozen state ($t = 35$ min). For the liquid reference the calibration reproduces the measurement, confirming the correct implementation of the two-point calibration procedure.

The model reproduces the measured voltage gain magnitude reasonably well over most of the investigated frequency range for ice. A pronounced maximum occurs in the measured voltage gain magnitude at approximately 8 kHz, which is also reproduced by the model and is in good agreement with the dielectric relaxation behavior of ice [5]. Below approximately 30 kHz, the model follows the measured data relatively closely, capturing both the rising slope and the position of the maximum. The average deviation between the modeled and measured gain magnitude for ice is approximately 30 %, with the largest deviations occurring above 30 kHz. In this frequency range, the measured gain continues to decrease, whereas the model approaches a plateau. The observed deviations can be attributed to the use of a magnitude-based two-point calibration in combination with a single-relaxation Debye model. Despite these limitations, the model provides a reasonable representation of the measured response and establishes a promising basis for future automated fitting approaches. By further refining the calibration and model formulation, such approaches could enable the estimation of relevant material parameters from magnitude-only impedance spectroscopy data for the characterization of dielectric changes during phase transitions.




Fig. 3. Measured and modeled voltage gain magnitude for the liquid reference (top) and for ice (bottom).

3.3. Maximum Acquisition Speed

To evaluate the maximum acquisition speed of the measurement electronics, a simple RC circuit was characterized over 100 frequency points between 10 kHz and 1 MHz. A complete frequency sweep required only 12.75 ms. The measurement was repeated ten times and compared with the theoretical transfer function of the RC circuit. As shown in Fig. 4,

the measured mean closely agrees with the theoretical response over the entire frequency range. Furthermore, the results demonstrate excellent repeatability, with a median relative standard deviation of only 0.058 %. Compared with conventional impedance spectroscopy, the achieved short acquisition time demonstrates the potential of the developed system for high speed measurements [1].




Fig. 4. Measured voltage gain of the RC circuit as a function of frequency. The upper plot shows the mean of 10 measurements, the corresponding standard deviation, and the theoretical response. The lower plot shows the relative standard deviation.

4. Conclusion

This work successfully establishes the foundation for high speed acquisition of transient processes using magnitude-based impedance spectroscopy. The water to ice phase transition and the main features of the frequency-dependent electrical response were successfully captured without phase information. In addition, the RC circuit validation demonstrated the acquisition of 100 frequency points from 10 kHz to 1 MHz within 12.75 ms, demonstrating the high speed acquisition capability of the developed system. However, a systematic deviation remained in the predicted ice gain magnitude. Future work will therefore focus on refining the calibration routine to reduce the remaining residual error and on developing automated fitting approaches for the estimation of relevant material parameters. In addition, the system architecture will be extended to a multichannel measurement system to enable the simultaneous investigation of faster transient processes at several points.

References

- [1]. P. Boškoski, A. Debenjak, B. Mileva Boshkoska, Fast Electrochemical Impedance Spectroscopy, *Springer*, Cham, 2017.
- [2]. J. Padilha Leitzke, H. Zangl, A review on electrical impedance tomography spectroscopy, *Sensors*, Vol. 20, Issue 18, 2020, 5160.
- [3]. C. Tan, S. Huang, G. Liang, M. J. Da Silva, et al., Electrical resistance/capacitance dual-mode tomography based on dual-frequency response, *IEEE Transactions on Instrumentation and Measurement*, Vol. 72, 2023, 4506311.
- [4]. D. C. Elton, The origin of the Debye relaxation in liquid water and fitting the high frequency excess response, *Physical Chemistry Chemical Physics*, Vol. 19, Issue 28, 2017, pp. 18739-18749.
- [5]. R. P. Auty, R. H. Cole, Dielectric properties of ice and solid D₂O, *The Journal of Chemical Physics*, Vol. 20, Issue 8, 1952, pp. 1309-1314.
- [6]. P. A. Sorichetti, C. L. Matteo, Low-frequency dielectric measurements of complex fluids using high-frequency coaxial sample cells, *Measurement*, Vol. 40, Issue 4, 2007, pp. 437-449.
- [7]. P. Debye, Polar Molecules, *Chemical Catalog Co.*, New York, 1929.
- [8]. L. G. Hector, H. L. Schultz, The dielectric constant of air at radiofrequencies, *Physics*, Vol. 7, Issue 4, 1936, pp. 133-136.
- [9]. V. G. Artemov, A. A. Volkov, Water and ice dielectric spectra scaling at 0 °C, *Ferroelectrics*, Vol. 466, Issue 1, 2014, pp. 158-165.

Predictive Anemometric Surveillance: Forecasting the IEC 61400-12-1 Annex K Wind-Speed Discrepancy for Proactive Sensor Fault Detection

Máximo Paz Iaconis^{1,2}, Lucas Carmusciano¹ and Juan Manuel Gonzalez Kapnik¹

¹ University of Palermo, Engineering Department, Mario Bravo 1050,
Ciudad Autónoma de Buenos Aires, Argentina

² Polytechnic University of Madrid, Higher Technical School of Industrial Engineers,
Department of Mechanical Engineering – Abascal 2 – Madrid, Spain
Tel.: + 54 9 11 6563-1030
E-mail: iaconis@live.com.ar

Summary: Annex K of the IEC 61400-12-1 standard quantifies the discrepancy between a control and a reference anemometer mounted at the same height on a meteorological mast. It is normally used reactively, to confirm sensor degradation once a warning threshold has been crossed. This work turns the metric into a forecasting target: the growth of the combined discrepancy is predicted three days ahead, so that maintenance can be scheduled before the alarm appears in daily operation. Three families of machine-learning models are compared on a five-height met-mast dataset sampled at ten minutes over 29 months: multiscale temporal windows, multi-height spatial fusion, and a combination of the two. All of them are evaluated under a blocked temporal split with purging and a 35-day embargo, so that no Annex K window is shared between training and test. The 28-day window of the standard introduces a temporal inertia that limits early diagnosis. The best model is a joint multi-output ridge regression over the 28-day discrepancy at all five heights, with a median coefficient of determination of 0.630; the multiscale windows add nothing, because of the multicollinearity between windows of the same signal. Removing the window overlap between partitions costs between 0.01 and 0.19 of that coefficient, depending on the model, and does not alter the ranking.

Keywords: Anemometer fault detection, IEC 61400-12-1 Annex K, Wind-speed discrepancy, Predictive maintenance, Ridge regression, Multi-height fusion, Condition monitoring, Blocked temporal validation.

1. Introduction

Annex K of the IEC 61400-12-1 standard [1] evaluates the consistency between a control anemometer and a reference anemometer placed at the same height of a meteorological mast. In its conventional use the metric is computed reactively, to verify sensor calibration and behaviour, flagging an anomaly only after the discrepancy has already crossed a warning threshold. The combined discrepancy δ is decomposed into a systematic deviation γ and a statistical deviation σ :

$$\gamma = \frac{1}{n} \sum_{i=1}^n (V_{\text{est},i} - V_i), \quad (1)$$

$$\sigma = \frac{\text{std}(V_{\text{est}} - V_i)}{\sqrt{n}}, \quad (2)$$

$$\delta = \sqrt{\gamma^2 + \sigma^2}, \quad (3)$$

where V_{est} is the primary anemometer speed estimated by a linear transfer relation, defined for this study in Section 3.3. According to the IEC power-performance standard, a combined discrepancy above 0.1 m/s denotes a potential drift or early bearing degradation, while a value above 0.2 m/s invalidates the measurements.

This work reformulates Annex K as a proactive task: forecasting the growth of δ at a short horizon so

that maintenance (bearing replacement, soiling correction or de-icing) can be planned before the alarm materialises in daily operation. A companion study by the same authors [7] computed features over a 28-day moving window with an 8-hour stride and trained a Ridge model to predict δ five days ahead at a single height, reporting a coefficient of determination of 0.866 on a two-month test set. Here the five heights are modelled jointly, several multiscale window sets are tested against that single-scale baseline, and the chronological hold-out is replaced by a protocol that controls for leakage. The horizon was also shortened to three days, which improves predictability and is still early enough for preventive scheduling.

Sensor-level fault detection in wind energy has developed along two lines. Model-based diagnosis, built around fault-tolerant control benchmarks, uses analytical redundancy relations and interval observers to isolate actuator and sensor faults [8, 9]. Data-driven condition monitoring grew instead around SCADA records and learned models of normal behaviour [10, 11]. Both lines address the turbine and not the measurement mast that feeds it. Anemometry has its own literature on calibration and uncertainty, including the error introduced by the linear transfer function of IEC 61400-12-1 [12], and wind-resource practice supplies quality-control procedures for mast records [13]. None of this work tries to anticipate the drift, which is the step added here on top of the standard.

2. The Temporal-Inertia Problem

Although the coefficient of determination obtained by the baseline is statistically high, the practical predictive value for early diagnosis is limited. The cause is the temporal scale of the descriptor windows. With a 28-day moving block and an 8-hour stride, the data overlap between consecutive measurements is 98.8 %. This has three critical consequences for the time series: (i) the resulting metric varies extremely slowly, so daily and weekly behaviour is buried under the average of the four preceding weeks; (ii) a severe mechanical failure or abrupt icing at day t takes weeks to be reflected in the discrepancy, so the model predicts an inertial behaviour that lags the real wear; and (iii) the regression has no information about short-term noise dynamics or the velocity of the drift, making rapid ruptures impossible to anticipate.

3. Methods and Alternatives Evaluated

3.1. Dataset and Experimental Setup

All alternatives use a single met mast (tower D231004) with ten cup anemometers mounted in redundant pairs at five heights (40, 60, 80, 100 and 120 m) and sampled every ten minutes. The record covers 1 September 2023 to 2 February 2026, that is 885 days and 127 420 rows. The site is cold and windy: air temperature ranges from -30.5 to 26.1 °C and the median wind speed at 120 m is 11.3 m/s. The location cannot be named for confidentiality reasons.

Annex K metrics are evaluated only at the stride boundaries and not at every 10-minute row: at each 8-hour timestamp they are computed over the raw records contained in the preceding window. This is 48 times cheaper than the row-by-row scheme of the reference implementation, and the values it produces differ only marginally. The aggregation also sets the size of the supervised dataset, which is worth making explicit because that size is much smaller than the raw record count. Three 8-hour blocks per day over 885 days give 2 655 candidate timestamps. Of these, 2 533 survive the 28-day burn-in, the autoregressive lags, the forecast horizon and the removal of windows that fail the Annex K minimum-count criteria. Those 2 533 samples are the dataset seen by every model; the LSTM included; they are not a subsample of the 127 420 raw rows. The horizon is three days in every case.

3.2. Data Quality, Missing Records and Cold-Climate Conditions

The mast records reached us after the provider's own validation, which limits what further quality control can add. No anemometer channel contains null values or zero readings, so the guard that discards rows with a stalled sensor removes nothing on this tower; we kept it because the same pipeline has to run on

masts where it does fire. Missing data appear as absent timestamps and not as gaps to be filled, and no interpolation is performed at any stage: 164 of the 127 584 expected records are missing (0.13 %), spread over twelve days, with a single contiguous outage of 19.5 hours. Since every window is defined by elapsed time and not by sample count, these gaps shorten the affected windows instead of displacing them, and any window left too sparse is discarded by the Annex K minimum-count criteria, which require ten retained samples per window and three per speed bin.

After that, two filters account for most of the data reduction. The IEC speed range of 4 to 12 m/s retains 47.6 % of the records at 120 m, and the discarded fraction is dominated by high winds rather than by calms, which follows from the wind climate of the site. Residuals are then aggregated per 1 m/s speed bin. Extreme residuals are not removed explicitly: 1.8 % of them exceed three standard deviations and 0.89 % exceed 1 m/s in absolute value. Bin averaging attenuates their influence, although it does not remove it, and a robust bin estimator would be a reasonable refinement.

Icing conditions are present in the record and their effect is measurable. Taking air temperature at or below 0 °C together with relative humidity at or above 95 % as a proxy, 1.42 % of the records fall in icing conditions, concentrated in the local winter. Under those conditions the standard deviation of the paired residual at 120 m rises from 0.165 m/s above 5 °C to 0.317 m/s, close to double. Icing therefore inflates the statistical component σ without necessarily indicating degradation, and a spell of rime accretion can look like an emerging fault. This mast carries no de-icing flag or ice-detection channel; dedicated ice sensors and power-curve-based detection are the usual instruments for that purpose [14-16]. We chose not to remove the icing periods, so the target the models are trained on still contains this contamination.

3.3. The Annex K Transfer Relation and Its Baseline Assumption

Annex K compares the measured primary speed against an estimate V_{est} from a linear transfer relation, $V_{\text{est}} = a \cdot V_{\text{ref}} + b$, calibrated over a period in which both instruments are known to be sound. In practice that calibration is hard to justify, because a fault-free baseline can rarely be verified in retrospect. This study uses the identity relation, $a = 1$ and $b = 0$. Each pair consists of two nominally identical cup anemometers at the same height on the same mast, so the expected relation between them is one to one, and no segment of the record could be certified fault-free on independent evidence. The assumption is not free, though. Fitting the relation per height over the first 90 days gives slopes between 0.978 and 0.996 and intercepts between 0.018 and 0.175 m/s (Table 1). At 120 m the fitted relation is almost the identity, while at 80 and 60 m the intercept reaches the same order as the 0.1 m/s warning threshold.

Under the identity assumption δ carries a constant component attributable to mounting geometry and flow distortion, and only its variation over time reflects degradation. This also accounts for a pattern discussed below: the mean discrepancy at 120 m is roughly four times smaller than at the intermediate heights, which says more about the transfer relation of that pair than about the state of its sensors. The forecasting task itself is not affected, since a constant offset shifts features and target alike, but the comparison of absolute values against the IEC thresholds is. Calibrating the relation per height over a documented baseline period would fix this, and it needs maintenance records that we did not have.

Table 1. Annex K transfer relation fitted per height over the first 90 days, against the identity relation actually used, with the mean discrepancy on the test set.

	Slope a	Intercept b (m/s)	Mean δ (m/s)
120 m	0.996	+0.018	0.051
100 m	0.988	+0.079	0.125
80 m	0.978	+0.151	0.138
60 m	0.980	+0.175	0.129
40 m	0.987	+0.115	0.093

3.4. Blocked Validation with Purging and Embargo

A chronological 80/20 split is not sufficient here. Each sample summarises the preceding 28 days and its target is read three days ahead, so consecutive samples share up to 98.8 % of their raw data and a plain split lets the partitions communicate through their windows. Under the naive split the earliest test sample draws on raw records from 7 July 2025, thirty-five days before the boundary of 11 August. Nothing is copied from one partition to the other, but the test features are partly determined by data the model was fitted on, and this is the mechanism that inflates apparent skill in temporally structured data [17, 18].

We therefore adopt a blocked split with purging and embargo [19]. The last nine training samples are purged, so that no training target reads raw data beyond the boundary, and the first 105 test samples are embargoed, so that no test feature reaches back across it. Those 105 steps of 8 hours cover the 28-day Annex K window plus the 6.7 days of sequence history consumed by the LSTM, which has the longest look-back of the models compared. The separation can then be checked directly: training draws on raw records up to 10 August 2025 at 16:00 and testing from 11 August at 08:00, so the two partitions are disjoint in raw data by construction, with a gap of 38.3 days in sample terms. Training retains 2 017 samples and the test set 402, covering 15 September 2025 to 30 January 2026.

A second leak, independent of the first, was also removed. The early-stopping checkpoint of the LSTM had been selected on the test partition; validation is now taken from the final 15 % of the training block. All linear and tree models use scikit-learn [5] and the

LSTM uses PyTorch. Table 2 reports both protocols so that the size of the gap between them can be seen.

3.5. Alternative A – Multiscale Single-Height Features

This alternative keeps a regularised linear estimator [3] or a random-forest estimator [4] but enriches it with windows of 3, 7, 14 and 28 days plus gradient features: discrepancy differentials ($\delta_{3D} - \delta_{28D}$), noise ratios (σ_{3D}/σ_{28D}) and short-term local regression slopes. The short windows are intended to act as an immediate-warning radar while the long ones provide the robust IEC baseline.

3.6. Alternative B – Dynamic LSTM Transfer Model (Proposed, not Evaluated)

Instead of the classical linear fit of Annex K, an LSTM would be trained on the raw 10-minute series to act as a non-linear, dynamic transfer operator: it receives a sequence of reference speeds and climate variables and predicts the instantaneous primary-sensor speed, learning aerodynamic inertia, wake and turbulence effects and thermal interactions [2]. Failure detection would then be performed on the dynamic residual:

$$e_t = V_{\text{real},t} - V_{\text{est},t} \quad (4)$$

Applying the Annex K windows to this residual should give a cleaner indicator, largely free of ambient noise. We describe the design here for completeness but report no results for it: the LSTM evaluated in this work operates on the aggregated feature series and not on the raw signal, and training a sequence model on the raw records is a separate exercise that we leave for future work.

3.7. Alternatives C and D – Multi-Height Fusion

Alternative C exploits the hypothesis that the sensor state at different heights carries orthogonal information: each anemometer pair operates in its own atmospheric regime, so the Annex K metrics of the five heights are genuinely independent predictors. The vertical wind profile follows a power law, so a degraded sensor alters the vertical ratio characteristically:

$$V(h) = V_{\text{ref}}(h/h_{\text{ref}})^\alpha \quad (5)$$

Alternative D combines both axes – four temporal windows per height across the five heights – predicting the five future discrepancies jointly (multi-output). Intra-height windows remain collinear, but inter-height features are nearly orthogonal (cross-correlation 0.14–0.84 versus $\rho \geq 0.97$ within a height), which anchors the regression and prevents it from collapsing onto the 28-day window.

4. Results and Discussion

Table 2 summarises the main configurations under both protocols. Every value was recomputed with a single unified pipeline, so the two columns can be compared directly. The strongest global model is the joint multi-output Ridge baseline over the 28-day discrepancy across the five heights (Alternative D, M1), with a median coefficient of determination of 0.630 on the leakage-controlled test set. Taken over the five heights it does better than the single-height multiscale models of Alternative A and than the per-height fit of Alternative C.

Table 2. Median coefficient of determination (R^2) by configuration under the naive and the blocked protocol (3-day horizon); the mean over the five heights is given for the multi-height models. The companion-study row is shown for reference only and uses a 5-day horizon on a two-month test set.

Model	Heights	naive	blocked	mean
Companion study [7]	120 m	0.866	n/a	n/a
Alt A · M1 Ridge base.	120 m	0.754	0.686	0.686
Alt A · M2 multiscale	120 m	0.632	0.438	0.438
Alt A · M3 Rand. Forest	120 m	0.613	0.602	0.602
Alt C · M4 RidgeCV/h.	5	0.653	0.584	0.595
Alt D · M1 Ridge base.	5	0.672	0.630	0.618
Alt D · M3 RidgeCV mult.	5	0.666	0.581	0.520
Alt D · M4 Rand. Forest	5	0.606	0.513	−0.202
Alt D · LSTM (1L/32u)	5	0.472	0.410	0.397

The multiscale hypothesis is not validated, and controlling for leakage refutes it more clearly than the naive split did. Adding short windows and gradients degrades performance consistently: at 120 m, Alternative A falls from 0.686 to 0.438 once the multiscale block is included. The cause is multicollinearity between windows of the same signal, whose variance-inflation factor tends to infinity and cancels the additional information. Regularised cross-validation drives the short windows almost to zero, so the enriched model converges towards the 28-day baseline while still paying the variance cost of the extra dimensions.

Removing the overlap costs between 0.01 and 0.19 of the coefficient depending on the model, and the configurations that lose most are the ones with most capacity: the Random Forest at 120 m drops from −0.76 to −3.05. The relative ordering of the configurations is not affected, so the inflation was real without changing which model should be preferred.

Table 3 details the per-height behaviour under the blocked protocol. The 120 m level is again the problematic one. Its discrepancy has a cross-correlation below 0.29 with every other height and a mean four times smaller, which follows from the near-identity transfer relation of Section 3.3 and not from a distinct sensor state, and it is also the height that suffers most when the overlap is removed: Alternative

D M1 falls there from 0.647 to 0.399. What keeps the level usable is per-height regularisation. With $\alpha^* = 2043$ the single-scale per-height fit of Alternative C reaches 0.500 at 120 m, above the joint baseline, so at that height the aggregate ordering of the two alternatives is reversed and 120 m is better modelled on its own. The Random Forest collapses there because of a scale divergence between heights. The earlier failure of the LSTM at 40 m turned out to be an artefact of test-set early stopping: with validation drawn from the training block the coefficient recovers from −0.09 to 0.18.

Table 3. Per-height coefficient of determination for the leading models under the blocked protocol, and target autocorrelation at lag 9.

Height	Alt C M4	Alt D M1	Alt D LSTM	ACF(9)
120 m	0.500	0.399	0.384	0.839
100 m	0.700	0.722	0.410	0.821
80 m	0.735	0.783	0.553	0.854
60 m	0.584	0.630	0.461	0.771
40 m	0.454	0.559	0.176	0.738

A coefficient between 0.56 and 0.78 at the four lower heights means the model explains most of the three-day variance of δ under a protocol that shares no data between partitions. For a signal this stationary, with target autocorrelation at lag 9 between 0.74 and 0.85, that is enough for maintenance prioritisation but not for high-confidence automatic alerts. The 80 m and 100 m levels are the best candidates for a semi-automatic early-warning system, while 120 m should be used only as a supporting indicator and preferably through a per-height model.

The comparison with the companion study (0.866 [7]) has to be qualified in two ways. That result came from a two-month test set, while ours covers four and a half months with greater target variance. It was also obtained under a chronological hold-out without purging or embargo, which is the protocol whose effect we quantify here at between 0.01 and 0.19. Part of the gap is therefore attributable to the evaluation protocol and not only to the models, and the two figures should not be read as directly comparable.

4.1. Limitations

Three constraints bound these results. The blocked test window runs from September to January and contains no icing season, so the behaviour of the models under rime accretion, which is the condition that most distorts σ , remains untested. The transfer relation is assumed to be the identity, and this makes the absolute level of δ depend on mounting geometry (Table 1), although the forecasting task itself is not affected. The third constraint is the single mast: the 120 m pattern may well be a property of this installation and not a general feature of upper-level pairs.

5. Conclusions

Reframing the IEC 61400-12-1 Annex K discrepancy as a short-horizon forecasting target is feasible and operationally useful for preventive anemometer maintenance. Across the architectures evaluated here the spatial dimension dominates the temporal one: five heights with a single 28-day window beat four windows at a single height. Four findings survive the leakage-controlled protocol. First, joint multi-output prediction exploits cross-height correlations that per-height training ignores, and gives the best result over the five heights. Second, multiscale windows add no value, because of the intra-height multicollinearity, and the penalty they carry grows once the window overlap between partitions is removed. Third, 120 m is the exception to the first point: per-height regularisation beats the joint model there, and that level should be modelled on its own. Fourth, blocked validation costs between 0.01 and 0.19 of the coefficient of determination without altering the ranking of the configurations. Future work includes validating the 120 m pattern on further towers, calibrating the transfer relation per height over a documented fault-free period, training the dynamic transfer operator of Alternative B on the raw signals, extending the test window over a full icing season, reformulating the task as threshold-crossing classification, and adding wind-vane readings as directional features, which the joint multi-height scheme now makes viable.

Acknowledgements

The authors thank the providers of the meteorological-mast dataset and the wind-resource assessment team for the discussions that motivated this study. We also thank the reviewers of this abstract, whose comments on temporal leakage and on data handling led to the methodology reported here.

References

- [1]. Wind energy generation systems – Part 12-1: Power performance measurements of electricity producing wind turbines, IEC 61400-12-1:2017, *International Electrotechnical Commission*, Geneva, 2017.
- [2]. S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural Computation*, Vol. 9, Issue 8, 1997, pp. 1735-1780.
- [3]. A. E. Hoerl, R. W. Kennard, Ridge regression: Biased estimation for nonorthogonal problems, *Technometrics*, Vol. 12, Issue 1, 1970, pp. 55-67.
- [4]. L. Breiman, Random forests, *Machine Learning*, Vol. 45, Issue 1, 2001, pp. 5-32.
- [5]. F. Pedregosa, G. Varoquaux, A. Gramfort, et al., Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research*, Vol. 12, 2011, pp. 2825-2830.
- [6]. G. Ke, Q. Meng, T. Finley, et al., LightGBM: A highly efficient gradient boosting decision tree, in *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS)*, 2017, pp. 3146-3154.
- [7]. M. P. Iaconis, L. Carmusciano, J. M. Gonzalez Kapnik, Predictive anemometry surveillance: An interpretable ridge regression framework based on IEC 61400-12-1 Annex K, submitted to *International Conference on Novel Intelligent Digital Systems (NiDS)*, 2026 (in print).
- [8]. P. F. Odgaard, J. Stoustrup, M. Kinnaert, Fault-tolerant control of wind turbines: A benchmark model, *IEEE Transactions on Control Systems Technology*, Vol. 21, Issue 4, 2013, pp. 1168-1182.
- [9]. H. Sanchez, T. Escobet, V. Puig, P. F. Odgaard, Fault diagnosis of an advanced wind turbine benchmark using interval-based ARR and observers, *IEEE Transactions on Industrial Electronics*, Vol. 62, Issue 6, 2015, pp. 3783-3793.
- [10]. J. Tautz-Weinert, S. J. Watson, Using SCADA data for wind turbine condition monitoring: A review, *IET Renewable Power Generation*, Vol. 11, Issue 4, 2017, pp. 382-394.
- [11]. A. Stetco, F. Dinmohammadi, X. Zhao, et al., Machine learning methods for wind turbine condition monitoring: A review, *Renewable Energy*, Vol. 133, 2019, pp. 620-635.
- [12]. R. V. Coquilla, J. Obermeier, B. R. White, Calibration procedures and uncertainty in wind power anemometers, *Wind Engineering*, Vol. 31, Issue 5, 2007, pp. 303-316.
- [13]. M. C. Brower, Wind Resource Assessment: A Practical Guide to Developing a Wind Project, *John Wiley & Sons*, Hoboken, 2012.
- [14]. M. C. Homola, P. J. Nicklasson, P. A. Sundsbø, Ice sensors for wind turbines, *Cold Regions Science and Technology*, Vol. 46, Issue 2, 2006, pp. 125-131.
- [15]. N. N. Davis, Ø. Byrkjedal, A. N. Hahmann, N.-E. Clausen, et al., Ice detection on wind turbines using the observed power curve, *Wind Energy*, Vol. 19, Issue 6, 2016, pp. 999-1010.
- [16]. V. Lehtomäki, S. Rissanen, S. Wadham-Gagnon, et al., Available Technologies for Wind Energy in Cold Climates (2nd Ed.), *International Energy Agency (IEA)*, 2018.
- [17]. C. Bergmeir, J. M. Benítez, On the use of cross-validation for time series predictor evaluation, *Information Sciences*, Vol. 191, 2012, pp. 192-213.
- [18]. D. R. Roberts, V. Bahn, S. Ciuti, et al., Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure, *Ecography*, Vol. 40, Issue 8, 2017, pp. 913-929.
- [19]. M. López de Prado, Chapter 7, in *Advances in Financial Machine Learning*, *John Wiley & Sons*, Hoboken, 2018.

Deep Latent Clustering of Kinematic Waveforms for Unbiased Functional Profiling

Shmuel Amour and Sara Rosenblum

Department of Occupational Therapy, University of Haifa, 199 Abba Khoushy Ave,
3498838 Haifa, Israel

E-mail: samour@campus.haifa.ac.il, rosens@univ.haifa.ac.il

Summary: Traditional human functional assessments rely on subjective ordinal scales susceptible to demographic confounders and hardware constraints. This paper presents an unsupervised, device-agnostic framework designed to extract deep latent representations and perform hierarchical clustering of raw kinematic waveforms from standardized physical tasks. Using a self-supervised pre-trained, frozen Time-Series Masked Autoencoder (MAE) Transformer-Encoder [4, 5], we map high-dimensional kinematic waveforms from heterogeneous Wacom and Samsung tablets into a shared, device-invariant 128-dimensional feature space. Crucially, the feature space excludes demographic priors, forcing the network to capture latent biomechanical patterns solely from movement dynamics. We evaluate our pipeline on a complete-case, single-record cohort of $N = 359$ unique participants (54.6 % female) under a fold-isolated preprocessing protocol. Unsupervised Gaussian Mixture Models (GMM) in the 128D space reveal four macro-clusters, validated by a multi-metric framework (Inertia, Silhouette, Davies–Bouldin), resolved via nested sub-clustering into stable, demographic- and hardware-invariant micro-phenotypes, establishing a highly generalizable digital biomarker for clinical stratification.

Keywords: Unsupervised clustering, Kinematic waveforms, Masked autoencoders, Hardware-agnostic, Patient stratification, Functional profiling.

1. Introduction

Traditional functional profiling in geriatric care and occupational therapy relies on manual, discrete, and subjective clinical scales that analyze physical, cognitive, and sensory systems in isolation [8]. These clinical batteries are static and treat functional systems as independent, while standard linear Principal Component Analysis (PCA) fails to capture complex, non-linear motor-compensatory mechanisms [3]. Grounded in the International Classification of Functioning, Disability and Health (ICF) framework [8], functional health is defined as an individual's capacity to execute activities (d-codes) in real-world contexts, yet current tools measure only binary completion, not the dynamic execution process [7].

Clinical stratification remains confounded by demographic variables – biological sex and chronological age – leading to cognitive diagnostic biases such as clinical ageism: clinicians frequently categorize older individuals into high-risk classes purely on age-associated norms, ignoring high-performing motor-compensatory patterns. In parallel, digital health assessments are hindered by hardware-dependent variations (e.g., sampling rate and pressure discrepancies between Wacom and Samsung tablets), introducing device-specific batch effects that prevent generalizable digital biomarkers.

The Cognitive-Motor Coupling framework [7] establishes that overt fine motor output – such as continuous handwriting – directly manifests underlying Executive Functions (EF) resources. Consequently, kinematic waveforms recorded during standardized tasks encode latent biomechanical signatures that reflect functional health status at a

resolution unavailable to traditional scales. Deep learning approaches applied to handwriting kinematic streams have demonstrated sensitivity to motor execution quality and task difficulty [10]. More broadly, self-supervised time-series representation learning has shown that compact latent encodings of raw sensor signals capture discriminative biomechanical structure without requiring labelled data [12, 4], while digital phenotyping pipelines have established the feasibility of deriving clinically meaningful patient subgroups from passively collected movement data [11]. However, existing approaches either depend on supervised labels, remain hardware-specific, or do not explicitly decouple demographic confounders from functional signal – limitations that motivate the present work.

To address these limitations, we propose the Human Functional Profiling (HFP) framework. The core research contribution is threefold: (1) a zero-leakage unsupervised pipeline that excludes all demographic and device priors from the learned representation; (2) a hierarchical clustering paradigm that resolves latent space into clinically meaningful micro-phenotypes; and (3) a formal operationalization of the Person-Environment-Occupation (PEO) model [14] as a continuous, demographic-invariant framework for representing occupational performance capacity.

2. Methods

2.1. Dataset and Kinematic Task

The clinical dataset consists of unique adults (age range 19–94 years; 54.6 % female, 45.4 % male)

performing a standardized continuous handwriting task mapped to ICF activity code **d170** (Writing) [8]. Participants were recorded on two heterogeneous digitization platforms: Wacom Intuos Pro ($n = 214$) and Samsung Galaxy S10 ($n = 145$), introducing systematic hardware batch effects in sampling rate and pressure resolution. Each participant contributes exactly one record, yielding a zero repeated-measures design that eliminates trial-level group leakage.

The dataset encompasses five structurally distinct writing tasks. The Samsung Tablet cohort ($n = 145$) performed a continuous Hebrew prose passage. Four Wacom task variants are included: **dic5** (identical prose, elderly participants aged 70–94 years, $n = 21$); **dic6** (a different prose passage of comparable length, $n = 80$); **dic7** (sequential inscription of all 22 Hebrew alphabet letters, alef-bet, $n = 80$); and **dic8** (prose from a validated Developmental Coordination Disorder (DCD) diagnostic instrument, young adults aged 19–42 years, $n = 62$). The inclusion of the alphabet task (dic7) is methodologically significant for representation learning: letter-by-letter inscription imposes obligatory pen lifts, directional reversals, and allographic switching – a motor programme qualitatively unlike continuous prose. Exposure to this diversity of kinematic regimes during pre-training widens the encoder’s coverage of the biomechanical manifold and prevents collapse to a prose-specific representation.

The cohort of $N = 359$ substantially exceeds the range typical of digitizer-based handwriting kinematic studies, where published cohorts commonly span $N = 50$ – 200 participants [15, 10]. For unsupervised deep representation learning, sample adequacy is governed by cluster structural validity rather than parametric effect-size power; the convergent multi-metric validation framework (silhouette = 0.10, consistent with a continuous functional manifold in real-world clinical data [1]) confirms stable separation at the identified $K = 4$ solution.

The raw kinematic signal is a 7-dimensional time-series capturing: spatial coordinates (x, y), pen pressure, tilt angle, azimuth (in-plane rotation), inter-sample time delta (Δt), and a binary on-surface indicator (*is_touching*) at each temporal index. Multi-device hardware disparities are normalized using local feature scaling (StandardScaler) prior to Transformer projection, establishing a device-invariant input space.

Kinematic Drift

Kinematic Drift is defined as the row-wise ℓ_2 Euclidean distance between a participant’s standardized latent vector $\tilde{x}_i$ and the centroid μ_k of their assigned macro-cluster k :

$$KD_i = \|\tilde{x}_i - \mu_k\|_2$$

This scalar metric operationalizes intra-cluster biomechanical instability: high Kinematic Drift

reflects atypical, irregular execution dynamics relative to the cluster’s functional phenotype, while low values indicate smooth, stable motor trajectories. The metric is computed exclusively in the 128D feature space, preserving full geometric fidelity.

2.2. Self-Supervised Pre-Training

To extract robust biomechanical structure without manual feature engineering, we employ a self-supervised, frozen Time-Series Masked Autoencoder (MAE) Transformer-Encoder architecture [4-6]. The Transformer Encoder consists of two stacked multi-head self-attention blocks (head size = 64, $H = 4$ attention heads, feed-forward dimension = 128, dropout = 0.1), projecting to a 32-dimensional latent bottleneck via a linear dense layer.

During self-supervised pre-training, raw kinematic inputs are masked at a stochastic 40 % ratio (random temporal indices resampled per iteration), and the network is trained to reconstruct original physical values by minimizing a Mean Squared Error (MSE) loss on masked indices:

$$L_{pretrain} = \frac{1}{|M|} \sum_{t \in M} \left| x_t - \hat{x}_t \right|^2,$$

where M is the set of masked temporal indices, x_t is the original physical vector, and $\hat{x}_t$ is the reconstructed value. The architecture itself was not modified from the reference design [4]; only its parameters were optimized via the Adam optimizer (default $\eta = 10^{-3}$, batch size = 64). Pre-training operated on 67978 temporal windows extracted from the $N = 359$ cohort recordings (50-timestep window ≈ 5.5 seconds; 80/20 split yielding 54330 train / 13648 test sequences), with no demographic or device labels entering the loss function. The sole training signal is temporal self-reconstruction of masked kinematic values, establishing a zero-leakage barrier with respect to functional labels: the encoder learns biomechanical dynamics, not participant identity. Training converged at epoch 38 of 50, achieving a final champion MSE of **0.0311** (see Fig. 2), with multi-seed robustness confirmed over 5 independent random initializations ($S = \{42, 101, 777, 2026, 9999\}$).

Ablation: Transformer vs. LSTM Baseline

To justify the architectural choice, a comparative ablation was conducted against a deep stacked LSTM Autoencoder across latent dimensions $\text{Dim} \in \{16, 32, 64\}$ (Fig. 1). The Transformer achieves a 21.5 % reduction in reconstruction MSE relative to the best LSTM baseline ($\text{Dim} = 16$: 0.03963 vs. 0.03110; Table 1).

2.3. Latent Space Clustering and K Selection

The frozen Transformer-Encoder maps all $N = 359$ participants into a 128-dimensional statistical feature

space. Specifically, the 32-dimensional latent bottleneck is summarised per participant across temporal windows using four per-dimension statistics (mean, Standard Deviation (SD), 90th percentile, 10th percentile), yielding a $32 \times 4 = 128$ feature matrix $X \in \mathbb{R}^{359 \times 128}$. To avoid the “Uniform Manifold Approximation and Projection (UMAP) clustering antipattern” – where clustering is applied to non-linearly compressed coordinates that distort global distances – we enforce a strict mathematical separation: all clustering, density estimation, and validation are computed exclusively in the original 128D feature space (X_{scaled}). UMAP is deployed strictly downstream as a passive 2D visualization layer without influencing any metric. The optimal number of macro-clusters was determined via a multi-metric internal validation sweep across $K = 2$ to $K = 8$, computed directly in 128D (Fig. 3): (1) the **Inertia Elbow Curve** shows a clear inflection at $K = 4$; (2) the **Silhouette Structural Density Profile** peaks at $K = 2$ but plateaus with a secondary elbow at $K = 4$, yielding a mean score of 0.10 (weak structure, consistent with a continuous functional manifold – expected in real-world clinical motor data where phenotypic boundaries are gradational rather than discrete [1]); and (3) the **Davies–Bouldin Separation Index** [13] (lower is better) shows a local optimum at $K = 4$ before further marginal improvement. The convergent agreement of all three independent metrics unambiguously supports $K = 4$ [2]. We implement a Gaussian Mixture Model (GMM) with a full covariance matrix Σ_k to estimate soft membership probabilities in the 128D space. To uncover fine-grained profiles, a nested, energy-driven hierarchical sub-clustering algorithm is applied to each of the four locked macro-clusters. This nested phase operates on locally standardized latent coordinates $\tilde{X}_k = \frac{(X_k - \mu_k)}{\sigma_k}$, reducing between-cluster demographic confounds: by re-centering and re-scaling within each cluster’s local coordinate system, the residual variance is driven primarily by biomechanical structure rather than cohort-level demographic composition.

Table 1. Ablation: Transformer vs. LSTM Reconstruction Error.

Architecture	Config	Val. MSE	Δ vs. Champion
LSTM	Dim=64	0.04491	+44.4%
LSTM	Dim=32	0.04195	+35.0%
LSTM	Dim=16	0.03963	+27.4%
Transformer	Dim=32	0.03110	—

2.4. Post-Hoc Clinical Label Assignment

Clinical micro-phenotype labels (e.g., *Successful Aging*, *Acute Breakdown*, *bradykinesia*) were assigned through a post-hoc interpretive protocol combining three criteria: (1) the mean and distribution of Kinematic Drift within each sub-cluster; (2) the age

and sex composition relative to the full cohort; and (3) established clinical literature on motor signatures in occupational therapy [7]. Crucially, these labels are *descriptive interpretations* of the data-driven clusters and carry no causal clinical diagnosis – they represent hypotheses for future external validation.




Fig. 1. LSTM baseline sterile evaluation profile (Dim = 16, 32, 64). All LSTM configurations exhibit higher validation MSE and greater instability than the champion Transformer (MSE = 0.0311).




Fig. 2. Champion Transformer Autoencoder convergence profile. Training and cross-validation MSE losses across epochs. The locked champion target MSE = 0.0311 is achieved at epoch 38, demonstrating stable convergence.

3. Results

3.1. Cohort Characterization and Writing Task Distribution

Fig. 4 presents the full post-hoc characterization of the four macro-clusters. Sex distribution is near-uniform (50.0–57.8 % female) and mean ages are tightly clustered (50.7–54.5 years), confirming that neither demographic variable drove cluster assignment. Writing task composition varies systematically – an emergent kinematic consequence – as detailed in Section 3.5.

3.2. Demographic Invariance

As confirmed by Fig. 4 (top panel), age ($\mu = 50.7$ –54.8 years, overlapping interquartile ranges) and sex (50.0–57.8 % female) are uniformly

distributed across all four macro-clusters. This finding directly contradicts the prevailing clinical OT paradigm, in which age and sex serve as first-order stratifiers: established normative benchmarks for handwriting kinematics are stratified by age and sex, with older adults expected to show lower writing speed and higher stroke variability [15]. By excluding all demographic signal from the reconstruction loss, the present encoder forces the 128D feature space to represent pure biomechanical structure, enabling

cross-age-boundary phenotyping consistent with recent unsupervised motor analysis [16]. The most striking illustration is Sub-ID 1_3 (Successful Aging, $N = 59$, median age 66.0 years, $KD = 7.82 \pm 2.27$), which achieves the lowest Kinematic Drift in the cohort despite being the oldest sub-group; Sub-ID 1_1 (Acute Breakdown, $N = 6$, mid-life) exhibits the highest instability ($KD = 15.50 \pm 4.81$) – a pattern wholly invisible to age-stratified normative scoring.




Fig. 3. Multi-metric validation framework in the 128D feature space. All three metrics (Inertia Elbow, Silhouette Profile, Davies–Bouldin Index) converge on $K = 4$ (red dashed line) as the optimal macro-cluster solution.

3.3. Micro-Phenotype Characterization

Table 2 presents the complete characterization of the 12 micro-phenotypes discovered by the nested sub-clustering pipeline, demonstrating complete demographic and hardware invariance.

Table 2. Kinematic Micro-Phenotypes within Discovered Macro-Clusters ($N = 359$).

Macro	Sub-ID	N	Age ($\mu \pm SD$)	Med. Age	KD ($\mu \pm SD$)
0	0.0	26	47.15 $\pm$ 16.53	46.5	9.74 $\pm$ 2.87
	0.1	30	49.60 $\pm$ 23.43	44.5	8.26 $\pm$ 2.09
	0.2	28	57.42 $\pm$ 21.65	62.0	9.37 $\pm$ 2.40
1	1.0	12	50.50 $\pm$ 22.49	50.5	10.75 $\pm$ 2.38
	1.1	6	55.50 $\pm$ 22.39	52.5	15.50 $\pm$ 4.81
	1.2	25	49.80 $\pm$ 24.34	41.0	9.92 $\pm$ 2.45
	1.3	59	57.27 $\pm$ 21.31	66.0	7.82 $\pm$ 2.27
2	2.0	35	52.34 $\pm$ 21.02	52.0	10.30 $\pm$ 3.62
	2.1	57	49.66 $\pm$ 22.62	46.0	9.30 $\pm$ 2.90
3	3.0	23	50.52 $\pm$ 18.79	42.0	9.47 $\pm$ 2.90
	3.1	17	45.65 $\pm$ 17.30	45.0	11.12 $\pm$ 3.05
	3.2	41	53.88 $\pm$ 20.02	66.0	8.02 $\pm$ 1.91

Cluster 0 sub-clusters must be interpreted with caution. Post-hoc task analysis reveals that 38.1 % of Cluster 0 participants performed the dic7 alphabet task – 1.8 times the cohort-level rate of 20.9 % – compared to only 11.1 % in Cluster 3. Because the alef-bet task imposes a structurally discrete motor programme (letter-by-letter strokes, frequent pen lifts, obligatory directional reversals) qualitatively different from continuous prose, the moderately elevated Kinematic Drift observed in **Sub-IDs 0_0** and **0_2** ($KD = 9.74 \pm 2.87$ and 9.37 ± 2.40 , respectively) may reflect

task-specific motor complexity rather than functional impairment; these sub-phenotypes are reported as preliminary hypotheses requiring task-controlled replication.

The most structurally robust sub-clusters are those with balanced task representation and cross-diagnostic clinical composition. **Sub-ID 1_3** (Successful Aging, $N = 59$, Median Age = 66.0 years) draws participants from all five writing tasks and achieves the lowest Kinematic Drift in the cohort ($KD = 7.82 \pm 2.27$), indicating preserved motor resilience in late adulthood. Post-hoc labeled participants in this sub-cluster span MCI, mild Alzheimer’s disease, and DCD – all exhibiting low instability – confirming that the model groups by functional kinematic similarity rather than diagnosis. Cluster 2 provides strong evidence for the Cognitive-Motor Coupling hypothesis [7]: **Sub-IDs 2_0** ($N = 35$, $KD = 10.30 \pm 3.62$) and **2_1** ($N = 57$, $KD = 9.30 \pm 2.90$) co-locate participants from the elderly cognitive decline spectrum (dic5: MCI/Mild AD) with young DCD adults (dic8) – two clinically distant populations converging on overlapping kinematic execution profiles in the latent space. **Sub-ID 1_1** (Acute Breakdown, $N = 6$, $KD = 15.50 \pm 4.81$) isolates the highest motor instability at mid-life but requires replication given its small size. **Sub-ID 3_1** (Early-Onset Volatility, $N = 17$, $KD = 11.12 \pm 3.05$) captures premature motor breakdown in younger adults.

3.4. Diagnostic Value of the Latent Space

The four macro-clusters do not represent diagnostic labels. Rather, they constitute a continuous functional

manifold aligned with the PEO model [14, 9]: each coordinate in the 128D space encodes occupational performance capacity as an intrinsic biomechanical property, free from demographic or hardware confounders. The clinical interpretability of micro phenotypes – supported by the Kinematic Drift metri – provides a scalable and objective reference framework for longitudinal monitoring and early detection of functional decline. External validation against established clinical assessment tools remains the focus of subsequent dissertation stages scalable and objective reference framework for longitudinal monitoring and early detection of functional decline. External validation against established clinical assessment tools remains the focus of subsequent dissertation stages.

3.5. Post-Hoc Zero-Knowledge Characterization

The entire clustering pipeline operated without access to any metadata; Fig. 4 visualises the resulting post-hoc distributions. Demographic uniformity is confirmed: age and sex overlap across all four clusters (see figure, top panel). Task composition varies as an emergent imprint of kinematic structure – Cluster 0 is enriched 1.8× for the structurally discrete Hebrew alphabet task (38.1 % vs. 20.9 % expected) while

Clusters 1 and 2 approximate the cohort-level task mixture, reflecting balanced, task-agnostic functional groupings. Among the 97 clinically labeled participants (27 %; labels withheld during training), DCD diagnoses span all clusters (54–70 %), and Cluster 2 uniquely co-locates cognitive decline labels (MCI/Mild AD, 11.1 %) with the DCD population – consistent with the Cognitive-Motor Coupling hypothesis that distinct aetiological pathways can converge on similar kinematic execution profiles. These distributions are post-hoc and interpretive; they are presented to illustrate the functional signal recoverable from a purely biomechanical latent space.

4. Conclusions

The main contribution of this study is an unsupervised, leakage-free pipeline that establishes a generalizable, objective, and unbiased digital biomarker of human functional status. This framework uniquely integrates heterogeneous tablet hardware through device-invariant latent projection, and delivers fine-grained clinical interpretability through nested hierarchical sub-clustering [3]. The Kinematic Drift metric provides a formal, mathematically grounded scalar index of motor instability within a functional phenotype group.




Fig. 4. Post-hoc characterization of the four macro-clusters (N = 359). Top: Cohort demographics – age and sex are near-uniform across clusters, confirming neither demographic drove assignment. Middle: Writing task composition; Cluster 0 is enriched 1.8× for the alphabet task (dic7: 38.1 % vs. 20.9 % cohort rate), an emergent kinematic signature. Bottom: Clinical label distribution among the 97 labeled participants (27 %; labels withheld during training). Cluster 2 uniquely co-locates cognitive decline (MCI/Mild AD) and DCD, consistent with Cognitive-Motor Coupling. All distributions are post-hoc observations from the zero-leakage pipeline.

Limitations include the cross-sectional, task-specific design, restricting inference to a single hand-writing task without longitudinal tracking. While

N = 359 exceeds domain precedents for digitizer-based kinematic studies, multi-site replication with larger independent cohorts is

warranted to test the generalizability of the discovered micro-phenotype structure. The post-hoc nature of clinical labels (e.g., bradykinesia, Successful Aging) constitutes interpretive hypotheses requiring prospective clinical validation. Sub-ID 1_1 (N = 6) is reported for structural completeness but should be interpreted with caution given its very small sample; stability of this sub-cluster requires replication in larger cohorts. Future work will focus on: (1) external validation of micro-phenotypes against established clinical and patient-reported assessment tools; (2) extension to multi-task longitudinal kinematic protocols; and (3) enrichment of the latent space with complementary data streams within the broader HFP research programme.

References

- [1]. P. J. Rousseeuw, Silhouettes: A graphical aid to the interpretation and validation of cluster analysis, *Journal of Computational and Applied Mathematics*, Vol. 20, 1987, pp. 53-65.
- [2]. L. McInnes, J. Healy, N. Saul, L. Großberger, UMAP: Uniform manifold approximation and projection, *Journal of Open Source Software*, Vol. 3, Issue 29, 2018, 861.
- [3]. K. Merkelbach, S. Schaper, C. Diedrich, S. J. Fritsch, et al., Novel architecture for gated recurrent unit autoencoder trained on time series from electronic health records enables detection of ICU patient subgroups, *Scientific Reports*, Vol. 13, Issue 1, 2023, 4053.
- [4]. G. Zerveas, S. Jayaraman, D. Patel, A. Bhamidipaty, et al., A transformer-based framework for multivariate time series representation learning, in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2021, pp. 2114-2124.
- [5]. K. He, X. Chen, S. Xie, Y. Li, et al., Masked autoencoders are scalable vision learners, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 16000-16009.
- [6]. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, et al., Attention is all you need, in *Advances in Neural Information Processing Systems 30* (I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, et al., Eds.), *Curran Associates, Inc.*, Red Hook, 2017, pp. 5998-6008.
- [7]. S. Rosenblum, Handwriting measures as reflectors of executive functions among adults with developmental coordination disorders (DCD), *Frontiers in Psychology*, Vol. 4, 2013, 357.
- [8]. International Classification of Functioning, Disability and Health: ICF, *World Health Organization*, Geneva, 2001.
- [9]. S. Rosenblum, Towards daily function homeostasis: A conceptual health framework and keys for action, *Journal of Family Medicine and Disease Prevention*, Vol. 3, Issue 2, 2017, 053.
- [10]. V. Babushkin, H. Alsuradi, M. H. Jamil, M. O. AlKhalil, et al., Assessing handwriting task difficulty levels through kinematic features: A deep-learning approach, *Frontiers in Robotics and AI*, Vol. 10, 2023, 1193388.
- [11]. S. D. Dlima, S. Shevade, S. R. Menezes, A. Ganju, Digital phenotyping in health using machine learning approaches: Scoping review, *JMIR Bioinformatics and Biotechnology*, Vol. 3, Issue 1, 2022, e39618.
- [12]. E. Eldele, M. Ragab, Z. Chen, M. Wu, et al., Time-series representation learning via temporal and contextual contrasting, in *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 2021, pp. 2352-2359.
- [13]. D. L. Davies, D. W. Bouldin, A cluster separation measure, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 1, Issue 2, 1979, pp. 224-227.
- [14]. M. Law, B. Cooper, S. Strong, D. Stewart, et al., The person-environment-occupation model: A transactive approach to occupational performance, *Canadian Journal of Occupational Therapy*, Vol. 63, Issue 1, 1996, pp. 9-23.
- [15]. N. van Drempt, A. McCluskey, N. A. Lannin, A review of factors that influence adult handwriting performance, *Australian Occupational Therapy Journal*, Vol. 58, Issue 5, 2011, pp. 321-328.
- [16]. M. Pfaff, M. Casteran, Identification of whole-body reaching movement phenotypes in young and older active adults: An unsupervised machine learning approach, *bioRxiv*, 2024, 10.1101/2024.07.09.600023.

Simulation-Based Study of a GMR Sensor Device for Event-Driven Tracking of Magnetic Microparticles

A. Shahryari¹, F. Pardo¹, J.A. Boluda¹, C. Reig¹, R. Afonso² and S. Cardoso²

¹ Universitat de València, ETSE, 46100 Burjassot, Spain

² INESC-MN, 1000-029 Lisboa, Portugal

E-mail: candid.reig@uv.es

Summary: We report an event-driven magnetic particle tracking system combining Giant Magnetoresistance (GMR) sensor arrays with microfluidic channels and a Complementary Metal-Oxide Semiconductor (CMOS) circuit implementing Selective Change Driven-Address Event Representation (SCD-AER) logic. An 8×1 GMR array, fabricated with standard clean-room techniques, acts as a linear magnetic pixel array. Simulations using synthetic stimuli and scaled real particle-flow signals confirm robust detection of particle-transit events with inherent noise rejection. This work demonstrates the feasibility of event-driven GMR-based cytometry for high-throughput, low-power point-of-care diagnostics.

Keywords: Giant magnetoresistance sensors, Magnetic particle tracking, Event-driven approach, Address event representation, Microfluidics, Point-of-care diagnostics.

1. Introduction

Tracking of functionalized magnetic particles is a common technique in point-of-care diagnostics. Magnetic nanoparticles or microparticles are coated with biological recognition molecules (antibodies, aptamers, or nucleic acids) that selectively bind to a target analyte (a pathogen, biomarker, or protein) [1]. A widely adopted strategy is to detect, quantify, or manipulate these particles using magnetic fields and sensors integrated within microfluidic circuits. Due to the required integration scale and the low magnetic fields involved, solid-state-based magnetic sensors are preferred in these applications [2, 3].

When the number of particles increases, conventional acquisition schemes saturate, and different approaches need to be considered. In this sense, event-based vision sensors have solved equivalent problems related to generic particle tracking [4].

With relation to previous work, the article builds on our SCD-GMR work along two complementary axes. It covers a finite element model of a magnetic micro-particle tracking device, based on a linear array of GMR sensors attached to a microfluidic channel, although the extraction of particle trajectories from the sensor array's response was left as an open problem [5]. This work focused on the design, fabrication, and circuit-level characterization of individual GMR sensor topologies and their CMOS interfaces, while our 2023 modelling paper [6] provided a Verilog-A dynamic compact electrical model of the GMR sensor for standard IC design tools; neither addressed the sensor as part of a complete particle-tracking system.

The present paper extends that sensor-level characterization to a complete particle-tracking system, demonstrating its response to both synthetic square-pulse stimuli and pseudo-real particle-flow signals derived from experimental GMR-cytometer

measurements, and quantitatively evaluating the SCD-AER logic's ability to suppress noise while correctly recording particle-transit events.

We proposed the application of address representation-based algorithms to the tracking of magnetic particles, based on linear arrays of giant magnetoresistance (GMR) sensors attached to microfluidic chips. These sensors act as magnetic pixels, connected to a CMOS integrated circuit implementing the event-driven logic. The present work is, to our knowledge, the first to combine these elements into a characterized event-driven device, demonstrating the capability of tracking magnetic microparticles with a functional GMR-based system inspired by SCD-AER pixels.

2. Methodology

GMR sensors, developed with standard clean-room techniques, are linearly distributed (8×1) in a substrate, and then attached to specifically fabricated microchannels. These sensors should be externally connected to the ASIC [7], including the IC with the event-driven based design [8], as indicated in Fig. 1. For the system's conversion device, we have considered a particular approach of the address event representation (AER) protocol denominated selective change driven (SCD). An SCD device gives, when required, the coordinates and level of the pixel with a higher variation. Initially developed for vision sensors, this architecture has recently been suggested for magnetic imaging by replacing photodiodes by elemental magnetic sensors [5, 9].

As described in the title, this work features a simulation-based study, making use of Python's Magpylib library. This tool is very convenient for our research. It gives us the benefit of building our magnetic cytometry system, implementing the particle

dynamics and interactions with the array with all the formulas in a very precise and controlled environment, avoiding the computational weight carried by more powerful Multiphysics software packages. All simulation parameters are chosen to mimic real flow cytometry environments. By doing so, we can compare our simulated data with experimental results.




Fig. 1. Scheme of the proposed system. 8-element GMR sensor connected to the event-driven IC CMOS.

At first, we modelled a Dynabeads M-280 streptavidin-coated bead – with parameters taken from [13] (2.8 μm diameter, magnetic moment 1.31×10^{-13} A·m² for applied fields exceeding 200 mT, pp 194) – as a proxy for a labelled bacterium, represented as a uniformly magnetized sphere (for

which the exterior field is exactly that of a point dipole at its centre) swept past the linear 8-element GMR array. For the target microfluidic chip (Fluidic 138, 200 μm -deep channel), we take the minimum and maximum sensor-particle distances as the particle radius and the channel height minus that radius, i.e. the particle resting at the channel floor or at the opposite wall of the microfluidic channel.

3. Simulation Results

This section presents the simulation results obtained to validate the proposed event-driven GMR sensing chain, complementing the models and setup described in Section 2.

3.1. Response to Synthetic Stimuli


To test the proper working principle of our ASIC, we prepared a preliminary test against synthetic square-pulse stimuli with amplitudes of 10 mV and 1 mV. The 8 signals are generated in Python and saved as data files. These are then fed consecutively to each of the GMR sensors as custom inputs. All configurations are done in an “ocean” (.ocn) script using Cadence’s “SKILL” programming language and loaded via “CIW”, Cadence’s Interpreter Window.

As observed in Fig. 2, our simulated ASIC can properly detect and generate the associated events as a response of the square stimuli inputs.

a) Clock



b) 10 mV - pk



c) 1 mV - pk




Fig. 2. Device response under square-pulse excitation.

3.2. Response to Pseudo-Real Particle-Flow Signals

After preliminary verification through simulation, we then tested the SCD-AER processing logic against real particle-flow signals extracted from the GMR array-based cytometer developed at INESC Microsistemas e Nanotecnologias (INESC-MN), Lisbon, as reported in [11]. Since these signals originate from a different acquisition setup, we applied a software gain to bring their amplitude within the input range of our simulated system; this test therefore validates the correct functional behavior of the event-driven detection and data-reduction logic under realistic event waveforms, rather than the absolute signal levels expected from the physical device.

Obtained simulated results are shown in Fig. 3. As observed, the events coming from particle flow are properly processed, discarding the unwanted data, while the noise components are naturally filtered as a result of the ED architecture use.

3.3. Data-Volume Efficiency

A key motivation for the SCD-AER approach is the reduction of the data volume that must be transmitted and processed, compared with conventional continuous-sampling acquisition. We compared the data volume generated by our event-driven system against three recently reported GMR- and impedance-based particle-counting systems [10-12], by calculating the data volume each system generates over the 9.6 ms time window. Table 1 summarizes the comparison, together with the real number of parallel measurement channels reported for each system, so that the structural differences between the architectures being compared are explicit rather than hidden behind a normalization.

This comparison is useful to characterize the capabilities of our system, but the actual data volume generated in a given test will depend on the real particle flux, since the number of particles crossing the array within the fixed measuring window cannot be controlled between iterations; this window is set so that each particle transit yields a single recorded event.

It should be stressed that this comparison is scoped specifically to the particle-counting task addressed in this article, and that it compares raw generated data volume, not overall system capability. The three reference systems perform substantially more than particle counting: the platform in [11] additionally implements automated bacteria screening on clinical samples; Caetano et al. [12] apply an on-board artificial neural network to classify the acquired magnetic signatures; and Chou et al. [10] integrate real-time impedance-based size classification directly on-chip. Still, the comparison remains a meaningful way to quantify the structural advantage of event-driven acquisition over classical, continuous sampling. Even using each reference system's own reported bit depth rather than a uniform assumption, the

event-driven architecture reduces the generated data volume by close to three orders of magnitude relative to the most modest reference system [12], close to four orders of magnitude relative to [11], and by more than five orders of magnitude relative to the highest-throughput reference system [10].

The number of channels listed for each reference system is taken directly from its source publication rather than assumed uniformly, since the architectures differ substantially: [11], for instance, uses six real channels (three GMR sensors replicated across two microfluidic channels, operated simultaneously at 200 kSps/channel), compared with the eight channels of our own GMR array. From this channel count, the samples per channel are obtained directly from each system's reported acquisition rate over the 9.6 ms measuring time, independently of the channel count; the total samples or events per row then follows as the product of samples per channel and the number of channels. Our own event-driven system does not have a fixed per-channel sample rate, since a sensing element only produces an event when a particle transit is detected, so its samples-per-channel entry is instead reported as an average across the array (8 events over 8 channels).

The 9.6 ms measuring time is the same acquisition window used throughout this comparison, already introduced in Fig. 3, and is listed explicitly so that the total-bytes column can be read directly as the volume of data each system would generate over that interval, calculated from its reported acquisition rate. Total bytes are obtained by converting each row's total sample or event count at the bit depth reported for that system (bytes = samples $\times$ bits/sample $\div$ 8): 8 bits for [10], and 16 bits for [11] and [12]. Our own system's 4 bytes follow from its 4-bit-wide event code (colbits + 1): 8 events, each encoded with 4 bits, give $8 \times 4 = 32$ bits, or 4 bytes in total. Those 4 bits suffice to address the 9 possible states of the array.

4. Conclusions and Future Work

We have presented a simulation-based study of an event-driven device based on GMR sensor arrays for tracking magnetic microparticles in a microfluidic environment. The SCD-AER architecture correctly records particle-transit events from synthetic stimuli down to hundreds of μ V and from pseudo-real particle-flow signals derived from experimental GMR-cytometer measurements, while naturally suppressing noise. The data-volume comparison (Table 1) further shows that this behavior translates into a substantial reduction – between three and five orders of magnitude, depending on the reference system, using each system's own reported bit depth – in the amount of raw acquisition data that needs to be generated, transmitted and processed for the same particle-counting task, relative to continuous-sampling and impedance-based systems reported in the literature.

Future work will focus on full experimental validation with the complete integrated system, including a hardware implementation of the SCD-AER interface characterized in [7], optimization of sensor

geometry, extension towards two-dimensional arrays for multi-analyte point-of-care platforms, and a systematic study of the sensitivity limits of the complete detection chain.



Fig. 3. Device response to pseudo-real particle flow. Behavior of ED logic from a simulated 8×1 GMR-based cytometry array.

Table 1. Comparison of data-volume between similar references over a 9.6 ms window. Using real channel count reported for each system.

System	Channels	Measuring time	Samples/ events	Samples/ channel	Bits/ sample	Total bytes
Chou et al. 2024 [10] (60 MHz, channel count not stated)	1*	9.6 ms	576.000	576.000	8 (reported)	576.000
R. Afonso et al. 2023 [11] (INESC-MN, 200 kSps/channel)	6	9.6 ms	11.520	1.920	16	23.040
Caetano et al. 2023 [12] (INESC-MN, 200 kSps/channel)	1	9.6 ms	1.920	1.920	16	3.840
Event-Driven Tracking System	8	9.6 ms	8	1	4	4

* Channel count is not stated in the publicly available material for [10]; a conservative single-channel configuration is therefore assumed for the total-sample calculation, which, if anything, would be the best-case scenario regarding data volume and does not inflate the comparison in our favor. Channel counts for [11] and [12] are taken directly from the respective publications (six real channels, and a single-sensor GMR configuration, respectively).

Acknowledgements

This work has been partially supported by projects with references: TSI-069100-2023-0012 and PID2023-149071NB-C53.

References

- [1]. B. Rezaei, P. Yari, S. Sanders, et al., Magnetic nanoparticles: A review on synthesis, characterization, functionalization, and biomedical applications, *Small*, Vol. 20, Issue 15, 2024, 2304848.
- [2]. C. Reig, S. Cardoso, S. C. Mukhopadhyay (Eds.), Giant Magnetoresistance (GMR) Sensors: From Basis to State-of-the-Art Applications, *Springer*, Berlin, 2013.
- [3]. K. Wu, et al., Giant magnetoresistance biosensors in biomedical applications, *ACS Applied Materials & Interfaces*, Vol. 14, Issue 8, 2022, pp. 9945-9969.
- [4]. J. Howell, T. Hammarton, Y. Altmann, High-speed particle detection and tracking in microfluidic devices using event-based sensing, *Lab on a Chip*, Vol. 20, Issue 17, 2020, pp. 3024-3035.
- [5]. C. Reig, F. Pardo, J. A. Boluda, et al., Advanced giant magnetoresistance (GMR) sensors for selective-change driven (SCD) circuits, in *Proceedings of the 13th Spanish Conference on Electron Devices (CDE)*, 2021, pp. 85-88.
- [6]. C. Reig, M. D. Cubells-Beltrán, F. Pardo, et al., Dynamic electrical compact model of GMR sensors for neuromorphically inspired applications, *AIP Advances*, Vol. 13, Issue 2, 2023, 025225.
- [7]. A. Shahryari, F. Pardo, J. A. Boluda, C. Reig, Event-driven based CMOS interface for magnetic sensing, in *Proceedings of the 14th IEEE Spanish Conference on Electron Devices (CDE)*, 2024, pp. 1-4.
- [8]. T. Sen, A. C. Sreekantan, S. Sen, A magnetic feedback-based Δ - Σ digitizing interface for giant magnetoresistance sensors, *IEEE Transactions on*

- Circuits and Systems II: Express Briefs*, Vol. 70, Issue 1, 2023, pp. 36-40.
- [9]. J. Gómez, C. Reig, M. D. Cubells, et al., Elemental giant magnetoresistance (GMR) sensors for neuromorphical applications, in *Proceedings of the 11th International Conference on Sensor Device Technologies and Applications (SENSORDEVICES)*, 2020, pp. 1-5.
- [10]. T.-H. Chou, S. Yu, C. Wilson, J. Dawes, et al., An impedance measurement system-on-chip for flow cytometry with on-chip size classification, in *Proceedings of the IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2024, pp. 1-5.
- [11]. R. Afonso, D. M. Caetano, R. Lorena, et al., Multi-sensor low-noise modular magnetic flow cytometer for bacteria detection, in *Proceedings of the IEEE Sensors Conference*, 2023, pp. 1-4.
- [12]. D. M. Caetano, R. Afonso, A. R. Soares, et al., Artificial neural networks for GMR-based magnetic cytometry, *IEEE Transactions on Instrumentation and Measurement*, Vol. 72, 2023, 2002311.
- [13]. R. F. Afonso, Electronic platforms and signal processing for dynamic magnetic biochips, PhD Thesis, Instituto Superior Técnico, *Universidade de Lisboa*, Lisbon, 2024.

A Low-Cost SDR-Based Telecommunications Laboratory Using Repurposed TV Boxes

**Arquimedes Lopes da Silva, Diego F. C. Siqueira, Leandro L. Alves, Maria Clara R. N.
and Eliézer Alves Teixeira**

Instituto Federal de Goiás, Goiânia-GO, Brasil

Tel.: + 55 62 3227-2730

E-mail: arquimedes.silva@ifg.edu.br, siqueiradiego2005@gmail.com, alves.leal@estudantes.ifg.edu.br,
claranog.eng@gmail.com, eliezerteixeira@ifg.edu.br

Summary: This article investigates the reuse of processors from commonly discarded or seized Android-based TV Box devices, transforming them into general-purpose System-on-Chip (SoC) computing platforms. The conversion process consists of the complete removal of illicit software and the reconfiguration of the hardware through the installation of free Linux-based operating systems. Furthermore, the initiative acts as a vector for digital inclusion and fosters the creation of low-cost academic laboratories. In this work, the proposed platform integrates the ARM/Linux (Advanced RISC Machine) hardware of the TV Box with an accessible radio frequency (RF) front-end and the free GNU Radio software, enabling the reception and processing of AM and FM signals. This set of low-cost tools is used pedagogically as an alternative to traditional and expensive benchtop equipment in teaching fundamental engineering subjects, such as Modulation Techniques and Digital Signal Processing (DSP). As a result, the methodology adopted allows students to perform interactive practical experiments, visualizing, manipulating, and analyzing real signals in both the time and frequency domains. In this way, the technical and pedagogical viability of reusing these computational systems to improve higher technological education is validated.

Keywords: Analog modulation AM, Analog modulation FM, ARM/Linux, SDR, GNU Radio, TV Box.

1. Introduction

Embedded systems, increasingly present in today's world, have become popular due to the growing processing power, low cost, and features offered by microcontrollers. Among microcontroller architectures, the one that has grown the most in recent years is ARM (Advanced RISC Machine) [1].

The popularization of TV Boxes in Brazil and worldwide is due to an economical solution to transform old TVs into "Smart TVs," centralizing streaming. However, the main driver of this sales explosion is the ease of use of these devices as IPTV systems and third-party applications for free or very cheap access to paid channels.

Educational institutions often lack adequate laboratory infrastructure because professional RF instruments, such as oscilloscopes, spectrum analyzers and signal generators, are expensive.

The conversion of an Android-based TV Box into a general-purpose ARM/Linux SoC platform (System-on-Chip), similar to an IoT (Internet of Things) device applied in the educational field. According to [2] performance, these TV Boxes converted with free operating systems are viable for use in public schools. The author [3] detailed the process of converting these devices into minicomputers, highlighting the removal of the DTB (Device Tree Blob) and the installation of Linux Armbian as a low-cost alternative to the Raspberry Pi.

This work presents a methodology for converting seized Android TV Boxes into low-cost ARM/Linux platforms, as described by author [4], for telecommunications teaching laboratories.

The main objective is to show that a platform, formed by RTL-SDR (Realtek – Software Defined Radio) hardware and GNU Radio software, is a low-cost, flexible and efficient solution for creating experimental scripts in telecommunications classes.

2. Fundamentals of Analog Modulation

2.1. Signal Flow in Amplitude Modulation

AM (Amplitude Modulation) is a telecommunications technique that transmits information (such as audio) by varying the amplitude of a radio carrier wave, while its frequency remains constant [5].

Fig. 1 illustrates how AM modulation works. A high-frequency wave (the carrier) is generated to carry the information (the modulating signal). These two signals are "mixed", resulting in a high-frequency modulated signal ready for transmission!




Fig. 1. Signal flow representing AM modulation.

The entire equation for AM modulation is given by:

$$m(t) = A_m \cos(2\pi f_m t), \quad (1)$$

$$c(t) = A_c \cos(2\pi f_c t), \quad (2)$$

$$s(t) = A_c [1 + m \cos(2\pi f_m t)] \quad (3)$$

where $m(t)$ is the modulating signal, $c(t)$ is the carrier signal, and $s(t)$ is the AM signal. The respective AM signals are represented graphically in the time domain in Fig. 2.




Fig. 2. AM Modulation Graphics.

The transmitted modulated signal can be viewed in the frequency domain, revealing the spectrum (amplitude and frequency) of the information. Fig. 3 illustrates the AM modulation spectrum, where the message spectrum is shifted to the carrier frequency f_c .




Fig. 3. AM Modulation Spectrum.

2.2. Frequency-Modulated Signal Flow

In FM modulation (see Fig. 4), the carrier frequency is altered according to the modulating signal. The phase angle also changes; therefore, Frequency Modulation is also referred to as Angle Modulation. To modulate a signal in terms of frequency, the carrier frequency is varied using the low-frequency signal such that the instantaneous value of the high-frequency signal is proportional to the instantaneous amplitude of the modulating signal [6].

The mathematical expression for a Frequency Modulation (FM) signal is defined as a cosine whose angle includes the integral of the message signal. The general equation for the FM-modulated signal $s(t)$ is given by:

$$c(t) = A_c \cos(2\pi f_c t), \quad (4)$$

$$s(t) = A_c \cos(2\pi f_c t + 2k_f \pi \int_0^t f(x) dx), \quad (5)$$

where, in this modulation, $s(t)$ is the FM signal. All these time-domain signals are plotted in Fig. 5.




Fig. 4. Signal flow representing FM modulation.




Fig. 5. FM Modulation Graphics.

Note in Fig. 6 that the message spectra are shifted relative to the carrier frequency f_c .




Fig. 6. FM modulation spectrum.

3. Proposed Platform

In this work, the proposed platform employs a device as a low-cost RF front-end for AM and FM experiments, an ARM/Linux TV Box, and the GNU Radio software, see Fig. 7.




Fig. 7. Proposed platform architecture [Author].

RTL-SDR is the front-end that converts the RF signal captured by the antenna into a sampled $x[n]$ signal compatible with the ARM processor of the repurposed TV Box, this digital signal is processed according to the signal flowchart present in the design of each educational experiment carried out in GNU Radio [7].

3.1. RTL-SDR

This RTL-SDR device is a low-cost SDR receiver based on Realtek chips, originally designed for DVB-T digital TV reception. The term RTL comes from the chipset:

These devices, the RTL2832U (Realtek), see Fig. 8 – were "hacked" by the user community to function as generic SDR receivers.



Fig. 8. RTL device.

This is a software-defined radio receiver featuring an RTL2832U ADC chip, R820T2 tuner, 1PPM TCXO, SMA-F connector, and an aluminum case with passive cooling.

In a GNU Radio signal flow project, the RTL-SDR appears as a software layer known as a "Source block" (RF signal source), enabling the capture of real-world signals from the spectrum.

The RTL-SDR Source block allows the user to:

- set the center frequency;
- sample rate;

- gain;
- frequency correction;
- Typical range.

3.2. ARM/Linux TV Box

The TV Box used on this platform is from manufacturer BTV, model B11, featuring the X3BL46 V2.0 2020.03.21 motherboard, DTB m1 ac213 2g board, Amlogic S905X3 CPU (Cortex-A55 family, speed 100–1908 GHz), Mali-G31 MC1 GPU, AP6212 Wi-Fi, 1GB of memory, and 16GB of storage.

The processor used was the ARMv8, as shown in Fig. 9.




Fig. 9. ARM processor version.

3.3. GNU Radio Companion

The GNU Radio software features a GUI (Graphical User Interface) called GRC (GNU Radio Companion), which allows for the creation of signal processing applications using blocks that can be "dragged" across the graphical area.

Fig. 10 presents the GRC main screen, highlighting:

- (1) the menu bar;
- (2) the button bar;
- (3) the block library;
- (4) the terminal;
- (5) the graphical area with two initial blocks.

The platform supports experiments involving AM modulation and demodulation, FM reception, spectrum visualization, signal filtering, and introductory DSP concepts. Students can interactively modify GNU Radio flowgraphs and immediately observe the effects in both the time and frequency domains.

3. Educational Experiments

Fig. 11 shows a flowchart of signals through connected functional blocks forming an AM Modulator; note the signals explained as the message (audio signal) and the carrier.




Fig. 10. GRC – GNU Radio's official GUI [Author].




Fig. 11. Signal Flow for the AM Modulator [7].

4. Preliminary Results

The experiments, such as AM and FM, cited in [5], were tested and applied to students in accordance with the theory found in [6]. Fig. 12 summarizes the experiment where the signal from a specific FM radio station in the city of Goiânia is monitored. In addition to the student following the entire signal flow, they can visualize the signal over time and its frequency spectrum at strategic points along the signal. They can even hear the sound of the signal through the speaker output of the TV Box device.

Fig. 12 presents the components of the platform used for this experiment, which is of great practical importance for the student to understand the entire process of receiving an FM-modulated signal: (1) the antenna for capturing FM signals, (2) the RTL-SDR receiver, (3) the TV Box converted into an ARM SoC,

(4) the monitor connected to the TV Box's HDMI output, (5) the Armbian/Linux operating system, (6) the GNU Radio software that enables signal flow processing from the FM receiver, and (7) the GNU Radio output - in this case, a spectrum analyzer.

The platform successfully executed GNU Radio flowgraphs for AM and FM reception using an RTL-SDR receiver operating on the ARM-based TV Box. Real-time spectrum visualization and audio demodulation were achieved without external computing resources.

5. Conclusions

The contribution of this article lies in the use of TV Boxes transformed into ARM devices due to their SoC reconfiguration in terms of operating system. These

devices come from material seized by the Federal Revenue Service, and are therefore low-cost. Thus, these new devices are used pedagogically, in a telecommunications laboratory, for teaching Telecommunications subjects such as Modulation

Techniques. The use of these devices will extend practices to subjects such as Digital Signal Processing. Future work includes the implementation of digital modulation experiments, DSP laboratories, and remote access for distance learning environments.




Fig. 12. FM Receiver Experiment [Author].

References

- [1]. L. S. Tavares, P. R. O. Valim, Placa de desenvolvimento com ARM para ensino de microcontroladores, *Anais do Computer on the Beach*, Vol. 9, 2018, pp. 21-29 (in Portuguese).
- [2]. F. T. de Souza, K. M. Lacerda, C. R. M. Mauricio, R. B. da Silva, et al., Análise de desempenho de aparelhos TV box descaracterizados para uso educacional, in *Proceedings of the 21st Congresso Latino-Americano de Software Livre e Tecnologias Abertas (Latinoware)*, 2024, pp. 364-367 (in Portuguese).
- [3]. H. G. Júnior, Re-caracterização de aparelhos de TV box para uso como minicomputadores: Procedimentos para investigação e formatação, *Anais dos Encontros Nacionais de Engenharia e Desenvolvimento Social*, Vol. 20, Issue 1, 2024, 28 (in Portuguese).
- [4]. V. A. S. Lopes, Uma proposta de computadores de baixo custo utilizando TV boxes apreendidas, Bachelor's Thesis, *Universidade de Brasília*, Brasília, 2023 (in Portuguese).
- [5]. S. Haykin, M. Moher, *Sistemas de Comunicação* (5th Ed.), *Bookman Editora*, Porto Alegre, 2011 (in Portuguese).
- [6]. B. P. Lathi, Z. Ding, *Modern Digital and Analog Communication Systems*, Vol. 3, *Oxford University Press*, New York, 2010.
- [7]. Exploring GNU Radio, <https://www.gnuradio.org>

Computer Vision Analysis of Sanding Marks in Wood

M. Colina¹, J. Ramos², L. Marín¹ and C. Reig²

¹ AIDIMME, Parque Tecnológico de Valencia, 46980, Paterna, Spain

² University of Valencia, ETSE, 46100, Burjassot, Spain

E-mail: mcolina@aidimme.es

Summary: This work presents a microscope-based, multi-illumination computer-vision pipeline for the quantitative analysis of individual sanding marks on wood surfaces. Wood panels sanded under controlled orbital-sanding conditions were imaged with an industrial microscope under automated sequential LED-ring illumination, producing twelve correlated images at different grazing-light angles. These images are fused by a temporal Discrete Fourier Transform (DFT), evaluated pixel by pixel across the illumination sequence, which condenses the shadow and highlight signature of each V-shaped groove into a single enhanced image. A dedicated ridge-following algorithm then detects, reconstructs, tracks the individual grooves doing a recentering of each trace on the groove crest, bridging fragmented segments, and measuring the geometry. Unlike prior optical methods to estimate aggregate roughness parameters, the proposed pipeline extracts per-groove descriptors such as length, width, and curvature radius, providing groove-level information for the quality control of automated and robotic wood-sanding processes.

Keywords: Computer vision, Wood surface, Sanding marks, Image processing, Groove detection, Surface quality, Microscopy.

1. Introduction

The analysis of sanding marks on wood surfaces is crucial for quality control, defect detection and anatomical diagnostics in wood science and industry [1]. The applied sanding process significantly affects the surface roughness, wettability and color of the wood [2, 3]. Nowadays, these analyses have gained renewed interest because of the ongoing migration from traditional manual sanding schemes to automated processes based on robots. The proper configuration of these systems requires a detailed correlation between the applied sanding parameters (force, velocity, sandpaper grit) and the features of the marks produced by the process, which in turn demands measurement techniques able to resolve and quantify individual sanding marks rather than only global surface indicators.

Traditionally, sanded-wood surface quality has been assessed with mechanical stylus profilometry. Although profilometry remains a reference technique for roughness parameters, it is intrinsically limited for generating areal, statistically representative data: it is a slow, line-based contact measurement, it can be distorted by the wood anatomy, and it does not preserve the trajectory of individual grooves [1, 3]. To overcome these limitations, non-contact optical approaches have been progressively adopted, including high-resolution digital microscopy [4], image-processing techniques [5] and, more recently, computer-vision and deep-learning methods [6].

Computer-vision and optical approaches to sanded-wood surface quality have already been investigated extensively. Early work showed that computer vision could estimate wood surface roughness in a non-contact way, correlating well with stylus tracing [7]. More recent studies estimate global

roughness parameters from surface images. The work presented in [8] applied computer-vision descriptors to predict surface roughness on wood-based sanded workpieces and, specifically, on parts processed by robotic sanding [9]; machine-vision systems based on reflectance statistics have been reported for abrasive-blasted surfaces [10]; and machine-learning and deep-learning regressors have been trained to predict roughness from micrographs of sanded wood [11, 12]. In parallel, data-driven pipelines classify surface defects or control color on wood using convolutional neural networks [6, 13], although these require large labelled datasets and typically return class labels or aggregate scores rather than the geometry of individual marks. A related family of techniques, developed mainly for industrial metal inspection, exploits multi-illumination and photometric-stereo acquisition to reveal shallow topographic features: multi-angle lighting has been combined with deep networks [14, 15] and, notably, with Fourier-Gabor filtering to enhance scratches against a textured background [16].

Despite this rich body of work, the reviewed methods either estimate aggregate roughness parameters or detect and classify defects; to the best of our knowledge, none of them explicitly detects and tracks individual sanding grooves and reports their per-groove geometric descriptors. So, in the present work, we propose a microscopy-based, multi-illumination image-processing pipeline for detecting, reconstructing, and tracking individual sanding grooves on wood surfaces, and for extracting geometric descriptors (average length, width and curvature radius) of each groove.

The pipeline combines three elements that, taken together, have not previously been reported for this problem. We combine a twelve-angle

grazing-illumination ring that modulates the shadow and highlight the signature of every V-shaped groove, a temporal Discrete Fourier Transform (DFT) evaluated across the illumination sequence, which isolates that modulation without blurring the original spatial features and an oriented-kernel ridge-following algorithm with a perpendicular-recentering and segment-bridging stage that reconstructs continuous groove trajectories.

Compared with optical roughness estimation [8-12], which condenses the surface into scalar parameters, our method preserves and measures the individual grooves; compared with contact profilometry [1, 3], it is a non-contact, areal and better suited to statistical analysis; and compared with deep-learning surface-analysis approaches [6, 13], it is interpretable, requires no labelled training set, and returns explicit groove geometry rather than a classification. This groove-level, statistical characterization is precisely the information required to correlate sanding parameters with the resulting marks in automated and robotic sanding.

The remainder of this article is as follows. Section 2 presents the complete pipeline detailing the microscopy and multi-illumination setup, the Fourier-based groove-enhancement stage and the ridge-following detection and tracking algorithm. Section 3 reports the qualitative and quantitative results and Section 4 summarizes the conclusions.

2. Methodology

The characterization method proposed in this work is organized as the sequential pipeline outlined in the block diagram of Fig. 1. Conceptually, it is built from two complementary stages. The first stage takes the set of micrographs captured under the automated multi-angle illumination and condenses them into a single enhanced image, in which the sanding marks are made to stand out from the underlying wood texture. The second stage takes that enhanced image and turns it into quantitative information, isolating each mark, following its path along the surface, and measuring its geometry. The process therefore advances from a stack of raw grazing-light images to a population of individually characterized grooves, with no manual intervention in between. The blocks that make up this diagram are described in detail in the following subsections.

2.1. Sample Preparation and Sanding Process

The study was conducted using wood block samples (10 cm × 10 cm × 1.6 cm) extracted from a larger panel typically used for the manufacture of modular furniture for bedrooms. Surface preparation was performed using a MIRKA DEROS random-orbital sander (150 mm plate diameter, 5.0 mm orbit).




Fig. 1. Image processing workflow.

Although the sander was hand-guided, the process was kept under controlled parameters: a single abrasive grade (ceramic-grain paper, grit P240), a fixed orbit and a constant plate speed were used throughout, so that the resulting marks are representative of a well-defined sanding condition rather than of an arbitrary manual operation.

2.2. Microscope Setup and Image Acquisition

To capture the surface topography, an industrial microscopy system was developed. The experimental setup, shown in Fig. 2, is based on a NIKON Eclipse LV100 industrial microscope equipped with a NIKON LU Plan Fluor 5X/0.15 lens and a NIKON DS-Ri1 digital camera.




Fig. 2. Experimental setup showing a) the industrial microscope and b) the custom automated LED ring for sequential grazing illumination.

The acquisition process was managed via NIKON NiS-Elements BR 3.10 64-bit software, running on a personal computer with the Windows 7 Professional operating system. Images of 1280 × 1024 pixels were acquired at a scale of 1.82 μm/pixel.

To enhance the visibility of the sanding grooves, a custom illumination system consisting of a ring with 12 equally spaced LEDs was implemented. This system is electronically automated via a micro-

controller, which activates a single LED at a time, holding the illumination long enough for a single micrograph to be captured. This sequential scheme allows the generation of twelve correlated images per field of view, each captured under a different grazing-light angle. This multi-angle approach is essential for effectively isolating the shadow and highlighting signatures of the V-shaped grooves against the natural wood texture.

2.3. Multi-illumination Fourier Enhancement

The twelve images are not processed independently, nor is a single best-illuminated image selected. The images are fused into one enhanced image by means of a temporal Discrete Fourier Transform (DFT) evaluated pixel by pixel across the illumination sequence. This stage constitutes the first block of the processing pipeline described in Fig. 1, and produces the enhanced image shown in Fig. 3.




Fig. 3. Sanding marks enhanced after applying the DFT.

The justification is that a conventional two-dimensional FFT of a single micrograph is of little use here, because the sanding lines are not periodic and therefore do not concentrate energy at characteristic spatial frequencies; likewise, simply adding the twelve images does not create any periodic pattern. Instead, we treat the twelve images as a temporal sequence. For every Cartesian position (x , y) a vector of twelve intensity samples is formed, one from each image, along an axis orthogonal to the image plane. Assuming a sampling rate of one image per second, a full revolution of the light source around the ring takes twelve seconds, so the shadow modulation produced by a V-shaped groove as the grazing light rotates appears at the fundamental frequency of $1/12$ Hz (equivalent to $2\pi/12$ rad/s).

Accordingly, the DFT of each twelve-sample vector is evaluated only at $1/12$ Hz, and its magnitude (the square root of the sum of the squares of the real and imaginary parts of that single coefficient) is stored in a magnitude map. For the 1280×1024 image, this

requires 1 310 720 single-frequency DFT evaluations. Because the operation acts along the illumination (temporal) axis, orthogonal to the image plane, it fuses the twelve views into a single map while retaining the original two-dimensional spatial features without any blurring. The resulting magnitude map is finally rescaled by an intensity-level normalization (histogram stretching to the 0-255 range) to obtain the enhanced image of Fig. 3. As a possible extension, the phase of the same DFT coefficient could be exploited, since pixels belonging to the same groove are expected to share a common phase angle.

2.4. Detection Pipeline

The enhanced image is then processed by a dedicated algorithm that detects, reconstructs and tracks the individual grooves following an iterative ridge-following approach as shown in Fig. 1. The building blocks rely on well-established image-processing operators: morphological black top-hat filtering, a standard operation for isolating thin bright/dark features against an uneven background, is combined with adaptive contrast enhancement in the preprocessing stage; and a bank of oriented Gaussian kernels is used to estimate the local orientation and the ridge (line) response, in the spirit of classical curvilinear-structure and ridge detectors.

Our contribution does not consist of these individual operators, but rather of how they are coupled with the multi-illumination DFT enhancement. From seed points placed on a grid (rejecting ridge crossings), a path-tracing routine follows each ridge and applies a perpendicular-recentering mechanism that keeps the trace on the ridge crest, together with a segment-bridging step that reconnects fragments belonging to the same groove across gaps and intersections. This yields robust, continuous groove trajectories even on challenging backgrounds, from which the geometric descriptors (length, width and curvature radius) are measured. For visualization, the detected grooves are colorized on top of the original image to highlight the algorithm's findings.

3. Results

In this section, the capabilities of the detection framework are demonstrated using the wood samples described in the methodology. The results are organized to reflect the processing stages, starting with the qualitative detection of individual sanding marks and concluding with the statistical characterization of groove geometry.

3.1. Qualitative Evaluation of the Detection Pipeline

The application of the detection pipeline to the enhanced images yields both visual reconstructions

and geometric data. Fig. 4 illustrates the qualitative performance of the tracking algorithm. The overlaid colored lines demonstrate that the method effectively reconstructs continuous grooves, successfully bridging fragmented segments while filtering out background noise and irregular, highly tortuous shapes.




Fig. 4. Sanding grooves tracking output.

Building upon this, Fig. 5 shows the ability of our algorithm to precisely identify the groove boundaries (highlighted in green) and the groove centerline (highlighted in red). This extraction is essential for characterizing groove width and provides a technical foundation for future research on depth estimation. These geometric descriptors are critical for training autonomous robotic systems, enabling real-time parameter adjustments in automated wood-sanding processes.




Fig. 5. Detection of ridge boundaries (green) and centerlines (red).

3.2. Statistical Analysis of Detected Grooves

From a quantitative perspective, the geometric descriptors of every retained groove were evaluated. The statistical distributions of groove length, width and

curvature radius are summarized in histograms of Fig. 6.

Groove lengths span from a few tens up to a few thousand microns; groove widths concentrate in the single-digit-micron range (around 4-6 μm), almost an order of magnitude smaller than the nominal P240 grain size being $\approx 58.5 \mu\text{m}$, which is consistent with each groove being plowed by the sharp cutting edge of an individual abrasive particle rather than by its full diameter; and curvature radii, shown on a logarithmic scale, are concentrated between 10-1000 microns range reflecting a mixture of markedly curved and nearly straight grooves.




Fig. 6. Quantitative analysis of the detected grooves geometry.

4. Conclusions

A complete pipeline for the characterization of individual sanding marks in wood has been developed. Unlike previous optical approaches that estimate aggregate roughness parameters, the method detects

and tracks each groove and quantifies its geometry (average length, width and curvature radius). The combination of twelve-angle grazing illumination, a temporal DFT groove-enhancement stage and an oriented-kernel ridge-following tracer provides a non-contact, interpretable and statistically representative tool, well suited to correlating sanding parameters with the resulting surface in automated and robotic sanding.

Acknowledgements

This work has been partially supported by the Ministry of Digital Transformation and Civil Service under the Recovery, Transformation and Resilience Plan, funded by the European Union – NextGenerationEU, through the PERTE Chip Chairs program (Project Ref: TSI-069100-2023-0012).

References

- [1]. P. L. Tan, S. Sharif, I. Sudin, Roughness models for sanded wood surfaces, *Wood Science and Technology*, Vol. 46, Issue 1, 2012, pp. 129-142.
- [2]. A. Laskowska, A. Piwek, K. Lipska, T. Kłosińska, et al., Evaluation of the selected surface properties of European oak and Norway maple wood sanded with aluminum oxide sandpapers of different grits, *Coatings*, Vol. 15, Issue 7, 2025, 813.
- [3]. L. Adamčík, L. Dzurenda, A. Banski, R. Kminiak, Comparison of surface roughness of beech wood after sanding with an eccentric and belt sander, *Forests*, Vol. 15, Issue 1, 2024, 45.
- [4]. P. Kitin, J. C. Hermanson, H. Abe, S. Nakaba, et al., Light microscopy of wood using sanded surface instead of slides, *IAWA Journal*, Vol. 42, Issue 3, 2021, pp. 322-335.
- [5]. C. Łukianowicz, Identification of machining marks on abrasive treated surfaces, using image processing techniques, *Mechanik*, Issue 11, 2018, pp. 1020-1022 (in Polish).
- [6]. J. Li, S. Wu, Image processing and analysis for color control of wood surfaces based on convolutional neural networks, in *Proceedings of the 2nd International Conference on Data Science and Information System (ICDSIS)*, 2024, pp. 1-5.
- [7]. X. Zhao, Study on measurement of wood surface roughness by computer vision, *Journal of Northeast Forestry University*, Vol. 20, Issue 3, 1992, pp. 69-74 (in Chinese).
- [8]. F. Iglesias, A. Aguilera, A. Padilla, A. Vizán, et al., Application of computer vision techniques to estimate surface roughness on wood-based sanded workpieces, *Measurement*, Vol. 223, 2023, 113917.
- [9]. F. Iglesias, A. Aguilera, A. Padilla, M. Madaf, et al., Optical evaluation of surface roughness in wood parts processed by robotic sanding, in *Proceedings of the 15th Ibero-American Congress of Mechanical Engineering (IACME)*, 2023, pp. 487-495.
- [10]. K. Kim, Development of a machine vision system for the average roughness measurement of shot- and sand-blasted surfaces, *Lubricants*, Vol. 12, Issue 10, 2024, 339.
- [11]. C. Shang, Z. Lieping, K. A. Gepreel, H. Yi, Surface roughness measurement using microscopic vision and deep learning, *Frontiers in Physics*, Vol. 12, 2024, 1444266.
- [12]. M. Güneş, Machine learning-based prediction of surface roughness in heat-treated and sanded Siberian pine wood, *European Journal of Wood and Wood Products*, Vol. 84, 2026.
- [13]. P. Kodytek, A. Bodzas, P. Bilik, A large-scale image dataset of wood surface defects for automated vision-based quality control processes, *F1000Research*, Vol. 10, 2021, 581.
- [14]. Y. Cao, et al., Photometric-stereo-based defect detection system for metal parts, *Sensors*, Vol. 22, Issue 21, 2022, 8374.
- [15]. X. Fang, Q. Luo, B. Zhou, C. Li, et al., Research progress of automated visual surface defect detection for industrial metal planar materials, *Sensors*, Vol. 20, Issue 18, 2020, 5136.
- [16]. Y. Yue, W. Sang, K. Zhai, M. Lin, Enhanced scratch detection for textured materials based on optimized photometric stereo vision and Fast Fourier Transform–Gabor filtering, *Applied Sciences*, Vol. 14, Issue 17, 2024, 7812.

Deterministic Modeling and Advanced Micro-Doppler Analysis of Quadrotor Drones in the X-Band

A. Le Dréo¹, A. Khenchaf¹, J.-C. Cexus¹ and P. Pouliguen²

¹ Lab-STICC, UMR CNRS 6285, ENSTA, Institut Polytechnique de Paris, 29200 Brest, France

² Agence de l'Innovation de Défense, DGA, France

E-mail: aymeric.le_dreo@ensta.fr, ali.khenchaf@ensta.fr, jean-christophe.cexus@ensta.fr,
philipe.pouliguen@intra.def.gouv.fr

Summary: The detection and classification of micro-drones is a major challenge for radar surveillance. Faithful modeling of their micro-Doppler (μD) signature is limited by classical point-scatterer models, which neglect two physical and geometric blade properties that are rarely integrated jointly: *Twist* (angular blade pitch variation) and *Taper* (chord tapering, i.e. the geometric characteristic of an aircraft wing or a propeller blade whereby the chord width decreases from the blade root to the blade tip). Xu et al. [1] model the blade geometry in the near field but with a uniform chord profile and no angular twist. Al-Nuaimi et al. [2] integrate material properties via a Method-of-Moments (MoM) approach but on a constant rectangular geometry. Our contribution jointly integrates, for the first time in a deterministic model dedicated to micro-UAVs, *Twist* as a phase-decorrelation mechanism and *Taper* as a spatial apodization, acting as a natural windowing function that weights the contribution of the elementary scatterers according to the local chord evolution ((r)). A multi-material polarimetric analysis (carbon fiber, ABS) complements this study by quantifying the impact of blade material composition on RCS dynamics.

Keywords: Radar, Micro-Doppler, Drone, Deterministic simulation, Spectral signature, Twist, Taper.

1. Introduction

The proliferation of civilian and military drones presents significant surveillance challenges due to their low Radar Cross-Section (RCS) and low-altitude operation, which often masks them within ground clutter or mimics the kinematic signatures of birds. The radar return of a drone is characterized by the micro-Doppler (μD) effect – a unique "fingerprint" induced by rotor rotation that is vital for detection and classification [3].

Accurate modeling of this effect remains a bottleneck. Classical point-scatterer models following Chen [3-5] represent each blade by a single equivalent reflector, yielding closed-form spectra but no spectral fill-in. Thin-wire models extend this representation to a one-dimensional line of scatterers along the blade span, favoring computational efficiency for large-scale, waveform-oriented studies such as recent FMCW/ISAC micro-Doppler frameworks [6], but at the cost of a uniform, geometry-agnostic reflectivity along the span. Conversely, full 3D component-based frameworks relying on Method-of-Moments or physicaloptics meshes [2] capture arbitrary blade geometry and material composition, at a computational cost prohibitive for large-scale synthetic dataset generation. We propose a deterministic segmental model that occupies an intermediate position: it retains the closed-form, low-cost structure of thin-wire/point-scatterer formulations while incorporating the radius-dependent *Twist* and *Taper* neglected by both, without resorting to a full 3D mesh.

We demonstrate that the spectral fill-in primarily originates from two factors: "Twist" (angular pitch variation), which causes continuous phase decorrelation along the blade, and "Taper" (chord tapering), which modulates the reflecting surface and redistributes energy. Neglecting these properties leads to an underestimation of spectral spread and energy distribution inaccuracies. This work rigorously integrates these effects to provide a robust model suitable for training classification algorithms [7-9].

The remainder of this article is organized as follows: Section 2 establishes the theoretical μD signal model. Section 3 analyzes polarimetric and material impacts. Section 4 quantifies the contributions of *Twist* and *Taper*, and Section 5 validates the model against reference data. Section 6 concludes and discusses future work.

2. μD Signature of a Target

The Doppler phenomenon corresponds to a frequency shift induced by the relative motion between a target and a radar. In radar theory, this effect is exploited to estimate a target's radial velocity.

By extension, the micro-Doppler (μD) effect refers to additional frequency modulations generated by the internal or oscillatory motion of a target's moving parts (rotations, vibrations). These contributions provide detailed information on the structure and kinematics of the observed object.

2.1. Modeling the μ D Effect: Chen's Model

Chen's model [3] is a reference approach for micro-Doppler signature modeling. It relies on a physical-optics formulation that represents the moving parts of the target by point scatterers, and derives the micro-Doppler frequency $f_{\mu D}$ as the time derivative of the phase (ϕ) of the backscattered wave [14]. It rests on three assumptions: a physical-optics approximation ($\lambda \ll a$, a the target's characteristic dimension), a far-field hypothesis, and a point-reflector representation of the scattering surfaces. Motion is described across a radar frame, a body reference frame, and a local frame attached to each moving element. The model is widely used for its simplicity, despite neglecting diffraction, surface waves, and multiple interactions between structural parts.

Consider a continuous-wave (CW) radar emitting a signal $(t) = S_0 \exp(-j2\pi f t)$ toward an object located at a distance $R_0 \gg a$. Neglecting the antenna radiation pattern, the received signal $s_r(t)$ is expressed as:

$$s_r(t) = \rho(x, y, z) \exp[-j(2\pi f t + \phi(t))], \quad (1)$$

where $\rho(x, y, z)$ is the reflectivity function of the point scatterers.

A point P of the object undergoes a micro-motion consisting of a translation from P_0 to P_1 at velocity $\vec{v}$, followed by a rotation from P_1 to P_2 about a vector $\vec{\omega}$ at angular velocity $\Omega = \|\vec{\omega}\|$. The theoretical micro-Doppler frequency is obtained by differentiating the phase:

$$f_{mD} = \frac{1}{2\pi} \frac{d\phi(t)}{dt} = \frac{2f}{c} \vec{v} \cdot \vec{n} + \frac{2f}{c} (\vec{\omega} \times \vec{r}_1^i) \cdot \vec{n}, \quad (2)$$

where $\vec{r}_1^i = e^{\hat{\omega}t} \vec{r}_1$ is the instantaneous position of P_1 under the rotation matrix $R_t = e^{\omega t}$ (obtained via Rodrigues' formula), and $\vec{n} = \frac{\vec{R}_0 + \vec{v}t + R_t \vec{r}_1}{|\vec{R}_0 + \vec{v}t + R_t \vec{r}_1|}$ is the radar line of sight.

The study of this signal relies on time-frequency analysis (Fourier transform and spectrograms) to capture the fast time variations induced by the target's oscillatory motion.

2.2. μ D Signature of a Quadrotor Drone

To evaluate the fidelity of the deterministic model, the simulations are based on a Class C2 quadrotor drone, representative of civilian micro-UAVs. The physical and kinematic parameters are detailed in Table 1. Unless otherwise stated, the results shown throughout Sections 2 and 4 correspond to the co-polarized HH channel; Section 3 additionally exploits the VV channel and its systematic exploitation (in particular for material discrimination, see Section 3) is left to future work.

We simulate the micro-Doppler (MD) signal of a hovering multi-rotor drone located at an initial distance

$R_0 = 0.1$ m (reference distance for phase analysis), with azimuth angle $\gamma = 0^\circ$ and elevation angle $\beta = 0$ in the frame (X, Y, Z) (see Fig. 1). The value $R_0 = 10$ m is set as the reference distance used to normalize the phase of the received signal. This deliberately small value makes it possible to isolate the μ D signature from propagation attenuation ($L_{total} \propto R_0^{-2}$) in order to analyze the spectral structure under near-maximal signal-to-noise conditions. The results are independent of this normalization convention and remain valid for any operational range $R \gg \lambda$ satisfying the far-field assumption. The continuous-wave (CW) radar emits an analytic signal $(t) = \exp(j2\pi f_c t)$ at a carrier frequency $f_c = 9.81$ GHz (X-band), corresponding to a wavelength $\lambda \approx 3.06$ cm.

Table 1. Drone Parameters and Radar Configuration.

Parameter	Symbol	Value
<i>Drone Properties</i>		
Number of blades per rotor	N_b	2
Propeller radius	R_{pale}	12.7 cm
Hub radius	r_{hub}	0.45 cm
Blade length	$L = R_{pale} - r_{hub}$	12.25 cm
Chord at root / tip	C_{root} / C_{tip}	3.04 cm / 1.25 cm
Nominal rotation speed	Ω	3414.35 rpm
Radar/Drone distance	R_0	10 m
<i>Radar Configuration</i>		
Carrier frequency	f_c	9.81 GHz (X-band)
Pulse repetition frequency	PRF	20 kHz
Integration time	T_{int}	100 ms
Studied polarizations	-	HH, VV

The drone is modeled as a central body and $N_h = 4$ motors, each driving $N_b = 2$ blades of length $L = 0.1225$ m. The rotation speed is set to 3414.35 RPM (approximately 56.9 rps). In contrast to models found in the literature, each blade is here discretized into a set of M elementary scatterers along its radial axis, in order to capture local geometric variations such as *twist* and chord taper.

The robustness of the segmental model is assessed by varying the number of scatterers M per blade from 5 to 50. Fig. 3 shows that the μ D descriptors (spectral envelope, spectrogram structure, and CVD periodicity) converge from $M = 15$ segments onward. Below this threshold ($M \leq 10$), insufficient discretization produces an underestimation of the spectral spread and amplitude artifacts in the time-domain signal. Above it ($M \geq 20$), the descriptors are stable to within 1 % variation, confirming that $M = 15$ offers the optimal trade-off between physical fidelity and computational cost, since computation time grows linearly with M .




Fig. 1. Radar-rotor system modeling a hovering quadrotor drone using a segmental mesh of elementary scatterers.




Fig. 2. Detail of the propeller discretization with Taper and Twist parameters.




Fig. 3. Convergence of the received signal $s_r(t)$ with the number of segments per blade M .

The received signal $s_r(t)$ is the coherent sum of the contributions of the airframe and of the $N_h \times N_b$ blades. Neglecting higher-order couplings, the signal is expressed as:

$$s_r(t) = A_c e^{j\frac{4\pi}{\lambda} R_c(t)} + \sum_{h=0}^{N_h-1} \sum_{b=0}^{N_b-1} \int_{r_{hub}}^L \rho(r) e^{j\frac{4\pi}{\lambda} R_{h,b}(r,t)} dr \quad (3) + n(t),$$

where (x_m, y_m) are the motor positions and $\phi_{h,b,t} = \Omega t + \phi_0 + 2\pi b / N_b$ is the instantaneous rotation angle. The term (r) incorporates the material reflectivity and the polarization gain, computed from the projections of the surface normal vectors.

(t) is the distance to the drone's center, including micro-vibrations, and $R_{h,b}(r,t)$ is the instantaneous distance of the blade segment located at radial position r . This distance projects the rotational motion onto the radar line of sight according to:

$$R_{h,b}(r,t) = d_c + x_m \cos \beta \cos \gamma + y_m \cos \beta \sin \gamma + r \cos \beta \cos(\phi_{h,b,t} - \gamma) + z_p \sin \beta \quad (4)$$

From this modeling, we obtain the time-domain and frequency-domain signatures shown in Figs. 4a, 4b, and 4c. In Fig. 4a, the signal exhibits characteristic

amplitude peaks, called "blade flashes", occurring when the blades are perpendicular to the incident beam, thereby maximizing the instantaneous radar cross-section (RCS).

The time-domain signal shows $2 \times N_b$ amplitude peaks per full rotation (i.e. 8 peaks for this quadrotor equipped with two-bladed rotors), occurring when the radar line of sight (LOS) is normal to the longitudinal axis of the blades.

In Fig. 4b, the Doppler spectrum spans approximately $[-4.0, 4.0]$ kHz, centered at zero (since $v_{\text{drone}} = 0$), with amplitude extrema at its edges corresponding to the maximum tangential velocity at

the blade tip, $v_{\text{tip}} = \Omega \cdot L$. The positive-frequency side corresponds to blade segments approaching the receiver, while the negative-frequency side corresponds to segments moving away from it.

In Fig. 4c (spectrogram), the observed "flashes" confirm that most of the incident energy is backscattered when the blades cross the specular position facing the radar. The sinusoidal structures visible over $[-4.0, 4.0]$ kHz reflect the continuous phase modulation induced by rotation around the central Doppler frequency $f_D = 0$ Hz, since the drone is hovering.




Fig. 4. Time-domain signal, Fourier transform, and spectrogram of the hovering drone's μD signature.

The value $T_{\text{int}} = 100$ ms (Table 1) is the total observation duration used to generate the CVD, not the spectrogram's analysis window. At $\Omega_0 = 3414.35$ rpm, the rotor period is $T_{\text{rot}} \approx 17.6$ ms ($f_{\text{rot}} \approx 56.9$ Hz) and the blade-passage cadence is $f_{\text{cad}} = N_b f_{\text{rot}} \approx 113.8$ Hz.

$T_{\text{int}} \approx 5.7 T_{\text{rot}}$ gives a CVD resolution of $1/T_{\text{int}} = 10$ Hz, sufficient to resolve the ± 2.6 Hz cadence broadening from motor jitter p_r , while accumulating ≈ 11 blade-passage cycles for a stable cadence spectrum. The spectrogram itself (Fig. 4c) uses a shorter, 12.8 ms sliding window (50 % overlap) – inside the $[1/f_{\text{cad}}, t] = [8.8, 17.6]$ ms range required to resolve individual blade flashes while preserving quasi-stationarity of the instantaneous Doppler frequency.

To further characterize the drone, a Cadence-Velocity Diagram (CVD) is generated by applying a second Fourier transform along the time axis of the spectrogram, extracting the intrinsic periodicities of the μD signal as modulation frequencies (cadences). The CVD reveals a symmetric structure centered on zero modulation frequency; intensity peaks at specific cadences (here over $[-400, 400]$ Hz) correspond to the blade-flash repetition frequencies, isolating the fundamental rotation frequency Ω and its harmonics independently of instantaneous phase. This facilitates discrimination between platform types (e.g., drone vs. bird) by highlighting the mechanical regularity of the rotors.

The parameter p_r models the angular-velocity jitter of the brushless electric motors, inherent to load variations and imperfect ESC regulation. Its value of 2.25 % is representative of variations measured on 250 mm-class motors under hovering conditions [9]. For $\Omega = 3414.35$ rpm and $N_b = 2$, this dispersion induces $\Delta f_{mod} \approx \pm 2.6$ Hz around the nominal chopping frequency, consistent with the peak width observed in the CVD (Fig. 5) and distinct from numerical instability, whose signature would be aperiodic and unbounded.

While the Fourier transform bounds the frequency envelope, only the spectrogram (STFT) isolates the time dynamics of rotation; the CVD then extracts structural periodicities, enabling precise identification of the number of motors and their rotation regime. Together, these time frequency patterns form a biometric fingerprint of the drone, tied to Ω , L , N_b , and attitude relative to the sensor. Section 3 details the impact of polarimetry and material diversity on the radar signature.




Fig. 5. Cadence-Velocity Diagram (CVD) of the signal (Fig. 4a).

3. Polarimetric Analysis and Spectral Decomposition

The study of the μD signature is deepened through a comparative analysis of the HH and VV polarization channels, considering blades and airframe made of ABS polymer or carbon fiber. Polarization affects the segment-level RCS according to the instantaneous blade orientation and the electromagnetic properties of the material – complex permittivity $\epsilon_r = \epsilon_r' - j\epsilon_r''$ and conductivity σ . The received signal is extended by a scattering matrix incorporating the polarimetric gains and propagation losses [10]:

$$s_{r,pol}(t)s_{r,pol}(t) = L_{total} \left[s_{c,pol}(t) + \sum_{h=0}^{N_h-1} \sum_{b=0}^{N_b-1} \sum_{i=1}^M G_{pol}(\theta_{h,b,i}, \text{mat}) e^{j\frac{4\pi}{\lambda} R_{h,b,i}(t)} \right] + n(t) \quad (5)$$

The polarimetric gain G_{pol} projects the incident field vector onto the local normal $\hat{n}_{h,i}$ of each blade segment:

$$G_{HH} = (\hat{h} \cdot \hat{n})^2, G_{VV} = (\hat{v} \cdot \hat{n})^2, \quad (6)$$

where $\hat{h}$ and $\hat{v}$ are the horizontal/vertical polarization unit vectors in the radar frame, valid under physical optics for locally planar surfaces with $a \gg \lambda$ [3].

Thermal noise P_n and total attenuation L_{total} are modeled following [11]:

$$P_n = k_B T B F, L_{total} = \left(\frac{\lambda}{4\pi R_0} \right)^2 \cdot 10^{-\frac{\gamma_R R_0}{2000}}, \quad (7)$$

where k_B is the Boltzmann constant, $T = 290$ K is the system temperature, B is the bandwidth, F is the receiver noise figure, and γ_R is the specific rain attenuation (dB/km), computed at 9.81 GHz following ITU-R P.838 for a rain rate $R_r = 10$ mm/h.

3.1. Material Modeling

Materials are characterized by their electromagnetic properties in the X-band. Carbon fiber is treated as a good conductor ($\sigma_c = 6.25 \times 10^4$ S/m), such that its reflectivity is dominated by the conductive term [12]. ABS polymer is a nearly transparent dielectric ($\sigma_{ABS} \approx 10^{-15}$ S/m), whose reflection results from the impedance discontinuity at the air/material interface: $\epsilon_r' \approx 2.7$, $\tan \delta \approx 0.003$ at 10 GHz [13]. The normal-incidence reflection coefficient is then:

$$\Gamma = \frac{1 - \sqrt{\epsilon_r'}}{1 + \sqrt{\epsilon_r'}} \approx -0.24, \quad (8)$$

i.e., an intrinsic RCS for ABS about 12 dB lower than that of a perfect conductor of the same surface area – which explains, without fully accounting for, the observed difference between the carbon and ABS configurations.

3.2. Comparative HH and VV Analysis

As illustrated in Fig. 6, the HH channel shows a systematically higher amplitude than the VV channel at the reflection peaks. This dominance is explained physically by Eq. (6): when the blade is close to the specular position ($\beta \approx 0$, azimuth $\gamma \approx 90$), the local normal $\hat{n}$ preferentially aligns with the horizontal plane, maximizing G_{HH} and minimizing G_{VV} .

The fine structure of the VV channel shows more choppy modulations, attributable to variations of $(\hat{v} \cdot \hat{n})$ during rotation – this term changes sign twice per revolution, introducing projection nulls that are absent from the HH channel. This differential behavior constitutes a potential descriptor for classification: the

instantaneous ratio G_{HH}/G_{VV} encodes the blade tilt angle.

3.3. Material Impact: Carbon Fiber vs. ABS

Fig. 7 compares the two extreme blade/airframe configurations (all-carbon and all-ABS behave similarly to case a and follow the same trend, and are omitted for space). The high conductivity of carbon

fiber produces a blade RCS about 12–15 dB higher than that of ABS (consistent with $|\Gamma|^2$ computed in (8)), which translates into sharp Doppler flashes and a clear emergence of the μD signal above the noise floor P_n (case a). In the ABS/ABS configuration (case b), the signature is instead dominated by the continuous peak at $f_D = 0$ Hz from the airframe: the SNR of the μD components is then reduced by about 10 dB, making their extraction sensitive to the thermal noise level.




Fig. 6. Decomposition of the μD spectrum by polarization: global signature (left) and zoom on the fine structure (right). HH channel in blue, VV channel in black.




Fig. 7. HH/VV spectral response as a function of material (carbon- and ABS-dominated cases; the mixed configurations follow the same trend as (a) and are omitted for space). The RCS gap between carbon and ABS (≈ 12 –15 dB) is consistent with the Fresnel coefficient computed in (8).

4. Signature Analysis and Results

4.1. Smeared Effect Metric

To quantify the spectral fill-in, we define the spectral fill-in ratio τ_s as the ratio between the power contained in the inter-line band and the total power of the μD spectrum:

$$\tau_s = \frac{\int_{BW} \mathcal{S}(f) df - \sum_k P_k}{\int_{BW} \mathcal{S}(f) df}, \quad (9)$$

where $\mathcal{S}(f)$ is the normalized power spectral density, P_k is the power of the discrete lines at the chopping

frequency $k \cdot f_{mod}$, and BW is the -10 dB bandwidth. An ideal point-scatterer model tends toward $\tau_s \rightarrow 0$ (pure line spectrum), while a physically filled spectrum tends toward $\tau_s \rightarrow 1$.

4.2. Impact of Twist on the Smeared Effect

Comparing Fig. 8a and 8b, we observe that *Twist* generates the continuous spectral "fill-in." In the literature, this effect is often attributed to multiple reflections; our model suggests that it intrinsically arises from the phase decorrelation induced by blade twist.

4.3. Analytical Analysis of Phase Decorrelation

For a blade segment located at radial position r with a local Twist angle $\theta_t(r)$, the phase contribution to the received signal is:

$$\phi(r, t) = \frac{4\pi}{\lambda} r \cos \beta \cos(\Omega t + \phi_0) \cos \theta_t(r) \quad (10)$$

Without Twist ($\theta_t = \text{const}$), all segments contribute with the same factor $\cos \theta_t$, producing discrete spectral lines at the harmonics of f_{mod} . With a variable Twist:

$$\theta_t(r) = \theta_{root} + (\theta_{tip} - \theta_{root}) \frac{r - r_{hub}}{R_{pale} - r_{hub}}, \quad (11)$$

coherent integration over the blade introduces a continuous phase modulation:

$$s_{pale}(t) \propto \int_{r_{hub}}^{R_{pale}} \cos \theta_t(r) \cdot e^{j\frac{4\pi}{\lambda} r \cos \beta \cos(\Omega t)} dr \quad (12)$$

Table 2. Validation of μD descriptors – comparison between simulation and reference.

Descriptor	Ref.	Sim.	ϵ_{rel}	Source
f_{mod} (Hz)	116.67	113.81	2.45 %	[2] (measurement)
$BW_{-10 \text{ dB}}$ (Hz)	7235.7	7142.85	1.28 %	[1] (simulation)
Optimal polarization	HH	HH	—	[2] (measurement)

This integral is analogous to the Fourier transform of a non-rectangular window – Twist acts as an angular apodization that continuously broadens the line spectrum into a dense spectrum. Taper (r) then acts as an additional spatial weighting, redistributing energy from the blade tip toward the blade center.




Fig. 8. Stepwise validation of the Twist/Taper apodization: from the baseline point-scatterer spectrum (a) to the full segmental model (d).

4.4. Role of Taper and Full Model

Adding *Taper* redistributes power toward the center of the blade (Fig. 8c); combining both effects (Fig. 8d) yields the full segmental T+T model used throughout this paper.

5. Validation by Comparison with the Literature

The fidelity of the model is assessed by comparison with the reference scenarios of Xu *et al.* [1] and Al-Nuaimi *et al.* [2], selected for the availability of their numerical and experimental data on geometries close to our configuration. The propeller parameters are rescaled to $\Omega = 3414.35$ rpm and $R_{pale} = 15$ cm to allow direct comparison.

5.1. Modulation Frequency and Temporal Coherence

The chopping frequency $f_{mod} = N_b \Omega / 60$ is the primary descriptor for identifying the number of blades. Our model predicts $f_{mod} = 113.81$ Hz, a deviation of 2.45 % from the value measured experimentally by Al-Nuaimi [2] on a rotor with comparable geometry in an anechoic chamber ($f_{mod,exp} = 116.67$ Hz). This deviation, below the 5 % threshold used as a convergence criterion [1], is attributable to the difference in effective radius between the two configurations ($\Delta R = 2.3$ cm).

The Doppler bandwidth BW is evaluated as:

$$\epsilon_{rel} = \frac{|BW_{simu} - BW_{lit}|}{BW_{lit} \times 100} \quad (13)$$

It should be noted that reference [1] relies on a MoM simulation, whereas [2] includes experimental measurements in an anechoic chamber: the two validation levels are therefore distinct and complementary. A measurement campaign on our own configuration is planned to address this limitation (see Section 6).

5.2. Energy Distribution and Smeared Effect

The joint integration of *Twist* and *Taper* produces a power spectral density (PSD) that is qualitatively consistent with the distributions published by Xu *et al.* [1]: energy is redistributed from the discrete harmonic lines toward the inter-line continuum, in agreement with the *Smeared Effect* observed experimentally. Our spectral fill-in ratio $\tau_s \approx 0.82$ (T+T model) is consistent with the bandwidth $BW = 7235$ Hz reported for Type 4 in Kang *et al.* [9], although the definition of τ_s differs from their metric – a cross-normalization remains to be formalized.

In a near-field configuration, Xu *et al.* [1] document a temporal tilt of the blade flashes (*slanting effect*), qualitatively reproduced by our spectrogram

(Fig. 4c); a direct figure-to-figure comparison would require a full geometric recalibration, left for future work.

Taken together, these results – relative errors below 3 % on the key descriptors, qualitative consistency of the PSD, and consistency of the time-domain structure – validate the robustness of the segmental T+T model against classical point-scatterer approaches, while clearly identifying its current limitations.

6. Conclusion and Future Work

This paper presented a segmental deterministic model of the micro-Doppler (μD) signature of micro-UAV drones in the X-band, jointly integrating *Twist* (angular pitch variation) and *Taper* (chord tapering) – two geometric parameters often neglected in analytical models. We showed that a classical point-scatterer model underestimates the spectral spread, whereas the full T+T model converges toward literature reference data with a relative error below 3 %. The polarimetric analysis further showed that the HH channel offers better sensitivity to Doppler flashes, and that blade/airframe material (carbon fiber vs. ABS) radically changes the segment-level RCS dynamics. As a fast, physically-grounded simulator, this model is well suited to generating large synthetic datasets for training UAV/bird classification algorithms (CNN/RNN), which rely on accurate fine spectral structure.

Future work will integrate aerodynamic effects such as blade *flapping* and motor-regime variations during unstable hover, and an anechoic-chamber measurement campaign is planned to validate the model against our own configuration and to quantify the impact of material composition on signal depolarization.

Acknowledgment

The authors would like to thank the AID-DGA (Agence de l'Innovation de Défense – Direction Générale de l'Armement) for their support in this research work.

References

- [1]. S. Xu, G. Luo, L. Wang, Z. Chen, Modelling and analysis of micro-Doppler signatures of rotating blades in near-field for high-resolution radars, *Remote Sensing*, Vol. 15, Issue 12, 2023, 3060.
- [2]. M. S. Al-Nuaimi, I. F. Al-Shaikhli, J. S. Mandeep, Characterization of radar micro-Doppler signature of small UAVs rotating blades based on MoM simulation and measurement, *IEEE Access*, Vol. 7, 2019, pp. 138136-138150.
- [3]. V. C. Chen, The Micro-Doppler Effect in Radar (2nd Ed.), *Artech House*, Boston, 2019.
- [4]. F. Zhou, Y. Yang, S. Wu, Analysis of micro-Doppler signature of helicopter rotating blades, in *Proceedings*

- of the *IEEE International Conference on Mechatronics and Automation (ICMA)*, 2006, pp. 2110-2114.
- [5]. M. A. Richards, *Fundamentals of Radar Signal Processing* (2nd Ed.), *McGraw-Hill Education*, New York, 2014.
 - [6]. H. C. Alves Costa, S. J. Myint, C. Andrich, S. W. Giehl, et al., Modeling micro-Doppler signature of multi-propeller drones in distributed ISAC, *IEEE Journal of Selected Topics in Electromagnetics, Antennas and Propagation*, Vol. 1, Issue 1, 2025, pp. 208-222.
 - [7]. J. Yan, H. Hu, J. Gong, D. Kong, et al., Exploring radar micro-Doppler signatures for recognition of drone types, *Drones*, Vol. 7, Issue 4, 2023, 280.
 - [8]. K. B. Kang, J. H. Choi, B. L. Cho, J. S. Lee, et al., Analysis of micro-Doppler signatures of small UAVs based on Doppler spectrum, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 57, Issue 5, 2021, pp. 3252-3267.
 - [9]. T. Li, B. Wen, Y. Tian, S. Wang, et al., Optimized radar waveform parameter design for small drone detection based on echo modeling and experimental analysis, *IEEE Access*, Vol. 7, 2019, pp. 101527-101538.
 - [10]. M. Alae-Kerahroodi, P. Babu, M. Soltanian, M. B. Shankar, *Signal Design for Modern Radar Systems*, *Artech House*, Norwood, 2022.
 - [11]. B. R. Mahafza, *Radar Signal Analysis and Processing Using MATLAB*, *Chapman and Hall/CRC*, Boca Raton, 2016.
 - [12]. Z. Li, A. Haigh, C. Soutis, A. Gibson, Principles and applications of microwave testing for woven and non-woven carbon fibre-reinforced polymer composites: A topical review, *Advanced Composites Letters*, Vol. 25, Issue 4, 2018, pp. 965-982.
 - [13]. B. Riddle, J. Baker-Jarvis, J. Krupka, Complex permittivity measurements of common plastics over variable temperatures, *IEEE Transactions on Microwave Theory and Techniques*, Vol. 51, Issue 3, 2003, pp. 727-733.
 - [14]. S. Borderieux, J.-C. Cexus, A. Khenchaf, P. Pouliguen, Influence of the propeller blade orientation on the micro-Doppler signatures of UAVs, in *Proceedings of the 20th European Conference on Antennas and Propagation (EuCAP)*, Dublin, Ireland, 2026, pp. 1-5.

Impact of Sparsity and Algorithm Parameters on SFT-based Wideband Spectrum Sensing

N. Luberne^{1,2}, A. Khenchaf¹, F. Comblet¹, B. Gorisse² and M. Bauche²

¹ENSTA, Lab-STICC UMR CNRS 6285, 2 r. François Verny, 29200 Brest, France

²THALES, 10 av 1ere DFL, 29200, Brest, France

Tel.: + 33 783387125

E-mail: nicolas.luberne@ensta.fr

Summary: This paper addresses wideband sensing and spectrum monitoring by targeting efficient estimation of spectral occupancy in increasingly crowded wireless environments. Specifically, we consider sparse spectrum recovery: for a complex signal of length N , only $S \ll N$ coefficients of its Discrete Fourier Transform (DFT) are assumed to carry significant energy, while the remaining coefficients are assumed negligible. Efficiently identifying these S active frequency bins is a key requirement in wideband sensing, cognitive radio, and signal compression. We study a specific Sparse Fourier Transform (SFT) voting scheme, a non-iterative method that estimates significant spectral components using pseudo-random permutations, sub-sampling, and short FFTs. Unlike iterative compressed-sensing approaches, the SFT achieves favorable computational complexity, making it suitable for real-time spectrum monitoring. A central contribution is an analytical characterization of the SFT voting stage. Under the assumptions that the active buckets have been correctly recovered and that inactive candidates are uniformly distributed, we derive a closed-form expression for the false-alarm probability P_f as an explicit function of the algorithm parameters: the number of branches L , the sub-sampling factor B , the number of selected peaks per branch M , and the voting threshold G . We then empirically investigate how the peak-selection parameter M influences both the probability of correct detection P_d and the false-alarm rate P_f revealing the resulting detection–false-alarm trade-off. The analytical predictions for P_f are validated through 10000 Monte Carlo trials, showing strong agreement between model and experiments. The analysis is carried out in the noiseless setting with integer-valued frequency bins, providing a rigorous baseline for future extensions to noisy and non-integer frequency scenarios.

Keywords: Sparse Fourier transform, Sparse spectrum, Wideband sensing, Detection and false alarm rates.

1. Introduction

Wideband sensing and spectrum monitoring are critical challenges in modern Cognitive Radio (CR) systems, where the growing density of wireless users demands efficient and reliable surveillance of wide frequency ranges. Identifying unoccupied spectral resources (commonly referred to as ‘white spaces’) requires both high spectral resolution and low computational overhead, two objectives that are often in tension.

Two major approaches are commonly considered for wideband spectrum sensing: sub-Nyquist sampling, such as Compressed Sensing (CS) [1-3], and direct sampling (DS) method.

Compressed Sensing enables wideband spectrum recovery using fewer samples than the traditional Nyquist rate, thereby reducing the requirements for high-speed and high-cost analog-to-digital converters (ADCs), at the expense of relying on computationally intensive, iterative recovery algorithms. In contrast, direct sampling (DS) leverages conventional spectrum analysis techniques, including the Discrete Fourier Transform (DFT) and the Wigner-Ville Distribution (WVD). For DFT computation, the Fast Fourier Transform (FFT) algorithm [4] is widely used due to its high efficiency.

However, as high sampling rates and fine spectral resolution become increasingly feasible, the FFT processes a large number of frequency bins, even

though, in practice, only a small subset may contain significant energy i.e., the signal spectrum is sparse [5]. This leads to inefficiencies in both computation and resource utilization.

The Sparse Fourier Transform (SFT) [5-11] addresses this limitation by efficiently estimating the significant frequency bins through a sequence of shorter FFTs and targeted processing. Several SFT variants exist and are briefly reviewed in Section 2; this paper focuses on a specific voting scheme [6].

Despite its practical relevance, the performance of the SFT voting scheme remains insufficiently characterized, in particular its false-alarm behavior as a function of the algorithm parameters. While detection performance has been studied for other SFT-based algorithms [8, 9], the present paper focuses on the voting stage of the scheme in a noiseless setting with integer-valued frequency bins. Under the assumption of correct recovery of the active buckets and uniformly distributed inactive candidates, we derive a closed-form expression for the false-alarm probability. We then empirically validate the theoretical results through extensive Monte Carlo simulations. This provides a first step toward a more general performance characterization.

The remainder of this paper is organized as follows. Section 2 presents the problem statement. Section 3 offers a detailed overview of the SFT algorithm, while Section 4 describes the proposed theoretical model and the main performance metrics. Section 5 discusses the

experimental results, and the paper concludes with final remarks and perspectives for future research.

2. Problem Statement & Sparse Fourier Transform Principle

In this paper, we consider a discrete signal x of length N . The signal is assumed to be sparse in the Fourier domain, modeled as the sum of S pure frequency components, with $S \ll N$. Specifically, x is given by [4]:

$$x(t) = \sum_{k=1}^S A_k e^{\frac{j2\pi n_k t}{N}}, \quad (1)$$

where A_k represents the amplitude of the k -th frequency component and n_k is the corresponding frequency bin (centered within one FFT bin for simplicity). Our objective is to accurately recover the indices n_k without the knowledge that only S are present out of the N possible.

3. SFT Algorithm Overview

The term Sparse Fourier Transform (SFT) covers distinct algorithm families for sparse spectrum recovery using short FFTs. These methods differ notably in their use of sub-sampling, sample permutation, reconstruction strategy, or decision rule. For example, some approaches rely on different sub-sampling factors with collision detection [5] and direct resolution of the frequency bins [8, 10], while permutation-based methods may use summation with a voting scheme [6, 7] or, alternatively without voting [9]. The main advantage of the permutation-based approach is that it yields a “uniformized” spectrum after permutation, even when the original spectrum contains clusters [11]; accordingly, this paper focuses on the permutation-based SFT method proposed in [6, 7].

The permutation-based SFT method proceeds as follows for each branch, where the number L of branches is one of the parameters which will be studied and commonly set to $L = O(\log N)$ [7], before aggregating the results in the voting stage:

1. Pseudo-random permutation:
The N samples are permuted according to a pseudo-random index σ coprime with N . This mixing ensures each active frequency appears at a different position for each permutation [6].
2. Windowing:
A weighting window is applied, which preserves information through the following sub-sampling step [7].
3. Sub-sampling and short FFTs (“Buckets”):
The permuted signal is sub-sampled by a factor B , via summation [6]. A reduced-size FFT is then computed on the resulting bucket.
4. Peak detection:

In each bucket, local maxima are identified. The positions J of the M largest peaks are retained as potential active frequencies.

5. Spectrum unfolding:
Using the permutation and sub-sampling parameters, the detected peaks are mapped back to candidate frequency bins in the original N -point DFT via the hashing function h_σ .
6. Voting and aggregation:
After L iterations, frequency bins that receive at least G votes (out of a maximum of L) are declared active and correspond to true signal components.

In this paper, we investigate the relationship between S , L , B , M , and G , and analyze the probability of correctly recovering the occupied frequency bins, even in the noiseless case.

4. Theoretical Model and Performance Metrics

The proposed analytical model relies on two simplifying assumptions:

- H1: All expected (active) frequencies are correctly recovered. The scope and validity of this hypothesis are specifically investigated;
- H2: Unexpected (inactive) frequency bins, i.e., those not corresponding to true signal components, are assumed to be uniformly distributed across the N -point DFT. This uniformity is supported by the empirical distribution of the inactive voted bins, illustrated in Fig. 1. In addition, a chi-square (χ^2) goodness-of-fit test was performed, and it does not allow rejection of the uniformity assumption (p-value of 1).




Fig. 1. Histogram of the voted inactive bins for $N = 1024$, $B = 32$, $S = 7$ and $M = 12$.

It is important to note that H1 may not always be valid in practice. Two common scenarios lead to its failure:

- $M < S$: The number of selected peaks per branch is insufficient compared to the actual number of active frequencies. As a result, some components cannot be captured or resolved;

- $M \geq S$, but the frequency index n_k is not exactly aligned with a measurement bucket. In this case, a component's energy may be spread across multiple buckets, potentially leading to missed detections, especially when M is too small to ensure all active buckets are captured.

The impact of these two scenarios will be highlighted in Section 5 through experimental results.

If the active buckets are truly identified, it is easy to verify that:

$$\{n, |X(n)| > 0\} \subset \cap_{\sigma} I_{\sigma}, \quad (2)$$

where

$$I_{\sigma} = \{n \mid h_{\sigma}(n) \in J, n \in \llbracket 0, N-1 \rrbracket\} \quad (3)$$

Assumption H2 holds as long as the selected inactive buckets result from random computational noise and the random selection of permutation indices, ensuring that any "false" detections remain uniformly distributed over all possible frequency bins.

Assuming both hypotheses H1 and H2 are valid, the metrics are defined as follows:

- P_d is the probability that an expected frequency bin is declared active. When H1 holds, this implies that $P_d = 1$;
- P_f is the probability that an inactive frequency bin is incorrectly declared active.

For one branch, the probability that an inactive frequency bin is voted is

$$P(n \in I_{\sigma} \mid n \notin \{n_k\}) = \frac{MB-S}{N-S} \quad (4)$$

Although votes coming from different branches are not strictly independent for finite N (because they are computed from the same underlying signal and differ only through the random permutations/sub-samplings), the analytical model treats these branch-wise votes as independent. Under this standard approximation, the number of votes received by an inactive bin is modeled as a Bernoulli process, and the resulting false alarm probability P_f is given by a binomial tail (5). Empirically, this approximation is validated by the agreement between the predicted and observed vote statistics.

$$P_f = \sum_{k=G}^L \binom{L}{k} \left(\frac{MB-S}{N-S}\right)^k \left(1 - \frac{MB-S}{N-S}\right)^{L-k} \quad (5)$$

This expression therefore quantifies the probability that an inactive bin is incorrectly declared active by receiving at least G votes across the L branches.

5. Experimental Results

To validate the theoretical performance analysis, we conducted a set of experiments with the following parameter settings: $N = 1024$ and $L = \log_2(N)$. The value of L is chosen according to [6]. The performance

of the SFT algorithm is primarily governed by the interplay between the signal sparsity and the number of selected peaks per branch, i.e., the validity of hypothesis H1. Therefore, our experiments investigate how these two quantities affect both the probability of correct detection and the false alarm probability. In addition, we select different values of B and G , and we run 10000 independent trials for each configuration. For each trial, the set of active frequency bins and the permutation indices are drawn randomly to ensure statistical generality. By systematically analyzing the probability of correct detection and the false alarm rate for these different configurations, we assess the practical efficacy of the SFT algorithm and compare empirical results with theoretical expectations.

5.1. Effect of Sparsity (with a Fixed Number of Selected Peaks)

The sparsity level S is a crucial parameter in SFT-based recovery. In most practical scenarios, S is not known a priori, which makes it necessary to assess how performances changes with S while controlling the per-branch selection budget. In this section, we study the impact of the sparsity parameter S . To this end, we keep M fixed and vary S across different experimental configurations, defined by the subsampling factor B and the voting threshold G . This allows us to isolate the impact of the number of true active components on both detection reliability and the false alarm behavior.

Fig. 2. depicts the measured probability of correct detection P_d as a function of the number of active frequency bins S . As anticipated, for large values of S (i.e. $S > M$), P_d falls progressively below 1. In that regime, reducing S improves the probability of recovering all active buckets, which progressively supports hypothesis H1.

This observation also suggests that the voting threshold G plays a key role. Indeed, when $S > M$, not all S active components can be detected on every branch; in this case, employing a lower threshold G mitigates missed detections across branches, leading to better recovery.




Fig. 2. Evolution of the detection probability P_d as a function of S for $M = 16$.

Fig. 3. compares the measured false alarm probability P_f as a function of the number of active frequency bins S with the theoretical prediction given by (5). Firstly, for fixed values of B, G and M, P_f remains nearly constant as S varies. This behavior is explained by the assumption $S \ll N$, under which the dependence of the per-branch vote in (5) on S becomes negligible. Secondly, when $B = 32, M = 16$ and $G = 5$, we observe that $P_f > 0.5$. In that configuration, the total number of frequency bins that are voted across all branches is $BM=N/2$, meaning for each branch half of the bins are selected as candidates. As a consequence, a large fraction of the voted frequencies corresponds to inactive bins, which leads to high false alarm probability.




Fig. 3. Comparison of theoretical and measured false alarm probabilities as a function of S for $M = 16$ across different (B, G) settings.

5.2. Effect of the Number of Selected Peaks (with a Fixed Number of Input Frequencies)

After studying how the sparsity level S affects the recovery performance (section 5.1), we now turn to the complementary factor that drives the SFT voting-scheme decision process: the number of selected peaks per branch, denoted by M .

In practice, S is generally unknown, and cannot be chosen by the receiver, whereas M is a key algorithm parameter that the designer can set.

Accordingly, we fix the sparsity level S to a worst-case situation (arbitrary value in the example), corresponding to the maximum number of input active frequency bins that the system may encounter. This analysis is performed for multiple (B, G) pairs, similarly to Section 5.1, in order to verify how well the theoretical predictions match the empirical behavior under different algorithm settings.

Fig. 4. shows the measured probability of correct detection P_d as a function of the number of detected peaks per branch M . As expected, for small M, P_d is strictly below 1. Increasing M improves the algorithm's ability to recover all active buckets, which progressively corroborates hypothesis H1. The missed detections observed at a low M are therefore attributed to occasional failures of H1, meaning that the algorithm may not be able to capture all active

frequency bins when too few peaks are selected. Compared with Fig. 2., the detection curves clearly reveal the second failure scenario of H1: even when $M \geq S$, the probability of detection remains below 1.

Finally, as in Fig. 2., the voting threshold G has a significant impact on the detection behavior: larger values of G lead to a more pronounced (steeper) decline of P_d in the regime where $M < S$.




Fig. 4. Evolution of the detection probability P_d as a function of M for $S = 7$.

Fig. 5. compares the measured false alarm probability with the expected theoretical values derived earlier. As expected, for small values of M , hypothesis H1 is sometimes invalid, which results in an underestimation of P_f relative to the theoretical predictions. As M increases, the measured false alarm probability converges to the theoretical values.

In comparison with Fig. 3, the false alarm rate increases with M . This behavior is anticipated because M plays a central role in (5), particularly through the product MB . It is also worth noting that the alarm rate approaches 1 as MB tends to N . Since $N = MB$, this condition represents an upper bound on either M or B : it implies that every bin is voted in each branch, so all bins are considered active.

These observations lead to the conclusion that achieving a good trade-off between detection performance and false alarm rate requires selecting competitive values of M and B such that $MB \approx S$, and in particular $MB > S$. Finally, for the configuration $B = 16$ and $G = 8$, the number of Monte Carlo trials is insufficient to reliably estimate false alarm rates at low probabilities.




Fig. 5. Comparison of theoretical and measured false alarm probabilities as a function of M for $S = 7$ across different (B, G) settings.

6. Conclusion and Future Work

In this paper, we analyze a spectrum recovery algorithm under a sparse spectrum hypothesis: the permutation- and-voting-based SFT. Even in a noiseless setting, the voting mechanism can lead to both detection failure and false alarm, motivating the need for a quantitative performance model.

We derive a closed-form analytical expression for the false alarm probability P_f as a function of the SFT parameters (N, B, M, G) and the sparsity parameter S . This model provides a practical framework for tuning the algorithm to achieve a desired trade-off between detection reliability and false alarm control. The theoretical predictions are validated through extensive numerical experiments (10000 Monte Carlo trials).

The numerical results show that the sparsity parameter S has a negligible impact on P_f as long as $S \ll N$, while the product MB plays a key role in determining the false alarm behavior. Specifically, increasing MB increases the number of candidate voted bins, which improves detection reliability but can rapidly raise P_f when the candidate set becomes too large. Consequently, reliable performances require selecting algorithm parameters such that MB remains sufficiently close to (and typically larger than) the effective sparsity parameter. A desirable criterion is $P_f < 1/N$, which corresponds to an average of fewer than one false alarm coefficient per spectrum recovery trial. In the simulation setting ($N = 1024, L = 10, S \approx 10$), using $B = 16, M = 16$ and $G = 8$ gives $P_f \approx 3.3 \times 10^{-4}$. This corresponds to roughly one false alarm coefficient per 10 spectrum recovery trials.

Finally, our study assumes integer-valued frequency components and a noiseless observation model. Future work will extend the proposed permutation- and-voting-based SFT framework to handle non-integer (off-grid) frequencies and more realistic noisy scenarios, while incorporating the corresponding impact on both P_d and P_f .

Acknowledgements

This work was carried out within the framework of the CORMORANT GIS (Scientific Interest Group) in the context of a partnership between Thales and Lab-STICC CNRS UMR 6285.

References

- [1]. M. Mishali, Y. C. Eldar, From theory to practice: Sub-Nyquist sampling of sparse wideband analog signals, *IEEE Journal of Selected Topics in Signal Processing*, Vol. 4, Issue 2, 2010, pp. 375-391.
- [2]. R. Sharma, R. Shrestha, S. K. Sharma, Hardware-efficient and short sensing-time multicoset-sampling based wideband spectrum sensor for cognitive radio network, *IEEE Transactions on Circuits and Systems I: Regular Papers*, Vol. 70, Issue 3, 2023, pp. 1298-1310.
- [3]. J. C. Martin, Analysis of the Nyquist folding receiver (NYFR), PhD Thesis, *University of Oklahoma*, Norman, 2018.
- [4]. J. W. Cooley, J. W. Tukey, An algorithm for the machine calculation of complex Fourier series, *Mathematics of Computation*, Vol. 19, Issue 90, 1965, pp. 297-301.
- [5]. H. Hassanieh, L. Shi, GHz-wide sensing and decoding using the sparse Fourier transform, in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM)*, 2014, pp. 2256-2264.
- [6]. J. Hou, J. Lv, Y. Zhang, Z. Ma, et al., A spectrum monitoring and analyzing algorithm based on sparse Fourier transform, in *Proceedings of the 7th International Conference on Computer and Communications (ICCC)*, 2021, pp. 2308-2312.
- [7]. H. Hassanieh, P. Indyk, D. Katabi, E. Price, Simple and practical algorithm for sparse Fourier transform, in *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2012, pp. 1183-1194.
- [8]. S. Pawar, K. Ramchandran, R-FFAST: A robust sub-linear time algorithm for computing a sparse DFT, *IEEE Transactions on Information Theory*, Vol. 64, Issue 1, 2018, pp. 451-466.
- [9]. S. Wang, V. M. Patel, A. Petropulu, The robust sparse Fourier transform (RSFT) and its application in radar signal processing, *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 53, Issue 6, 2017, pp. 2735-2755.
- [10]. S. Pawar, K. Ramchandran, FFAST: An algorithm for computing an exactly k-sparse DFT in $O(k \log k)$ time, *IEEE Transactions on Information Theory*, Vol. 64, Issue 1, 2018, pp. 429-450.
- [11]. A. Rauh, G. R. Arce, Optimized spectrum permutation for the multidimensional sparse FFT, *IEEE Transactions on Signal Processing*, Vol. 65, Issue 1, 2017, pp. 162-172.

Design and Evaluation of an ESP32-Based Wireless Sensor Network for Subsurface Temperature Monitoring in Road Energy-Harvesting Applications

T. Bača, P. Šarašin, O. Karpiš and M. Formanek

University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1, 010 26 Žilina, Slovakia
Tel.: + 421 41/513 4365
E-mail: tomas.baca@fri.uniza.sk

Summary: Thermoelectric energy harvesting is a promising approach for powering low-power sensor nodes in locations with stable temperature gradients. Road infrastructure represents an ideal application because the road surface and deeper subgrade layers exhibit distinct thermal behavior during daily and seasonal variations. This paper presents the design and experimental evaluation of a wireless sensor network developed to collect subsurface temperature data for evaluating road-based thermoelectric energy harvesting. The proposed system consists of a battery-powered ESP32-C6 sensor node and a central node that forwards measured data to cloud storage. Local data buffering is implemented in both nodes to eliminate data loss during temporary wireless or Internet outages. The prototype was evaluated in an outdoor environment over a two-week period. The results verify the functionality of the measurement system, including cloud data transfer and robust recovery after simulated communication outages.

Keywords: Thermoelectric energy harvesting, Wireless sensor network, Subsurface temperature monitoring, ESP32-C6, Road infrastructure.

1. Introduction

With the increasing number of battery-powered devices, methods to decrease energy consumption and prolong the regular battery replacement period have become important design challenges in embedded and wireless sensor systems. The operational lifetime of battery-powered sensor nodes can be extended by reducing both hardware and software energy consumption. Hardware level approaches include the selection of low-power microcontrollers, sensors, and communication modules, as well as disconnecting inactive peripherals from the power supply. Software level approaches include duty-cycled operation, where the control unit spends most of its time in sleep mode and wakes up only for measurement, processing, and wireless communication.

However, reducing energy consumption represents only a partial solution for prolonging the operation time of battery-powered devices, since the battery will still eventually discharge. Energy harvesting can significantly extend the operational lifetime of battery-powered devices and enable maintenance-free operation over a defined period. Photovoltaic harvesting is commonly used for outdoor low-power systems. However, it is not suitable in all locations or operating conditions. Thermoelectric harvesting represents an alternative when a usable thermal gradient is available.

The main contribution of this paper is the design and preliminary outdoor validation of a battery-powered wireless sensor network for collecting subsurface temperature data required for the evaluation of road-based thermoelectric energy harvesting. The proposed system combines multi-depth temperature

measurement, low-power sensor node operation, local data buffering, and cloud synchronization. The paper focuses on the measurement infrastructure rather than on the final thermoelectric generator prototype.

Unlike our previous work [6], which reviewed thermoelectric transducer principles and general smart-road applicability, this paper does not analyze thermoelectric harvesting itself. Its contribution lies in the data-acquisition layer: the design, implementation, and outdoor validation of a low-cost ESP32-C6 wireless sensor network with dual-node local buffering and automatic recovery, developed to support future thermoelectric road-harvesting studies with long-term subsurface thermal profiling.

2. Thermoelectric Energy-Harvesting Systems

Thermoelectric energy harvesting systems in the form of Seebeck generators are capable of generating electrical energy when exposed to a thermal gradient, which means a temperature difference between the two sides of the thermoelectric generator [1-3]. The generated open-circuit voltage V_{OC} can be calculated as:

$$V_{OC} = S \cdot \Delta T, \quad (1)$$

where S is the Seebeck coefficient, which depends on the used thermoelectric generator, and ΔT is the thermal gradient. The thermoelectric generator can be electrically modeled as a real voltage source supplying the open-circuit voltage V_{OC} with an internal resistance R_{IN} . The current I passing through a load with resistance R_{LOAD} can be calculated as:

$$I = \frac{V_{OC}}{R_{IN} + R_{LOAD}} = \frac{S \cdot \Delta T}{R_{IN} + R_{LOAD}} \quad (2)$$

Based on the maximum power transfer theorem, the voltage source delivers the maximum possible power to the load when the load resistance R_{LOAD} equals the internal resistance R_{IN} of the voltage source:

$$R = R_{LOAD} = R_{IN} \quad (3)$$

The maximal delivered current can, under these conditions, be calculated as:

$$I_{MAX} = \frac{V_{OC}}{2 \cdot R} = \frac{S \cdot \Delta T}{2 \cdot R} \quad (4)$$

The maximal delivered power can, under this condition, be calculated as:

$$P_{MAX} = \frac{V_{OC}^2}{4 \cdot R} = \frac{(S \cdot \Delta T)^2}{4 \cdot R} \quad (5)$$

Thermoelectric energy harvesting systems represent an interesting alternative for applications with naturally occurring thermal gradients where one object can act as a passive heat source and another object as a passive cooler. Naturally occurring thermal gradients can reverse their polarity over time. For example, the side heated by solar radiation during the day may become the colder side at night. This reversal of the thermal gradient results in a change in the output voltage polarity of the thermoelectric generator. Consequently, the energy management subsystem connected directly to the output of the thermoelectric generator must be capable of keeping the system close to its optimal operating point throughout the entire operation. At the same time, it must convert the generated energy to the voltage level required by the subsequent components, even when the output voltage polarity reverses.

These conditions are met by several integrated circuits designed to serve as thermoelectric generator energy management systems. For systems in which the polarity of the temperature gradient can reverse, suitable power management circuits must support operation with very low input voltages and bipolar input conditions. An example of such a circuit is the LTC3109, which is designed for energy harvesting from low-voltage thermoelectric generators.

3. Road-Based Thermoelectric Energy-Harvesting Applications

Our research group focuses on applications of thermoelectric energy harvesting systems in transportation, an area that offers multiple possibilities where naturally occurring thermal gradients are present. In the project described in this article, we focused on the thermal gradient that occurs between the road surface and the subgrade layers [4-7]. The subgrade or subsurface soil layers have a significantly

higher effective thermal inertia than the road surface. Therefore, their temperature changes more slowly during daily ambient temperature variations. The road surface has a much smaller thermal capacity. As a consequence, its temperature is more strongly affected by changes in ambient temperature and the level of solar radiation to which it is exposed.

During the summer months, the road surface represents the hot side of the thermoelectric generator, and the subgrade layer represents the cold side. During the winter months, the situation reverses. The road surface becomes the cold side of the thermoelectric generator, and the subgrade layer becomes the hot side. The transitional seasons are especially important because the magnitude and polarity of the temperature gradient can vary more rapidly. During these periods, the available temperature difference may become too small for continuous operation of a thermoelectric energy harvesting system, which can result in interruptions in system operation.

One approach is to construct a complete thermoelectric harvesting prototype, place one side of the thermoelectric generator near the road surface, thermally connect the other side to a selected subgrade depth, and measure the generated energy over a defined period. Although this approach provides direct experimental results, the outcome strongly depends on the selected generator, heat conductor, insulation, and power management circuit. The total energy output is limited by the thermal resistance of the conductive rod, the level of thermal insulation, the Seebeck coefficient, and the conversion efficiency of the power management circuit. Consequently, this approach can provide only the amount of generated electrical energy under a specific set of conditions.

A second approach is the measurement of ambient air temperature, road surface temperature, and soil temperature at different depths. Based on the measured datasets and mathematical models describing heat conduction in materials, thermal losses, and the efficiency of the circuits, the total amount of generated energy at different depths can be predicted by calculation. This allows evaluation of many different scenarios where the properties of the components can be varied. In the current prototype evaluation, subsurface and ambient temperatures were measured. In the long-term deployment, road surface temperature will be obtained either from an additional surface sensor or from a road weather dataset and will be used together with the subsurface measurements for energy output estimation. However, the results of measurements and mathematical models are still only approximations of real-world behavior. Scenarios that appear optimal must still be practically validated by constructing the complete system according to the first approach.

4. Wireless Sensor Network

A wireless sensor network (WSN) consists of sensor nodes that acquire physical quantities and

transmit measured data to a central node or gateway. In the proposed system, the sensor node performs temperature measurement and local buffering, while the central node receives the data and forwards them to cloud storage [8-17]. Both the sensor node and the central node are based on the ESP32-C6 microcontroller. The sensor node was designed as a standalone printed circuit board tailored for outdoor temperature monitoring, as shown in Fig. 1.

The board integrates an ESP32-C6 module, connectors for external temperature sensors, battery supply circuitry, and voltage regulation within a waterproof enclosure. The sensor node is powered by two Li-Ion batteries connected in series, providing a voltage in the range of 7.0 V to 8.4 V. This voltage is converted to the 3.3 V required by the microcontroller and sensors by a high-efficiency switching power supply.



Fig. 1. Prototype of the battery-powered ESP32-C6 sensor node installed in the outdoor environment.

Since the sensor node is battery-powered, it spends the majority of its time in deep sleep mode. Once every 10 minutes, the sensor node wakes up, measures the current temperature from all connected sensors, stores the measured data in a local database stored in flash memory, and attempts to transfer the data to the central node using the low-power ESP-NOW wireless communication protocol. After completing these steps, the sensor node returns to deep sleep mode.

The central node is powered from a mains-connected 5 V power supply. As a consequence, the central node can remain constantly in the active state and receive new data from the sensor node at any time. This approach removes the need for complex time synchronization between nodes. The central node is implemented as two standalone development boards containing the ESP32-C6 microcontroller,

interconnected using the wired UART communication protocol. The first development board receives data from the sensor node through ESP-NOW and stores them in its own local database stored in flash memory. Afterwards, the received data are transferred to the second development board via UART, which attempts to send them over Wi-Fi to the cloud data storage in the ThingSpeak application.

Since both the sensor node and the central node maintain local data buffers, temporary communication failures do not result in data loss. At the beginning of communication, the central node requests the sample number ID of the last stored measurement in the sensor node database and compares it with its own database to download all missing data. Simultaneously, the central node keeps track of the ID of the last dataset successfully transferred over Wi-Fi to ThingSpeak. Therefore, even when Wi-Fi becomes temporarily unavailable, the central node can later transfer all missing data to the cloud storage in a burst. The database is implemented using binary files to reduce memory usage. Each record contains one measurement sample, including measured temperatures and battery voltage. The sampling interval of 10 minutes results in 144 records per day.

5. Experimental Evaluation of the Wireless Sensor Network

The designed wireless sensor network was experimentally evaluated over a two-week period in an outdoor environment. Two DS18B20 temperature sensors were connected to the sensor node and installed in the subsurface soil layer at depths of 20 cm and 35 cm below the surface. One DHT22 sensor was used to measure ambient air temperature and humidity. The DS18B20 sensors communicate with the sensor node using the 1-Wire protocol. The measured temperature data collected during this evaluation period are illustrated in Fig. 2.




Fig. 2. Subsurface soil temperatures at depths of 20 cm and 35 cm plotted against ambient air temperature over the two-week evaluation period.

The DS18B20 sensors used for subsurface temperature measurements have a manufacturer-specified accuracy of ± 0.5 °C over the range from -10 °C to $+85$ °C and an overall measurement range from -55 °C to $+125$ °C. In the present system, the sensors were operated at 12-bit resolution.

The DS18B20 devices were installed in commercial plastic housings with sealed cable entries, providing protection against prolonged exposure to moist soil. The probes were directly buried, and native soil was carefully backfilled and compacted around each probe tip to improve thermal contact. This installation method minimizes, but does not completely eliminate, measurement lag and errors caused by residual air gaps between the probe housing and the surrounding soil.

The DHT22 sensor used to measure ambient air temperature and relative humidity has a specified temperature accuracy of ± 0.5 °C over its operating range of -40 °C to $+80$ °C. Its specified relative-humidity accuracy is typically ± 2 % RH, with a maximum stated deviation of ± 5 % RH under the manufacturer's specified conditions.

During the experimental evaluation, it was verified that the designed WSN can successfully transfer data to the cloud data storage. The WSN prevented data loss when a simulated shutdown of the Wi-Fi network was performed for a period of 5 hours. After the Wi-Fi network was restored, all missing data were transferred to the cloud in a burst. The WSN also prevented data loss when a simulated shutdown of the central node was performed for a period of 4 hours. After the central node was reconnected to the power supply, it successfully synchronized the missing data from the sensor node database and uploaded them to the cloud.

The required functionality of the WSN was successfully verified. To address the uncertainty associated with the battery-capacity-based estimate, the current consumption of the sensor node was subsequently measured using an INA226 current monitor. Measurements were acquired at intervals of approximately 4 ms during representative one-hour operating cycles.

The sensor node operated in three principal states. In deep-sleep mode, the measured current was approximately 4.1 mA. During sensor measurement and data processing, the current was approximately 40 mA. During ESP-NOW communication, the maximum observed current reached 216 mA. In a representative hourly cycle, the node was active for approximately 30 s, of which approximately 18 s involved ESP-NOW communication. The remaining 3570 s were spent in deep-sleep mode. The duration of the communication interval varied slightly among individual cycles because of the asynchronous nature of the wireless transmission.

Because 216 mA represents the maximum observed current rather than the mean current throughout the complete communication interval, using this value for the entire 18 s interval provides a conservative estimate of the average current consumption. Assuming that the communication

interval is included within the 30 s active interval, the upper-bound average current is

$$I_{AVG,UB} = \frac{I_{TX}t_{TX} + I_{ACTIVE}(t_{ACTIVE} - t_{TX}) + I_{SLEEP}(3600 - t_{ACTIVE})}{3600} \quad (6)$$

Substituting the measured values:

$$I_{AVG,UB} = \frac{(216 \text{ mA})(18 \text{ s}) + (40 \text{ mA})(12 \text{ s}) + (4.1 \text{ mA})(3570 \text{ s})}{3600 \text{ s}} = 5.28 \text{ mA}$$

The corresponding charge consumption over 336 h is

$$Q = I_{AVG,UB}t = (5.28 \text{ mA})(336 \text{ h}) \approx 1774 \text{ mAh} \quad (7)$$

This result is close to the previously obtained battery-capacity-based estimate of approximately 1800 mAh, corresponding to an average current of 5.36 mA. However, the battery-based value is retained only as an independent consistency check, whereas the revised estimate is based on direct current measurements and the observed operating duty cycle.

Owing to the approximately 4 ms measurement interval and the variability of the ESP-NOW transmission duration, the short-duration current waveform could not be fully resolved. Therefore, the reported value should be interpreted as a conservative estimate of the average current consumption rather than an exact measurement of individual transmission peaks. The operational time could be further prolonged by more optimized software that shortens the time the board spends in active mode, or by choosing temperature sensors with shorter startup times.

6. Conclusions

This paper presented the design and experimental evaluation of an ESP32-C6-based wireless sensor network for subsurface temperature monitoring in road thermoelectric energy harvesting applications. The prototype sensor node measured ambient air temperature and subsurface temperatures at several depths and transmitted the data to a central node, which forwarded them to cloud storage. Local buffering in both nodes enabled complete recovery after simulated wireless and Internet communication outages.

The prototype was evaluated in an outdoor environment over a two-week period. The results verified the basic functionality of the measurement system, including sensor acquisition, local storage, wireless data transfer, cloud synchronization, and battery-powered operation. The system is therefore suitable for long-term deployment aimed at collecting seasonal temperature profiles in subsurface layers. Future work will focus on long-term measurements, thermal modeling of the road surface/subgrade temperature difference, and estimation of the potential energy output of thermoelectric generators.

References

- [1]. C. Gould, Thermoelectric energy harvesting, in *A Guide to Small-Scale Energy Harvesting Techniques* (R. Manyala, Ed.), *IntechOpen*, London, 2020, pp. 1-19.
- [2]. G. J. Snyder, Thermoelectric energy harvesting, in *Energy Harvesting Technologies* (S. Priya, D. J. Inman, Eds.), *Springer*, Boston, 2009, pp. 325-336.
- [3]. Y. Kadohiro, S. Cheng, J. S. Cross, All-day energy harvesting power system utilizing a thermoelectric generator with water-based heat storage, *Sustainability*, Vol. 12, Issue 9, 2020, 3659.
- [4]. S. A. Tahami, M. Gholikhani, S. Dessouky, Thermoelectric energy harvesting system for roadway sustainability, *Transportation Research Record*, Vol. 2674, Issue 2, 2020, pp. 135-145.
- [5]. M. A. M. Lallmamode, A. S. M. Al-Obaidi, Harvesting energy from vehicle transportation on highways using piezoelectric and thermoelectric technologies, *Journal of Physics: Conference Series*, Vol. 2120, Issue 1, 2021, 012016.
- [6]. T. Baca, P. Sarafin, M. Chochul, M. Kubascik, Thermoelectric transducers: A promising method of energy generation for smart roads, *Applied Sciences*, Vol. 16, Issue 3, 2026, 1662.
- [7]. R. De Fazio, M. De Giorgi, D. Cafagna, C. Del-Valle-Soto, et al., Energy harvesting technologies and devices from vehicular transit and natural sources on roads for a sustainable transport: State-of-the-art analysis and commercial solutions, *Energies*, Vol. 16, Issue 7, 2023, 3016.
- [8]. M. Hodon, M. Chovanec, L. Cechovic, M. Hudik, et al., Maximizing performance of low-power WSN node on the basis of event-driven-programming approach: Minimization of operational energy costs of WSN node control unit, in *Proceedings of the IEEE Symposium on Computers and Communication (ISCC)*, 2015, pp. 204-209.
- [9]. M. Hodon, P. Sarafin, P. Sevcik, Monitoring and recognition of bird population in protected bird territory, in *Proceedings of the IEEE Symposium on Computers and Communication (ISCC)*, 2015, pp. 198-203.
- [10]. P. Sevcik, S. Zak, M. Hodon, Wireless sensor network for smart power metering, *Concurrency and Computation: Practice and Experience*, Vol. 29, Issue 24, 2017, e4247.
- [11]. R. Zalman, V. Olesnanikova, P. Sevcik, P. Sarafin, Monitoring of CO2 amount in closed objects via WSN, in *Proceedings of the Federated Conference on Computer Science and Information Systems (FedCSIS)*, 2015, pp. 1257-1260.
- [12]. M. Chochul, P. Sevcik, A survey of low power wide area network technologies, in *Proceedings of the 18th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, 2020, pp. 69-73.
- [13]. M. Hodon, J. Micek, O. Karpis, P. Sevcik, Robotic chassis for Android phones, in *Proceedings of the 9th International Conference on Education and New Learning Technologies (EDULEARN)*, 2017, pp. 3692-3696.
- [14]. O. Karpis, J. Micek, V. Olesnanikova, MYrobot – Mobile educational platform, in *Advances in Service and Industrial Robotics* (N. Aspragathos, P. Koustoumpardis, V. Moulانيتis, Eds.), *Springer*, Cham, 2019, pp. 687-694.
- [15]. O. Karpis, V. Olesnanikova, L. Formanek, M. Hodon, et al., Compressed sensing - a way to spare energy in WSN for UAV, *IFAC-PapersOnLine*, Vol. 55, Issue 4, 2022, pp. 170-176.
- [16]. P. Sarafin, L. Formanek, M. Chochul, Dynamic system parameter identification based on the acceleration data, in *Proceedings of the 18th International Conference on Emerging eLearning Technologies and Applications (ICETA)*, 2020, pp. 570-573.
- [17]. L. Formanek, O. Karpis, V. Olesnanikova, Reducing number of nodes in WSN with neural network usage, in *Proceedings of the 4th International Conference on Intelligent Computing in Data Sciences (ICDS)*, 2020, pp. 1-4.

AI-Based Virtual Sensors as Cybersecurity Mechanisms in Industrial IoT Systems: A Systematic Literature Review, Taxonomy, and Cyber-Resilience Framework

I. Gabrišová, G. Koman and M. Kubina

Department of Management Theories, Faculty of Management Science and Informatics,
University of Žilina, Univerzitná 8215/1, 010 26 Žilina, Slovakia
E-mail: ivana.gabrisova@stud.uniza.sk

Summary: The convergence of artificial intelligence (AI), the Industrial Internet of Things (IIoT), and cyber-physical systems (CPS) has strengthened industrial monitoring while increasing exposure to attacks on sensing and measurement integrity. This study systematically reviews cybersecurity-oriented virtual sensing in industrial environments, with particular emphasis on AI-based approaches, and reports the review in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines. Searches of Web of Science, IEEE Xplore, and Scopus identified 72 records; after duplicate removal and screening, 21 studies were included in the comparative synthesis. Virtual sensor is used as an umbrella term covering inferential/process soft sensing, analytical-redundancy/state-estimation virtual sensing, and contextual/behavioural virtual sensing. An evidence-derived four-dimensional taxonomy distinguishes sensing paradigm, modelling approach, cybersecurity/resilience role, and deployment purpose. A four-layer cyber-resilience framework links sensing and data acquisition, virtual sensing and state estimation, attack detection and isolation, and resilience and recovery.

Keywords: Virtual sensors, Artificial intelligence, Cybersecurity, Industrial internet of things, Cyber-resilience.

1. Introduction

The rapid advancement of Industry 4.0 technologies has accelerated the adoption of Industrial Internet of Things (IIoT) architectures, cyber-physical systems (CPS), cloud computing, edge intelligence, and artificial intelligence (AI). Sensors constitute the core data-acquisition layer of IIoT and CPS environments, providing the information required for monitoring, control, optimisation, and decision-making [1, 2]. Industrial organisations can achieve higher levels of automation, operational efficiency, and data-driven decision-making by applying advanced sensing technologies alongside intelligent analytics [1].

The increasing interconnectivity of industrial systems, however, has significantly expanded the cyberattack surface. Modern industrial control systems (ICS) depend on numerous foundational components, including sensors and distributed control architectures, creating multiple attack vectors that can compromise data integrity and confidentiality. Recent studies identify false-data injection (FDI), denial-of-service (DoS), replay, spoofing, and sensor manipulation attacks as major threats to industrial and cyber-physical environments [3-6]. Such attacks may manipulate process measurements, conceal operational failures, disrupt production processes, or trigger unsafe operating conditions.

Traditional cybersecurity approaches in industrial environments often focus on network monitoring, access control, and system-level threat detection [6-8]. However, these mechanisms may offer limited protection against attacks that directly target sensor

measurements and the sensing layer's integrity. Consequently, virtual sensing has attracted increasing attention as a complementary mechanism for enhancing cybersecurity and cyber-resilience in industrial environments [9-12].

The terminology surrounding virtual sensing is not fully standardised across the literature [13, 14]. In this study, the term "virtual sensor" is therefore used as an umbrella term for software- or model-based sensing mechanisms that infer, estimate, reconstruct, or validate variables or system states using information from physical sensors, other virtual sensors, process models, or historical data [13-16]. The term soft sensor is used more specifically to refer to the inferential process-sensing tradition, in which difficult-to-measure process or quality variables are estimated from readily measurable auxiliary variables [14, 15], [17-19]. Software sensor is treated as a related designation occurring in the literature rather than as a separate technical category or a strictly interchangeable synonym [13, 14].

Within this umbrella concept, the review distinguishes three virtual-sensing paradigms according to the type of information that is inferred or provided: (1) inferential/process soft sensing, in which difficult-to-measure process or quality variables are estimated from available auxiliary measurements [15], [17-19]; (2) analytical-redundancy/state-estimation virtual sensing, in which independent estimates of physical measurements or system states are generated for validation, residual generation, secure estimation, reconstruction, or substitution [9-11, 20]; and (3) contextual/behavioural virtual sensing, in which sensor, system, or communication data are processed

to infer abstract system conditions, behavioural states, or contextual indicators [21]. These paradigms are distinguished by their predominant role in sensing rather than by the computational algorithm used to implement them. Consequently, machine-learning and deep-learning techniques may support different virtual-sensing paradigms and are treated separately as modelling approaches in the proposed taxonomy.

Digital-twin-integrated virtual sensing was also considered during the comparative coding but did not emerge as a distinct evidence-supported paradigm in the included corpus; digital-twin integration is therefore treated as an emerging research direction rather than a separate taxonomy category. Although the review places particular emphasis on AI-based implementations, non-AI model- and observer-based approaches are retained where they provide relevant evidence on cybersecurity and cyber-resilience functions of virtual sensing.

Although inferential soft sensors were originally developed primarily for process optimisation, quality monitoring, and indirect measurement [14, 15, 17, 19], virtual sensing concepts have subsequently been applied to analytical redundancy, sensor validation, state estimation, attack detection, and resilient control in cyber-physical systems [9-11, 20]. Across these paradigms, virtual sensing can provide complementary or analytically redundant information for reconstructing missing or compromised measurements, validating physical-sensor outputs, and estimating process states [10-12, 22]. In analytical redundancy and state-estimation applications, discrepancies between estimated and measured values can generate residuals or other indicators of abnormal behaviour, thereby supporting the detection of cyberattacks, sensor faults, communication failures, or process disturbances [3, 10, 20, 23].

Recent research has demonstrated the cybersecurity potential of analytical-redundancy/state-estimation virtual sensing and AI-enabled virtual sensing implementations for secure state estimation, attack detection, measurement reconstruction, and resilient control of cyber-physical systems [3, 9-11, 24]. Virtual-sensor-based approaches have been applied to detect false-data injection attacks, identify or compensate for compromised sensors, reconstruct affected measurements, and preserve operational continuity in critical infrastructures and industrial environments [9-11, 20]. Advances in machine learning and deep learning have further extended AI-based virtual sensing capabilities, supporting nonlinear estimation, measurement reconstruction, adaptive modelling, and anomaly detection in industrial environments [6, 8, 17, 25].

However, AI-driven soft sensors can themselves become targets of adversarial attacks. Recent research shows that neural-network-based soft sensors can be manipulated by carefully crafted adversarial inputs, posing new challenges for reliability, robustness, and trustworthiness in industrial environments [26, 27]. This introduces a complementary research problem: virtual sensing mechanisms may serve as cybersecurity

tools while AI-based implementations may themselves require protection against adversarial manipulation [12, 26, 27].

Against this background, the objective of this study is threefold: (1) to systematically review the literature on cybersecurity-oriented virtual sensing in industrial environments, with particular emphasis on AI-based approaches; (2) to develop a four-dimensional taxonomy that distinguishes virtual sensing solutions according to virtual-sensing paradigm, modelling approach, cybersecurity/resilience role, and deployment purpose; and (3) to propose a cyber-resilience framework showing how virtual sensing can contribute to the detection, response, and recovery capabilities of IIoT and CPS environments.

Existing reviews have examined AI-driven soft sensors in process industries [17], machine-learning and deep-learning approaches to IoT and network intrusion/anomaly detection [6, 7], machine-learning-based cybersecurity in power systems [8], false-data injection and cyber-physical attacks affecting state estimation [4, 5], secure and resilient networked or cyber-physical control [3, 24, 28], and domain-specific applications of virtual sensing [16]. In contrast, this review focuses specifically on the intersection of virtual sensing, industrial cybersecurity, and cyber-resilience. Its contribution is to integrate inferential/process, analytical-redundancy/state-estimation, and contextual/behavioural virtual-sensing paradigms and to relate them systematically to their computational implementation, cybersecurity/resilience roles, and operational purposes. Rather than classifying algorithms or attack types in isolation, the proposed taxonomy separates four complementary dimensions: (1) virtual-sensing paradigm, (2) modelling approach, (3) cybersecurity/resilience role, and (4) deployment purpose, thereby revealing how similar modelling techniques can support different sensing paradigms and how similar resilience functions can be achieved through different implementations. The four-layer framework subsequently maps these mechanisms to detection, response, and recovery capabilities. The review also considers the dual role of AI-based virtual sensing as both a cybersecurity mechanism and a potential target of adversarial manipulation.

2. Methodology

This study employed a systematic literature review methodology and was reported in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines [29]. Literature searches were conducted on 14 June 2026 in Web of Science, IEEE Xplore, and Scopus. The same complete search expression was applied in all three databases: ("virtual sensor" OR "soft sensor" OR "digital twin") AND ("cybersecurity" OR "cyberattack" OR "anomaly detection") AND ("industrial IoT" OR "CPS" OR "ICS"). Searches were restricted to open-access publications published from

2015 up to the search date. After these database-level filters, 72 records were identified: 54 from Web of Science, 17 from IEEE Xplore, and 1 from Scopus.

Studies were eligible if they addressed virtual sensing or a related estimation mechanism in industrial, IoT, CPS, or control-system contexts and provided either direct evidence of cybersecurity/resilience or foundational evidence relevant to virtual-sensing concepts, modelling, functions, or deployment. Records outside this scope or lacking sufficient methodological or substantive information for reliable classification were excluded.

Records retrieved from the three databases were merged in Zotero, and duplicate records were identified using DOIs, titles, and bibliographic metadata and removed before screening. All three authors independently screened titles and abstracts and subsequently assessed potentially eligible reports in full text against the predefined eligibility criteria; disagreements were resolved through discussion and consensus. After removal of 7 duplicate records, 65 unique records underwent title and abstract screening. Of these, 32 were excluded, and 33 reports proceeded to full-text assessment.

Of the 33 full-text reports assessed, 21 met the criteria for inclusion in the comparative synthesis. Twelve reports were not included in the synthesis: nine were retained only as contextual/supportive references, while three were excluded during full-text assessment because they fell outside the virtual sensing/estimation scope ($n = 2$) or did not provide sufficient evidence relevant to industrial cybersecurity or resilience ($n = 1$).

The complete study-selection process is summarised in Fig. 1. The 21 included studies comprised nine original studies and twelve secondary reviews; the nine contextual references were excluded from comparative coding.

A qualitative appraisal of study quality and relevance was conducted during full-text assessment, considering methodological transparency, clarity of the virtual-sensing or estimation approach, adequacy of the reported evaluation or synthesis, relevance to the review scope, and traceability of the reported results. Because the corpus comprised both original research publications and secondary reviews, the criteria were interpreted according to study type rather than combined into a single performance score.

Data extraction covered publication characteristics, study type, industrial application domain, virtual-sensing terminology and paradigm, modelling approach and specific model or method, attack or fault type, dataset or evaluation environment, cybersecurity/resilience role, deployment purpose, and resilience-related capability. These extracted characteristics formed the basis of the comparative study matrix presented in Table 1.

The taxonomy was developed through comparative coding of the 21 included studies across four analytical dimensions: virtual-sensing paradigm, modelling approach, cybersecurity/resilience role, and deployment purpose. Taxonomy categories were

retained only when supported by evidence in the included corpus; original research publications provided direct evidence of implemented virtual-sensing mechanisms, while secondary reviews were used to corroborate and contextualise the identified research strands. Studies could be assigned to multiple cybersecurity/resilience roles or deployment purposes were supported by the reported implementation.




Fig. 1. PRISMA 2020 flow diagram of the study selection process.

The cyber-resilience framework was derived by grouping the recurrent mechanisms and resilience outcomes identified during coding into a sequential functional structure spanning sensing/data acquisition, virtual sensing/state estimation, attack detection/isolation, and resilience/recovery.

3. Findings

Among the nine original studies, three represented inferential/process soft sensing [26, 27, 30]; five represented analytical-redundancy/state-estimation virtual sensing [9-11, 20, 22]; and one represented contextual/behavioural virtual sensing [21]. Cross-study comparison showed that the modelling technique does not determine sensing paradigm: machine-learning and deep-learning methods span multiple paradigms. At the same time, observer- and filter-based approaches are prominent in secure state estimation and resilient control. The identified functions range from measurement validation and anomaly/attack detection to secure state estimation, reconstruction, resilient control, and adversarial robustness.

Table 1. Evidence matrix of the 21 included studies.

Study / evidence group	Domain	Paradigm & model	Threat / evaluation	Cybersecurity/ resilience contribution
Vallejo [13], Abeykoon [14], Curreri [15], Paepae [16], Perera [17], Luo [18], Zhu [19] (R)	Process industries; soft metrology; water quality; blast furnace; fermentation	P1; model-based, statistical, ML/DL, fuzzy and hybrid soft sensors	No direct cyberattack focus; reviewed studies and process datasets	Indirect measurement, process monitoring, validation/fault context; D1/D3
Belgacem & Chihi [31] (R)	Intelligent sensor systems	P2; observer-, model- and AI-based methods	Sensor faults; reviewed studies	C2/C4/C6/C7; D3/D4
Rafique et al. [6] (R)	IoT/IIoT networks	P0; ML/DL IDS	DDoS, spoofing, replay, injection; multiple datasets	C3; D2
Sandberg [3], Xing & Shen [24], Wan & Cao [28] (R)	Networked/CPS control	P2/P0; state estimation, observers, ML/DL, watermarking	FDI, replay, DoS, hybrid attacks; reviewed studies	C3–C7; D2/D4/D5
Chen [26]; Guo [27] (O)	Sulfur recovery; air preheater	P1; DNN/DL soft sensors + adversarial methods	Adversarial manipulation; industrial datasets	C8; D1/D3/D4
Kordestani et al. [20] (O)	Steam turbine	P2; MLP soft sensor + residual + SVM	DoS, deception, replay; simulation	C1/C3/C4; D2
Paridari et al. [9] (O)	ICS/energy system	P2; knowledge detection + SVM + Kalman filter	Measurement manipulation; simulation/real measurements	C3/C5/C6/C7; D2/D4/D5
Madichetty & Mishra [10] (O)	DC microgrid	P2; adaptive observer + residual	FDI; hardware test bench	C1/C3/C5/C6/C7; D2–D5
Xin et al. [11] (O)	Multi-sensor CPS/GNSS	P2; VSN + DRL + Kalman fusion	FDI + DoS; KITTI/MS2	C5–C7; D3–D5
Ludwig et al. [22] (O)	Industrial process plant	P2; LSTM autoencoder	Missing/faulty sensors; simulated + real data	C6; D3/D4
Ströbel et al. [30] (O)	CNC milling	P1; physics-informed ML/DL soft sensing	Process/tool anomalies; experimental data	C2; D1
Trouli & Kornaros [21] (O)	Automotive CAN	P3; k-NN + behavioural profiling	Impersonation/frame insertion; real CAN data	C3; D2

Note: O – original research; R – secondary review; P0 – no direct virtual-sensor implementation; P1 – inferential/process soft sensing; P2 – analytical-redundancy/state-estimation virtual sensing; P3 – contextual/behavioural virtual sensing. C1 – validation; C2 – anomaly/fault detection; C3 – cyberattack/intrusion detection; C4 – localisation/isolation; C5 – secure state estimation; C6 – reconstruction/substitution; C7 – resilient control/recovery; C8 – adversarial robustness. D1 – process monitoring; D2 – security monitoring; D3 – sensor backup; D4 – operational continuity; D5 – mitigation/recovery. ML – machine learning; DL – deep learning; IDS – intrusion detection system; DNN – deep neural network; MLP – multilayer perceptron; SVM – support vector machine; VSN – Virtual SensorNet; DRL – deep reinforcement learning; GNSS – Global Navigation Satellite System; LSTM – long short-term memory; CAN – Controller Area Network; k-NN – k-nearest neighbours; MS2 – Multi-Spectral Stereo; DDoS – distributed denial-of-service; KITTI – Karlsruhe Institute of Technology and Toyota Technological Institute dataset. Grouped rows aggregate studies with closely related scope and summarise their shared evidence strands; individual references remain explicitly identified within each group.

Based on Table 1, the taxonomy comprises four dimensions: (1) virtual-sensing paradigm – inferential/process, analytical-redundancy/state-estimation, or contextual/behavioural sensing; (2) modelling approach—physics/model-based, observer/filter/statistical, machine learning, deep learning/reinforcement learning, or hybrid/knowledge-informed methods; (3) cybersecurity/resilience role – validation, anomaly/fault or cyberattack detection, localisation/isolation, secure state estimation, reconstruction, resilient control/recovery, or adversarial robustness; and (4) deployment purpose – process or security monitoring, sensor backup, operational continuity, or mitigation/recovery.

The taxonomy separates dimensions that are often conflated: a modelling technique does not uniquely determine the sensing paradigm or security role. Machine-learning and deep-learning models support both inferential and analytical-redundancy applications, while reconstruction may be learning-

based or observer-based. Solutions can therefore be compared by what is virtually sensed, how the estimate is generated, what resilience function it performs, and why it is deployed. For example, deep learning supports adversarially targeted inferential soft sensors [26, 27], whereas Xin et al. [11] combine a learned virtual sensor with state estimation in the presence of sensor attacks.

Based on the recurrent cybersecurity and resilience functions identified across the included studies, a four-layer cyber-resilience framework is proposed comprising: (1) sensing and data acquisition, (2) virtual sensing and state estimation, (3) attack detection and isolation, and (4) resilience and recovery. The framework illustrates how virtual sensors can support secure monitoring, detection of compromised measurements, reconstruction of unavailable or manipulated sensor data, and continued system operation under adverse cyber conditions in Industry 4.0 environments [9-11, 22].

4. Conclusions

This study systematically examined cybersecurity-oriented virtual sensing in Industrial Internet of Things (IIoT) and cyber-physical systems (CPS) environments, with particular emphasis on AI-based approaches. By treating virtual sensing as an umbrella concept, the review distinguishes three evidence-supported virtual-sensing paradigms: inferential/process soft sensing, analytical-redundancy/state-estimation virtual sensing, and contextual/behavioural virtual sensing. The reviewed literature indicates that these paradigms provide complementary capabilities for attack detection, measurement validation, secure state estimation, measurement reconstruction, resilient control, and operational continuity in industrial environments [9-11].

The proposed taxonomy distinguishes four complementary aspects of cybersecurity-oriented virtual sensing: the virtual-sensing paradigm, the modelling approach, the cybersecurity/resilience role, and the deployment purpose. Its added value lies in exposing relationships across these dimensions rather than grouping studies solely by algorithms or attack types. The synthesis shows that similar computational methods can implement different virtual-sensing paradigms and that comparable resilience functions can be achieved through different modelling approaches. In parallel, the four-layer cyber-resilience framework describes the progression from sensing and state estimation through attack detection and isolation to reconstruction, mitigation, and continued operation. Together, the taxonomy and framework provide complementary structures for analysing and designing virtual-sensor-supported cybersecurity and resilience in IIoT and CPS environments.

The reviewed studies indicate that these mechanisms can improve resilience against false-data injection, spoofing, denial-of-service attacks, and sensor failures while supporting operational continuity in Industry 4.0 environments [10, 11, 22, 23].

The review also identifies an emerging research stream focused on protecting virtual sensors themselves against adversarial manipulation. As AI-powered virtual sensors become increasingly integrated into critical industrial processes, guaranteeing their robustness, reliability, and trustworthiness becomes essential. Recent studies demonstrate the need to improve the robustness and trustworthiness of AI-based virtual sensing models under adversarial or anomalous conditions [12, 26, 27].

From a managerial perspective, the synthesis suggests that organisations should treat virtual sensing technologies as strategic cybersecurity assets rather than merely as process-monitoring tools. Integrating virtual sensors into cybersecurity governance, resilience engineering, and digital transformation strategies may improve operational continuity, reduce dependence on costly hardware redundancy, and enhance cyber situational awareness across industrial systems.

The review is limited by the open-access restriction and the defined scope of the database and search terms, which may have excluded relevant studies outside these boundaries.

Future research ought to focus on explainable and robust artificial intelligence, federated and distributed learning, digital-twin integration, real-time deployment architectures, and standardisation frameworks that support secure, interoperable, and trustworthy virtual sensing solutions for critical industrial infrastructures [1, 8, 12].

Acknowledgements

Co-financed by the European Union through the Slovakia Programme under project No. NFP401101C360: Research and development of advanced artificial intelligence solutions for the detection of cyber threats and defense against sophisticated attacks.

References

- [1]. H. Xu, W. Yu, D. Griffith, N. Golmie, A survey on Industrial Internet of Things: A cyber-physical systems perspective, *IEEE Access*, Vol. 6, 2018, pp. 78238-78259.
- [2]. P. A. Rezvy, V. L. N. Komanapalli, Investigation on the role of artificial intelligence in measurement system, *IEEE Access*, Vol. 13, 2025, pp. 96483-96502.
- [3]. H. Sandberg, V. Gupta, K. H. Johansson, Secure networked control systems, *Annual Review of Control, Robotics, and Autonomous Systems*, Vol. 5, 2022, pp. 445-464.
- [4]. F. Liberati, E. Garone, A. Di Giorgio, Review of cyber-physical attacks in smart grids: A system-theoretic perspective, *Electronics*, Vol. 10, Issue 10, 2021, 1153.
- [5]. M. Irfan, A. Sadighian, A. Tanveer, S. J. Al-Naimi, et al., A survey on detection and localisation of false data injection attacks in smart grids, *IET Cyber-Physical Systems: Theory & Applications*, Vol. 9, Issue 4, 2024, pp. 313-333.
- [6]. S. H. Rafique, A. Abdallah, N. S. Musa, T. Murugan, Machine learning and deep learning techniques for Internet of Things network anomaly detection: Current research trends and challenges, *Sensors*, Vol. 24, Issue 6, 2024, 1968.
- [7]. J. Lansky, et al., Deep learning-based intrusion detection systems: A systematic review, *IEEE Access*, Vol. 9, 2021, pp. 101574-101599.
- [8]. M. A. Adewusi, M. M. Mustafa, Machine learning for power-grid cybersecurity: A review, *Discover Computing*, Vol. 29, Issue 1, 2026, 282.
- [9]. K. Paridari, N. O'Mahony, A. El-Din Mady, R. Chabukswar, et al., A framework for attack-resilient industrial control systems: Attack detection and controller reconfiguration, *Proceedings of the IEEE*, Vol. 106, Issue 1, 2018, pp. 113-128.
- [10]. S. Madichetty, S. Mishra, Cyber attack detection and correction mechanisms in a distributed DC microgrid, *IEEE Transactions on Power Electronics*, Vol. 37, Issue 2, 2022, pp. 1476-1485.

- [11]. L. Xin, G. He, Z. Long, Secure state estimation for multi-sensor cyber-physical systems using virtual sensor and deep reinforcement learning under multiple attacks on major sensor, *IEEE Transactions on Network Science and Engineering*, Vol. 12, Issue 3, 2025, pp. 1470-1481.
- [12]. G. Sternharz, J. Skackauskas, A. Elhalwagy, A. J. Grichnik, et al., Self-protected virtual sensor network for microcontroller fault detection, *Sensors*, Vol. 22, Issue 2, 2022, 454.
- [13]. M. Vallejo, C. de la Espriella, J. Gómez-Santamaría, A. F. Ramírez-Barrera, et al., Soft metrology based on machine learning: A review, *Measurement Science and Technology*, Vol. 31, Issue 3, 2020, 032001.
- [14]. C. Abeykoon, Design and applications of soft sensors in polymer processing: A review, *IEEE Sensors Journal*, Vol. 19, Issue 8, 2019, pp. 2801-2813.
- [15]. F. Curreri, G. Fiumara, M. G. Xibilia, Input selection methods for soft sensor design: A survey, *Future Internet*, Vol. 12, Issue 6, 2020, 97.
- [16]. T. Paepae, P. N. Bokoro, K. Kyamakyia, From fully physical to virtual sensing for water quality assessment: A comprehensive review of the relevant state-of-the-art, *Sensors*, Vol. 21, Issue 21, 2021, 6971.
- [17]. Y. S. Perera, D. A. A. C. Ratnaweera, C. H. Dasanayaka, C. Abeykoon, The role of artificial intelligence-driven soft sensors in advanced sustainable process industries: A critical review, *Engineering Applications of Artificial Intelligence*, Vol. 121, 2023, 105988.
- [18]. Y. Luo, X. Zhang, M. Kano, L. Deng, et al., Data-driven soft sensors in blast furnace ironmaking: A survey, *Frontiers of Information Technology & Electronic Engineering*, Vol. 24, Issue 3, 2023, pp. 327-354.
- [19]. X. Zhu, K. U. Rehman, B. Wang, M. Shahzad, Modern soft-sensing modeling methods for fermentation processes, *Sensors*, Vol. 20, Issue 6, 2020, 1771.
- [20]. M. Kordestani, A. Chaibakhsh, M. Saif, SMS – A security management system for steam turbines using a multisensor array, *IEEE Systems Journal*, Vol. 14, Issue 3, 2020, pp. 3813-3824.
- [21]. G.-I. Trouli, G. Kornaros, Automotive virtual in-sensor analytics for securing vehicular communication, *IEEE Design & Test*, Vol. 37, Issue 3, 2020, pp. 91-98.
- [22]. B. Ludwig, J. Ehrhardt, O. Niggemann, Creating virtual sensors using neural networks, in *Proceedings of the 30th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2025, pp. 1-8.
- [23]. K. Yang, R. Wang, Y. Jiang, C. Luo, et al., Enhanced resilient sensor attack detection using fusion interval and measurement history, in *Proceedings of the International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, 2018, pp. 1-3.
- [24]. W. Xing, J. Shen, Security control of cyber-physical systems under cyber attacks: A survey, *Sensors*, Vol. 24, Issue 12, 2024, 3815.
- [25]. X. Kong, X. Jiang, B. Zhang, J. Yuan, et al., Latent variable models in the era of industrial big data: Extension and beyond, *Annual Reviews in Control*, Vol. 54, 2022, pp. 167-199.
- [26]. L. Chen, Q.-X. Zhu, Y.-L. He, Adversarial attacks for neural network-based industrial soft sensors: Mirror output attack and translation mirror output attack, *IEEE Transactions on Industrial Informatics*, Vol. 20, Issue 2, 2024, pp. 2378-2386.
- [27]. R. Guo, H. Liu, D. Liu, When deep learning-based soft sensors encounter reliability challenges: A practical knowledge-guided adversarial attack and its defense, *IEEE Transactions on Industrial Informatics*, Vol. 20, Issue 2, 2024, pp. 2702-2714.
- [28]. Y. Wan, J. Cao, A brief survey of recent advances and methodologies for the security control of complex cyber-physical networks, *Sensors*, Vol. 23, Issue 8, 2023, 4013.
- [29]. M. J. Page, et al., The PRISMA 2020 statement: An updated guideline for reporting systematic reviews, *BMJ*, Vol. 372, 2021, n71.
- [30]. R. Ströbel, S. Deucker, H. Zhou, H. Kader, et al., Hybrid machine learning for CNC process monitoring, *IEEE Access*, Vol. 13, 2025, pp. 91875-91888.
- [31]. H. Belgacem, I. Chihi, Toward reliable and intelligent sensor systems: A comprehensive study of fault diagnosis and mitigation, *IEEE Sensors Reviews*, Vol. 2, 2025, pp. 511-536.

Utilisation of Ultrasonic Anemometers for Acquiring Airflow Data in Tunnel Fire Spread Models

M. Hodoň¹, L. Formanek¹ and P. Danišovič²

¹ University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1,
010 26 Žilina, Slovakia

² University of Žilina, Department of Construction Management, Univerzitná 8215/1,
010 26 Žilina, Slovakia
Tel.: +421 41 513 4358
E-mail: michal.hodon@fri.uniza.sk

Summary: Accurate characterisation of airflow is essential for reliable modelling of smoke propagation during tunnel fires. This paper presents a compact measurement platform based on WindSonic ultrasonic anemometers for long-term monitoring at road-tunnel portals and inside tunnel tubes. The system combines an ESP32-S3 microcontroller, an RS485 interface, microSD storage with flash-memory fallback, battery-voltage monitoring, and a local Wi-Fi access point. Embedded control software written in MicroPython performs sensor-frame validation, buffered data logging, storage management, and provides a browser-based interface for live monitoring and access to historical data. A 12 h field experiment with two sensors installed at different heights produced paired wind-speed and wind-direction records. The wind-speed series showed a correlation coefficient of 0.90 and a root-mean-square inter-sensor difference of 0.19 m/s; the wind-direction series showed a correlation coefficient of 0.93 and a root-mean-square circular difference of 24.30°. The platform provides datasets suitable for investigating portal-flow variability and supporting CFD/FDS model validation.

Keywords: Ultrasonic anemometry, Tunnel airflow monitoring, Road tunnel safety, Fire dynamics simulator, Embedded measurement systems.

1. Introduction

Road tunnels are critical components of modern transport infrastructure, enabling safe and efficient transit through geographically challenging regions. During a tunnel fire, smoke can rapidly reduce visibility, increase exposure to toxic combustion products, and complicate evacuation and emergency response. Smoke movement is strongly coupled to longitudinal airflow and ventilation conditions; therefore, reliable characterisation of tunnel airflow is essential for the design of ventilation strategies and for fire-safety modelling [1, 2].

Computational Fluid Dynamics (CFD) tools, including the Fire Dynamics Simulator (FDS), are widely used to analyse airflow, heat transfer, and smoke propagation in tunnel environments [3]. Their predictive accuracy depends on appropriate modelling assumptions, representative boundary conditions, and experimental data acquired under real operating conditions. Although tunnel-fire modelling has advanced substantially, field datasets describing airflow under varying environmental conditions remain comparatively difficult to obtain.

Meteorological conditions at tunnel portals can influence the airflow distribution within tunnel tubes, particularly during low-velocity states relevant to evacuation and smoke-control scenarios [1, 2]. These external influences may affect emergency-ventilation performance and alter smoke propagation during a fire. Experimental investigation therefore requires measurements that capture both tunnel airflow and environmental conditions in the portal area.

To address this need, a dedicated platform based on ultrasonic anemometers was developed. The proposed system combines custom embedded electronics, industrial RS485 communication, local data storage, and software for sensor management and data acquisition. It is intended for long-term deployment in tunnel environments and for acquiring airflow data suitable for CFD/FDS model validation. This paper describes the measurement system, the embedded hardware and MicroPython software, the web-based user interface, and preliminary field results from two sensors installed at different heights.

2. Airflow Measurement System

Airflow measurements were performed using WindSonic ultrasonic anemometers (Gill Instruments, UK) configured with an RS485 communication interface and a 5-30 V DC supply range. The sensors provide simultaneous wind-speed and wind-direction measurements without moving mechanical parts, which makes them suitable for long-term outdoor deployment and critical infrastructure monitoring.

The ultrasonic measurement principle is based on the influence of airflow on the transit time of acoustic pulses travelling between pairs of transducers. By comparing propagation times along opposing acoustic paths, the sensor determines wind-speed components and reconstructs the horizontal wind vector [4]. Compared with cup or vane anemometers, ultrasonic instruments contain no moving mechanical parts and therefore require less maintenance during long-term

environmental monitoring. The WindSonic unit used in this work provides wind speed and direction through its serial interface and is designed for outdoor operation [5].

Measurement stations were deployed at selected locations near tunnel portals and within the tunnel environment. The sensor arrangement was designed to capture the influence of external meteorological conditions on airflow in the tunnel tube. For the preliminary comparison presented in Section 5, two WindSonic anemometers were operated simultaneously in the same portal area at different installation heights. The paired 12 h measurement enabled a comparison of wind speed and direction while limiting differences caused by horizontal separation of the instruments. This configuration is useful for identifying height-dependent variability in the approaching flow near the tunnel portal.



Fig. 1. Field deployment near a motorway tunnel portal with two WindSonic ultrasonic anemometers installed at different heights.

3. Data Acquisition Platform and Embedded Software

A dedicated data-acquisition platform was developed for long-term field measurements. The hardware comprises an ESP32-S3 microcontroller [6], an RS485 communication interface, a microSD card storage subsystem, battery-voltage monitoring, and power-management circuitry. The ESP32-S3 performs serial acquisition from the WindSonic sensor, protocol validation, timestamping, buffering, storage management, battery monitoring, and operation of the embedded web server. The measurement node is self-contained and can be accessed locally without an external network connection. Fig. 2 summarises the functional architecture of the device.




Fig. 2. Functional block diagram of the ESP32-S3-based measurement system with Gill ASCII communication over RS485.

The firmware separates sensor acquisition, storage handling, and user-interface functions into dedicated MicroPython modules. This modular structure allows the node to continue logging measurements while simultaneously serving the local web interface and responding to changes in storage status.

3.1. Hardware Implementation

The WindSonic sensor is connected to UART2 of the ESP32-S3 through a TP8485E RS485 transceiver. Communication uses the Gill ASCII serial protocol over the RS485 physical layer at 9600 bit/s [5], and a dedicated driver-enable signal controls the transceiver direction. A separate analogue input is used for supply-voltage monitoring. Battery voltage is reconstructed from the ADC reading using the divider gain, corrected by an empirically determined calibration offset, and smoothed with a ten-sample moving average. The battery value is updated once per second and is available to both the data logger and the web interface.

For unattended operation, the storage subsystem prioritises the microSD card. The storage manager verifies the mounted file system and writes measurement records to the card when it is available. If a write operation fails, the firmware automatically switches to the ESP32-S3 internal flash memory and repeats the write. This failover mechanism reduces the risk of losing an entire measurement campaign because of a removable-storage fault, while the active storage location is reported through the user interface. The assembled measurement-node prototype is shown in Fig. 3.




Fig. 3. Prototype of the ESP32-S3 measurement electronics with RS485 interface and microSD storage.

3.2. MicroPython Firmware and Sensor Protocol Handling

The control software is implemented in MicroPython and uses the asynchronous `uasyncio` framework for the main application loop and embedded HTTP server [7]. WindSonic messages are processed by an interrupt-driven UART parser implementing the Gill ASCII framing used by the sensor [5]. The parser recognises STX/ETX frame delimiters, reads the two hexadecimal checksum characters, verifies the terminating CR/LF sequence, and validates the payload using an XOR checksum. The incoming sensor payload contains comma-separated ASCII fields; only a complete, checksum-valid telegram is converted into a measurement record containing the sensor identifier, wind direction, wind speed, engineering units, and status code.

Validated measurements are appended to an in-memory byte buffer and written to non-volatile storage in batches. For data logging, the validated quantities are serialised as semicolon-delimited records in CSV files, clearly distinguishing the storage format from the comma-separated fields of the incoming Gill ASCII telegram. In the current configuration, the buffer is flushed every 30 s, or earlier when required by the buffer-size limit or a daily file rollover. Daily files follow the pattern

YYYY-MM-DD-ws<ID>.csv. Before rollover, pending records are committed to the previous file and the first record belonging to the new date is preserved, preventing records from different days from being mixed.

The firmware also propagates the WindSonic diagnostic status to the web interface. A status value of 0 denotes normal operation, while non-zero values are retained in the CSV record and displayed as diagnostic information. Storing the raw status together with the measured variables supports subsequent screening of questionable data during long-term campaigns.

3.3. Local Wireless Web Interface

At start-up, the measurement node creates a WPA/WPA2-protected Wi-Fi access point with the static local address 192.168.1.1. The embedded HTTP server listens on port 80, allowing a laptop, tablet, or smartphone to configure and inspect the node through a standard web browser. The interface is self-contained, and the charting library is stored locally on the device; therefore, Internet access is not required during field operation.

The measurement page displays current wind speed, wind direction, battery voltage, sensor status, and the active CSV file. Live measurements are delivered to the browser using server-sent events with a nominal 1 s update period. The history view allows selection of a stored CSV file and an optional From/To time interval. The file-management view lists data stored in internal flash memory and, when available, on the microSD card; files can be viewed, downloaded, or deleted. A settings page allows the real-time clock to be synchronised from the browser or set manually. Fig. 4 shows the mobile layout of the user interface.

4. Data Collection, Visualisation and Processing

Each accepted sensor telegram is stored as a semicolon-delimited record in a CSV file. The file header contains five fields: time, sensor status code, battery voltage, wind speed, and wind direction. Time is stored as HH:MM:SS using the ESP32-S3 real-time clock, while the file name contains the calendar date and measurement-node identifier. This simple text representation enables direct inspection and straightforward import into numerical data-processing tools without a proprietary conversion step.

The web application reads historical records from either active storage location and converts them into a compact data structure for plotting. For interactive use, the firmware supports an optional time-range filter and limits the number of returned history points to protect the limited RAM of the embedded platform. The live chart retains a shorter rolling history, whereas explicitly loaded files can contain substantially larger datasets. Fig. 5 illustrates the time-filtered history view.

Several measures improve data integrity during unattended operation. Corrupted serial frames are rejected by checksum validation, sensor diagnostic codes are stored with each measurement, buffered writes reduce non-volatile-memory access, daily rollover prevents unlimited growth of a single file, and automatic SD-to-flash failover provides an alternative logging destination if removable storage becomes unavailable. Battery telemetry provides an additional indicator of node operating condition.




Fig. 4. Mobile web interface for live monitoring of wind speed, wind direction, battery voltage, sensor status, and recorded data.

The resulting time series serve two related purposes: characterisation of natural and ventilation-driven airflow at selected tunnel locations, and provision of experimental data for CFD/FDS model validation. Measurements at portal and in-tunnel stations can be used to determine whether changes in external meteorological conditions coincide with changes in longitudinal tunnel airflow. For model validation, measured velocity and direction can be aligned with simulation outputs at corresponding sensor locations and time intervals, while diagnostic and battery channels help identify periods that should be excluded from quantitative comparison.




Fig. 5. Historical-data view with CSV-file selection and time-range filtering.

5. Preliminary Field Measurement Results

A preliminary 12 h field experiment was conducted on 19 April 2026 using two WindSonic sensors positioned at different heights in the portal area shown in Fig. 1. Both sensors were sampled over the same observation interval, producing paired time series of wind speed and direction. Because the instruments were intentionally mounted at different heights, the differences reported below should be interpreted as inter-sensor differences in the local flow rather than as measurement error alone. The arrangement was intended to reveal vertical variability associated with the atmospheric boundary layer, terrain, roadside structures, and short-term turbulence.

Fig. 6 compares the wind-speed records from the two installation heights. Over the 12 h period, the two series had a correlation coefficient of 0.90, a mean signed difference of 0.02 m/s, a mean absolute inter-sensor difference of 0.14 m/s, and a root-mean-square difference of 0.19 m/s. The close correspondence of the large-scale variations indicates that both sensors captured the same dominant airflow evolution, while short-term deviations may reflect local turbulence and vertical shear.

Fig. 7 compares wind direction over the same interval. The directional comparison yielded a

correlation coefficient of 0.93. Circular-distance metrics gave a mean signed circular difference of 3.28° , a mean absolute circular difference of 15.76° , and a root-mean-square circular difference of 24.30° . Circular-statistical measures [8] were applied to account for the directional nature of the data and to

avoid artificial discontinuities when observations crossed the $0^\circ/360^\circ$ boundary. The close correspondence of the directional trends indicates that both sensors captured the same dominant changes in the approaching flow, while short-lived deviations may reflect height-dependent local variability.




Fig. 6. Twelve-hour comparison of wind speed measured by two ultrasonic anemometers installed at different heights.

The paired experiment demonstrates the value of simultaneous multi-height measurements for tunnel-airflow studies. The results show that the developed platform can acquire long-duration records from independent measurement nodes while preserving the temporal variations needed for inter-station comparison. The observed differences between installation heights also suggest that a

single-point measurement may not fully represent the approaching flow field used as a boundary condition in CFD/FDS simulations. Field intercomparison studies of sonic anemometers [9] and full-scale road-tunnel measurements demonstrating the influence of naturally induced airflow [10] provide complementary context for these observations.




Fig. 7. Twelve-hour comparison of wind direction measured by two ultrasonic anemometers installed at different heights.

6. Conclusions

This paper presents a compact WindSonic-based platform for road-tunnel airflow monitoring. It integrates an ESP32-S3 controller, an RS485 interface, microSD storage with flash-memory fallback, battery monitoring, and a local web interface. The MicroPython firmware validates sensor frames, buffers and stores data, manages daily file rollover, and exposes diagnostics for unattended operation. The 12 h two-sensor experiment demonstrated simultaneous acquisition at different heights and close correspondence in the dominant wind-speed and wind-direction trends.

The architecture can be replicated at tunnel portals and inside tunnel tubes to create spatially distributed airflow datasets. Future work will address measurement uncertainty and completeness, longer-term circular-statistical assessment, inter-station comparison, and direct CFD/FDS validation under representative operating conditions.

Acknowledgements

This work was supported by the Slovak Research and Development Agency (APVV) under project No. APVV-24-0429, Experimental Investigation and Computer Modelling of Airflow during Road Tunnel Fire.

References

- [1]. J. Gehandler, Road tunnel fire safety and risk: A review, *Fire Science Reviews*, Vol. 4, 2015, 2.
- [2]. H. Ingason, Y. Z. Li, A. Lönnemark, Tunnel Fire Dynamics (2nd Ed.), *Springer*, Cham, 2024.
- [3]. K. McGrattan, R. McDermott, C. Weinschenk, G. Forney, Fire dynamics simulator, technical reference guide (6th Ed.), NIST Special Publication 1018, *National Institute of Standards and Technology*, 2013.
- [4]. A. Cuerva, A. Sanz-Andrés, On sonic anemometer measurement theory, *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 88, Issue 1, 2000, pp. 25-55.
- [5]. WindSonic ultrasonic anemometer user manual, Doc. No. 1405-PS-0019, Issue 31, *Gill Instruments*, 2024.
- [6]. ESP32-S3 series datasheet, Version 2.2, *Espressif Systems*, 2026.
- [7]. MicroPython documentation, Release v1.20.0, *MicroPython Project*, 2023.
- [8]. N. I. Fisher, Statistical Analysis of Circular Data, *Cambridge University Press*, Cambridge, 1993.
- [9]. M. Mauder, M. J. Zeeman, Field intercomparison of prevailing sonic anemometers, *Atmospheric Measurement Techniques*, Vol. 11, Issue 1, 2018, pp. 249-263.
- [10]. M. Król, A. Król, P. Koper, P. Wrona, Full scale measurements of the operation of fire ventilation in a road tunnel, *Tunnelling and Underground Space Technology*, Vol. 70, 2017, pp. 204-213.

Distributed VOC and Environmental Monitoring System for Road Tunnel Applications

M. Hodoň¹, L. Formanek¹ and P. Danišovič²

¹ University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1,
010 26 Žilina, Slovakia

² University of Žilina, Department of Construction Management, Univerzitná 8215/1,
010 26 Žilina, Slovakia

Tel.: + 421415134358

E-mail: michal.hodon@fri.uniza.sk

Summary: Reliable wireless communication is essential for the deployment of distributed monitoring systems in road tunnels, where conventional wired infrastructure can be costly and difficult to extend. This paper presents the design and experimental evaluation of a distributed environmental monitoring platform combining compact sensor nodes with LoRa-based wireless communication. The nodes integrate temperature, humidity, and volatile organic compound (VOC) sensing together with communication performance monitoring. The communication subsystem is based on the ESP32-S3 and Semtech SX1262 platform using the Meshtastic communication framework and is intended to support multi-hop networking in future distributed deployments.

The communication performance was experimentally evaluated in three real road tunnels in Slovakia with lengths of approximately 420 m, 900 m, and 3.6 km. The experiments were intentionally performed as direct point-to-point links using a single transmitter and receiver without intermediate relay nodes. Successful communication was achieved over the complete lengths of the 420 m and 900 m tunnels and at distances of up to 2.2 km in the 3.6 km Čebrať tunnel. The average SNR decreased from 3.28 dB at 420 m to -19.40 dB at 2.2 km, while no gaps were observed in the sequential application payload during the successfully established measurement intervals.

The results demonstrate the potential of sub-GHz LoRa communication for distributed sensing in elongated road tunnel environments and provide an experimental basis for subsequent multi-hop deployment in tunnels where direct communication cannot provide complete coverage.

Keywords: LoRa mesh, Wireless sensor networks, Tunnel monitoring, Environmental monitoring, Air quality monitoring, Volatile organic compounds, Embedded systems, Internet of things.

1. Introduction

Road tunnels are critical components of modern transportation infrastructure, enabling efficient transport through geographically challenging regions while reducing environmental impact and improving traffic safety. To ensure safe operation, tunnels are equipped with numerous monitoring and control systems responsible for supervising environmental conditions, ventilation performance, traffic flow, and emergency situations. Continuous acquisition of operational data is therefore essential for maintaining tunnel safety and supporting informed decision-making by tunnel operators. Traditional monitoring systems are predominantly based on wired communication infrastructure. Although such solutions provide high reliability, their deployment and maintenance in tunnel environments are often associated with significant installation costs, limited flexibility, and challenges related to future expansion. Wireless sensor networks have emerged as a promising alternative capable of reducing deployment complexity while enabling distributed monitoring of environmental and operational parameters. Among available wireless technologies, LoRa has gained considerable attention due to its long communication range, low power consumption, and suitability for

industrial Internet of Things (IoT) applications. However, conventional LoRaWAN architectures rely on direct communication between sensor nodes and gateways, which may be problematic in elongated underground structures such as road tunnels. The use of a LoRa Mesh topology offers a potential solution by enabling multi-hop communication between neighboring nodes, thereby extending network coverage and improving communication resilience.

This paper presents the design and implementation of a distributed monitoring platform based on LoRa Mesh communication for tunnel environments. The proposed system combines wireless communication capabilities with environmental sensing, including temperature, humidity, and volatile organic compound (VOC) monitoring. The objective is to evaluate the applicability of LoRa-based wireless communication for reliable data acquisition in tunnel infrastructures while providing a basis for subsequent multi-hop deployment.

Previous studies by the authors addressed airflow measurements, fire scenarios, emergency events, and numerical modelling in road tunnels [1-5]. In contrast, the present work focuses on the development and experimental evaluation of a distributed wireless sensing infrastructure. The proposed platform integrates LoRa-based communication with compact

autonomous environmental sensor nodes specifically intended for deployment along elongated tunnel structures. Its applicability was experimentally evaluated in three real road tunnels in Slovakia with lengths ranging from 420 m to 3.6 km. Successful communication was achieved over the complete lengths of the 420 m and 900 m tunnels, while in the 3.6 km Čebrať tunnel a communication distance of up to 2.2 km was demonstrated. Previous studies have demonstrated the feasibility of wireless sensor networks for tunnel monitoring, while also highlighting the strong influence of tunnel geometry, surface materials, and sensor-node placement on radio communication performance [6, 7]. Linear and multi-hop network architectures have therefore been investigated as suitable approaches for elongated underground environments where direct communication with a central gateway may not always be possible [8, 9].

2. Materials and Methods

The proposed monitoring platform is based on autonomous wireless sensor nodes designed for experimental deployment in tunnel environments. Each node integrates environmental sensing, local data processing, and wireless communication capabilities within a compact battery-powered device.

The hardware platform utilizes an ESP32-WROOM-32D microcontroller as the main processing unit. The microcontroller is responsible for sensor data acquisition, local data processing, communication management, and power control. Environmental monitoring is provided by a DHT11 sensor for temperature and relative humidity measurements and an SGP40 sensor for volatile organic compound (VOC) monitoring. The DHT11 was selected for the prototype implementation as a low-cost sensor sufficient for demonstrating the functionality of the distributed sensing and communication platform. Due to its limited measurement range and accuracy, it is not intended as an industrial-grade sensor for permanent tunnel monitoring. The modular architecture of the developed node allows it to be replaced by a more accurate and robust temperature and humidity sensor in future deployments without modification of the overall system architecture. In addition, a photoresistor-based light intensity sensor is incorporated to provide supplementary information about environmental conditions. The sensor node includes several auxiliary functions such as battery voltage monitoring, RGB status indication, an acoustic buzzer for local signaling, and a user button for device control and configuration. Communication interfaces including UART and I2C are available for future expansion and integration of additional sensing modules. The complete device is powered by four AAA batteries and enclosed within a compact housing designed for deployment in transportation infrastructure environments. The current prototype is battery-

powered; however, a detailed power consumption characterization was not performed within the scope of the presented experiments. The primary objective was to evaluate the functionality of the sensing platform and the achievable LoRa communication range under real tunnel conditions. Power consumption in different operating modes, including sensing, transmission, reception, and sleep states, will be evaluated as part of future long-term deployment studies. The sensors used in the current prototype were not individually calibrated against reference measurement instruments, as the primary objective of this stage was to evaluate the functionality of the distributed sensing and communication platform rather than to provide metrological-grade environmental measurements. The measured values are therefore used primarily for monitoring relative and temporal changes in environmental conditions. Future tunnel deployments intended for quantitative environmental assessment will require calibrated sensors with accuracy appropriate to the monitored parameters. The overall architecture of the developed sensor node is illustrated in Fig. 1, where the hardware implementation is shown.



Fig. 1. Developed sensor node.

3. VOC Monitoring

Air quality monitoring within road tunnels represents an important component of tunnel safety and operational management. Vehicle emissions, maintenance activities, and limited natural air exchange may contribute to the accumulation of various gaseous pollutants, affecting both user comfort and ventilation system performance. For this reason, the developed monitoring platform incorporates a volatile organic compound (VOC) sensor to provide continuous information about air quality conditions. Road tunnels represent specific environments in which traffic-related emissions can result in elevated concentrations of gaseous and particulate pollutants [10]. Recent studies have also demonstrated the potential of distributed low-cost sensing systems for monitoring spatial and temporal variations of air quality parameters in road tunnels [11].

The proposed sensor node utilizes the SGP40 air quality sensor, which communicates through an I2C interface and provides raw measurement data that are processed using the manufacturer's VOC algorithm. Instead of directly measuring the concentration of individual gases, the sensor outputs a dimensionless VOC Index ranging from 1 to 500, where higher values indicate deteriorating air quality conditions. This approach enables long-term monitoring of air quality trends while maintaining low power consumption and compact hardware requirements.

Although the VOC Index does not represent the concentration of a specific pollutant, it provides a useful indicator of relative changes in VOC levels and the presence of volatile organic compounds generated by vehicle traffic and other environmental sources. Consequently, changes in the measured VOC Index can be used to identify periods of increased traffic activity, evaluate ventilation effectiveness, and detect abnormal environmental conditions within the monitored tunnel environment. It should be emphasized that the VOC Index is used in the proposed platform as a supplementary indicator of temporal changes in the tunnel environment rather than as a direct measure of regulated tunnel pollutants. The SGP40 does not provide concentrations of individual gases and is therefore not intended to replace dedicated monitoring of safety-critical parameters such as CO, NO_x, particulate matter, or visibility. Its role in the proposed system is primarily to provide additional information on relative changes in environmental conditions and to support the evaluation of temporal trends together with the other measured parameters.

The integration of VOC monitoring with the LoRa Mesh communication infrastructure enables distributed air quality sensing at multiple locations throughout the tunnel. Such an approach allows environmental data to be collected simultaneously with communication performance metrics, providing a comprehensive dataset for future studies focused on tunnel operation, environmental monitoring, and intelligent infrastructure management.

4. LoRa Communication

Reliable wireless communication represents one of the key challenges in road tunnel environments. The elongated geometry of tunnel structures, concrete surfaces, technological equipment, and the absence of conventional communication infrastructure can significantly affect radio signal propagation. At the same time, the tunnel geometry may provide specific propagation conditions in which multiple reflections from the tunnel surfaces contribute to signal propagation over relatively long distances.

The developed monitoring platform uses the XIAO ESP32S3 & Wio-SX1262 hardware platform, combining an ESP32-S3 microcontroller with a Semtech SX1262 LoRa transceiver. Communication is implemented using the Meshtastic platform in the European 868 MHz ISM band. For the tunnel

experiments, the radio was configured to operate at a center frequency of 869.5 MHz with a bandwidth of 125 kHz, spreading factor SF12, coding rate 4/6, and a transmit power of 22 dBm. These parameters were selected to provide high receiver sensitivity and robust communication over long distances. Meshtastic additionally provides mesh networking capabilities, allowing individual nodes to forward messages and thus enabling multi-hop extension of the monitoring network where direct communication is not available.

Meshtastic provides the communication and networking layer used by the developed platform. Messages are transmitted as LoRa packets and can be retransmitted by neighboring nodes in a mesh deployment, allowing the communication path to be extended beyond the range of a single radio link. The platform also supports encrypted communication between nodes sharing the same channel configuration. In the experiments presented in this work, however, the mesh forwarding functionality was intentionally not utilized. The communication topology consisted exclusively of one transmitter and one receiver, and therefore no routing decisions or intermediate hops were involved. This configuration allowed the radio propagation and packet delivery performance of a single LoRa link to be evaluated independently of the effects introduced by mesh routing.

For the experimental evaluation, the measurements were intentionally performed without intermediate relay nodes in order to determine the achievable direct LoRa communication range under real tunnel conditions. Each measurement involved only one transmitting and one receiving device. The devices were positioned at predefined locations near the side of the tunnel rather than along its central axis, corresponding more closely to the expected placement of monitoring equipment in practical tunnel installations.

During the experiments, communication performance was evaluated using parameters reported by the receiving node, particularly Received Signal Strength Indicator (RSSI) and Signal-to-Noise Ratio (SNR). The experiments were complemented by radio-frequency spectrum observations using an SDR receiver, allowing the LoRa transmissions to be observed independently in the frequency domain. The obtained measurements were subsequently used to evaluate communication performance as a function of distance and tunnel geometry.

5. Experimental Evaluation in Road Tunnels

The communication performance of the developed platform was experimentally evaluated in three real road tunnels in Slovakia: Svrčinovec, Poľana, and Čebrať. The tunnels have approximate lengths of 420 m, 900 m, and 3.6 km, respectively, providing different propagation distances for the evaluation of LoRa communication under realistic underground conditions.

For each measurement, only two LoRa devices were used, operating as a transmitter and a receiver. No intermediate relay or mesh nodes were deployed. The devices were positioned near the side of the tunnel at predefined locations rather than in the central axis of the tunnel tube. This configuration was selected to represent a realistic installation scenario for distributed monitoring nodes. The same radio configuration was maintained during the experiments: a center frequency of 869.5 MHz, bandwidth of 125 kHz, spreading factor SF12, coding rate 4/6, and transmit power of 22 dBm.

The experimental campaign was motivated by the specific propagation conditions of elongated tunnel structures. Although reinforced concrete and installed infrastructure can introduce attenuation and multipath propagation, the approximately tubular geometry may also support radio propagation through multiple reflections from tunnel surfaces. The objective of the measurements was therefore to determine the achievable direct LoRa communication distance under real tunnel conditions rather than to establish a detailed electromagnetic propagation model. During the experimental campaign, the communication devices were installed at predefined positions along the tunnel. Both communication devices were equipped with vertically oriented omnidirectional whip antennas designed for the 868 MHz ISM band. Rather than placing the antennas in the central axis of the tunnel tube, the nodes were positioned close to the tunnel side wall, as illustrated in Fig. 2. This arrangement was selected to approximate a practical installation scenario, where sensor and communication nodes would typically be mounted along the tunnel infrastructure. The measurement positions were determined using the tunnel distance markers, enabling the transmitter–receiver separation to be defined for individual experiments.



Fig. 2. LoRa communication devices based on the XIAO ESP32S3 and Wio-SX1262 platform used during the tunnel experiments and their placement near the tunnel wall.

Reliable packet reception was achieved over the complete lengths of the Svrčinovec and Poľana tunnels, corresponding to communication distances of approximately 420 m and 900 m, respectively. At 420 m, the recorded packets exhibited an average RSSI of approximately -111.3 dBm and an average SNR of 3.3 dB. At 900 m, the average reported RSSI decreased

to approximately -127.5 dBm, while the average SNR decreased to approximately -10.1 dB.

Longer-distance measurements were subsequently performed in the 3.6 km Čebrať tunnel. Successful communication was confirmed at distances of 1.5 km and 2.2 km. At 1.5 km, an average SNR of approximately -15.3 dB was recorded, while at 2.2 km the average SNR decreased to approximately -19.4 dB. Communication over the complete 3.6 km length of the Čebrať tunnel was not achieved. Therefore, 2.2 km represents the maximum successfully demonstrated point-to-point communication distance during the experimental campaign.

Notably, the 2.2 km communication distance was achieved using a single direct LoRa link without any intermediate relay nodes. Although the experiment did not provide complete coverage of the Čebrať tunnel, the results demonstrate that a substantial portion of a multi-kilometre road tunnel can be covered by a single LoRa hop. This finding also provides practical motivation for the mesh architecture of the proposed platform, where intermediate nodes could be used to extend communication beyond the achievable direct-link range.

Packet delivery was evaluated using a sequential counter included in the application payload. Each newly generated message incremented the counter by one, allowing missing messages to be identified directly from the received data. Duplicate receptions were excluded from the packet delivery calculation. No gaps in the payload sequence were observed at the successfully established measurement points, corresponding to a 100 % delivery rate for unique application messages during the evaluated intervals.

Communication quality was evaluated primarily using the SNR values reported by the receiving LoRa node. Fig. 3 shows the average SNR obtained at the successfully tested transmitter–receiver distances. A progressive decrease in SNR can be observed with increasing communication distance, from positive values at 420 m to approximately -19.4 dB at the maximum successfully demonstrated distance of 2.2 km.




Fig. 3. Average SNR measured for successful point-to-point LoRa communication at different transmitter–receiver distances.

The obtained results are consistent with recent observations demonstrating that LoRa communication can achieve kilometre-scale coverage under favorable tunnel propagation conditions [7]. In contrast, experiments performed in more constrained underground environments reported substantially shorter single-hop communication ranges and consequently employed multi-hop LoRa architectures [8]. These observations highlight the strong dependence of the achievable LoRa communication range on the geometry and propagation characteristics of the particular underground environment.

In addition to the average values, the SNR of individual received messages was analyzed to evaluate the stability of the established links. As shown in Fig. 4, the SNR remained relatively stable during each measurement interval despite the substantial differences in communication distance. Even at 2.2 km, where the received signal operated close to the sensitivity limits of the LoRa link, successful packet reception was maintained with SNR values concentrated around -19.4 dB.




Fig. 4. SNR of individual received LoRa messages during successful point-to-point communication tests at different transmitter–receiver distances.

The LoRa transmissions were additionally monitored using an SDR receiver. The recorded spectrum confirmed the presence of the characteristic LoRa chirp modulation around the configured center frequency of 869.5 MHz. An example of the received signal spectrum during the long-distance tunnel experiment is presented in Fig. 5.

6. Conclusion

This work presented a distributed environmental monitoring platform designed for road tunnel applications, combining compact sensor nodes with LoRa-based wireless communication. The communication performance of the developed platform was experimentally evaluated under real conditions in three road tunnels in Slovakia.

Point-to-point communication was successfully established over the complete lengths of the Svrčinovec (420 m) and Poľana (900 m) tunnels. In the

3.6 km Čebrať tunnel, successful communication was demonstrated at distances of 1.5 km and up to 2.2 km, while direct communication over the complete tunnel length was not achieved. The average SNR decreased from 3.28 dB at 420 m to -19.40 dB at 2.2 km. Importantly, these results were obtained using only a single transmitter and receiver without intermediate relay nodes.




Fig. 5. SDR spectrum and waterfall representation of the LoRa transmission observed during the tunnel communication experiment at 869.5 MHz.

The experiments demonstrate the potential of sub-GHz LoRa communication for distributed sensing in elongated road tunnel environments. The observed communication distances may be influenced by the specific tunnel geometry and multipath propagation; however, a detailed characterization of the propagation mechanisms was outside the scope of this study.

The demonstrated 2.2 km single-hop communication range also provides practical motivation for the mesh architecture of the proposed platform. In longer tunnels where a direct link is insufficient, intermediate nodes can potentially extend the communication coverage through multi-hop transmission. Future work will therefore focus on multi-hop experiments, optimization of relay-node placement, long-term network reliability, and the integration of additional tunnel-relevant environmental sensors.

Acknowledgement

This work was supported by the Slovak Research and Development Agency (APVV) under project No. APVV-24-0429, Experimental Investigation and Computer Modelling of Airflow during Road Tunnel Fire.

References

- [1]. P. Danišovič, J. Šrámek, M. Hodoň, J. Glasa, et al., Measurements and fire simulation models in road

- tunnels, *MATEC Web of Conferences*, Vol. 196, 2018, 04077.
- [2]. P. Danišovič, J. Šrámek, M. Hodoň, Testing measurements of airflow velocity in road tunnels, *MATEC Web of Conferences*, Vol. 117, 2017, 00035.
 - [3]. P. Danišovič, J. Šrámek, M. Hodoň, 3D and 2D models of smoke spread in road tunnel as didactic tool, in *Proceedings of the 10th International Conference on Education and New Learning Technologies (EDULEARN)*, 2018, pp. 6817-6824.
 - [4]. P. Danišovič, J. Šrámek, M. Hodoň, Simulation of emergency events with fire in virtual road tunnels, *MATEC Web of Conferences*, Vol. 251, 2018, 02033.
 - [5]. J. Šrámek, P. Danišovič, M. Hodoň, Evaluation and analysis of emergency events in road tunnels, *Transportation Research Procedia*, Vol. 40, 2019, pp. 1571-1575.
 - [6]. X. Li, Z. Ji, H. Zhu, et al., A feasibility study of the measuring accuracy and capability of wireless sensor networks in tunnel monitoring, *Frontiers of Structural and Civil Engineering*, Vol. 6, Issue 2, 2012, pp. 111-120.
 - [7]. M. Meier, S. Gorenflo, N. Canzani, Y. S. Femenias, Measurement of 868 MHz LoRa radio propagation in a tunnel, *Hochschule für Technik und Umwelt FHNW*, 2026, 10.26041/fhnw-15373.
 - [8]. A. Abrardo, A. Pozzebon, A multi-hop LoRa linear sensor network for the monitoring of underground environments: The case of the medieval aqueducts in Siena, Italy, *Sensors*, Vol. 19, Issue 2, 2019, 402.
 - [9]. H. Yang, H. Guo, J. Jia, Z. Jia, et al., Self-organizing and routing approach for condition monitoring of railway tunnels based on linear wireless sensor network, *Sensors*, Vol. 24, Issue 20, 2024, 6502.
 - [10]. S. Marinello, F. Lolli, R. Gamberini, Roadway tunnels: A critical review of air pollutant concentrations and vehicular emissions, *Transportation Research Part D: Transport and Environment*, Vol. 86, 2020, 102478.
 - [11]. S. De Vito, A. Del Giudice, G. D'Elia, E. Esposito, et al., Correlating air pollution concentrations and vehicular emissions in an Italian roadway tunnel by means of low cost sensors, *Atmosphere*, Vol. 14, Issue 4, 2023, 679.

Wind Feature Engineering for Tree-Sway Energy Harvesting Assessment in Forest IoT Nodes

P. Kolok, J. Kapitulík, M. Chochul and J. Šumský

University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1, 010 26 Žilina, Slovakia

Tel.: +421904390571

E-mail: peter.kolok@fri.uniza.sk

Summary: Wind-induced tree sway is a potential kinetic energy source for ultra-low-power sensing nodes deployed in forest environments. This paper presents a wind time-series processing and feature-extraction workflow intended to provide the wind-side input representation for subsequent model-based evaluation of tree-sway energy harvesting. Field-measured wind-speed samples are calibrated, quality-control flagged, filtered with a short median filter, and divided into fixed 10 min windows. Each window is represented by descriptive and excitation-related features, including mean and median wind speed, variability, high-percentile and maximum values, turbulence intensity, gust factor, threshold exceedance ratio, and nonlinear proxies based on squared and cubed wind speed. The workflow is demonstrated using one 24 h record sampled at 10 samples/s, corresponding to 864000 nominal samples and 144 nominal analysis windows. The record had a mean wind speed of 2.62 m/s, a 95th-percentile value of 6.8 m/s, and a maximum of 13.87 m/s. These values show that the daily mean alone does not describe the upper tail of the measured wind regime. Squared and cubed terms emphasize high-speed intervals but are treated as physical opportunity indicators rather than direct measurements of harvested electrical energy. The resulting feature vector is intended for validation against synchronized inertial tree-sway measurements and downstream harvester models.

Keywords: Wind feature extraction, Wind time-series analysis, Tree sway, Energy harvesting, Forest IoT, Gust Indicators, Inertial sensing.

1. Introduction

Wireless sensor networks and IoT systems are increasingly deployed in outdoor and remote environments to monitor environmental conditions and structural behaviour. In forest deployments, long-term energy supply remains a critical limitation because access to individual nodes is difficult and regular battery replacement increases maintenance requirements and environmental impact.

Energy harvesting can reduce this dependence by converting locally available ambient energy into electrical energy. Solar harvesting is widely used, but dense vegetation, shading, seasonal variability, and installation constraints can limit its reliability. Wind-induced mechanical motion therefore represents a complementary energy source that is naturally present in forest environments and spatially co-located with many sensing tasks [1, 2].

Wind excitation in forests is strongly time-varying and includes short gusts, turbulent fluctuations, and slower changes associated with weather evolution. The resulting tree response also depends on crown geometry, structural stiffness, damping, and root-soil interaction. Consequently, similar mean wind speeds can be associated with different motion histories and different opportunities for kinetic-energy conversion [3, 4].

This distinction is relevant because mechanically important wind-related quantities do not scale linearly with wind speed. Aerodynamic loading is approximately related to the square of wind speed, whereas wind-power density is proportional to its cube [5]. A descriptor based only on the mean can therefore

conceal short high-speed events. Window-based feature extraction offers a practical compromise by retaining information about variability, gustiness, and nonlinear excitation while reducing the volume of data supplied to subsequent motion and harvester models.

This study is deliberately limited to characterization of the measured wind-speed time series. It does not claim to measure harvested electrical energy directly. Instead, it extracts window-based descriptors that can later be compared with synchronized inertial tree-motion features and used as inputs to physically based or data-driven harvester models.

The contribution of the paper is a focused preprocessing and feature-engineering workflow that converts raw wind measurements into a compact set of statistically and physically interpretable predictors. These descriptors form the wind-side input layer of the planned model-based evaluation: their value must ultimately be established by comparison with synchronized tree-motion features and electrical harvester output. The workflow is demonstrated on one 24 h field record, and its present scope and validation limits are stated explicitly.

2. Measurement Platform and Dataset

The measurement concept combines wind-speed sensing with inertial measurement of tree motion. The present paper reports the wind-measurement channel as the primary input for feature extraction. The inertial channel defines the downstream mechanical response against which the extracted wind descriptors can be validated in subsequent work.

Measurements were acquired using the C6-Datalog field platform developed for synchronized wind and inertial sensing [6]. The device is based on an ESP32-C6 module and integrates analog acquisition, I²C sensor interfaces, a microSD card, and a low-power operating mode. Tree motion is measured using a six-axis MPU-6050 inertial measurement unit, while wind speed is measured using an analog cup anemometer. The anemometer has a nominal range of 0-30 m/s represented by a 0-5 V output signal [7].

For the analysed campaign, the anemometer was mounted on a steel support near the monitored tree at approximately half of the tree height. The logger recorded wind speed at 10 samples/s. This configuration provides fine temporal sampling of the sensor output while remaining compatible with long-duration local logging; the dynamic response of the cup anemometer itself was not independently characterized. The campaign consisted of 24 h records collected over one week, and the present paper analyses one selected 24 h record without claiming that it represents seasonal wind conditions.

At startup, the logger obtains an NTP-synchronized timestamp and subsequently disables wireless communication during acquisition. Each wind and IMU sample is stored with both an absolute timestamp and a relative time from the beginning of the record. Sampling and storage are separated using double buffering to reduce the effect of variable microSD write latency on acquisition timing. The resulting CSV files can therefore be aligned during offline processing without manual synchronization.

Calibration was performed in two stages. First, a fixed offset of approximately 0.25 V introduced by the analog front end was removed. Second, comparison with a UNI-T UT362 reference anemometer indicated that the cup-anemometer channel underestimated wind speed by approximately 15 %. A multiplicative correction factor was therefore applied before feature extraction [6]. These corrections are particularly important for the upper tail of the distribution because calibration errors are amplified when squared and cubed wind-speed terms are evaluated.

The platform stores synchronized measurements locally and enables offline processing of wind and motion time series. Fig. 1 separates the wind-processing and IMU-processing branches and shows where their feature vectors enter the downstream model-based evaluation. Only the wind-processing and wind-feature-extraction branch is quantitatively evaluated in this paper.

3. Wind Signal Processing and Feature Extraction

Raw wind-speed samples are first calibrated and subjected to quality-control (QC) flagging to identify sensor malfunctions and missing segments. A short odd-length median filter, restricted to a kernel of 3 or 5 samples, is then applied to suppress isolated measurement spikes while retaining short wind-speed

peaks. The cleaned signal is subsequently divided into fixed analysis windows W_k , with a primary duration of 10 min.




Fig. 1. Separate wind and IMU processing branches for subsequent model-based evaluation of wind-driven tree-sway energy harvesting.

The raw record is represented by the absolute timestamp, relative timestamp, and 12-bit ADC code for each sample. The ADC code is first converted to voltage using the 3.3 V reference. After subtraction of the measured front-end offset, the compensated voltage is mapped linearly to wind speed using the nominal gain of $6 \text{ m s}^{-1} \text{ V}^{-1}$ and a calibration factor of approximately $1/0.85$ [6]. Negative compensated values are clipped to zero to avoid non-physical wind-speed estimates.

Quality control is retained as an explicit stage rather than being hidden inside the filter. Separate flags identify non-monotonic timestamps, irregular sampling, missing intervals, values outside the nominal sensor range, and nearly constant segments that may indicate a stuck sensor. The selected 3- or 5-sample kernel is held fixed throughout a processed record. The kernel is intentionally short because aggressive smoothing could attenuate genuine gust peaks that are relevant to mechanical excitation.

The cleaned signal is divided into fixed 10 min analysis windows. At a sampling rate of 10 Hz, each complete window contains approximately 6000 samples; a complete 24 h record therefore contains 144 nominal windows and 864000 nominal samples. QC status is retained alongside each window so that incomplete or suspect intervals remain identifiable in downstream analysis. A secondary 60 min aggregation can be used for longer-term summaries. The median and the 10th-90th percentile envelope are additionally calculated for visualization of the within-window distribution.

For each window W_k containing N_k valid samples, the statistical features include the mean wind speed μ_v , k , standard deviation σ_v , k , 95th percentile $v_{95, k}$, and maximum wind speed $v_{\max, k}$. Gust-related descriptors are represented by turbulence intensity and gust factor, which distinguish a relatively steady interval from a highly intermittent interval having the same mean speed [8]:

$$TI_k = \frac{\sigma_{v,k}}{\bar{v}_k}, \quad GF_k = \frac{v_{max,k}}{\bar{v}_k}$$

The threshold-exceedance duty cycle is defined as $Dk(v > v_0) = N_k^{-1} \sum I(v_{i,k} > v_0)$, where v_0 is an application-specific operating threshold and $I(\cdot)$ is the indicator function. The threshold remains configurable and is not used to infer electrical energy directly.

Since aerodynamic loading approximately scales with v^2 , while wind-power availability scales with v^3 [5], nonlinear harvesting-related proxies are computed as:

$$\overline{v^2}_k = \frac{1}{N_k} \sum_{i=1}^{N_k} v_{i,k}^2, \quad \overline{v^3}_k = \frac{1}{N_k} \sum_{i=1}^{N_k} v_{i,k}^3$$

The resulting descriptors represent the background wind level, short-term variability, and concentration of high-speed events. The corresponding integrated cubic index is $Ev^3_k = N_k \Delta t \cdot \text{mean}(v^3)_k$, where Δt is the sampling interval. Squared and cubed terms are used as physically motivated proxies for aerodynamic loading and wind-energy opportunity; they are not interpreted as direct measurements of tree-sway energy or electrical harvester output. In near-calm windows where μv_k approaches zero, turbulence intensity and gust factor are treated as undefined rather than forcing numerically large values.

Table 1. Wind features extracted for each analysis window.

Feature	Physical interpretation
Mean (μv)	Background wind level and reference operating condition.
Variability (σv)	Short-term variability and unsteadiness of the airflow.
P95 (v_{95})	Robust high-wind indicator less sensitive to isolated spikes.
Maximum (v_{max})	Largest gust observed within the analysis window.
Turbulence intensity (TI)	Variability normalized by the mean wind speed.
Gust facto (GF)	Peak-to-mean ratio describing intermittent gust severity.
Duty cycle ($D(v > v_0)$)	Fraction of samples exceeding an application threshold.
Drag proxy ($\overline{v^2}$)	Approximate indicator of aerodynamic loading on the tree.
Power proxy ($\overline{v^3}$)	Nonlinear indicator of instantaneous wind-energy availability.
Energy index (E_v)	Time-integrated opportunity for energy accumulation in a window.

4. Observed Wind-Feature Behaviour and Role in Model-Based Evaluation

The extracted features provide a compact wind-side representation for subsequent comparison with tree-sway descriptors and downstream model-based energy-yield analysis. The wind signal is not represented only by its mean because this would conceal information about variability and short

high-speed events. Each window is therefore described using central-tendency, dispersion, upper-tail, gust, threshold, and nonlinear proxy features.

Fig. 2 presents examples of the extracted features over the analysed 24 h period. Fig. 2a shows the wind-speed envelope, where the median represents the central tendency and the 10-90 % band describes the within-window spread. Fig. 2b shows squared and cubed wind-speed proxies. These transformations increase the visual emphasis assigned to high-speed intervals, as expected from their nonlinear definitions.




Fig. 2a. Wind-speed median and 10-90 % envelope over the analysed 24 h record.




Fig. 2b. Squared and cubed wind-speed proxies over the analysed 24 h record.

The analysed record represents a low-to-moderate background wind regime interrupted by several short high-wind episodes. Across the complete 24 h period, the mean wind speed was 2.62 m/s, the 95th percentile was 6.8 m/s, and the maximum recorded value was 13.87 m/s [6]. The 95th-percentile speed was approximately 2.60 times the daily mean, while the maximum was approximately 5.29 times the mean. The separation between these values shows quantitatively that the daily average alone does not fully describe the measured wind regime.

Three elevated intervals are visible in the time-domain features. During these intervals, the median and high-percentile wind speeds increase together with within-window variability. Turbulence intensity and gust factor provide complementary normalized descriptions: the former expresses variability relative to the mean, whereas the latter expresses the largest valid sample relative to the mean. Values from near-calm windows are omitted when the normalization is not numerically meaningful.

The nonlinear proxies concentrate most strongly around the same elevated intervals. Squaring wind speed increases the contrast of periods associated with larger aerodynamic loading, while cubing it produces sharper peaks for high-speed events. The

time-integrated cubic index additionally distinguishes a brief isolated peak from a sustained interval. This behaviour follows directly from the mathematical transformations and supports their use as candidate predictors; it does not prove that they predict tree motion or electrical output better than mean wind speed.

The empirical cumulative distribution in Fig. 3 is strongly right-skewed, with most samples concentrated at low-to-moderate wind speeds and a comparatively sparse upper tail [6, 9]. From a harvester-design perspective, such a distribution suggests long intervals of limited excitation separated by rarer high-speed events. Threshold duty-cycle features quantify this availability by reporting the fraction of valid samples in each window above a selected operating threshold.




Fig. 3. Empirical cumulative distribution function of wind speed in the analysed record.

For subsequent wind-sway analysis, the extracted wind-feature vector can be compared with inertial descriptors such as angular range, angular-velocity RMS, dominant sway frequency, and spectral energy. High-percentile and nonlinear wind features are plausible candidates for this comparison, but their predictive value must be established using synchronized wind and IMU measurements rather than assumed from the wind signal alone.

Previous work has demonstrated the practical feasibility of powering a low-power wireless sensor node from tree movement [10], while recent field research shows that energy transfer during tree motion depends on wind conditions and forest configuration [11]. These studies support synchronized wind-motion analysis, but they do not remove the need to validate the present feature set on the measured platform.

The present study therefore stops at feature extraction and descriptive analysis. It does not compare electromagnetic harvester architectures, estimate conversion efficiency, or report electrical energy yield. Those stages require explicit mechanical and electrical models or synchronized prototype measurements and are treated as downstream work.

5. Limitations and Future Work

The current evaluation is based on one 24 h record from a single measurement site. The selected record is not claimed to represent the seasonal wind climatology of the site. It does not capture seasonal changes in

canopy density, storm conditions, differences among tree species, or changes in surrounding terrain. The anemometer was installed near the monitored tree at approximately half-tree height, so the measurements represent local excitation and should not be interpreted as a complete vertical wind profile.

The cup anemometer measures wind-speed magnitude only, and wind direction was not included. Calibration was performed against a handheld reference instrument, while uncertainty over a wider speed range was not quantified in the present analysis. In addition, the dynamic response of the anemometer was not independently characterized, and only fixed 10 min windows were evaluated. Future experiments should compare multiple window lengths, document the treatment of missing and near-calm windows explicitly, repeat calibration over a wider range, and evaluate direction-sensitive or multi-point wind sensing.

Most importantly, the squared wind-speed term, cubed wind-speed term, and integrated cubic index are physical proxies rather than direct measurements of tree-sway energy or harvested electrical energy. Complete validation requires synchronized wind, IMU, mechanical displacement, and harvester-output measurements. Longer datasets can then be used to compare feature subsets and quantify how accurately the selected wind descriptors explain tree motion and usable energy delivered through the power-management circuit.

6. Conclusions

This paper presented a wind time-series feature-extraction workflow that supplies the wind-side input representation for preliminary model-based assessment of wind-driven tree-sway energy-harvesting opportunities. The workflow combines calibration, explicit quality-control flagging, short-window spike suppression, fixed-window segmentation, and extraction of statistical, gust-related, threshold-based, and nonlinear wind descriptors. The analysed 24 h record contained 864000 nominal samples and 144 nominal 10 min windows. Its mean wind speed was 2.62 m/s, the 95th-percentile value was 6.8 m/s, and the maximum was 13.87 m/s. The 95th percentile and maximum were approximately 2.60 and 5.29 times the mean, respectively, confirming that the daily mean does not fully describe the upper tail of the measured wind regime.

Squared and cubed wind-speed proxies provide a compact means of emphasizing high-speed intervals, but they remain opportunity indicators rather than direct measurements of harvested energy. The extracted feature vector is suitable for subsequent validation against synchronized inertial tree-sway features and electrical output. Future work should extend the dataset across seasons and sites, compare window lengths and feature subsets, and quantify whether the selected descriptors improve prediction of

measured tree motion and usable prototype-harvester output over simpler mean-wind baselines.

References

- [1]. S. Roundy, P. K. Wright, J. Rabaey, Energy Scavenging for Wireless Sensor Networks: With Special Focus on Vibrations, *Springer*, Boston, 2004.
- [2]. P. D. Mitcheson, E. M. Yeatman, G. K. Rao, A. S. Holmes, et al., Energy harvesting from human and machine motion for wireless electronic devices, *Proceedings of the IEEE*, Vol. 96, Issue 9, 2008, pp. 1457-1486.
- [3]. R. B. Stull, An Introduction to Boundary Layer Meteorology, *Springer*, Dordrecht, 1988.
- [4]. D. Sellier, T. Fourcaud, Crown structure and wood properties: Influence on tree sway and response to high winds, *American Journal of Botany*, Vol. 96, Issue 5, 2009, pp. 885-896.
- [5]. J. F. Manwell, J. G. McGowan, A. L. Rogers, Wind Energy Explained: Theory, Design and Application (2nd Ed.), *John Wiley & Sons*, Chichester, 2009.
- [6]. J. Šumský, Energy harvesting systems, PhD Thesis, *University of Žilina*, Žilina, 2026.
- [7]. Guide to Instruments and Methods of Observation (WMO-No. 8) (2023 Ed.), *World Meteorological Organization*, Geneva, 2023.
- [8]. B. Yu, A. G. Chowdhury, Gust factors and turbulence intensities for the tropical cyclone environment, *Journal of Applied Meteorology and Climatology*, Vol. 48, Issue 3, 2009, pp. 534-552.
- [9]. C. G. Justus, W. R. Hargraves, A. Mikhail, D. Graber, Methods for estimating wind speed frequency distributions, *Journal of Applied Meteorology*, Vol. 17, Issue 3, 1978, pp. 350-353.
- [10]. S. McGarry, C. Knight, Development and successful application of a tree movement energy harvesting device, to power a wireless sensor node, *Sensors*, Vol. 12, Issue 9, 2012, pp. 12110-12125.
- [11]. K. Kamimura, K. Nanko, A. Matsumoto, S. Ueno, et al., Energy transfer during tree movement for different wind conditions and forest configurations, *Forest Ecology and Management*, Vol. 571, 2024, 122223.

Structural Optimization of a Circular Restoring Spring for a Low-Frequency Electromagnetic Standard Vibrator

X. F. Zhang^{1,2}, **T. Xu**¹, **T. T. Du**¹ and **M. C. Xu**¹

¹ College of Mechanical Engineering, Taiyuan University of Technology, No. 79 Yingze Street, Taiyuan 030024, Shanxi, China

² MOE Key Laboratory of Advanced Transducers and Intelligent Control System, Taiyuan University of Technology, Taiyuan 030024, Shanxi, China
Tel.: +86 18535127169
E-mail: zhangxufei@tyut.edu.cn

Summary: Low-frequency electromagnetic standard vibrators are used for vibration sensor calibration, where stable sinusoidal motion and low waveform distortion are required. Under low-frequency and large-displacement operation, nonlinear restoring-spring behavior can degrade waveform quality. This study investigates how the gap width and thickness of a circular restoring spring affect the first-order natural frequency and displacement waveform distortion and performs coordinated multi-objective optimization. Finite element analysis, range analysis, equivalent stiffness modeling, response surface modeling and a multi-objective gray wolf optimizer were combined. Spring thickness was found to dominate both responses. Relative to the unfavorable initial design, the optimized design increased the simulated first-order natural frequency by 467.33 % and reduced the simulated waveform distortion by 89.15 %. Finite element verification yielded 21.419 Hz and 0.38 % distortion, whereas experiments on the optimized structure yielded 26 Hz and 3.15 %, respectively. The results demonstrate the effectiveness of structural optimization for improving intrinsic vibrator dynamics.

Keywords: Circular restoring spring, Electromagnetic vibrator, Waveform distortion, Finite element analysis, Multi-objective optimization.

1. Introduction

Electromagnetic vibrators provide controllable sinusoidal excitation for vibration-sensor calibration [1, 2]. At low frequency and large displacement, nonlinear support stiffness, damping and magnetic-field nonuniformity can introduce harmonic components into the output signal [3]. Support-spring optimization is not new: Zhou and Han optimized a laminated composite support spring using experimental modal analysis and finite element analysis to improve low-frequency shaker response [4]. Waveform distortion has also been addressed mainly through signal-level compensation and harmonic-suppression methods [5-7]. These studies establish the importance of both support design and waveform control, but do not address coordinated geometry optimization of a circular restoring spring for natural frequency and displacement distortion.

Unlike previous work focused mainly on signal-level waveform control or general support-spring design, this study investigates how the gap width D and thickness h affect both the first-order natural frequency and displacement waveform distortion of a circular restoring spring. A response-surface-based multi-objective gray wolf optimizer (MOGWO) is then used to obtain a compromise structure verified by finite element analysis and experiment.

2. Method

The circular restoring spring consists of an inner ring, an outer ring and several support arms. The outer ring is fixed to the base, while the inner ring is connected to the moving part of the vibrator. Considering the spatial constraints of the existing vibrator, the maximum radius of the support arms was limited to 75 mm. The gap width D and thickness h were selected as the design variables. The tested levels were $D = 4, 6$ and 8 mm, and $h = 1, 1.5, 2, 2.5$ and 3 mm, giving 15 parameter combinations.

The system can be reasonably simplified for simulation purposes. Specifically, the vibrator is modeled as an equivalent vibration system composed of the lumped mass of the moving parts and the circular restoring spring. Based on this simplification, the model was imported into ANSYS for meshing, and the resulting simplified simulation model is shown in Fig. 1.




Fig. 1. Simplified simulation model.

Since the deformation is mainly concentrated in the support arms, the support arms are idealized as thin plates subjected to transverse loading. The equivalent stiffness K_s and the first-order natural frequency f can be approximately expressed as:

$$\begin{cases} K_s = C \cdot \frac{Eh^3}{D^2} \\ f = \frac{1}{2\pi} \sqrt{\frac{K_s}{m}} \end{cases} \quad (1)$$

where C is a constant determined by the structural configuration and boundary conditions; E is the Young's modulus of the material. Eq. (1) indicates that the stiffness is strongly dependent on the circular spring thickness and tends to decrease with increasing gap width. This provides the theoretical basis for explaining the dominant influence of h observed in the finite element results.

A second-order polynomial response surface model was constructed to approximate the mapping between the structural parameters and performance indices:

$$f(x) = \beta_0 + \sum_{i=1}^n \beta_i x_i + \sum_{i=1}^n \beta_{ii} x_i^2 + \sum_{i=1}^{n-1} \sum_{j=i+1}^n \beta_{ij} x_i x_j, \quad (2)$$

where $f(x)$ denotes the predicted response; x_i and x_j are the design variables; n represents the number of design variables; and β_i and β_{ij} are the polynomial coefficients.

The steady-state displacement histories obtained from finite element analysis were exported and processed in MATLAB using the fast Fourier transform (FFT). Displacement waveform distortion was quantified as the total harmonic distortion (THD), calculated as

$$D = \frac{\sqrt{\sum_{n=2}^N A_n^2}}{A_1} \times 100\%, \quad (3)$$

where A_1 is the amplitude of the fundamental displacement component, A_n is the amplitude of the n th harmonic component, and N is the highest harmonic order included in the calculation. The experimental displacement data were processed using the same FFT-based procedure.

The optimization was formulated to minimize displacement THD and maximize the first-order natural frequency. In a minimization framework, it can be expressed as:

$$\begin{cases} x = [D, h]^T = [x_1, x_2]^T \\ \min f(x) = [y_1(x), -y_2(x)]^T \end{cases} \quad (4)$$

3. Results and Validation

The finite element results show that increasing h substantially raises the first-order natural frequency and generally reduces displacement THD, whereas increasing D decreases the first-order natural frequency. Range analysis confirms that h is the

dominant parameter, with ranges of 14.5103 Hz for frequency and 2.24 % for distortion, both larger than those of D . Fig. 2 shows the response-surface fits for the first-order natural frequency and displacement THD. The natural-frequency model achieved $R^2 = 0.99$, while the distortion model achieved $R^2 = 0.89$.

The multi-objective gray wolf optimizer (MOGWO) produced a non-dominated solution set, from which $D = 5$ mm and $h = 3$ mm was selected as a representative compromise solution, as shown in Fig. 3.




Fig. 2. Main numerical results: (a) first-order natural frequency response surface, (b) displacement THD response surface.

For the unfavorable initial design, Combination 11 ($D = 8$ mm and $h = 1$ mm), the simulated first-order natural frequency was 3.7166 Hz and the displacement THD was 3.78 %. MOGWO predicted 21.0856 Hz and 0.41 %, while finite element verification yielded 21.419 Hz and 0.38 %, respectively. The corresponding prediction errors were below 10 %.




Fig. 3. Non-dominated solution set obtained by the multi-objective gray wolf optimizer.

Experimental validation was carried out using a signal generator, power amplifier, single-axis electromagnetic vibrator, data acquisition card, PC and eddy-current displacement sensor. The optimized structure achieved an experimental first-order natural frequency of 26 Hz and a displacement THD of 3.15 %. Compared with Combination 11 and Combination 8, the first-order natural frequency increased by 271.43 % and 52.94 %, respectively. The experimental displacement THD of 3.15 % was substantially higher than the finite element value of 0.38 %. This discrepancy is attributed to effects not

fully represented in the simplified structural model, including manufacturing tolerances, assembly constraints, nonlinear damping, nonuniform electromagnetic force and measurement disturbances. Accordingly, the simulated distortion is used primarily as a relative indicator for structural comparison rather than an exact prediction of the complete experimental system. The frequency difference is additionally affected by simplified boundary conditions and deviations in the actual moving mass.

4. Conclusions

A multi-parameter structural optimization method was developed for the circular restoring spring of a low-frequency electromagnetic vibrator. Spring thickness was the dominant parameter affecting both first-order natural frequency and displacement THD. The optimized design with $D = 5$ mm and $h = 3$ mm increased the simulated first-order natural frequency and reduced the simulated displacement THD, while the experiment yielded a first-order natural frequency of 26 Hz and a THD of 3.15 %. Although the experimental THD remained higher than the numerical prediction, the results support structural optimization as a complementary approach to signal-level waveform-control methods for improving the intrinsic dynamic performance of vibration-calibration systems.

References

- [1]. S. Hoffait, F. Marin, D. Simon, B. Peeters, et al., Measured-based shaker model to virtually simulate vibration sine test, *Case Studies in Mechanical Systems and Signal Processing*, Vol. 4, 2016, pp. 1-7.
- [2]. K. P. S. Rana, V. Kumar, Fuzzy control of an electrodynamic shaker for automotive and aerospace vibration testing, *Expert Systems with Applications*, Vol. 38, Issue 9, 2011, pp. 11335-11346.
- [3]. A. Saraswat, N. Tiwari, Modeling and study of nonlinear effects in electrodynamic shakers, *Mechanical Systems and Signal Processing*, Vol. 85, 2017, pp. 162-176.
- [4]. L. Zhou, J. Han, Optimum design of electrodynamic shaker's support spring to improve low frequency performance, *Journal of Vibroengineering*, Vol. 17, Issue 1, 2015, pp. 346-356.
- [5]. Y. Gao, J. Zhou, W. He, A harmonic suppression method for low-frequency electromagnetic vibrators, *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, Vol. 238, Issue 7, 2024, pp. 2616-2627.
- [6]. X. F. Zhang, W. He, C. Y. Wang, Self-sensing waveform control for a low-frequency electromagnetic vibrator, *IEEE/ASME Transactions on Mechatronics*, Vol. 22, Issue 2, 2017, pp. 785-793.
- [7]. P. Hippold, G. Kleyman, L. Woiwode, et al., An iteration-free approach to excitation harmonization, *Mechanical Systems and Signal Processing*, Vol. 233, 2025, 112732.

Hydrostat – Real-Time Monitoring of Water Quality Parameters

M. Kubaščík, A. Tupý, M. Tebeľák and P. Ševčík

University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1, 010 26 Žilina, Slovakia
Tel.: + 421415134368
E-mail: michal.kubascik@fri.uniza.sk

Summary: Hydrostat is an integrated system for real-time monitoring of key water-quality and environmental parameters in aquatic systems. The platform comprises a distributed sensor node based on the ESP32-C6 microcontroller, equipped to measure water temperature, dissolved oxygen, water level, and ambient air conditions. Sensor data are transmitted wirelessly to a central unit using the ESP-NOW communication protocol. The central unit integrates an ESP32-C6 with an nRF9161 Development Kit, enabling reliable data aggregation and long-range communication. Collected measurements are transferred from the ESP32-C6 to the nRF9161 via a UART interface and subsequently transmitted over LTE-M networks to a remote server for storage in a database. The acquired data are further processed and visualized through a web-based interface, providing intuitive access to time-series graphs and system status. The proposed architecture ensures efficient data acquisition, low-power communication, and scalable deployment.

Keywords: Water quality monitoring, ESP-NOW, Environmental sensing, Internet of things, LTE-M connectivity, Dissolved oxygen monitoring.

1. Introduction

The monitoring of water quality parameters is a fundamental requirement for ensuring the stability, safety, and sustainability of aquatic environments. Water quality directly affects ecological balance, industrial processes, and human health, making its continuous assessment a critical task in modern environmental management. Key physicochemical parameters, including temperature, dissolved oxygen (DO), and water level, play a decisive role in regulating biological activity and chemical reactions within water systems.

In particular, dissolved oxygen is widely recognized as a primary indicator of water quality and ecosystem health [1], as it directly influences the survival, growth, and metabolic activity of aquatic organisms. Temperature is equally important, as it governs oxygen solubility, reaction kinetics, and biological processes, thereby strongly interacting with DO dynamics [2].

Recent studies emphasize that inadequate monitoring of these parameters can result in severe environmental and economic consequences, particularly in applications such as aquaculture and water resource management [3, 4]. Low oxygen levels, for instance, can lead to hypoxic conditions, significantly reducing organism health and productivity. Furthermore, increasing industrialization, urbanization, and climate variability have intensified water pollution and resource degradation, highlighting the need for more efficient and continuous monitoring solutions.

Traditional water quality assessment methods are typically based on manual sampling and laboratory analysis, which are time-consuming, costly, and incapable of providing real-time data. These

limitations hinder rapid response to dynamic environmental changes [5, 6]. Consequently, recent research has focused on leveraging Internet of Things (IoT) technologies and embedded systems to enable automated, real-time monitoring and remote data access. IoT-based systems offer significant advantages, including continuous data acquisition, reduced human intervention, and improved accuracy, making them suitable for large-scale and distributed deployments.

Recent reviews have highlighted the rapid growth of IoT-based water quality monitoring systems employing distributed sensors, wireless communications, cloud databases, and web-based visualization platforms [5, 6, 10, 13]. Commonly monitored variables include water temperature, dissolved oxygen, pH, conductivity, turbidity, and water level. ESP32-based platforms have become particularly popular due to their low cost, low power consumption, and integrated wireless communication capabilities [3, 4, 9, 10]. Recent ESP32-based implementations have demonstrated the feasibility of real-time monitoring of water quality parameters using cloud-connected IoT architectures, further highlighting the suitability of this platform for environmental sensing applications. While these systems demonstrate the effectiveness of IoT technologies for environmental monitoring, the Hydrostat platform focuses on a communication architecture that combines ESP-NOW-based local sensor networking with LTE-M cloud connectivity through an ESP32-C6/nRF9161 gateway, enabling deployment in locations where conventional network infrastructure is unavailable.

In response to these challenges, the Hydrostat system is proposed as an integrated solution for real-time monitoring of water-quality and

environmental parameters. While numerous IoT-based water-quality monitoring platforms have been reported, the main contribution of Hydrostat is its distributed communication architecture. The system combines ESP32-C6 sensor nodes communicating through the ESP-NOW protocol with a central ESP32-C6/nRF9161 gateway providing LTE-M connectivity for cloud communication. This approach enables low-overhead local wireless communication together with reliable long-range data transmission in locations where Wi-Fi infrastructure may not be available.

2. Water Parameters and Use Cases

Water quality monitoring systems are typically designed to measure a set of essential physicochemical parameters that provide insight into the state of aquatic environments. Among these, temperature, dissolved oxygen, and water level are considered key variables due to their strong influence on biological activity and system stability [4]. These parameters are commonly used to assess environmental conditions and detect changes that may indicate contamination or system imbalance.

Although numerous water-quality indicators can be monitored, the selected parameters provide complementary information concerning aquatic conditions. Water temperature influences biological activity and oxygen solubility, dissolved oxygen reflects ecosystem health and water quality, while water level provides insight into hydrological variations and operational conditions. Simultaneous monitoring of these variables enables a more comprehensive assessment of aquatic environments than observation of a single parameter alone.

Dissolved oxygen plays a crucial role in supporting aerobic life and maintaining ecological equilibrium, while temperature governs reaction rates and oxygen solubility in water systems [1, 2]. The importance of accurate dissolved oxygen sensing and calibration procedures has also been highlighted in previous studies of environmental monitoring systems [11]. Water level monitoring is also important for controlling system operation, detecting leakage or overflow conditions, and ensuring safe and efficient process management.

The monitoring of these parameters is essential in several application domains. In aquaculture, maintaining stable oxygen and temperature levels is critical for ensuring optimal growth conditions and preventing stress or mortality of aquatic organisms. In environmental monitoring, continuous observation of water parameters enables early detection of pollution events and long-term assessment of ecosystem health.

In industrial and laboratory settings, water quality monitoring supports process control, regulatory compliance, and system optimization. Modern solutions increasingly rely on automated measurement systems capable of real-time data acquisition and

transmission, allowing operators to respond promptly to changes in environmental conditions [5, 6].

The Hydrostat system is designed with these application scenarios in mind, providing a flexible platform capable of monitoring key parameters and supporting decision-making processes across diverse water-related use cases.

3. Proposed System

The Hydrostat system was designed as a distributed monitoring platform whose key feature is the integration of ESP-NOW-based local sensor networking with LTE-M cloud connectivity. The architecture consists of ESP32-C6 sensor nodes for data acquisition, an ESP32-C6/nRF9161 gateway for aggregation and long-range communication, and a cloud-based backend for storage and visualization.

3.1. Sensor Node Design

The sensing subsystem (Fig. 1) is implemented as a compact sensor node built around the ESP32-C6 microcontroller. The node acquires water temperature using a DS18B20 digital temperature sensor, dissolved oxygen concentration using a DFRobot Gravity Analog Dissolved Oxygen Sensor (SEN0237), water level using a Bosch BMP585 pressure sensor, and ambient temperature and humidity using a DHT22 sensor. These sensing elements are connected to the microcontroller through 1-Wire, analog, I²C, and GPIO interfaces according to their respective communication requirements. The selected sensors provide a cost-effective solution for continuous environmental monitoring while maintaining sufficient measurement accuracy for aquatic applications. Prior to deployment, the dissolved oxygen sensor was calibrated according to the manufacturer's recommended procedure. Sensors intended for direct contact with water were installed using waterproof housings suitable for long-term operation in aquatic environments.




Fig. 1. Sensor node architecture.

The DS18B20 sensor was selected due to its digital interface, robustness, and suitability for submerged operation. Water-level estimation is performed using the BMP585 pressure sensor, which enables indirect measurement of water depth through hydrostatic pressure evaluation. The DFRobot SEN0237 sensor provides continuous dissolved oxygen monitoring and supports calibration procedures required for aquatic applications. Together, these sensing elements allow simultaneous observation of the most important parameters affecting water quality and ecosystem stability. The selected sensor configuration provides a balance between measurement capability, deployment simplicity, and overall system cost, making it suitable for both laboratory and field applications.

The ESP32-C6 platform was selected due to its low power consumption, integrated wireless capabilities, and suitability for IoT applications. Sensor data are periodically sampled and preprocessed locally to reduce noise and ensure measurement consistency before transmission. The modular design also enables straightforward integration of additional sensors for future system expansion.

3.2. Central Unit Architecture

The central unit (Fig. 2) serves as a bridge between local sensor nodes and the remote server infrastructure. It consists of two main components: an ESP32-C6 microcontroller and an nRF9161 Development Kit. The ESP32-C6 receives measurement data from sensor nodes via ESP-NOW and performs initial data aggregation.




Fig. 2. Central unit architecture.

In addition to data aggregation, the ESP32-C6 is responsible for validating incoming packets and organizing the acquired measurements into a structured format suitable for transmission to the LTE-M communication module. This separation of local communication and wide-area communication tasks simplifies the system architecture and improves modularity. The gateway design also allows multiple sensor nodes to be integrated into a single monitoring deployment.

The use of a dedicated gateway also improves system scalability by decoupling local sensing functions from wide-area communication tasks. As the number of deployed sensor nodes increases, the gateway can continue to aggregate measurements and forward them to the cloud infrastructure without requiring modifications to the individual sensing units.

This architecture supports gradual system expansion and facilitates deployment in distributed monitoring scenarios.

Subsequently, the data are transmitted from the ESP32-C6 to the nRF9161 module using a UART interface. The nRF9161, equipped with LTE-M connectivity, is responsible for secure wide-area communication. This architecture supports communication between geographically distributed sensing nodes and cloud services even in locations where conventional network infrastructure is unavailable.

3.3. Wireless Communication

Data acquired by the sensor node are transmitted to the central unit using the ESP-NOW protocol. ESP-NOW is a lightweight, connectionless wireless communication protocol that enables low-latency data exchange between ESP-based devices without the need for traditional Wi-Fi network infrastructure. This approach reduces communication overhead and power consumption, making it suitable for distributed sensing scenarios.

The connectionless nature of ESP-NOW eliminates the need for conventional Wi-Fi access points and simplifies deployment in outdoor environments. Furthermore, the protocol enables direct communication between sensor nodes and the gateway, reducing network complexity and supporting low-latency transfer of measurement data. These characteristics make ESP-NOW particularly attractive for environmental monitoring applications where energy efficiency and deployment flexibility are important considerations.

Compared to conventional Wi-Fi communication, ESP-NOW eliminates the need for connection establishment and network management procedures, reducing protocol overhead and simplifying operation in isolated environments. The lightweight communication model contributes to lower energy consumption and facilitates efficient exchange of measurement data between distributed sensing nodes and the central gateway.

Each sensor node periodically samples the monitored parameters and transmits structured data packets containing a node identifier, timestamp, water temperature, dissolved oxygen concentration, water level, ambient temperature, and humidity measurements to the central unit. The sampling interval is set to 30 s, enabling regular monitoring while minimizing communication overhead.

3.4. Cloud Communication and Data Storage

The nRF9161 module forwards the collected data to a remote server over LTE-M networks. This cellular communication technology ensures long-range connectivity with low power consumption, making it suitable for IoT-based environmental monitoring

applications. LTE-M has become an attractive solution for environmental monitoring applications because it combines wide-area coverage with energy-efficient operation suitable for distributed IoT deployments [12].

On the server side, incoming data are received using the HTTP protocol and stored in a PostgreSQL database for further processing. Time-series data organization allows efficient storage and retrieval of historical measurements, enabling long-term analysis and system evaluation.

The centralized storage architecture enables continuous collection of historical records that can subsequently be used for trend analysis, anomaly detection, and environmental assessment. Furthermore, the PostgreSQL database structure facilitates future integration of analytical tools and predictive models without requiring modifications to the sensing or communication layers. The modular backend architecture also enables future integration with alerting systems and external monitoring services. This layered design contributes to the overall scalability of the Hydrostat platform.




Fig. 3. Proposed system architecture.

3.5. Data Visualization

The stored data are processed and presented through a web-based interface. The visualization layer provides graphical representations of measured parameters, including time-series plots and system status indicators. This enables users to remotely monitor water conditions, detect anomalies, and analyze trends in real time.

The separation of sensing, communication, and visualization layers ensures system flexibility and scalability. The proposed methodology supports deployment in diverse scenarios, including aquaculture systems, laboratory environments, and field monitoring applications.

The web interface provides a unified view of measurements collected from all deployed sensor nodes. Historical data can be examined over different time intervals, enabling users to identify seasonal variations and gradual environmental changes. Such visualization capabilities support informed decision-making in aquaculture, environmental monitoring, and laboratory applications.

The visualization layer also enables comparison of multiple parameters within the same time interval, facilitating identification of relationships between environmental conditions and water-quality indicators. Such functionality is particularly useful for long-term observation and operational decision support.

The dashboard can be accessed remotely through a standard web browser, enabling supervision of monitored locations without requiring direct physical access. This capability is particularly valuable for geographically distributed deployments and long-term environmental observation applications.

4. Conclusions

This paper presented Hydrostat, a distributed IoT platform for monitoring water-quality and environmental parameters. The primary contribution of the proposed system is the integration of ESP32-C6 sensor nodes communicating via ESP-NOW with an ESP32-C6/nRF9161 gateway providing LTE-M connectivity for cloud-based data transmission. This architecture enables efficient local data collection and supports deployment in remote environments where conventional network infrastructure may be unavailable.

The developed platform integrates environmental sensing, low-overhead wireless communication, cellular connectivity, cloud-based storage, and web-based data visualization into a unified monitoring framework. By combining ESP-NOW local networking with LTE-M wide-area communication, Hydrostat differs from many existing IoT monitoring systems that rely on conventional Wi-Fi infrastructure, while maintaining flexibility and scalability for a wide range of deployment scenarios.

Overall, Hydrostat represents a flexible and cost-effective solution for continuous monitoring of key water-quality indicators, including dissolved oxygen, temperature, and water level. Future work will focus on extending the platform with additional sensing capabilities, advanced data analytics, and actuator-based control mechanisms for autonomous regulation of aquatic environments. Further research will also include a detailed evaluation of communication reliability, LTE-M performance, and power consumption under long-term deployment conditions.

References

- [1]. Hach Company, Dissolved oxygen measurement and its role in water quality monitoring, 2026, <https://www.hach.com>
- [2]. S. Larance, J. Wang, M. A. Delavar, M. Fahs, Assessing water temperature and dissolved oxygen and their potential effects on aquatic ecosystems, *Environments*, Vol. 12, Issue 1, 2025, 25.
- [3]. V. Sirisha, et al., IoT-based real-time water quality monitoring and predictive alert system using ESP32,

- American Journal of AI Cyber Computing Management*, Vol. 6, Issue 1, 2026, pp. 340-346.
- [4]. H. Gowda, M. A. K. N. Jailani, IoT-based water quality monitoring system, *International Journal of Research Publication and Reviews*, Vol. 4, Issue 6, 2023, pp. 4466-4470.
- [5]. M. Miller, A. Kisiel, D. Cembrowska-Lech, T. Miller, IoT in water quality monitoring – Are we really here?, *Sensors*, Vol. 23, Issue 2, 2023, 960.
- [6]. E. C. Vicente, L. A. Silva, A. M. R. Fernandes, W. D. Parreira, A structured review of IoT-based embedded systems and machine learning for water quality monitoring, *Applied Sciences*, Vol. 15, Issue 21, 2025, 11719.
- [7]. YSI (Xylem Inc.), Water quality monitoring instruments and sensors, 2026, <https://www.ysi.com>
- [8]. Atlas Scientific, Water quality sensors for industrial and environmental monitoring, 2026, <https://atlas-scientific.com>
- [9]. D. Patel, K. Koradia, P. Sharma, IoT-based water quality monitoring system using ESP32, in *Proceedings of the 7th International Conference on Energy, Power and Environment (ICEPE)*, 2025, pp. 1-5.
- [10]. Y. Singh, T. Walingo, Smart water quality monitoring with IoT wireless sensor networks, *Sensors*, Vol. 24, Issue 9, 2024, 2871.
- [11]. Y. Wei, Y. Jiao, D. An, D. Li, et al., Review of dissolved oxygen detection technology: From laboratory analysis to online intelligent detection, *Sensors*, Vol. 19, Issue 18, 2019, 3995.
- [12]. E.-C. Trînc, V. Niță, C. Stolojescu-Crișan, C. Ancuți, et al., Cost-effective TinyML-ready design and field deployment of a solar-powered environmental monitoring data collector using LTE-M communication, *Applied Sciences*, Vol. 16, Issue 7, 2026, 3237.
- [13]. G. A. López-Ramírez, A. Aragón-Zavala, Wireless sensor networks for water quality monitoring: A comprehensive review, *IEEE Access*, Vol. 11, 2023, pp. 95120-95142.

Measurement Error Analysis of Eddy-Current Sensors under Tilted Probe Installation

X. F. Zhang ^{1,2}, T. Xu ¹ and L. K. Wu ¹

¹ College of Mechanical Engineering, Taiyuan University of Technology, No. 79 Yingze Street,
Taiyuan 030024, Shanxi, China

² MOE Key Laboratory of Advanced Transducers and Intelligent Control System,
Taiyuan University of Technology, Taiyuan 030024, Shanxi, China
Tel.: +86 18535127169
E-mail: zhangxufei@tyut.edu.cn

Summary: Eddy-current sensors are attractive non-contact references for vibration displacement measurement in accelerometer calibration, but probe tilt caused by installation error changes electromagnetic coupling and biases the measured vibration amplitude. Unlike prior tilt studies in general displacement or nondestructive-testing settings, this work considers the dynamic lift-off encountered in single-axis vibration calibration and develops an analytical model using an equivalent rectangular-section eddy-current ring, Taylor segmentation, and the Neumann integral. A three-dimensional COMSOL model validates the predicted equivalent-inductance variation, which is further propagated through the sensor measurement-circuit model to quantify output-voltage error. Tilt breaks the axial symmetry of the eddy-current field and lowers equivalent inductance, with a stronger influence at smaller lift-off. For a 5 mm vibration amplitude, the voltage-amplitude errors are 0.554 %, 2.368 %, and 5.349 % at 5°, 10°, and 15°, respectively, while the maximum relative error between the analytical and finite element method (FEM) inductance results is 0.832 %.

Keywords: Eddy current sensor, Probe tilt, Vibration displacement, Equivalent inductance, Finite element method.

1. Introduction

Eddy-current displacement sensors are widely used in precision displacement and vibration measurement because of their non-contact operation, high resolution, and fast response [1, 2]. In accelerometer calibration systems, they can be employed to measure the vibration displacement of the moving target. In single-axis vibration calibration, installation errors may cause the probe to tilt relative to the target surface, altering the electromagnetic coupling between the probe coil and the target and thereby affecting the coil impedance and displacement-related output signal. This effect is further complicated by the dynamically varying lift-off during vibration.

Probe tilt has been investigated in previous studies. Le Bihan analyzed lift-off and tilt effects using three-dimensional finite elements [3]. Li et al. investigated the relationships among tilt angle, lift-off, and equivalent inductance in eddy-current displacement measurement and reported that increasing the tilt angle reduces the measured response [4]. Lu et al. studied tilt-induced impedance and voltage variations in a nondestructive-testing application [5]. These studies established the general influence of probe tilt on eddy-current sensor response, but mainly considered static displacement measurement or nondestructive-testing conditions.

The present study focuses on the accelerometer-calibration scenario with dynamically varying lift-off. An equivalent eddy-current ring combined with a Taylor-segmented Neumann-integral mutual-inductance model is used to analyze tilted-probe errors. The analytical results are compared with a

three-dimensional COMSOL model, and the predicted equivalent-inductance variation is further converted into output-voltage amplitude error. Thus, the contribution of this work lies not in identifying the general tilt effect itself, but in applying the proposed analytical formulation to quantify tilt-induced measurement error under dynamic vibration-calibration conditions.

2. Methodology

In practical installation, a probe tilt (Fig. 1) changes the electromagnetic coupling between the probe coil and the disc. With probe parameters R_p and L_p and equivalent eddy-current-ring parameters R_e and L_e , Kirchhoff's law gives the sensor equivalent impedance components:

$$\begin{cases} R_{eq} = R_p + \frac{\omega^2 M^2 R_e}{R_e^2 + (\omega L_e)^2} \\ L_{eq} = L_p - \frac{\omega^2 M^2 L_e}{R_e^2 + (\omega L_e)^2} \end{cases}, \quad (1)$$

where L_p and R_p are the self-inductance and resistance of the probe coil without eddy-current effects, respectively; L_e and R_e are the equivalent self-inductance and resistance of the eddy-current ring in the target disc, respectively; M is the mutual inductance between the probe coil and the equivalent eddy-current ring; L_{eq} and R_{eq} are the equivalent inductance and resistance of the probe considering the eddy-current effect; and ω is the angular frequency of the alternating current applied to the probe coil.




Fig. 1. Vibration displacement measurement model.

The tilted probe and the measured disc are modeled as two magnetically coupled systems. The eddy current induced in the disc is approximated as an equivalent ring with a rectangular cross-section, as shown in Fig. 2(a). The mutual-inductance model between the tilted probe coil and the equivalent eddy-current ring is illustrated in Fig. 2(b). To facilitate the analytical calculation, the two coils are subdivided using the Taylor segmentation scheme shown in Fig. 2(c). Based on the tilted geometric relationship in Fig. 2(d), the mutual inductance between the corresponding segments is calculated using the Neumann integral, and the total mutual inductance is obtained by summation.




Fig. 2. Analytical model for probe-tilt-induced equivalent inductance calculation.

Based on the skin-effect penetration depth and radial eddy-current distribution, the inner radius r_1 , outer radius r_2 , and central radius r_3 of the equivalent

eddy-current ring are determined. The corresponding equivalent-ring parameters are then calculated as follows:

$$\begin{cases} L_e = \mu_0 r_3 \left[\frac{\ln(8 \times r_3)}{g} - 2 \right] \\ R_e = \frac{2\pi}{\delta \sigma \ln(r_2/r_1)} \end{cases}, \quad (2)$$

where δ is the standard penetration depth, σ is the conductivity of the measured disc, and g is the average geometrical distance of the ring cross-section.

The mutual inductance is obtained by subdividing the probe coil and the equivalent eddy-current ring into ten segments. According to the Taylor segmentation scheme, the final mutual inductance is

$$M = \frac{(M_{01'} + M_{02'} + M_{03'} + M_{04'} + M_{01} + M_{02} + M_{03} + M_{04} - 2M_{00'})}{6} \quad (3)$$

Taking $M_{00'}$ as an example, the Neumann formula can be expressed as:

$$M_{00'} = \frac{N\mu_0}{4\pi} \int_0^{2\pi} d\theta \int_0^{2\pi} \frac{r_4 r_3 \cos(\beta - \theta)}{l} d\beta, \quad (4)$$

where N is the coil turn number, β and θ are angular coordinates of the tilted probe coil and the equivalent ring, respectively; r_4 depends on the probe-tilt projection, and l is the spatial distance between the two integration points.

3. Finite-Element Simulation

A three-dimensional COMSOL model was established using a copper probe coil, an aluminum disc, and a 200 mm × 200 mm × 100 mm air domain. The inner diameter, outer diameter, and thickness of the coil were 14.8 mm, 17.8 mm, and 1.5 mm, respectively. The disc radius and thickness were 60 mm and 3 mm, respectively. The probe excitation frequency was 500 kHz, while the target disc underwent sinusoidal vibration at 1 Hz. The initial lift-off was 11.8 mm, with vibration amplitudes of 3, 4, and 5 mm and probe tilt angles of 0°, 5°, 10°, and 15°. Therefore, the instantaneous lift-off varied periodically during each vibration cycle.

As shown in Fig. 3, L_{eq} decreases as the tilt angle increases. The difference is small at large instantaneous lift-off but grows near the displacement valley, consistent with stronger electromagnetic coupling at smaller lift-off. At $t = 0.75$ s for a 5 mm vibration amplitude, the comparison between the analytical and finite-element method (FEM) results remains within 0.832 % (Table 1). The FEM-derived time-varying equivalent inductance $L_{eq}(t)$ was introduced into the established sensor measurement-circuit model to obtain the corresponding output-voltage waveform, whose fundamental amplitude was extracted by FFT. For a 5 mm vibration amplitude, the resulting voltage-amplitude errors relative to the zero-tilt condition were 0.554 %, 2.368 %, and 5.349 % at tilt angles of 5°, 10°, and 15°, respectively.




Fig. 3. Equivalent inductance under a 5 mm vibration amplitude for different probe tilt angles.

Table 1. Analytical and finite-element equivalent inductance at $t = 0.75$ s for a 5 mm vibration amplitude.

Tilt	Analytical L_{eq} (uH)	Finite Element Method L_{eq} (uH)	Error (%)
0°	42.130	42.072	0.138
5°	42.061	41.983	0.186
10°	41.857	41.697	0.384
15°	41.484	41.139	0.832

4. Conclusions

Probe tilt in single-axis vibration calibration was analyzed using an equivalent-ring/Taylor-Neumann analytical model and three-dimensional FEM under dynamic lift-off. The model reproduces the FEM

equivalent-inductance response with a maximum selected-case error of 0.832 %. More importantly, the inductance variation produces a measurable vibration-output bias: at 5 mm amplitude, the voltage-amplitude error rises to 5.349 % at 15°. The results therefore quantify both the electromagnetic mechanism and its measurement consequence, providing a model-based basis for subsequent compensation.

References

- [1]. X. Xue, Y. Dong, X. Wu, Motion induced eddy current sensor for non-intrusive vibration measurement, *IEEE Sensors Journal*, Vol. 20, Issue 2, 2020, pp. 735-744.
- [2]. W. Liu, B. Liang, Z. Jia, D. Feng, et al., High-accuracy calibration based on linearity adjustment for eddy current displacement sensor, *Sensors*, Vol. 18, Issue 9, 2018, 2842.
- [3]. Y. Le Bihan, Lift-off and tilt effects on eddy current sensor measurements: A 3-D finite element study, *The European Physical Journal Applied Physics*, Vol. 17, Issue 1, 2002, pp. 25-28.
- [4]. Y. Li, X. Sheng, M. Lian, Y. Wang, Influence of tilt angle on eddy current displacement measurement, *Nondestructive Testing and Evaluation*, Vol. 31, Issue 4, 2016, pp. 289-302.
- [5]. M. Lu, X. Meng, R. Huang, A. Peyton, et al., Analysis of tilt effect on notch depth profiling using thin-skin regime of driver-pickup eddy-current sensor, *Sensors*, Vol. 21, Issue 16, 2021, 5536.

Ultra-Low-Power Odor Classification on the ESP32-C6 Using Integer-Based Models

M. Kubaščík, P. Kolok, L. Čechovič and M. Chochul

University of Žilina, Department of Technical Cybernetics, Univerzitná 8215/1, 010 26 Žilina, Slovakia

Tel.: + 421415134368

E-mail: michal.kubascik@fri.uniza.sk

Summary: This concept paper defines a reproducible implementation and evaluation protocol for event-triggered odor classification using the ESP32-C6 low-power (LP) core and the Bosch BME690 gas sensor. The work does not claim novelty for TinyML, quantization, or Bosch gas-scan training. Its proposed contribution is a system-level comparison of three transparent integer classifiers executed in the LP domain while the high-performance CPU remains asleep: prototype-based classification, a reduced k-nearest-neighbor classifier, and a single-layer linear neural classifier. The manuscript specifies an initial four-class benchmark, chamber and exposure timing, session-based validation, drift controls, quantization equations, analytical memory and operation budgets, event-confidence logic, and phase-resolved energy measurements. The custom workflow is explicitly distinguished from Bosch BME AI-Studio and BSEC, which provide an established sensor-configuration, training, and on-device gas-scan workflow. The term ultra-low-power is treated as an experimental hypothesis rather than an assumed property; the final implementation must report absolute energy per cycle and demonstrate a declared reduction against a main-CPU baseline before that designation is used.

Keywords: ESP32-C6, BME690, Electronic nose, Odor classification, Integer inference, LP core, Event-triggered sensing, TinyML, BME AI-Studio.

1. Introduction and Positioning

Electronic-nose systems combine gas-sensitive measurements with signal processing and pattern recognition to identify odor-related states. Portable implementations have already demonstrated machine-learning classification using metal-oxide sensors, while broad reviews describe established preprocessing, feature-extraction, and validation methods [1-3, 8, 9]. Likewise, quantized inference and TinyML deployment on constrained microcontrollers are mature research topics rather than new ideas [4-7]. The novelty of the present concept therefore cannot be based merely on using a gas sensor, a microcontroller, or integer arithmetic.

The proposed contribution is narrower and system-oriented. First, a custom classifier is intended to execute on the ESP32-C6 LP core while the high-performance CPU remains in a low-power state. Second, a confidence or rejection mechanism converts classification into an event trigger: only a target class, anomaly, persistent trend, or uncertain decision wakes the main CPU. Third, prototype-based, reduced-neighbor, and linear-neural classifiers would be compared under identical sensor data, features, integer precision, firmware, and sampling conditions. Fourth, the energy analysis would include sensor acquisition, heater activity, LP processing, main-CPU wake-up, and communication instead of reporting classifier execution in isolation.

Espressif documents that the ESP32-C6 LP core can remain active while the main CPU is in a low-power mode, can be periodically activated by the

LP timer, can access LP SRAM and low-power peripherals, and can wake the main CPU when a condition is met [12, 13]. These capabilities support the proposed partitioning, although direct BME690 acquisition through the LP I2C path, driver size, timing, and retained-memory use must still be verified experimentally. Consequently, this manuscript is presented as a concept and evaluation protocol, not as a completed performance study.

2. Relation to Bosch BME AI-Studio and Existing TinyML Work

Bosch already provides a complete BME688/BME690 software ecosystem. BME AI-Studio supports sensor configuration, data recording, labeling, algorithm training, optimization, and export, while BSEC runs on the host microcontroller and interprets sensor data using configuration strings generated by AI-Studio [15, 16]. The BME690 is also designed to retain the gas-sensing functionality of the BME688 while improving operation in high-condensation environments [14]. The proposed work does not replace or rebrand this workflow.

The distinction is that the planned classifier is deliberately small, transparent, and compiled as integer parameters for the ESP32-C6 LP core. BME AI-Studio/BSEC would serve as an established reference workflow on the main CPU where technically and legally applicable. The custom models would instead expose their features, distances,

weights, thresholds, memory use, and wake-up decisions for direct measurement. The scientific question is whether this restricted LP-core screening stage reduces complete-node energy without unacceptable false alarms, missed events, or uncertain outputs.

3. Proposed Event-Triggered Architecture

The architecture combines offline model preparation with two on-device execution levels. Labeled recordings are first transformed into compact features and used to train the three candidate classifiers. Scaling constants, prototypes, reduced reference samples, weights, biases, and decision thresholds are then converted into integer or fixed-point form and stored as constant arrays. During operation, the main CPU configures the sensor and LP program, loads parameters, and enters a low-power state. The LP core periodically updates a short feature

window, executes the selected classifier, and either returns to sleep or requests a main-CPU wake-up.

Table 1. Positioning of the proposed workflow relative to established solutions.

Aspect	Proposed distinction
TinyML / quantization	Established prior art; used as an implementation method, not claimed as novelty.
BME AI-Studio / BSEC	Reference ecosystem for gas-scan configuration, training, export, and host-MCU inference.
This concept	Custom integer screening on ESP32-C6 LP core with event-triggered main-CPU wake-up.
Primary comparison	Recognition quality, false wake-ups, memory, latency, and complete-cycle energy on the same hardware.




Fig. 1. Simplified workflow of LP-core odor classification with event-triggered main-CPU wake-up.

The initial feature vector is limited to eight values: gas resistance, relative gas-resistance change, short-window mean, minimum, maximum, variance, slope, and a normalized deviation from the recent baseline. Temperature and humidity are retained as compensation variables during data analysis; if they materially improve grouped validation, they can replace weaker gas-derived features without increasing the eight-feature budget. Features requiring long histories, floating-point filters, or large buffers are excluded unless their recognition benefit justifies their energy and memory cost.

A confidence rule accompanies every class decision. The prototype model uses the minimum distance and the gap to the second-best prototype; reduced k-nearest-neighbor uses neighbor agreement; and the linear classifier uses the difference between its two highest scores. A low-confidence sample is labeled unknown and wakes the main CPU rather than being silently classified as normal. Hysteresis and two consecutive detections are planned to suppress isolated noisy triggers. The main CPU may then record a longer sequence, execute BSEC or a more capable verification model, store raw data, activate Wi-Fi, or notify an external system.

4. Planned Dataset and Validation Protocol

The planned experimental dataset must be designed to support a reproducible and unbiased comparison of the proposed classifiers. Therefore, the data-collection protocol should define not only the considered odor classes, but also the sample preparation procedure, measurement chamber, exposure and recovery intervals, environmental conditions, number of recording sessions, and validation strategy. Particular attention must be paid to avoiding data leakage between temporally related samples and to evaluating the influence of temperature, humidity, sensor drift, and sensor-to-sensor variability. The following subsections describe the proposed dataset structure and the procedures that would be used for model training, validation, and independent testing.

4.1. Odor Classes, Preparation, and Chamber

The initial benchmark is defined before data collection to avoid an unspecified dataset. It contains four trained states: clean background air, ground coffee, ethanol vapor, and vinegar vapor. The benchmark is intended to test separability and embedded execution, not to claim universal chemical identification. An additional rejection set contains at least two untrained household odors and is used only to evaluate the unknown-class mechanism.

Measurements are planned in a 1 L glass chamber with a small mixing fan, fixed sensor position, specimen inlet, and filtered-air purge. The source is placed in an open 50 mL glass vial approximately 5 cm from the inlet. Each coffee exposure uses 10 g of

freshly ground coffee; ethanol and vinegar exposures use 5 mL of 70 % ethanol solution and 5 mL of approximately 5 % acetic-acid vinegar, respectively; the clean-air class uses an empty vial. Source order is randomized, the chamber is cleaned between sessions, and source mass or volume, room conditions, sensor serial number, and timestamps are recorded.

Before the first campaign, each sensor is powered for at least 24 h, consistent with Bosch guidance for initial stabilization [16]. Every session contains 5 min of clean-air baseline, 10 min of exposure, and at least 15 min of recovery. Recovery continues until gas resistance returns within 5 % of the pre-exposure median or until a documented maximum of 30 min is reached. The selected heater profile, scan duration, and output data rate are fixed for the complete campaign; one labeled observation corresponds to one completed gas-scan cycle, while classifier inputs are generated from non-overlapping short windows.

4.2. Number of Sessions and Data Split

The minimum dataset contains 80 independent labeled sessions, 20 per trained class, acquired across at least five days and two BME690 units. A further minimum of 10 sessions is reserved for untrained-odor rejection tests. Temperature and relative humidity are intentionally varied within approximately 18-30 °C and 30-70 % RH and are recorded with every scan. Session order is randomized so that class and environmental condition are not confounded. A supplementary leave-one-device-out analysis assesses whether calibration transfers between sensor units.

Sixteen labeled sessions (20 %) are locked as the final test set before feature or threshold selection. The remaining 64 sessions form the development set and are evaluated using five-fold grouped cross-validation, with all windows from one exposure session assigned to the same fold. No neighboring windows from the same recording may appear in both training and validation data. After model selection, the model is retrained on all development sessions and evaluated once on the held-out test sessions. Reported metrics include accuracy, macro F1-score, per-class precision and recall, confusion matrix, false-alarm rate, missed-event rate, unknown rate, and event-detection delay. Complete temporal sequences are analyzed to expose drift and repeated triggering [10, 11].

5. Integer Classifiers and Analytical Resource Budget

Each feature is mapped to a signed integer using a training-derived scale factor and saturation limits:

$$x_{Q,i} = \text{round}(S_i \cdot x_i)$$

The floating-point and quantized implementations are evaluated on identical validation sequences. Quantization is accepted only when the loss in macro

F1-score is no greater than two percentage points and no accumulator overflow occurs. This requirement follows the established principle that integer deployment must be validated after conversion rather than inferred from floating-point performance [6].

Table 2. Predefined dataset protocol for the planned benchmark.

Item	Planned specification
Classes	Clean air, coffee, ethanol, vinegar; separate untrained-odor rejection set.
Sessions	80 labeled (20/class) + at least 10 rejection sessions; ≥ 5 days; ≥ 2 sensors.
Timing	5 min baseline, 10 min exposure, 15-30 min recovery.
Validation	20 % locked test; grouped five-fold CV on remaining sessions; no window leakage.
Drift control	Randomized order; 18-30 °C and 30-70 % RH; device/day grouping and leave-one-device-out test.

5.1. Prototype and Reduced k-Nearest-Neighbor Models

For class k and feature dimension N , the squared integer distance to prototype c_k is

$$d_k = \sum_{i=1}^N (x_{Q,i} - c_{k,i})^2, \\ \hat{y} = \arg \min_k d_k$$

The class is therefore selected by an explicit minimization over valid class indices. Multiple prototypes may represent one class, while a maximum-distance threshold produces an unknown output. The reduced k-nearest-neighbor classifier applies the same distance to a compact reference set and predicts the majority label among the K nearest samples. K is selected only on the development set, and ties or weak neighbor agreement are returned as unknown.

5.2. Single-Layer Linear Neural Classifier

The linear classifier calculates one integer score per class:

$$z_j = b_j + \sum_{i=1}^N w_{j,i} x_{Q,i}, \\ \hat{y} = \arg \max_j z_j$$

The corrected decision rule explicitly maximizes the class score, in contrast to the distance models that minimize distance. Accumulators use a wider integer type than inputs and weights, and all intermediate ranges are checked using the training and validation data.

Table 3 gives an analytical, not measured, lower-bound budget for the initial configuration with $N = 8$ features, $C = 4$ classes, two prototypes per class, 32 retained k-nearest-neighbor references, 16-bit stored values, and 32-bit accumulators. Firmware code, stack, sensor buffers, and alignment are excluded and will be reported from the compiled binary.

Table 3. Analytical model-data and operation budget for the initial configuration.

Classifier	Stored model data	Per-decision core work
Prototype	≈ 136 B	64 feature-distance terms
Reduced 3-NN	≈ 544 B	256 feature-distance terms + neighbor vote
Linear neural	≈ 80 B	32 multiply-accumulate operations

6. Embedded Measurements and Decision Criteria

All three classifiers are deployed on the same ESP32-C6 board and compiled with the same ESP-IDF version and optimization settings. Flash use is obtained from the linker map, retained LP RAM from the LP binary and symbols, peak runtime RAM from instrumented buffers, and inference time from repeated timestamped executions. At least 1000 inference cycles are measured per classifier after warm-up. Floating-point reference outputs, saturation events, and class disagreements are logged.

Energy is measured with a calibrated shunt or dedicated power analyzer at a sampling rate sufficient to separate sensor acquisition, heater activity, LP inference, main-CPU start-up, and radio transmission. Three baselines are planned: a main-CPU implementation that wakes and classifies every cycle, a threshold-only LP trigger, and, where integration permits, the Bosch AI-Studio/BSEC workflow on the main CPU. Event-driven tests replay identical sequences at declared target-event rates of 0 %, 1 %, 5 %, and 10 %. Results report energy per cycle, average current, average power, number and duration of main-CPU wake-ups, number of transmissions, and estimated battery life for a stated battery and sampling interval.

The term ultra-low-power is not treated as self-evident. It will be used for the final prototype only if the complete non-radio monitoring path is below 1 mW on average for the declared duty cycle and the event-driven system reduces average energy by at least 20 % relative to the main-CPU baseline without violating the recognition criteria. Regardless of whether these thresholds are met, absolute measured values and uncertainty intervals will be reported. Sensor heating may dominate the budget, in which case a very small classifier would not justify the label

– a useful result, despite being less glamorous than the usual marketing séance.

Table 4 Planned measured outputs and preregistered acceptance criteria.

Measured output	Acceptance criterion / report
Recognition	Test accuracy $\geq 85\%$, macro F1 ≥ 0.80 , per-class metrics and confusion matrix.
Quantization	Macro-F1 loss ≤ 2 percentage points; zero observed overflow.
LP resources	Model + feature state target ≤ 8 kB; measured flash/RAM reported.
Latency	Feature preparation + inference target ≤ 10 ms per completed scan.
Event behavior	False-wake and missed-event rates reported; unknown output enabled.
Energy	Absolute energy per phase; $\geq 20\%$ average-energy reduction for ultra-low-power claim.

7. Limitations and Future Work

Metal-oxide gas sensors exhibit cross-sensitivity, environmental dependence, unit-to-unit variation, warm-up behavior, and long-term drift [10, 11]. The planned four-class benchmark therefore demonstrates only recognition of defined experimental states. It does not establish concentration measurement, chemical selectivity, or universal odor identification. The explicit unknown output is mandatory, especially if the architecture is later used for warning applications.

The principal implementation risk is that BME690 heater and acquisition energy may dominate the entire cycle or that LP-core driver and memory restrictions may require a short main-CPU acquisition phase. Both outcomes must be reported rather than hidden behind classifier-only current measurements. Future development will proceed through LP-I2C feasibility, dataset collection, grouped model comparison, integer deployment, phase-resolved energy measurement, and multi-day drift testing. Later work may evaluate adaptive prototypes, calibration transfer, or a two-stage classifier in which the LP model only rejects obviously normal samples.

8. Conclusions

This paper defines a concrete concept and evaluation protocol for event-triggered odor classification using the ESP32-C6 LP core and BME690 sensor. It positions the work relative to established TinyML, integer inference, portable electronic noses, and Bosch BME AI-Studio/BSEC; specifies the planned odor classes, chamber procedure, sample preparation, timing, dataset size, grouped validation, and drift controls; corrects the

minimization and maximization equations; and provides numerical analytical budgets and preregistered measurement criteria without fabricating experimental results.

The expected contribution is not a claim that one tiny classifier is universally best. It is a reproducible comparison showing when LP-core screening and event-triggered wake-up reduce complete-node energy, and when sensor or communication costs make the approach ineffective. Experimental measurements remain necessary before any accuracy, battery-life, or ultra-low-power claim can be accepted.

References

- [1]. Z. Ye, Y. Liu, Q. Li, Recent progress in smart electronic nose technologies enabled with machine learning methods, *Sensors*, Vol. 21, Issue 22, 2021, 7620.
- [2]. J. Palacín, E. Rubies, E. Clotet, D. Martínez, Classification of two volatiles using an eNose composed by an array of 16 single-type miniature micro-machined metal-oxide gas sensors, *Sensors*, Vol. 22, Issue 3, 2022, 1120.
- [3]. Y. Huang, I.-J. Doh, E. Bae, Design and validation of a portable machine learning-based electronic nose, *Sensors*, Vol. 21, Issue 11, 2021, 3923.
- [4]. L. Capogrosso, F. Cunico, D. S. Cheng, F. Fummi, et al., A machine learning-oriented survey on tiny machine learning, *IEEE Access*, Vol. 12, 2024, pp. 23406-23426.
- [5]. S. S. Saha, S. S. Sandha, M. Srivastava, Machine learning for microcontroller-class hardware: A review, *IEEE Sensors Journal*, Vol. 22, Issue 22, 2022, pp. 21362-21390.
- [6]. B. Jacob, S. Kligys, B. Chen, M. Zhu, et al., Quantization and training of neural networks for efficient integer-arithmetic-only inference, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 2704-2713.
- [7]. J. Lin, W.-M. Chen, Y. Lin, J. Cohn, et al., MCUNet: Tiny deep learning on IoT devices, in *Advances in Neural Information Processing Systems 33* (H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, et al., Eds.), *Curran Associates, Inc.*, Red Hook, 2020, pp. 11711-11722.
- [8]. S. Marco, A. Gutiérrez-Gálvez, Signal and data processing for machine olfaction and chemical sensing: A review, *IEEE Sensors Journal*, Vol. 12, Issue 11, 2012, pp. 3189-3214.
- [9]. R. Gutiérrez-Osuna, Pattern analysis for machine olfaction: A review, *IEEE Sensors Journal*, Vol. 2, Issue 3, 2002, pp. 189-202.
- [10]. A. Vergara, S. Vembu, T. Ayhan, M. A. Ryan, et al., Chemical gas sensor drift compensation using classifier ensembles, *Sensors and Actuators B: Chemical*, Vol. 166-167, 2012, pp. 320-329.
- [11]. A.-C. Romain, J. Nicolas, Long term stability of metal oxide-based gas sensors for e-nose environmental applications: An overview, *Sensors and Actuators B: Chemical*, Vol. 146, Issue 2, 2010, pp. 502-506.
- [12]. ULP LP core coprocessor programming – ESP32-C6, ESP-IDF Programming Guide, ver. 6.0.2, Espressif

- Systems, 2026, <https://docs.espressif.com/projects/esp-idf/en/stable/esp32c6/api-guides/ulp-lp-core.html>
- [13]. Sleep modes - ESP32-C6, ESP-IDF Programming Guide, ver. 6.0.2, Espressif Systems, 2026, https://docs.espressif.com/projects/esp-idf/en/stable/esp32c6/api-reference/system/sleep_modes.html
- [14]. BME690 – 4-in-1 environmental sensor with AI, Product Flyer BST-BME690-FL000-02, ver. 1.2, *Bosch Sensortec GmbH*, 2024.
- [15]. BME688 and BME690 software: BME AI-Studio Desktop and BSEC software, Software Documentation, BME AI-Studio v3.2.0 and BSEC v3.3.0.0, *Bosch Sensortec GmbH*, 2026.
- [16]. BME AI-Studio getting started: Data recording, training, evaluation, and algorithm export, Online Documentation, *Bosch Sensortec GmbH*, 2026.

Design and Preliminary Validation of a Modular ESP32-H2 Sensor Node

Andrej Brnušák, Andrej Šimkovič, Peter Šarašin, Matúš Formanek and Tomáš Bača

University of Žilina, Faculty of Management Science and Informatics, Department of Technical Cybernetics,
Univerzitná 8215/1, 010 26 Žilina, Slovakia

E-mail: brnusak@stud.uniza.sk, simkovic4@stud.uniza.sk, peter.sarasin@fri.uniza.sk,
matus.formanek@fri.uniza.sk, tomas.baca@fri.uniza.sk

Summary: This paper presents the design and preliminary functional validation of a modular battery-powered sensor node based on the ESP32-H2 microcontroller. The custom printed circuit board provides I2C, SPI and UART connectors for compatible sensor modules and includes configurable series-resistor and shunt-capacitor footprints intended for digital signal-integrity tuning. The proposed operating concept uses periodic wake-up, sensor acquisition and subsequent transition to deep sleep, but this sleep/wake sequence was not functionally verified. Basic prototype tests confirmed I2C communication with a BME680 sensor, SPI communication with an SD-card interface, battery-powered operation, and BLE advertising and device discoverability. These tests confirm the basic operation of the developed hardware platform. Quantitative characterization of energy consumption, signal integrity and wireless-link performance is outside the scope of this work.

Keywords: Sensor node, ESP32-H2, Modular hardware, Digital interfaces, Duty-cycled operation.

1. Introduction

Wireless sensor nodes represent a fundamental component in modern distributed data collection systems. Since these nodes are frequently deployed in remote environments without access to a continuous power grid, minimizing energy consumption is the primary challenge in their hardware and software design. Therefore, the operating lifespan of the device directly depends on the efficiency of its hardware architecture and the utilization of low-power sleep modes [1].

A significant limitation of many contemporary sensor platforms is their rigid architecture. Many nodes are developed as single-purpose devices with fixed sensor configurations, meaning that any modification to the system requires a complete redesign of the printed circuit board (PCB). Digital interconnects between a microcontroller and external sensor modules can be affected by fast signal edges, interconnect parasitics and electromagnetic interference. Under unsuitable interface conditions, these effects may cause ringing, timing violations or communication errors [2].

To address these limitations, this paper presents a modular sensor node based on the ESP32-H2 microcontroller. The hardware architecture provides I2C, SPI and UART connectors for compatible sensor modules. Optional series-resistor and shunt-capacitor footprints are provided to allow signal-edge behavior to be adjusted during hardware configuration.

The contribution of this work is the design and preliminary functional validation of a compact, modular ESP32-H2-based PCB with configurable I2C, SPI and UART sensor interfaces; the validation scope and results are summarized in Section 4. Comparable wireless sensor-node platforms illustrate the relevance of modular hardware architectures for adaptable monitoring applications [3-5].

2. Hardware Design and Interfaces

The core of the developed modular sensor node is the ESP32-H2 microcontroller, selected for its integrated wireless connectivity and low-power operating modes. According to the manufacturer, the microcontroller supports Bluetooth Low Energy (BLE) and IEEE 802.15.4, enabling Zigbee and Thread implementations, with Matter support possible over Thread [6].

The hardware platform is designed as a reusable mainboard for sensor modules that are compatible with its I2C, SPI or UART connectors and electrical interface requirements. The high-level topology and connection layout are illustrated in Fig. 1. In this work, modularity refers to the ability to exchange interface-compatible sensor modules without redesigning the mainboard; it does not imply compatibility with every device implementing the corresponding communication protocol.




Fig. 1. Block diagram of the modular sensor node architecture.

To support modular sensor connections, the PCB breaks out three standard digital communication interfaces: I2C, SPI and hardware UART. This enables compatible digital sensor modules to be connected

without redesigning the underlying PCB. The final physical layout and component placement of the developed PCB are shown in the 3D model in Fig. 2.




Fig. 2. 3D model of the designed printed circuit board (PCB).

The prototype firmware enabled Bluetooth Low Energy advertising. BLE advertising and device discoverability were confirmed using a mobile application. This test verified basic radio-interface and firmware operation but did not include end-to-end sensor-data transfer. The ESP32-H2 datasheet documents support for IEEE 802.15.4-based technologies, including Zigbee, Thread and Matter over Thread [6]. These technologies were not evaluated in the present prototype.

2.1. Configurable Interface Components and Design Constraints

The PCB provides local supply-decoupling components and optional footprints for adjusting the edge behavior of selected digital communication lines. At the ESP32-H2 supply pins, the documented capacitor set consists of 10 μ F, 100 nF, 1 nF, and 100 pF capacitors. At the sensor-connector supply pins, the corresponding set consists of 2.2 μ F, 100 nF, and 1 nF capacitors. These components are located close to the associated supply pins to provide local decoupling [4].

For the SPI interface, the documented design includes a 22 Ω series resistor and an optional 10 pF shunt capacitor to ground on each of the CS, MOSI, SCK, and MISO lines, as shown in Fig. 3. Other digital communication lines include series-resistor footprints intended for values in the range of 22–100 Ω .

Series resistance is commonly used to reduce transmission-line effects such as ringing and overshoot, while small shunt capacitances can be used to modify signal-edge behavior [2]. However, the effectiveness of specific component values depends on

the source and load impedances, interconnect geometry, parasitic capacitance, and operating frequency. The presented values should therefore be regarded as initial design choices rather than experimentally optimized parameters.




Fig. 3. SPI signal-conditioning circuit (22 Ω series resistors and 10 pF shunt capacitors); source: adapted from [4].

The 10 pF capacitors may be left unpopulated when their influence on signal-edge behavior is undesirable. Although the ESP32-H2 GP-SPI controller supports master clock rates of up to 48 MHz, this value represents a controller capability and not a validated operating speed for the complete assembled bus [6]. The proposed component network is therefore presented as a configurable design provision rather than a quantitatively validated signal-integrity improvement.

3. Intended Duty-Cycled Operation

The sensor node is designed to support timer-controlled duty-cycled operation. In the intended firmware sequence, the microcontroller enters deep sleep between acquisition cycles, wakes in response to an internal timer, initializes the connected sensor, acquires a measurement and subsequently returns to deep sleep. During deep sleep, the CPU core and wireless communication blocks are disabled while the functions required for timed wake-up remain active.

Duty-cycling is commonly used to reduce the average energy consumption of wireless sensor nodes by limiting the time spent in active measurement and communication states [7, 8]. Battery-powered operation of the assembled prototype was confirmed independently of the duty-cycled sleep/wake firmware sequence described above; the sleep/wake cycle itself, along with board-level current and timing measurements, was not functionally verified in this

work. Consequently, this work makes no quantitative claim regarding average current consumption, energy per acquisition cycle or battery lifetime.

4. Prototype Functional Validation

The assembled prototype was subjected to basic functional tests intended to detect assembly, power-supply and interface-connectivity faults. The tests were qualitative and were not designed to provide statistical or performance-related metrics. Table 1 summarizes the implemented functions and their validation status.

Table 1. Prototype implementation and validation scope.

Function	Implementation or test configuration	Validation status
I2C interface	BME680 environmental sensor	Basic communication confirmed
SPI interface	SD-card interface	Basic communication confirmed
UART interface	Hardware connector implemented	Not evaluated
Battery supply	Prototype powered from battery	Basic operation confirmed
BLE	Advertising monitored by mobile application	Advertising and discoverability confirmed
Interface tuning network	Optional series-resistor and shunt-capacitor footprints implemented	Not quantitatively characterized
Duty-cycled operation	Intended timer-driven firmware sequence (design concept)	Sleep/wake sequence not functionally verified

The validation confirms that the assembled PCB supports the tested interfaces and basic operating functions. End-to-end sensor-data transmission, wireless-link performance, signal integrity and

board-level energy consumption remain outside the experimental scope of the present work.

5. Conclusions

Overall, the prototype demonstrates the initial functional feasibility of the proposed modular hardware architecture through representative I2C and SPI peripheral tests, battery-powered operation and BLE advertising. The configurable interface networks provide a basis for subsequent optimization, but the present results do not constitute quantitative performance validation.

Further work should functionally verify the complete duty-cycled sleep/wake sequence and quantify current consumption, signal integrity and end-to-end wireless communication.

References

- [1]. H. J. Albeyboni, I. A. Ali, LPWAN technologies for IoT applications: A review, *Journal of University of Duhok*, Vol. 26, Issue 1, 2023, pp. 29-42.
- [2]. H. W. Johnson, M. Graham, High-Speed Digital Design: A Handbook of Black Magic, *Prentice Hall*, Englewood Cliffs, 1993.
- [3]. T. Polonelli, et al., Aerosense: A self-sustainable and long-range Bluetooth wireless sensor node for aerodynamic and aeroacoustic monitoring on wind turbines, *IEEE Sensors Journal*, Vol. 23, Issue 1, 2023, pp. 715-723.
- [4]. A. Simkovic, FRISStack – core and sensor module, Master's Thesis, *University of Žilina*, Žilina, 2026.
- [5]. D. Widhalm, K. M. Goeschka, W. Kastner, An open-source wireless sensor node platform with active node-level reliability for monitoring applications, *Sensors*, Vol. 21, Issue 22, 2021, 7613.
- [6]. ESP32-H2-MINI-1 and ESP32-H2-MINI-1U datasheet, Version 1.6, *Espressif Systems*, 2025.
- [7]. S. Khriji, R. Cheour, O. Kanoun, Dynamic voltage and frequency scaling and duty-cycling for ultra low-power wireless sensor nodes, *Electronics*, Vol. 11, Issue 24, 2022, 4071.
- [8]. M. Hodon, M. Chovanec, L. Cechovic, M. Hudik, et al., Maximizing performance of low-power WSN node on the basis of event-driven-programming approach: Minimization of operational energy costs of WSN node control unit, in *Proceedings of the IEEE Symposium on Computers and Communication (ISCC)*, 2015, pp. 204-209.

A Comparative Study of Variational Autoencoder Architectures for Anomaly-Based Fault Detection in an Industrial Control System

Ietezaz ul Hassan^{1,2}, Krishna Panduru^{1,2}, Daniel Riordan^{1,2} and Prof. Joseph Walsh^{1,2}

¹IMaR Research Centre, Munster Technological University, V92 HD4V, Co Kerry, Ireland.

² School of Science, Technology, Engineering, and Mathematics, Munster Technological University, V92 HD4V, Co Kerry, Ireland.

E-mail: Ietezaz.Ul.Hassan@mymtu.ie

Summary: Industrial Control Systems (ICS) require reliable fault detection methods to ensure operational safety and efficiency. However, the limited availability of labelled fault data in industrial environments makes supervised approaches difficult to implement. This study presents a comparative analysis of four unsupervised Variational Autoencoder (VAE)-based anomaly detection models: CNN-based-VAE, LSTM-based-VAE, BiLSTM-based-VAE, and GRU-based-VAE. The models were trained using healthy sensor data collected from a real-world industrial environment. Experimental results demonstrate that the CNN-based-VAE consistently outperforms the recurrent-based VAE models, achieving higher anomaly detection accuracy, higher true positive rates, and lower false positive rates across multiple industrial components. The results indicate that CNN-based-VAE is a robust and effective solution for anomaly detection in ICS, offering significant potential for condition monitoring and predictive maintenance applications.

Keywords: Industrial control systems (ICS), Anomaly detection, Fault detection, Condition monitoring, Predictive maintenance, Unsupervised learning, Deep learning.

1. Introduction

ICS are used to monitor industrial processes including manufacturing, power generation, transportation systems, and oil and gas operations. ICS is made up of hardware and software components to monitor, control, and automate the industrial processes [1]. Sensors, actuators, Programmable Logic Controllers (PLCs), Human Machine Interface (HMI), Supervisory Control and Data Acquisition (SCADA), and communication networks are some key components of an ICS that operate together to collect and process data related to industrial processes. ICS use a number of sensors to collect temperature, pressure, vibration, and other data, which could be processed and used to control processes using industrial computers known as PLCs [2].

The interconnected components of an ICS are vulnerable to equipment ageing, sensor failures, communication problems, software faults, human errors, environmental factors, process abnormalities, and so on. Any possible failure or anomaly in these components could result in a defect, breakdown, or operational inefficiency [3]. Therefore, it is important to stop ICS faults, as these systems have a responsibility to regulate significant operations, and could negatively affect profitability, productivity, safety, and equipment health. Condition monitoring is a maintenance method, that uses the collected data from sensors and monitoring devices to evaluate the operational conditions of machinery, equipment, and industrial processes in real time or at regular intervals [4]. In condition monitoring, the collected data could be processed using methods based on rule-based,

signal processing, statistical analysis, trend analysis, or artificial intelligence-based methods [5].

The increasing degree of complexity and interconnected nature of modern industrial machinery has promoted the transition from simple rule-based data processing methods to data-driven methods. The collected operational data could be processed for future fault detection and diagnosis using these data-driven methods. Data-driven methods could either be supervised, unsupervised, or semi-supervised learning based [6]. The supervised data-driven methods are very efficient, but they require labelled faulty data to develop, which is a challenging task in industrial environments [7]. Because machines are mostly operating normally, faults occur very rarely, and when they occur, the machines either stop working or the fault is repaired immediately. Therefore, there have been recent advancements in the application of unsupervised and semi-supervised learning-based methods for condition monitoring and fault detection [8].

In this article, a comparison of CNN, Long Short Term Memory (LSTM), Bi-Directional Long Short Term Memory (BiLSTM), and Gated Recurrent Unit (GRU) based VAE architectures on confidential real world industrial sensor data is carried out to evaluate their effectiveness, using healthy data for training and reconstruction error for anomaly detection.

2. Preliminaries and Background

This section summarises recent developments in fault and defect detection using traditional and artificial intelligence-based methods.

2.1. Fault Detection Methods

Fault detection involves identifying abnormalities in a system's operation. In general, fault detection methods are classified into three types: knowledge-based, model-based, and data-driven methods [8]. Knowledge-based fault detection methods can use expert knowledge and rule-based methods to detect faults based on specific algorithm capable of detecting defects relying on specific conditions or symptoms. Whereas, model-based methods use mathematical models to reflect the system or process underlying normal behavior, and they can detect defects by analyzing inconsistencies between model predictions and actual outputs. Whereas, data-driven methods rely on statistical and machine learning methods to process huge volumes of operational data so as to predict patterns which indicate defects [9]. Among all these methods, data-driven methods are suitable for fault detection, because knowledge-based and model-based methods require explicit knowledge descriptions, which are both challenging and resource-intensive [3].

2.2. Data-Driven Based Fault Detection Methods

Data-driven based fault detection methods employs artificial intelligence and machine learning methods to extract key insights directly from data [10]. It employs machine learning algorithms and historical data to detect faults. These methods learn the characteristics of faults and have made significant developments towards fault detection. Machine learning-based algorithms could either be supervised, unsupervised, or reinforcement learning-based, where supervised machine learning algorithms are trained on labelled data, which is very challenging in industrial environments [11]. Whereas, unsupervised machine learning-based algorithms could be trained on unlabeled data to detect patterns in the data and uses anomaly detection to predict a fault [12]. Lastly, reinforcement machine learning-based algorithms involve training model using a trial-and-error approach and are important for autonomous systems in a dynamic industrial setting [13]. Due to a lack of labelled data, machine learning algorithms used in a real-world industrial environment could be constrained to unsupervised or semi-supervised learning-based algorithms [14].

2.3. Anomaly Based Fault Detection

Anomaly detection has recently attracted a lot attention in condition monitoring and maintenance systems due to the limited availability of labelled data [15]. Deep learning has made significant advancements in anomaly detection by supporting data-driven modeling of complex patterns directly from raw sensors data [15]. For anomaly detection, a number of deep learning architectures, including CNNs, Recurrent Neural Network (RNNs), LSTM, GRUs, attention-based models, autoencoders, and

transformer-based frameworks, have demonstrate significant performance.

The employed deep learning-based method could determine how a model learns to distinguish between normal and abnormal behavior [16]. Deep learning-based methods for detecting anomaly could be classification-based, reconstruction-based, or prediction-based, depending on data availability, labelling conditions, and operational environments [17]. Classification-based anomaly detection methods apply supervised learning to assign a specified class (such as normal or abnormal) or a specific fault type to each data sample. This approach performs very effectively when fault labels are defined and labelled fault data is available. However, prediction-based anomaly detection methods learn the normal temporal dynamics of a system and could detect any observations that are different from the expected values. These methods rely on normal operational data, thus reducing the requirement for fault labels, which makes them suitable for real-world industrial environments. Lastly, reconstruction-based anomaly detection algorithms attempt to comprehend normal (healthy) system behavior. These methods detect anomaly based on the reconstruction error, which could be significant if the data deviates significantly from historically learnt normal patterns [18].

3. Proposed Framework

This section outlines the proposed framework employed to conduct this research study. The framework presented in recent article (currently accepted for publication) has been adopted here, as presented in Fig. 1. The proposed framework consists of key steps, including data collection, preprocessing, model development, training, and evaluation.




Fig. 1. Proposed framework

3.1. Data Collection

The data used in this article was sourced from a real-world industrial environment, and the name and machine specifics have been kept hidden for confidentiality purposes. The data in this industrial setting was collected using a number of sensors, including vibration, temperature, status-indicating sensors, and so on. These sensors send data to the PLC, which saves it in proprietary encoded files generated on an hourly basis.

3.2. Preprocessing

During the preprocessing step, the collected data are transformed from encoded representations to data

frames. The transformed data frames were then further processed to extract only the changing data (which did not remain constant over time) and scaled to fall within the same range. The scaled data was then used to generate segments that were used to train the models for anomaly detection. In this study, segments were generated using moving windows of size 100. Each generated segment served as a training sample for the model.

3.3. Model Development

This article evaluated several different reconstruction-based anomaly detection models including CNN-based-VAE, LSTM-based-VAE, BiLSTM-based-VAE, and GRU-based-VAE. The CNN-based-VAE model as presented in Fig. 2.

All of these models are built on an encoder-decoder architecture, with the encoder using the layers as a feature extractor to extract useful features from the data and the decoder using the layers to generate data of the same shape as the input. The CNN-based-VAE model as presented in Fig. 2, consists of two 1-dimensional convolutional layers followed by 1-dimensional max pooling layers in the encoder, and two 1-dimensional convolutional layers followed by 1-dimensional upsampling layers in the decoder. In this architecture, the convolutional layers are replaced with LSTM, BiLSTM, or GRU layers, depending on whether the VAE is based on LSTM, BiLSTM, or GRU. These VAEs (LSTM-based-VAE, BiLSTM-based-VAE, and GRU-based-VAE) bypass the pooling and upsampling layers because they process data sequentially across time rather than spatially across dimensions. Another reason for avoiding pooling and upsampling layers is their operational differences, which make them incompatible with their underlying processes.

The reconstruction error threshold for detecting anomalies was selected as the 99th percentile of the reconstruction errors obtained from normal (healthy) data.




Fig. 2. Proposed CNN-based-VAE.

3.4. Models Training

The developed VAE-based models for anomaly detection were trained for 10 epochs with a batch size of 128. During the training, 10 % of the training data

was used to evaluate the model's performance. These models were trained using Adam as the optimizer and a learning rate of 0.001.

3.5. Models Evaluation

In this article, the true negative rate, false negative rate, true positive rate, and false positive rate have been used to evaluate the performance of the VAE-based reconstruction models for anomaly detection.

4. Results and Discussion

This section presents the results of four VAE-based models for detecting anomaly in healthy and manipulated test data. The manipulated test data was obtained by manipulating the same healthy test data by adding Gaussian-based noise. The results indicate the percentage of the samples detected as anomalous by each model out of a total of 52,817 samples.

4.1. True Negative Rate

The true negative rate is the ratio of healthy samples identified as healthy by the model. The model's performance could be described as efficient, having a high true negative rate. The percentage of healthy samples identified as healthy samples by the four VAE-based models is presented in Table 1 below. This table demonstrates that CNN-based-VAE performs better than LSTM-based-VAE, BiLSTM-based-VAE, and GRU-based-VAE, with true negative rates of 99.34 %, 99.55 %, and 99 % for components 1, 3, and 4 data, respectively. However, with a true negative rate of 98.99 %, the BiLSTM-based-VAE has only outperformed other VAEs on component 2 data.

Table 1. Performance comparison of four VAE-based models using true negative rate.

Component	CNN-based-VAE	LSTM-based-VAE	BiLSTM-based-VAE	GRU-based-VAE
Component 1	99.34%	98.85%	91.18%	98.98%
Component 2	98.89%	98.9%	98.99%	98.98%
Component 3	99.55%	99.2%	98.84%	98.99%
Component 4	99%	98.81%	98.79%	98.72%

4.2. False Negative Rate

The false negative rate is the ratio of anomalous samples that the model incorrectly identifies as healthy samples. Table 2 presents the percentage of anomalous samples that were incorrectly identified as healthy samples (known as the false negative rate in this scenario) by the four VAE-based models. In terms of false negative rate, the optimum model is the one with

the lowest false negative rate, because a low false negative rate means that fewer anomalous samples are falsely labelled as healthy. As shown in the table, the CNN-based-VAE model achieves the lowest false negative rate for components 1, 3, and 4, which are 0.66 %, 0.45 %, and 1 %, respectively.

Table 2. Performance comparison of four VAE-based models using false negative rate.

Component	CNN-based-VAE	LSTM-based-VAE	BiLSTM-based-VAE	GRU-based-VAE
Component 1	0.66%	1.15%	8.82%	1.02%
Component 2	1.11%	1.1%	1.01%	1.02%
Component 3	0.45%	0.8%	1.16%	1.01%
Component 4	1%	1.19%	1.21%	1.28%

4.3. True Positive Rate

The true positive rate is the ratio of anomalous samples that the model correctly identifies as anomalous samples. Table 3 presents the percentage of anomalous samples that were correctly identified as anomalous samples (known as the true positive rate in this case) by the four VAE-based models. As a high true positive rate indicates that more anomalous samples are appropriately classified as anomalous, that's why the model with the highest true positive rate is considered the best one. Table 3 demonstrates that the CNN-based-VAE outperforms all other VAE models by reaching a 100% true positive rate for data from components 1, 2, and 4, and a 98.45 % true positive rate for data from component 3.

Table 3. Performance comparison of four VAE-based models using true positive rate.

Component	CNN-based-VAE	LSTM-based-VAE	BiLSTM-based-VAE	GRU-based-VAE
Component 1	100%	2.47%	31.34%	4.72%
Component 2	100%	1.52%	99.41%	7.18%
Component 3	98.45%	1.08%	2.11%	1.44%
Component 4	100%	1.36%	1.84%	1.43%

4.4. False Positive Rate

The ratio of healthy samples that the model incorrectly classifies as anomalous samples is known as the false positive rate. The percentage of healthy samples that four VAE-based models mistakenly identified as anomalous samples (referred to as the false positive rate in this case) is presented in Table 4. The best model could be selected as the one with the lowest false positive rate because fewer healthy samples are incorrectly classified as anomalous when there is a low

false positive rate. As shown in the table, the CNN-based-VAE model has the lowest false negative rate for components 1, 2, and 4, at 0 % and 1.55 % for data from component 3.

Table 4. Performance comparison of four VAE-based models using false positive rate.

Component	CNN-based-VAE	LSTM-based-VAE	BiLSTM-based-VAE	GRU-based-VAE
Component 1	0%	97.53%	68.66%	95.28%
Component 2	0%	98.48%	0.59%	92.82%
Component 3	1.55%	98.92%	97.89%	98.56%
Component 4	0%	98.64%	98.16%	98.57%

5. Conclusion and Future Work

A comparison of four VAE reconstruction-based models for anomaly detection, including CNN-based-VAE, LSTM-based-VAE, BiLSTM-based-VAE, and GRU-based-VAE, demonstrated that the CNN-based-VAE model could outperform the other VAE models. The CNN-based-VAE model demonstrated a high true negative rate, a low false negative rate, a high true positive rate, and a low false positive rate for anomaly detection in a real-world ICS. The CNN-based-VAE's superior performance could be attributed to its capacity to extract hierarchical spatial patterns more efficiently through parallel processing, as well as its ability to measure data variance probabilistically rather than relying on a set of historical sequences.

In Future, this work will be extended into a comprehensive journal article. The extended version will include a more detailed description of the experimental setup, additional benchmark datasets, extensive comparisons with state-of-the-art models, sensitivity analysis, and a deeper investigation of the proposed method. These additions will provide a more thorough evaluation of the approach and offer further insights into its strengths, limitations, and practical applicability.

Acknowledgements

The authors extend sincere thanks to the entire IMaR team at Munster Technological University in Co. Kerry, Ireland.

References

- [1]. C. Serôdio, P. Mestre, J. Cabral, M. Gomes, et al., Software and architecture orchestration for process control in Industry 4.0 enabled by cyber-physical systems technologies, *Applied Sciences*, Vol. 14, Issue 5, 2024, Article 2160.
- [2]. O. Akande, *Industrial Automation from Scratch: A Hands-on Guide to Using Sensors, Actuators, PLCs*,

- HMIs, and SCADA to Automate Industrial Processes*, Packt Publishing, 2023, 496 p.
- [3]. D. Leite, E. Andrade, D. Rativa, A. M. A. Maciel, Fault detection and diagnosis in Industry 4.0: A review on challenges and opportunities, *Sensors*, Vol. 25, Issue 1, 2025, Article 60.
 - [4]. I. U. Hassan, K. Panduru, J. Walsh, An in-depth study of vibration sensors for condition monitoring, *Sensors*, Vol. 24, Issue 3, 2024, Article 740.
 - [5]. I. U. Hassan, K. Panduru, J. Walsh, Review of data processing methods used in predictive maintenance for next generation heavy machinery, *Data*, Vol. 9, Issue 5, 2024, Article 69.
 - [6]. P. Jieyang, A. Kimmig, W. Dongkun, Z. Niu, et al., A systematic review of data-driven approaches to fault diagnosis and early warning, *Journal of Intelligent Manufacturing*, Vol. 34, Issue 8, 2023, pp. 3277–3304.
 - [7]. W. Strielkowski, A. Vlasov, K. Selivanov, K. Muraviev, et al., Prospects and challenges of the machine learning and data-driven methods for the predictive analysis of power systems: A review, *Energies*, Vol. 16, Issue 10, 2023, Article 4025.
 - [8]. K. Eltoumy, M. Gomaa, X. Liang, Unsupervised learning methods for data-driven vibration-based structural health monitoring: A review, *Sensors*, Vol. 23, Issue 6, 2023, Article 3290.
 - [9]. A. Melo, M. M. Câmara, J. C. Pinto, Data-driven process monitoring and fault diagnosis: A comprehensive survey, *Processes*, Vol. 12, Issue 2, 2024, Article 251.
 - [10]. L. Lei, W. Li, S. Zhang, C. Wu, et al., Research progress on data-driven industrial fault diagnosis methods, *Sensors*, Vol. 25, Issue 9, 2025, Article 2952.
 - [11]. T. Talaei Khoei, N. Kaabouch, Machine learning: Models, challenges, and research directions, *Future Internet*, Vol. 15, Issue 10, 2023, Article 332.
 - [12]. E. F. Agyemang, Anomaly detection using unsupervised machine learning algorithms: A simulation study, *Scientific African*, Vol. 26, 2024, Article e02386.
 - [13]. W. Zhao, Y. Lv, J. Liu, C. K. M. Lee, et al., Early fault diagnosis based on reinforcement learning optimized-SVM model with vibration-monitored signals, *Quality Engineering*, Vol. 35, Issue 4, 2023, pp. 696–711.
 - [14]. A. Saifoddin, N. Mirzaei, M. Allahrabbi Shirazi, H. Yousefi, Applications of supervised and unsupervised machine learning models in energy systems, *Journal of Energy Management and Technology*, Vol. 9, Issue 4, 2025, pp. 284–290.
 - [15]. A. Bouketta, S. Niar, H. Ouarnoughi, Deep learning for anomaly detection in railway systems: A structured survey, *Engineering Applications of Artificial Intelligence*, Vol. 181, Part 7, 2026, Article 115776.
 - [16]. H.-O. Choe, M.-H. Lee, Artificial intelligence-based anomaly detection technology for equipment condition monitoring in smart farms, *Applied Sciences*, Vol. 15, Issue 23, 2025, Article 12843.
 - [17]. Z. Zamanzadeh Darban, G. I. Webb, S. Pan, C. C. Aggarwal, et al., Deep learning for time series anomaly detection: A survey, *ACM Computing Surveys*, Vol. 57, Issue 1, 2024, Article 15, 42 p.
 - [18]. J. Miao, H. Tao, H. Xie, J. Sun, et al., Reconstruction-based anomaly detection for multivariate time series using contrastive generative adversarial networks, *Information Processing & Management*, Vol. 61, Issue 1, 2024, Article 103569.

Resistance-Based Thermal Sensing and Heat Propagation Modeling of a PTC-Heated Shape Memory Alloy Actuator

D. Arriaga-Ventura, Rogelio Bustamante-Bello and Javier Izquierdo Reyes

School of Engineering and Sciences, Tecnológico de Monterrey, Mexico City, Mexico

E-mail: a01338696@tec.mx, rbustama@tec.mx, javieriz@mit.edu

Summary: Shape Memory Alloy (SMA) actuators embedded in silicone offer compact and compliant motion for soft robotic systems, but their response depends strongly on temperature distribution and heat-transfer conditions. This work presents a resistance-based thermal characterization framework for a PTC-heated SMA actuator. A ceramic Positive Temperature Coefficient (PTC) heater was used as both the thermal excitation source and the electrical variable for indirect temperature measurement. The heater was excited using a pseudo-random binary sequence (PRBS) at voltage levels from 4 V to 12 V, while current and thermal-camera measurements were recorded. The PTC resistance was computed from voltage and current, and its logarithmic resistance-temperature behavior was fitted before and after the Curie region. Then, heat propagation along the SMA segment was modeled using a distance-dependent first-order differential equation. The PTC temperature reconstruction reached an average fit of 82.6 %, while the propagation model achieved 93.0 % at the farthest evaluated SMA point.

Keywords: Shape memory alloy, PTC heater, Thermal characterization, Resistance-based sensing, Heat propagation.

1. Introduction

Shape Memory Alloy (SMA) actuators have become an attractive alternative for soft robotic systems due to their compact geometry, high power-to-weight ratio, silent operation, and ability to generate compliant motion. In particular, NiTi wires embedded in silicone structures can produce deformation without bulky transmission mechanisms, making them suitable for compact soft actuators, artificial muscles, and bioinspired robotic platforms [1-3]. However, their actuation mechanism is inherently thermal, since motion is generated by the temperature-dependent martensite-austenite phase transformation [4, 5]. Therefore, accurate thermal characterization is essential to understand, predict, and control the response of SMA-based actuators.

Temperature measurement in silicone-embedded SMA actuators is challenging. Contact-based sensors can be placed near the heating element or directly on the SMA wire, but both configurations may introduce measurement uncertainty [6, 7]. When the sensor is placed near the heater, the measured value may not represent the temperature effectively transferred to the SMA. When the sensor is placed on the SMA wire, maintaining stable contact during operation is difficult because the wire and the silicone body deform during actuation. In addition, the temperature distribution along the wire is not necessarily uniform, which limits the use of a single-point temperature measurement as a representative thermal state of the actuator.

Direct Joule heating is one of the most common methods for activating SMA wires, since electrical current is applied directly through the actuator material [2, 3]. Nevertheless, this approach can require relatively high current levels and does not inherently

provide thermal self-regulation. As an alternative, ceramic Positive Temperature Coefficient (PTC) heaters offer a self-regulated heating mechanism associated with their resistance-temperature behavior [8, 9]. In the analyzed PTC element, the electrical resistance decreases as temperature increases below the Curie temperature, whereas after this point, the resistance increases sharply. This transition limits the heating process and allows the PTC element to be characterized not only as a thermal source, but also as a resistance-based thermal sensing variable.

The use of a PTC heater introduces an additional thermal modeling challenge. Unlike direct Joule heating, where heat is generated along the SMA wire, a PTC-heated configuration transfers heat from a localized ceramic element to the SMA through a contact region [10]. As a result, the temperature distribution depends on the thermal contact between the heater and the SMA, the distance from the anchoring point, the surrounding silicone, and the heat dissipation conditions. This produces a spatially distributed and time-dependent thermal response that cannot be fully described by the heater temperature alone. Therefore, relating the resistance-based temperature of the PTC heater to the distributed temperature response of the SMA wire is necessary to describe the thermal behavior of the actuator under localized heating [11].

This work presents a resistance-based thermal characterization framework for a silicone-embedded SMA actuator heated by a ceramic PTC element. The PTC heater is characterized through its resistance-temperature behavior, allowing its temperature to be obtained from electrical measurements. Then, the temperature propagation along the active SMA wire is characterized at different distances from the heating

point using a distance-dependent first-order differential equation obtained from experimental data. A pseudo-random binary sequence (PRBS) excitation signal is used to evaluate the dynamic thermal response of the PTC heater and the SMA segment under different voltage levels [12]. This approach provides a thermal self-sensing alternative for SMA actuators by exploiting the intrinsic electrical and thermal properties of the PTC-SMA system.

The main contribution of this work is the integration of the electrical resistance–temperature characterization of a ceramic PTC heater with a distance-dependent thermal propagation model for a silicone-embedded SMA actuator. Unlike conventional SMA thermal characterization approaches based on direct Joule heating or multiple temperature sensors, the proposed framework uses the PTC element simultaneously as a localized heat source and as an indirect temperature-sensing variable. This combination enables the PTC temperature to be reconstructed from electrical measurements and the SMA temperature response to be estimated at different distances from the heater.

2. Literature Review

SMA actuators are widely used in soft robotics because of their compact geometry, high power density, silent operation, and ability to generate compliant motion [1-3]. Their response is governed by a thermally induced phase transformation between martensite and austenite, which also introduces hysteresis, cooling limitations, and sensitivity to heat-transfer conditions [4, 5]. Consequently, accurate thermal characterization is necessary to improve the predictability and repeatability of SMA-based systems.

Direct Joule heating is the most common method for activating SMA wires, while their electrical resistance has also been investigated as a self-sensing variable. However, SMA resistance depends simultaneously on temperature, strain, stress, and phase transformation, which complicates the isolated estimation of temperature during operation. Contact-based sensors may provide direct measurements, but their installation can affect the actuator response or become unreliable when the SMA and surrounding silicone deform [6, 7].

Ceramic PTC heaters provide a self-regulated alternative because their resistance changes significantly near the Curie region [8, 9]. When a localized PTC heater is coupled to an SMA wire, heat propagation depends on the contact condition, distance from the heater, surrounding material, and heat dissipation [10, 11]. Therefore, combining the resistance–temperature characterization of the PTC element with a distance-dependent thermal model provides a suitable approach for estimating both the heater temperature and the distributed thermal response of the SMA actuator.

3. Methodology

This section describes the experimental setup and the procedures used to characterize the PTC heater and the heat propagation along the SMA segment. The methodology comprises electrical and thermal data acquisition, resistance–temperature modeling, and identification of the distance-dependent thermal model.

3.1. Experimental Setup

The experimental platform was designed to characterize the electrical and thermal behavior of the actuator under controlled excitation and environmental conditions. The measurement system included a thermal camera, which allowed for the collection of temperature data from five different measurement points. Additionally, a current sensor ACS712 with a measurement range from 0 to 5 A was incorporated, along with a MOSFET driver module to transmit the excitation signal to the heating element. The current sensor was connected in series between the output terminal of the MOSFET module and one terminal of the PTC heater, enabling the measurement of the current supplied to the heating element during operation. Finally, a Texas Instruments control board C2000 LaunchPad XL was used to program the excitation sequence and acquire data through MATLAB and Simulink.

The excitation input was defined as a PRBS, which was selected to obtain the thermal response of the system under different excitation frequency components. The PRBS signal was programmed in MATLAB and consisted of 20 binary states, each with a duration of 60 s. The binary values defined the activation state of the heater, allowing the PTC element to be switched on and off according to the generated sequence, as shown in Fig. 1. The characterization tests were performed under different excitation voltage levels, ranging from 4 V to 12 V with a step size of 2 V, to evaluate the thermal and electrical response of the actuator under different heating conditions.




Fig. 1. Binary PRBS excitation signal applied to the PTC heater through the MOSFET driver.

The PRBS signal was generated in MATLAB, implemented in Simulink, and executed using a Texas

Instruments C2000 LaunchPad XL. The controller switched the ceramic PTC heater through a MOSFET module powered by a GPS-3303 DC supply. Current and thermal-camera measurements were recorded simultaneously for excitation levels from 4 V to 12 V in 2 V increments.

The actuator consisted of a 12 W ceramic PTC heater mechanically attached to an active SMA segment. This segment was formed by two intertwined NiTi wires, each with a diameter of 1 mm and a length of 5 cm. The wires were twisted together three times using a drill to obtain a braided configuration. After twisting, the SMA segment was bent into a U-shape, and the central region of the U was mechanically anchored to the PTC heater. This configuration allowed the heater to act as a localized thermal source, while temperature measurements were acquired along the SMA legs to characterize heat propagation from the anchoring point. This configuration is illustrated in Fig. 2.




Fig. 2. Actuator configuration.

The experimental area was kept at a stable ambient temperature of 25 °C and shielded from air currents. To minimize disturbances from external temperature fluctuations, both the actuator and the thermal camera were placed inside an insulated box. This setup created a controlled environment for collecting thermal data and reduced the impact of convective effects during the characterization test.

3.2. Characterization of PTC Resistance

To characterize the thermal behavior of the ceramic PTC heater, temperature measurements were taken in the central region of the heating element during PRBS excitation tests. The temperature response obtained for each voltage level is shown in Fig. 3, where the thermal evolution of the PTC heater can be observed throughout the complete PRBS sequence.

The current supplied to the PTC heater was also measured during each test, as shown in Fig. 4. Using the measured voltage and current, the electrical resistance of the heater was calculated through Ohm's law.

Then, the natural logarithm of the resistance was used to describe the exponential behavior commonly associated with PTC elements in a simpler form. The

resulting resistance values were related to the corresponding temperature measurements of the heater. Before fitting the data, a moving-average filter with a window length of 10 samples was applied to reduce measurement noise, as shown in Fig. 5.




Fig. 3. Thermal response of the PTC heater under PRBS excitation at different voltage levels.




Fig. 4. Current response of the PTC heater under PRBS excitation at different voltage levels.




Fig. 5. Filtered $\ln(R_H)$ response as a function of temperature.

The Curie temperature was estimated from the filtered $\ln(R_H)$ -temperature data by fitting robust linear models to the decreasing region below 180 °C and the increasing region above 200 °C. Their intersection yielded $T_C = 193.4$ °C. For model formulation and implementation, the estimated value was rounded to 193 °C. A block-bootstrap uncertainty analysis produced a 95 % confidence interval of 192.0–195.1 °C, allowing the result to be reported as $T_C = 193 \pm 2$ °C. Varying the moving-average window from 5 to 12 bins changed the estimate by less than 0.2 °C, indicating low sensitivity to moderate variations in the filtering parameters. Based on this temperature, the dataset was divided into two regions: below and above T_C . In the pre-Curie

region, the data were fitted using a second-order polynomial, as indicated in Eq. (1). Conversely, the post-Curie region was fitted with a linear model, as represented in Eq. (2). Both models were derived using least-squares regression.

$$\ln(R_H) = 2.28 - 2.72 \times 10^{-4}(T_{PTC} - 193) + 1.22 \times 10^{-5}(T_{PTC} - 193)^2, \quad (1)$$

$$\ln(R_H) = 2.288 + 5.44 \times 10^{-2}(T_{PTC} - 193) \quad (2)$$

3.3. Heat Propagation Along the SMA Wire

To characterize heat propagation along the active SMA segment, temperature measurements were acquired at four different distances from the anchoring point between the PTC heater and the SMA wire. This anchoring point was considered the highest-temperature region because it remained in fixed contact with the ceramic heater during the test. The measurement positions were located at 1 mm, 1.2 mm, 4.6 mm, and 5.2 mm from the anchoring point. The temperature responses obtained at these positions during the PRBS excitation tests are shown in Fig. 6.




Fig. 6. Temperature response along the SMA segment at different distances from the anchoring point.

A lumped-parameter thermal model was used to describe heat propagation from the PTC heater to the SMA segment, as shown in Fig. 7. In this model, T_{PTC} represents the absolute temperature of the ceramic heater, while the thermal source $T_{PTC} - T_e$ represents the temperature rise of the heater with respect to the ambient temperature T_e . The parameter R_c represents the thermal contact resistance associated with the imperfect contact between the ceramic heater and the intertwined NiTi wires. The temperature at the SMA measurement point is represented by T_{SMA} , and C_{SMA} denotes the effective thermal capacitance of the corresponding SMA segment. The parameter R_{SMA} represents the distance-dependent thermal resistance between the SMA measurement point and the surrounding environment. The lower reference node corresponds to the ambient temperature (T_e).




Fig. 7. Lumped-parameter thermal model for heat propagation along the SMA segment.

By solving the lumped thermal circuit, the relationship between the PTC temperature and the SMA temperature at a distance x is expressed by the first-order differential equation in Eq. (3). Since the reference node represents the ambient temperature (T_e), the model is expressed in terms of the temperature rises of the PTC heater and the SMA segment with respect to the environment.

$$(T_{PTC} - T_e) = b(x) \frac{dT_{SMA}}{dt} + k(x)(T_{SMA} - T_e) \quad (3)$$

In Eq. (3), $b(x)$ is the dynamic thermal coefficient associated with the thermal inertia of the SMA segment and the thermal coupling between the PTC heater and the SMA measurement point. It has units of seconds because it represents the combined effect of thermal resistance and thermal capacitance. The coefficient $k(x)$ is a dimensionless static thermal coefficient associated with the heat-transfer and dissipation paths between the PTC heater, the SMA segment, and the surrounding environment. The distance x is measured from the anchoring point and is expressed in millimeters. Both coefficients were identified for each measurement distance and excitation voltage, averaged across the voltage levels, and fitted as functions of distance, resulting in the empirical expressions shown in Eq. (4) and Eq. (5).

$$b(x) = (11.22 \text{ s})e^{(0.0382 \text{ mm}^{-1} x)}, \quad (4)$$

$$k(x) = 1.63542 e^{(0.0044 \text{ mm}^{-1} x)} \quad (5)$$

4. Results

The same experimental datasets were used for model parameter estimation and fit calculation. Therefore, the reported fit values quantify the agreement between the identified models and the experimental data used during the fitting process. Future work will evaluate the performance of the proposed models using additional tests with different excitation voltage levels and alternative PRBS routines.

The fit percentages reported in this work were calculated using the MATLAB normalized fit criterion presented in Eq. (6).

$$\text{Fit}(\%) = 100 \left(1 - \frac{\|y - \hat{y}\|_2}{\|y - \bar{y}\|_2} \right), \quad (6)$$

where y is the measured temperature response, $\hat{y}$ is the estimated temperature response, and $\bar{y}$ is the mean of the measured data. According to Eq. (6), a fit of 100 % represents perfect agreement between the measured and estimated responses.

The resistance-temperature characterization of the PTC heater was performed using the transformed variable $\ln(R_H)$, where R_H was calculated from the measured voltage and current of the heating element. Fig. 8 shows the fitted $\ln(R_H)$ temperature models for the pre-Curie and post-Curie regions. The fitted models presented a correlation coefficient of 0.9856 and normalized root-mean-square error (NRMSE) of 5.64 %, indicating that the logarithmic resistance representation provides a suitable description of the thermal behavior of the ceramic heater.




Fig. 8. Fitted $\ln(R_H)$ -temperature models for the PTC heater.

As shown in Fig. 8, the fitted models follow the general resistance-temperature trend of the PTC heater. However, the experimental resistance values in the lower-temperature region present fluctuations that are not fully captured by the polynomial model. These variations may be associated with the limited resolution of the current sensor or with the presence of outliers in this operating region. As the temperature increases, the resistance behavior becomes less fluctuating, and the model follows the experimental data more closely. An important aspect for temperature reconstruction is the slope of the resistance-temperature curve, since similar resistance values may appear at different temperatures on opposite sides of the Curie point. Therefore, identifying whether the resistance is increasing or decreasing is necessary to determine the operating region of the PTC heater. An alternative is to restrict the operating range to one side of the Curie temperature to avoid ambiguity in the resistance-temperature relation.

The PTC temperature was reconstructed analytically by inverting the fitted $\ln(R_H)$ -temperature models. The quadratic formula was applied to the pre-Curie model, selecting the root satisfying $T_{\text{PTC}} < 193^\circ\text{C}$, whereas the post-Curie temperature was obtained by directly rearranging the linear model. The resulting inverse relation is given by

$$T_{\text{PTC}} = \begin{cases} 193 + \frac{2.72 \times 10^{-4} - \sqrt{D(R_H)}}{2.44 \times 10^{-5}}, & q = 0, \\ 193 + \frac{\ln(R_H) - 2.288}{5.44 \times 10^{-2}}, & q = 1, \end{cases} \quad (7)$$

$$D(R_H) = 4.88 \times 10^{-5} \ln(R_H) - 1.112 \times 10^{-4} \quad (8)$$

Eq. (7) provides the reconstructed PTC temperature according to the active operating region, where $q = 0$ represents the pre-Curie state, and $q = 1$ represents the post-Curie state. For the pre-Curie branch, Eq. (8) defines the discriminant obtained after rearranging the quadratic $\ln(R_H)$ -temperature model. The negative quadratic root was selected because it produces temperatures below 193°C , while the post-Curie branch was obtained directly from the inverse of the corresponding linear model.

Eq. (7) was implemented through a two-state machine initialized in the pre-Curie state. The transition to the post-Curie state was enabled when $V \geq 8\text{ V}$, since 8 V was the lowest tested excitation capable of reaching the Curie region under prolonged heating, whereas the lower voltage levels remained below 193°C . The transition was confirmed by detecting the minimum of $\ln(R_H)$ through a change in its discrete gradient from negative to positive. A subsequent minimum-crossing event while $q = 1$ returned the machine to the pre-Curie state.

During intervals in which the primary heating voltage was deactivated, 2 V sensing pulses with a duration of 20 ms were applied at 5 Hz . This provided one resistance measurement every 200 ms with a 10% duty cycle, reducing sensing-induced self-heating while maintaining sufficient temporal resolution. At each pulse, the resistance was calculated as $R_H = V_H/I_H$, and the temperature was updated using the branch of Eq. (7) selected by the state machine. Between consecutive measurements, the last estimated temperature was retained using a zero-order hold. When no previous temperature-state history was available, the sensing cycle was initialized only after the unpowered heater reached thermal equilibrium with the environment, ensuring a reliable initialization in the pre-Curie state.

Fig. 9 compares the measured PTC temperature with the temperature reconstructed from the measured resistance for the different excitation-voltage levels. Considering all cases, the average fit between the measured and estimated temperature responses was 82.6% . The lowest performance was obtained at 4 V , with a fit of 75.97% , which is consistent with the fluctuations observed in the low-temperature resistance data and their limited representation by the polynomial model. The second-lowest fit was obtained at 10 V , with 79.40% . In this case, the largest mismatch appeared near the final portion of the PRBS signal, where the PTC heater approached its saturation region above the Curie temperature. This difference may be associated with the transition between the pre-Curie and post-Curie fitting functions. Despite these deviations, the reconstructed temperature

followed the overall experimental trend, with better agreement observed for the 6 V and 12 V cases.




Fig. 9. Measured and resistance-estimated PTC temperature under PRBS excitation.

The heat propagation model was evaluated at the farthest measured point of the SMA segment, located 5.2 mm from the anchoring point. This position corresponds to the transition between the curved region of the U-shaped SMA segment and the straight SMA legs. This location was selected because it represents the region where the temperature transferred from the PTC reaches the portion of the SMA wire expected to participate in the shape-memory response. Fig. 10 shows the temperature response for the different excitation voltage levels. In each case, the dotted line represents the measured PTC temperature used as the input of the differential equation, the red line represents the temperature estimated by the propagation model, and the blue line represents the temperature measured directly with the thermal camera at the SMA point of interest.




Fig. 10. Measured and estimated SMA temperature at $x = 5.2$ mm using the heat propagation model.

The comparison between the measured and estimated SMA temperatures shows that the distance-dependent first-order model captures the main thermal behavior at 5.2 mm. Considering all voltage levels, the average fit between the measured and estimated propagated temperature responses was

93.03 %. In contrast to the PTC temperature reconstruction, the fit values of the propagation model did not vary markedly among the different voltage cases. Additionally, the steady-state temperature values were captured correctly by the model, including regions where the PTC temperature reconstruction presented larger discrepancies near saturation.

The propagation model achieved the highest fit but depends on the reconstructed PTC temperature; therefore, errors in the resistance-temperature model propagate to the SMA estimate. A lower-range, higher-resolution current sensor could improve reconstruction accuracy because the measured current remained below 1.2 A, whereas the sensor range was 5 A.

5. Conclusions

This work presented a resistance-based thermal characterization framework for a silicone-embedded SMA actuator heated by a ceramic PTC element. The PTC heater was characterized through its resistance-temperature behavior using the transformed variable $\ln(R_H)$, allowing the heater temperature to be reconstructed from electrical measurements. The resulting model achieved an average fit of 82.6 % across the different excitation voltage levels.

Heat propagation from the PTC heater to the SMA segment was characterized using temperature measurements acquired at different distances from the anchoring point. A distance-dependent first-order differential equation was obtained to describe the propagated temperature along the SMA wire. At 5.2 mm, the model achieved an average fit of 93.03 %, showing that the proposed formulation describes the main thermal dynamics of the SMA segment under localized PTC heating.

Overall, the PTC resistance-temperature characterization and the heat-propagation model provide a thermal-sensing alternative for silicone-embedded SMA actuators. This approach reduces the dependence on multiple contact-based temperature sensors along the SMA wire and establishes a basis for thermal monitoring in PTC-heated SMA systems.

References

- [1]. N. El-Atab, R. B. Mishra, F. Al-Modaf, L. Joharji, et al., Soft actuators for soft robotic applications: A review, *Advanced Intelligent Systems*, Vol. 2, Issue 10, 2020, 2000128.
- [2]. J. M. Jani, M. Leary, A. Subic, M. A. Gibson, A review of shape memory alloy research, applications and opportunities, *Materials & Design*, Vol. 56, 2014, pp. 1078-1113.
- [3]. M.-S. Kim, J.-K. Heo, H. Rodrigue, H.-T. Lee, et al., Shape memory alloy (SMA) actuators: The role of material, form, and scaling effects, *Advanced Materials*, Vol. 35, Issue 33, 2023, 2208517.
- [4]. C. Liang, C. A. Rogers, One-dimensional thermomechanical constitutive relations for shape

- memory materials, *Journal of Intelligent Material Systems and Structures*, Vol. 1, Issue 2, 1990, pp. 207-234.
- [5]. L. C. Brinson, One-dimensional constitutive behavior of shape memory alloys: Thermomechanical derivation with non-constant material functions and redefined martensite internal variable, *Journal of Intelligent Material Systems and Structures*, Vol. 4, Issue 2, 1993, pp. 229-242.
- [6]. S.-Y. Lee, G.-Y. Lee, Largely deformable torsional soft morphing actuator created by twisted shape memory alloy wire and its application to a soft morphing wing, *Scientific Reports*, Vol. 13, 2023, 17629.
- [7]. M. S. Kalairaj, H. Banerjee, C. M. Lim, P. Chen, et al., Hydrogel-matrix encapsulated Nitinol actuation with self-cooling mechanism, *RSC Advances*, Vol. 9, Issue 59, 2019, pp. 34244-34255.
- [8]. J. G. Bell, T. Graule, M. Stuer, Barium titanate-based thermistors: Past achievements, state of the art, and future perspectives, *Applied Physics Reviews*, Vol. 8, Issue 3, 2021, 031318.
- [9]. A. Yu, Q. Li, D. Fan, H. Zhang, Study on positive temperature coefficient of resistivity of co-doped BaTiO₃ with Curie temperature in room temperature region, *Science China Technological Sciences*, Vol. 62, Issue 5, 2019, pp. 811-819.
- [10]. D. S. Smith, N. Ghayoub, I. Charissou, O. Bellon, et al., Transient thermal gradients in barium titanate positive temperature coefficient (PTC) thermistors, *Journal of the American Ceramic Society*, Vol. 81, Issue 7, 1998, pp. 1789-1796.
- [11]. D. Silva, J. P. Araújo, A. M. Pereira, R. P. Silva, Modeling the transient response of thermal circuits, *Applied Sciences*, Vol. 12, Issue 24, 2022, 12555.
- [12]. J. N. Davidson, D. A. Stone, M. P. Foster, Real-time prediction of power electronic device temperatures using PRBS-generated frequency-domain thermal cross coupling characteristics, *IEEE Transactions on Power Electronics*, Vol. 30, Issue 6, 2015, pp. 2950-2961.

Cost-Effective Hybrid Raman-Infrared Non-Contact Thermometry for Industrial Process Monitoring

V. Multian, C. W. Rischau and J. Teyssier

Laboratoire de Technologie Avancée (LTA), Department of Quantum Matter Physics,
Optics Platform, Quai E. Ansermet 24, Geneva, Switzerland
Tel.: + 41 22 379 35 16
E-mail: jeremie.teyssier@unige.ch

Summary: Accurate non-contact temperature measurement is essential for process monitoring and quality control across a wide range of industrial applications. Raman-based thermometry has emerged as a promising alternative to conventional infrared pyrometry, overcoming limitations related to emissivity calibration and environmental sensitivity. This work aims to develop a compact and cost-effective Raman thermometry device suitable for industrial deployment while maintaining high measurement accuracy. A novel optical architecture has been designed to investigate different strategies for reducing system cost and footprint. In this article, we present the experimental setup developed for this purpose, describe the data processing and analysis methodology used for temperature determination, and discuss preliminary results obtained with the proposed system. These first findings highlight the potential of Raman thermometry as a scalable solution for advanced industrial monitoring and control.

Keywords: Non-contact thermometry, Calibration less optical thermometry, Raman thermometry.

1. Introduction

Accurate non-contact temperature measurement is essential for real-time process monitoring and quality control in industrial applications involving high-speed moving components, hazardous environments, contaminated or hygienically sensitive surfaces, and extreme temperature conditions such as molten or cryogenic materials. Infrared pyrometry, the current standard for non-contact thermometry, suffers from several limitations, including dependence on material-specific emissivity calibration, sensitivity to environmental conditions, and constraints on measurement speed and accuracy [1]. These challenges are particularly significant in laser-assisted manufacturing processes such as welding, surface texturing, polishing, and metal additive manufacturing [2].

This work presents a Raman-based non-contact thermometry approach that exploits the temperature dependence of Raman scattering generated during laser-material interaction. Unlike conventional pyrometry, the method does not rely on emissivity-based calibration and is especially well suited to laser-assisted processes, where the processing laser can simultaneously serve as the excitation source for Raman signal generation.

This method is already implemented in temperature measurement along optical fibers [3-5]. Also called Distributed Fiber Optic Temperature Sensing (DTS), this technology provides continuous, real-time temperature monitoring over long distances using optical fibers for both sensing and data transmission. However, analyzed light travels through an optical fiber and therefore cannot be used for remote, open space non-contact measurements. On the other hand, remote Raman systems have been developed for

different purpose. Remote collection on Arctic ice has shown the possibility of collecting high-quality spectra over a distance above 100 m. [6] Ultra-fast mineral composition measurement for lunar [7] and more generally in planet exploration [8] or detection of toxic compounds [9] was also performed in a range of a few meters distance. Solutions were also already tested to perform remote Raman analysis in industrial environment [10]. Applications combining remote measurement and Stokes/anti-Stokes analysis are rare. Nasa reported using Stokes/anti-Stokes Raman signal to measure nitrogen temperature in an operating gas turbine [11].

The feasibility of this approach has been demonstrated using a laboratory-grade Raman spectrometer. The present work focuses on translating this capability into a compact and cost-effective industrial device, targeting a tenfold reduction in system cost for pilot-scale deployment and up to a hundredfold reduction for full industrial implementation, while preserving measurement accuracy. We will present a recently developed optical architecture enabling the evaluation of multiple strategies for reducing system cost and footprint. Significant research and development efforts are still required, including advanced Raman signal processing, optimized hardware design, and dedicated software implementation. The goal is to establish a scalable Raman-based thermometry solution for next-generation industrial process monitoring and control.

2. Experimental Setup

The proposed thermometry approach is based on the simultaneous analysis of Raman scattering and thermal infrared emission. Temperature information is




Fig. 2. Broad range spectroscopic view of the different parts composing the spectrum collected on silicon at 873 K on a scale axis of Raman shift and absolute wavelength. Stokes and anti-Stokes regions are indicated in green and blue respectively. Experimental data and fit based on our unified model described in Eq. (2) and (3) are shown as symbols and red line respectively. Model is extrapolated to the region of maximal thermal radiance colored red.

4. Results

The hybrid Raman–infrared thermometry concept has been experimentally validated on silicon samples heated up to approximately 873 K (600 °C) using a homemade furnace. Stokes, anti-Stokes, and thermal emission background were collected continuously during heating with an exposure time of 5s. Model is already implemented in Speqle platform (www.speqle.ch) and the fit of experimental data is shown as red line in Fig. 2. The difference in temperature between the thermocouple reading and the Raman-fitted temperature is due to the overheating induced by the laser. In the future, the system will be optimized to reduce to the minimal the power of the laser and minimize this overheating.

The temperature extracted from the fit for the full set of measured temperatures is shown in Fig. 3 and compared to the temperature of the furnace (dashed line).




Fig. 3. Temperature extracted from the fit of spectra collected on silicon compared to the furnace temperature.

This preliminary test on silicon of our prototype module shows that in the region up to 600 K, the method achieves already temperature determination with an accuracy close to 2 % on silicon and without requiring material-specific calibration.

Although the temperature estimate remains acceptable at higher temperatures, its accuracy is degraded to about 10 % by the significant noise affecting the data at longer wavelengths. This limitation could be substantially mitigated by using a more recent spectrometer with a spectral range better suited to the measurement.

5. Discussion

The system was successfully tested on silicon which has a Raman response that is much lower than plastics or most oxides so that it could be quickly adapted to other applications where pyrometry temperature measurement can become unreliable. In several industrial applications, the surface properties or measurement conditions affect infrared radiation. For example, in semiconductor manufacturing, pyrometry fails when measuring thin wafers or highly reflective surfaces because emissivity can be low, variable, or affected by coatings and surface layers. Similarly, shiny metals such as aluminum, copper, and stainless steel can reflect surrounding radiation, causing inaccurate temperature readings. Measurements can also be affected in furnaces, boilers, and high-temperature processes where flames, dust, smoke, or steam interfere with the infrared signal. Transparent or semi-transparent materials, such as glass and some polymers, create additional challenges because the pyrometer may detect radiation from inside the material rather than the true surface temperature.




Fig. 4. Spectra, fit and extracted temperatures for different laser powers on a) stainless steel and b) Metal titanium in air. Appearance of oxide layer on titanium is visible for extremely small thicknesses.

Our method is not limited to Raman-active materials. For example, temperature extraction was performed on spectra acquired using our laboratory Raman spectrometer (Horiba Evolution HR) on materials that do not exhibit phonon peaks but show only electronic Raman contributions, such as stainless steel (Fig.4a) and titanium (Fig.4b). Fit curve used for temperature extraction is shown in the graph (orange curve) as well as laser power and temperature value. As oxides signal is much stronger than the electronic background, the appearance of oxide film on the material can be easily detected at the early stage of growth.

This method may also be particularly well suited to laser-assisted processes such as welding, additive manufacturing, and texturing, as the use of a high-power laser also generates intense Raman radiation that can be exploited thanks to our module.

6. Conclusions

This work demonstrates a universal, calibration-free optical thermometry method that combines Raman scattering and infrared thermal radiation within a single analytical framework. The approach provides accurate temperature measurements across a broad range of materials and temperatures, including systems with weak or absent vibrational Raman signatures. The development of a patented dual-channel optical architecture enables substantial simplification of the instrumentation while preserving the essential temperature information contained in the Raman Stokes/anti-Stokes ratio and thermal emission background. Experimental validation on graphite, stainless steel, and titanium confirms the feasibility of the method and highlights its potential for industrial applications such as laser-assisted manufacturing and standalone non-contact thermometry. Future work will focus on optimizing reduced-data acquisition strategies and implementing real-time temperature extraction in a compact, low-cost sensing device.

Acknowledgements

The authors gratefully acknowledge the financial support provided by the InnoGAP Innovation Grant of the University of Geneva and the FIF (FONGIT Innovation Fund) grant from the FONGIT startup incubator.

References

- [1]. X. P. V. Maldague, Theory and Practice of Infrared Technology for Nondestructive Testing, Wiley, New York, 2001.
- [2]. Z. Liu, et al., Recent advances and current developments of molten pool temperature measurement for laser additive manufacturing processes, *Recent Patents on Mechanical Engineering*, Vol. 13, Issue 1, 2020, pp. 13-23.
- [3]. J. Li, M. Zhang, Physics and applications of Raman distributed optical fiber sensing, *Light: Science & Applications*, Vol. 11, 2022, 128.
- [4]. J. Wu, et al., Distributed fiber sensors with high spatial resolution in extreme radiation environments in nuclear reactor cores, *Journal of Lightwave Technology*, Vol. 39, Issue 15, 2021, pp. 4873-4883.
- [5]. T. D. Vo, et al., Chalcogenide fiber-based distributed temperature sensor with sub-centimeter spatial resolution and enhanced accuracy, *Optics Express*, Vol. 22, Issue 2, 2014, pp. 1560-1568.
- [6]. F. Rull, A. Vegas, A. Sansano, P. Sobron, Analysis of Arctic ices by remote Raman spectroscopy, *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, Vol. 80, Issue 1, 2011, pp. 148-155.
- [7]. H. Zhao, et al., Development and testing of a compact remote time-gated Raman spectrometer for in situ lunar exploration, *Remote Sensing*, Vol. 17, Issue 5, 2025, 860.
- [8]. S. M. Angel, N. R. Gomer, S. K. Sharma, C. McKay, Remote Raman spectroscopy for planetary exploration: A review, *Applied Spectroscopy*, Vol. 66, Issue 2, 2012, pp. 137-150.
- [9]. W. Ortiz-Rivera, L. C. Pacheco-Londoño, S. P. Hernández-Rivera, Remote continuous wave and pulsed laser Raman detection of chemical warfare agents simulants and toxic industrial compounds, *Sensing and Imaging: An International Journal*, Vol. 11, Issue 4, 2010, pp. 131-145.
- [10]. G. Berlanga, et al., Remote Raman spectroscopy of natural rocks, *Applied Optics*, Vol. 58, Issue 32, 2019, pp. 8971-8980.
- [11]. Y. R. Hicks, W. A. DeGroot, R. J. Locke, R. C. Anderson, Combustion temperature measurement by spontaneous Raman scattering in a Jet-A fueled gas turbine combustor sector, Technical Report NASA-TM-62693, National Aeronautics and Space Administration, 2002, <https://ui.adsabs.harvard.edu/abs/2002ntrs.rept62693H>
- [12]. M. Cardona, Resonance phenomena, in *Light Scattering in Solids II: Basic Concepts and Instrumentation* (M. Cardona, G. Güntherodt, Eds.), Springer, Berlin, 1982, pp. 19-178.
- [13]. Multian, et al., Temperature measurement device and associated temperature measurement method and processing installation, Patent EP25198991.9, European Patent Office, 2025.

Ultra-Thin Inkjet-Printed Metal/Polymer Composite Strain Sensors for Wearable Health Monitoring

A. Khan¹, A. S. Nittala², S. Ali¹, A. Hassan² and A. Bermak¹

¹ Hamad Bin Khalifa University, College of Science and Engineering, 34110 Doha, Qatar

² University of Calgary, Department of Computer Science, 2500 University Drive NW, Calgary, Canada

Tel.: + 974 55377596

E-mail: arkhan4@hbku.edu.qa

Summary: This paper evaluates previously developed, fully inkjet-printed silver/PEDOT:PSS strain sensors on ultra-thin decal-transfer substrates for monitoring throat and cheek motion. The sensor consists of a silver-nanoparticle bottom layer and a PEDOT:PSS top layer printed using a piezoelectric drop-on-demand inkjet system. Following water-assisted transfer, the sensor conforms to curved skin surfaces. Proof-of-concept measurements were obtained from a neck-mounted sensor during voluntary coughing and saliva swallowing and from a cheek-mounted sensor during controlled chewing. The demonstrations produced distinct resistance-change waveforms for the tested activities. To establish suitability for wearable monitoring, further evaluation is required using multiple participants, independent event annotation, quantitative detection metrics, comparison with thicker substrate designs, and assessment of skin compatibility, drift, sweat exposure, and long-term adhesion.

Keywords: Wearable electronics, Inkjet printing, Strain sensors, Skin-conformable sensors, Metal/polymer composite, Swallowing monitoring, Soft electronics.

1. Introduction

Soft, skin-conformable wearable sensors have attracted growing interest for the continuous, unobtrusive monitoring of human physiological and physical activities, owing to their conformal, comfortable attachment to curved and dynamic body surfaces [1, 2]. Among the various classes of soft wearable sensors, strain sensors are particularly attractive because subtle, localised skin deformations can be directly translated into rich physiological information, such as swallowing, coughing, breathing, and facial or muscular activity [3, 4].

In our earlier work, we reported a fully inkjet-printed silver/PEDOT:PSS composite strain sensor on thermoplastic polyurethane (TPU) and textile-compatible substrates, showing decent sensitivity, a wide strain range, and outstanding cyclic stability for tracking large-deformation motions such as finger and elbow bending [5]. While effective, substrates in the tens-of-micrometer range still add a noticeable mechanical mismatch with the skin, which limits their ability to capture very subtle physiological motions such as the small displacements of the thyroid cartilage during swallowing or coughing.

To address this, we previously extended our metal/polymer composite sensing concept to commercial ultra-thin decal (temporary-tattoo-type) substrates, which are only a few micrometres thick, and reported the resulting sensor architecture and decal-transfer substrate [6]. This substrate reduces the sensor's bending stiffness by orders of magnitude relative to TPU or textile carriers, enabling truly conformal, imperceptible skin attachment while preserving a fully digital, mask-less, additive inkjet-printing fabrication route that is facile, affordable, energy-efficient and readily scalable. In the

present work, we evaluate this previously developed ultra-thin sensor for two applications not explored in [6]: real-time detection of swallowing and coughing via neck-mounted sensing over the thyroid cartilage, and monitoring of chewing/eating activity via cheek-mounted sensing, highlighting the broad applicability of our low-cost soft strain sensors for fine-grained human motion and activity tracking.

2. Materials and Methods

2.1. Materials

Silver nanoparticle ink (SmartInk S-CS01130) and water-based PEDOT:PSS ink (Orgacon™ IJ-1005), both formulated for piezoelectric inkjet printing, were procured from Sigma-Aldrich, consistent with our previous formulation [5, 6]. In this work, the composite sensing film (schematic shown in Fig. 1) was printed on a commercial ultra-thin decal (tattoo-paper) substrate, comprising a water-soluble release layer, a few-micrometer-thick polymer carrier film, and a protective top sheet, enabling a simple wet-transfer of the printed sensor onto the skin.




Fig. 1. Schematic of the fully inkjet-printed silver/PEDOT:PSS composite strain sensor on an ultra-thin decal substrate mounted onto the skin.

2.2. Fabrication

The sensing patches were fabricated using a Dimatix material inkjet printer (DMP-2800, Fujifilm) with a 10-pL cartridge (DMC-11610), using printing parameters (20 kHz firing rate, 20 μ m drop spacing) optimised in our earlier work to minimise the coffee-ring effect [5, 6]. The bottom silver film was printed first and thermally cured at 100 $^{\circ}$ C, followed by printing of the PEDOT:PSS top film, which was thermally dried/annealed at lower temperature (80 $^{\circ}$ C) to remove residual solvent and consolidate a continuous, conductive film, forming the metal/polymer composite sensing film illustrated schematically in Fig. 2. The printed patch, still on its decal carrier, was then conformally transferred onto the target skin site (neck or cheek) with water, after which the water-soluble release layer dissolved, leaving the ultra-thin composite sensor adhered to the skin via an adhesive film. This adhesive film electrically insulates the silver nanoparticle layer from the skin, while the decal carrier substrate forms the outward-facing, environment-exposed surface of the sensor; no additional encapsulation layer was applied.




Fig. 2. Schematic of the fabrication process of the fully inkjet-printed silver/PEDOT:PSS composite strain sensor on an ultra-thin decal substrate.

3. Results and Discussion

3.1. Electromechanical Performance

Owing to the hybrid microstructure combining the high electrical conductivity of the silver film with the intrinsic stretchability of the PEDOT:PSS top layer, the ultra-thin composite sensor exhibits a clear, repeatable relative resistance change, $\Delta R/R_0$ (where R_0 is the initial, unstrained resistance of the sensor and ΔR is the resulting change in resistance), under applied tensile strain, together with a wide operating strain range and reasonable gauge sensitivity, in line with the performance we previously reported for this composite system on thicker substrates [5]. Repeated stretch-release cycling confirms excellent mechanical and electrical durability of the ultra-thin sensor, with stable signal amplitude maintained over repeated loading, indicating that migrating to the ultra-thin decal substrate does not compromise the robustness of the sensing film while substantially improving skin conformability.

3.2. Wearable Sensing Demonstrations

The ultra-thin sensor's high skin conformability was exploited to capture subtle physiological motions

that are difficult to resolve with thicker sensor platforms. First, the sensor patch was conformally attached over the thyroid cartilage on the neck to monitor swallowing and coughing (Fig. 3). During swallowing, the upward-downward translation of the thyroid cartilage produces a distinct, relatively slow and larger-amplitude resistance-change waveform, whereas coughing generates a sharper, more transient resistance spike. The recorded real-time traces clearly and repeatably distinguish these two events, demonstrating the sensor's capability for throat-motion-based physiological monitoring.




Fig. 3. Real-time relative resistance response ($\Delta R/R_0$) of the neck-mounted ultra-thin strain sensor during repeated swallowing and coughing events, with the sensor placement over the thyroid cartilage shown in the schematic.

Second, the sensor was conformally attached to the cheek to monitor subtle muscle movements associated with chewing and eating (Fig. 4). The recorded real-time data clearly show a distinct, repeatable resistance-change pattern corresponding to each chewing cycle, confirming that the ultra-thin sensor can reliably capture the sensor wearer's eating activity. Such continuous, unobtrusive eating-pattern monitoring is relevant to applications in dietary tracking and the maintenance of a healthy lifestyle.




Fig. 4. Real-time relative resistance response ($\Delta R/R_0$) of the cheek-mounted ultra-thin strain sensor during repeated chewing/eating cycles, with the sensor placement on the cheek shown in the schematic.

Taken together, these results confirm that combining our metal/polymer composite sensing film with an ultra-thin decal substrate preserves the sensitivity, stretchability and durability of the sensing

film while substantially extending its applicability to the detection of fine, low-amplitude skin deformations that are relevant to everyday health and activity monitoring.

4. Conclusions

We demonstrated a fully inkjet-printed, ultra-thin, skin-conformable strain sensor based on a silver/PEDOT:PSS metal/polymer composite film, fabricated on a commercial decal substrate via a facile, fully digital, affordable, energy-efficient and scalable process. The sensor retains excellent sensitivity, stretchability and mechanical durability, while its improved conformability enables reliable detection of subtle physiological motions. Worn on the neck, it clearly and repeatably detected swallowing and coughing; worn on the cheek, it clearly captured chewing/eating patterns in real time. These results demonstrate the strong potential of our low-cost, ultra-thin soft strain sensors for wide-ranging, unobtrusive human physiological and activity monitoring.

Acknowledgements

We acknowledge the support provided by the Qatar Research Development and Innovation (QRDI) Council under Grant ARG01-0522-230274. The content is solely the responsibility of the authors and does not represent the official views of QRDI Council.

We also acknowledge the support provided by HBKU under the Joint Research Agreement in “*Digital Health*” with City University of Hong Kong (City UHK).

References

- [1]. M. Lin, H. Hu, S. Zhou, S. Xu, Soft wearable devices for deep-tissue sensing, *Nature Reviews Materials*, Vol. 7, Issue 11, 2022, pp. 850-869.
- [2]. A. S. Nittala, A. Khan, K. Kruttwig, T. Kraus, et al., PhysioSkin: Rapid fabrication of skin-conformal physiological interfaces, in *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, 2020, 287.
- [3]. S. Ali, S. Khan, A. Khan, A. Bermak, Developing conductive fabric threads for human respiratory rate monitoring, *IEEE Sensors Journal*, Vol. 21, Issue 4, 2021, pp. 4350-4356.
- [4]. Q. Liu, J. Chen, Y. Li, G. Shi, High-performance strain sensors with fish-scale-like graphene-sensing layers for full-range detection of human motions, *ACS Nano*, Vol. 10, Issue 8, 2016, pp. 7901-7906.
- [5]. A. Khan, S. Ali, A. S. Nittala, K. Riaz, et al., Fully inkjet-printed soft wearable strain sensors based on metal/polymer composite sensing films, in *Proceedings of the IEEE International Conference on Flexible and Printable Sensors and Systems (FLEPS)*, 2023, pp. 1-4.
- [6]. A. Khan, S. Ali, S. Khan, A. Bermak, Ultra-thin and skin-conformable strain sensors fabricated by inkjet printing for soft wearable electronics, in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, 2022, pp. 1759-1762.

Atmospheric Hydrometeor Monitoring Using 24 GHz Microwave Attenuation

B. H. Kwon and Z. Seo

Pukyong National University, Major of Environment Atmospheric Sciences,
45 Yongso-ro Nam-gu, Busan, Korea
Tel.: +82 10 64789733
E-mail: bhwon@pknu.ac.kr

Summary: This study investigates the use of a 24 GHz microwave link (ML) for path-integrated rainfall detection and evapotranspiration (ET) estimation. A 237.3 m link was operated at the Boseong Standard Weather Observatory, and its received signal level was evaluated using rainfall observations from automatic weather stations and an Automated Surface Observing System, together with ET derived from an eddy-covariance system. For rainfall detection, a clear-sky reference was established using a moving-window variance method, and variations in total path attenuation were compared with 1 min rainfall observations. During the event on 2 July 2024, the ML reproduced the observed rainfall duration and achieved a detection rate exceeding 80%. For ET estimation, specific residual attenuation during selected rain-free daytime periods was empirically related to eddy-covariance ET. Leave-one-day-out cross-validation yielded correlation coefficients above 0.97 and an RMSE of 0.05 mm h^{-1} , while a 3 h moving average improved the agreement to $R = 0.99$ and $\text{RMSE} = 0.03 \text{ mm h}^{-1}$. These results demonstrate the potential of a short 24 GHz ML to complement conventional observations of rainfall and land-atmosphere water exchange. Because gaseous absorption, refractivity fluctuations, antenna conditions, and equipment drift may also influence the signal, the ET retrieval should be regarded as a site-calibrated empirical relationship.

Keywords: Microwave link, Rainfall intensity, Evapotranspiration, Latent heat flux, Eddy correlation

1. Introduction

Localized heavy rainfall and urban flash floods cause substantial human and economic losses, creating a need for observations that complement radar and rain-gauge networks. Since Atlas and Ulbrich [1] demonstrated path-integrated rainfall measurement from microwave attenuation, microwave links (MLs) have been widely investigated as near-surface rainfall sensors. Their line-of-sight measurements provide path-averaged information at spatial scales between point gauges and weather radar. Recent studies have further developed statistical rainfall retrievals and quantified wet-antenna effects, which remain an important source of uncertainty [3,4].

Kim and Kwon [2] estimated rainfall using eight ML paths operated by Korea Telecom in Seoul. For Case 4 of that study, the rainfall detection rate was 83.5%, and the estimated rainfall intensity had a correlation coefficient of 0.90, an RMSE of 1.5 mm h^{-1} , and a mean bias of -0.5 mm h^{-1} ; the other validation case produced different performance statistics. These values therefore describe that specific earlier case and not the present Boseong experiment.

Microwave links are also sensitive to non-rainfall processes. A previous 24 GHz study investigated the feasibility of detecting fog and sea fog over an 11.4 km coastal path [5], while gaseous absorption — particularly water-vapour absorption near 22.235 GHz — varies with frequency, temperature, pressure, and humidity [6]. Microwave scintillometry can additionally provide path-integrated information related to turbulent heat and water-vapour fluxes, although separating temperature and humidity

contributions requires supplementary information and introduces uncertainty [7].

In contrast to the earlier fog-focused experiment [5], the present study uses a short 237.3 m link to evaluate two applications supported by the available validation data: rainfall detection and an empirical relationship between clear-day link attenuation and eddy-covariance ET. Fog classification is not evaluated in the present dataset.

2. Method

The experimental ML operated at 24 GHz with an initial transmitter output of 12 dBm. The transmitter was installed at a height of 20 m on the meteorological tower at the Boseong Standard Weather Observatory, and the receiver was located approximately 237.3 m southwest of the tower. Rainfall measurements from the Boseong and Beolgyo Automatic Weather Stations (AWSs) and the Automated Surface Observing System (ASOS) were used for comparison and validation.

Received signal level was converted to total path attenuation, $A_{\text{tot}}(t)$ [dB], relative to a reference signal level selected from stable rain-free periods using a moving-window variance method. Conceptually, $A_{\text{tot}}(t) = A_{\text{rain}}(t) + A_{\text{wet}}(t) + A_{\text{gas}}(t) + A_{\text{lw}}(t) + A_{\text{sys}}(t)$, where A_{rain} is rain-induced attenuation, A_{wet} is wet-antenna attenuation, A_{gas} is gaseous attenuation associated mainly with oxygen and water vapour, A_{lw} is attenuation by suspended liquid water such as fog or low cloud, and A_{sys} represents instrumental drift and other residual system effects; every term is expressed in decibels [4,6]. Rainfall

periods were identified using the local-variance statistic and checked against AWS/ASOS observations. Because the present analysis contains no independent visibility or liquid-water validation, A_{lw} is not retrieved separately and no fog-detection accuracy is reported.

For ET analysis, rainfall periods and unstable signal intervals were excluded, and a daily reference was defined from the upper 99th percentile of the received signal level. The resulting specific residual attenuation [dB km^{-1}] was related empirically to ET [mm h^{-1}] derived from latent heat flux measured by an eddy-covariance system. Linear and power-law regressions were evaluated using leave-one-day-out cross-validation. Water-vapour absorption, temperature- and humidity-dependent refractivity and scintillation, pressure, antenna effects, and equipment drift were not interpreted as individually separated ET components; instead, they form part of the calibrated clear-sky background and retrieval uncertainty [6,7].

3. Results

Microwave-link rainfall estimates were compared with 1 min ASOS rainfall observations. During the event on 2 July 2024, the ML-derived rainfall period closely followed the observed rainfall duration. The ratio of ML-detected rainfall intervals to ASOS-observed rainfall intervals exceeded 80%, demonstrating the ability of the short link to detect this rainfall event.




Fig. 1. Total path attenuation (black), local-variance test statistic (blue), and ML-detected rainfall periods (green).

Under selected rain-free daytime conditions, the 24 GHz link showed a repeatable diurnal signal cycle. Residual attenuation increased after sunrise, peaked near noon, and decreased after sunset. The upper 99th percentile of the daily received signal level defined the baseline, and specific residual attenuation was related to eddy-covariance ET using linear and power-law regression. This procedure yields an empirical, site-calibrated ET proxy and does not imply that ET is the sole physical cause of the measured attenuation.

Leave-one-day-out cross-validation showed strong agreement, with correlation coefficients above 0.97 and an RMSE of 0.05 mm h^{-1} . Applying a 3 h moving average improved the agreement to $R = 0.99$ and $\text{RMSE} = 0.03 \text{ mm h}^{-1}$. These statistics demonstrate predictive consistency within the observed dataset, while independent campaigns are required to test transferability across seasons, sites, and atmospheric conditions.




Fig. 2. Comparison of ET derived from the eddy-covariance system (ET_{EC}) and estimated from the microwave link (ET_{ML}).

4. Conclusions

The 24 GHz ML detected the duration of the selected rainfall event with a detection rate above 80% and provided a site-calibrated empirical estimate of clear-day ET that agreed closely with eddy-covariance observations. The results support the potential of short microwave links to supplement conventional near-surface hydrometeorological observations. However, the ET relationship may include contributions from gaseous absorption, refractivity fluctuations, antenna conditions, and instrumental drift. Future work should incorporate explicit gas-attenuation corrections, temperature-pressure-humidity controls, wet-antenna characterization, independent-site validation, and visibility/liquid-water observations before extending the method to fog detection.

Acknowledgements

This work was funded by the Korea Meteorological Administration Research and Development Program under Grant RS-2024-00404042.

References

- [1]. D. Atlas, C. W. Ulbrich, Path- and area-integrated rainfall measurement by microwave attenuation in the 1–3 cm band, *Journal of Applied Meteorology*, Vol. 16, Issue 12, 1977, pp. 1322–1331.

- [2]. M.-S. Kim, B. H. Kwon, Rainfall detection and rainfall rate estimation using microwave attenuation, *Atmosphere*, Vol. 9, Issue 8, 2018, Article 287.
- [3]. X. Zheng, H. Messer, Y. Qin, Y. Song, et al., A statistical approach to retrieve rainfall intensity by attenuation measurements from commercial microwave links, *Earth's Future*, Vol. 11, Issue 11, 2023, Article e2023EF003718.
- [4]. J. Tiede, C. Chwala, U. Siart, New insights into the dynamics of wet antenna attenuation based on in situ estimations provided by the dedicated field experiment ATTRRA2, *IEEE Geoscience and Remote Sensing Letters*, Vol. 20, 2023, Article 1002305, pp. 1–5.
- [5]. Z. Seo, B. H. Kwon, S. J. Kim, K. H. Lee, et al., Fog detection based on microwave signal attenuation, in *EMS Annual Meeting 2024*, Barcelona, Spain, 1–6 September 2024, EMS2024-624.
- [6]. International Telecommunication Union, Attenuation by atmospheric gases and related effects, Recommendation ITU-R P.676-13, *International Telecommunication Union*, Geneva, 2022.
- [7]. L. D. van der Valk, O. K. Hartogensis, M. Coenders-Gerrits, R. W. Hut, et al., Use of commercial microwave links as scintillometers: Potential and limitations towards evaporation estimation, *Atmospheric Measurement Techniques*, Vol. 18, Issue 21, 2025, pp. 6143–6165.

Wearable Heartbeat Sensors from Laser-Induced Graphene on a Poly(Dimethylsiloxane) Substrate with Poly(Ethylene Glycol) and Triton X Additives

M. Spasenović, A. Gavran, M. Tomić, I. Pešić, T. Vičentić and M. V. Pergal

Institute of Chemistry, Technology and Metallurgy, National Institute of the Republic of Serbia,
Belgrade, Serbia

Tel.: + 381 11 3640 232

E-mail: marko.spasenovic@ihtm.bg.ac.rs

Summary: Wearable sensors are a growing field of industry, with applications in sports and medical technology. Such sensors need to be thin, lightweight, and flexible, while reliably detecting physiological signals such as heartbeat, respiration rate, or other parameters. Laser-induced graphene is an excellent candidate active material for such sensors; however it needs to be manufactured on an appropriate polymeric substrate that conforms to demands of wearable technology. Poly(dimethylsiloxane) is an excellent polymer for wearable use; however it lacks the carbon structure that would make it susceptible to laser induction of graphene. We modify poly(dimethylsiloxane) with poly(ethylene glycol) and Triton X-100, subsequently inducing graphene on its surface and making heartbeat sensors. We perform physicochemical characterization of the devices, confirming excellent structural, mechanical and electrical performance, and heartbeat detection with a signal-to-noise ratio of over 26 dB.

Keywords: Graphene, Polymers, Wearable sensor, Heartbeat sensor, Poly(dimethylsiloxane).

1. Introduction

In recent years, substantial progress has been achieved in the field of wearable sensor technology, primarily driven by the escalating demand for portable electronic systems that are both stretchable and flexible, thereby offering viable alternatives to conventional rigid devices. These advancements have facilitated the development of epidermal patches capable of monitoring a wide range of physiological parameters, including cardiac activity, respiratory function, blood pressure, vocalization, and limb kinematics, with significant implications for medical diagnostics and athletic performance [1–3]. To satisfy the stringent requirements imposed on sensors designed for direct skin contact, constituent materials must exhibit minimal thickness, mechanical compliance, and favourable electrical characteristics, particularly high conductivity [4]. Among the candidate materials, laser-induced graphene (LIG) has emerged as a promising option. The laser induction process garnered considerable attention in 2014 when polyimide (PI) was successfully transformed into porous graphene via CO₂ laser irradiation [5]. This method provides a highly controlled approach for converting carbon-rich polymer domains into graphene, yielding a graphene/polymer heterostructure in which the graphene layer functions as the active sensing element, while the polymer substrate ensures mechanical flexibility and structural stability.

Commercial polyimide (PI) films represent the most widely employed substrates for the fabrication of high-quality LIG, primarily due to their excellent mechanical robustness and high thermal stability (up to 400 °C) [6–8]. PI has been extensively utilized in biomedical contexts, including implantable

microelectrodes, flexible neural interfaces, and microfluidic systems, owing to its favourable thermal endurance, mechanical integrity, and chemical resistance. Despite these advantages, PI often necessitates additional surface modifications or coatings to improve cellular compatibility and mitigate inflammatory responses. Furthermore, its intrinsic low elasticity limits applications requiring mechanical compliance [9–11]. Consequently, LIG is frequently transferred from PI substrates onto alternative polymers to enhance flexibility and stretchability [3, 12, 13]. Polymers with structural similarity to PI, such as polyetherimide and polysulfone, have also been reported as suitable matrices for LIG generation [5, 14–17]. Beyond synthetic polymers, LIG has been successfully produced from diverse carbon-rich materials including wood, textiles, paper, food products, and alginate [18–21]. Nevertheless, the development of LIG/polymer heterostructures that simultaneously achieve biocompatibility, electrical conductivity, and mechanical stretchability remains a significant challenge [21, 22].

Poly(dimethylsiloxane) (PDMS), a silicon-based elastomer, has emerged as a promising candidate in this regard. PDMS is inherently biocompatible and highly stretchable, making it particularly attractive for wearable and biomedical applications. Its advantageous properties include excellent mechanical flexibility, optical transparency, tunable surface chemistry, low water permeability, and remarkable stability under thermal, oxidative, and ultraviolet conditions, coupled with high gas permeability [23]. Owing to its low cost, biocompatibility, and unique surface characteristics, PDMS has been widely adopted in microfluidics, flexible electronics, and tissue engineering [24, 25]. Pure PDMS, however, is

not a good substrate for laser induction of graphene, due to low carbon content. Researchers are exploring avenues to enhance PDMS-based laser induction by adding carbon-containing compounds, such as poly(ether-ether-ketone) (PEEK), triethylene glycol (TEG), poly(ethylene glycol) (PEG), polyimide particles, and Triton X-100 [23, 26–29]. We have added Triton X-100 to PDMS in our earlier work [27], and although we have shown that this additive makes PDMS susceptible to laser induction of graphene, the sheet resistances of the resulting materials were very large. We have also added PEG to PDMS to make it susceptible to induction [29], however the LIG thus made had poor adhesion to the substrate. Here we introduce high-molar-mass PEG ($M_n = 20000$ g/mol) and Triton X-100 (TX) together as additives to PDMS to make the material susceptible to direct laser induction of graphene. The CO_2 laser used to prepare LIG was a DBK FL-350 with a wavelength of $10.6\ \mu m$. The maximum power output of the laser is 60 W, and the resolution was fixed at 1200 DPI. We used a reduced laser power of 9 W, with a scanning speed of 40 mm/s. The laser was kept in focus, in ambient air conditions. PEG and TX were added as carbon and aromatic ring sources, at different weight percentages. At least three separate samples were made for each weight percentage. LIG rectangles with dimensions 1×2 cm were made. We characterize the composite material with physicochemical methods, prior to applying the PDMS-PEG-TX-LIG composite as a wearable sensor of heartbeat.

2. Sample Characterization

2.1. Mechanical Properties

Mechanical properties are among the most important for wearable sensors. We perform full mechanical testing on a Universal Testing Machine (Shimadzu AGS-X, Japan), retrieving information such as tensile strength, elongation at break, and Young's modulus. Fig. 1 depicts results of tensile testing of devices with 30 % wt.% PEG in a PDMS matrix, with 10 % TX content wt.%. The figure represents results from three different samples produced in the same manner. Elongation at break lies between 129 % and 143 %, indicating a highly stretchable material, which is favourable for wearable devices. Table 1 presents data from the three samples, pointing to relatively consistent fabrication results across samples, and mechanical properties that satisfy requirements for wearable devices.

2.2. Raman Spectroscopy

The Raman spectra of the materials were recorded using a DXR Raman microscope (Thermo Fisher Scientific, Waltham, MA, USA) with a laser wavelength of $\lambda_i = 532$ nm and a power of 2 mW, focused onto a $2.1\ \mu m$ spot on the surface. A Raman

spectrum is shown in Fig. 2. All the characteristic peaks of graphene are present in the sample. The ratio of intensities of the D peak to the G peak is 0.38, which indicates graphene of moderate defect density. Such a defect content has been shown to have good performance in sensing [3].




Fig. 1. Tensile testing.

Table 1. Mechanical properties of three samples.

Measurement	Tensile Strength (MPa)	Elongation at break (%)	Young's modulus (MPa)
1.	1.01	129.04	0.747
2.	0.82	142.37	0.477
3.	0.9	136.45	0.591




Fig. 2. Raman spectrum.

2.3. Water Contact Angle

Hydrophobicity was investigated through surface tension measurements with a contact angle goniometer (Ossila, L2004A1, UK) using the sessile drop method to determine the water contact angle (Fig. 3). Contact angles were measured 3 times, and the median was found. Prior to inducing graphene on the substrate, the contact angle was 65° for this material. After induction of LIG, the contact angle was 127° . An angle of 65° indicates a hydrophilic surface, on which water can spread. After induction of graphene, the surface

becomes hydrophobic, i.e. the surface is not wettable, approaching the superhydrophobic limit (150°).




Fig. 3. Water contact angle measurements: a) without and b) with graphene.

2.4. X-Ray Diffractography

X-Ray diffractometry (XRD) was conducted using a Rigaku Ultima IV diffractometer (Japan). The X-ray beam was nickel-filtered $\text{CuK}_{\alpha 1}$ radiation ($\lambda = 0.154$ nm) operating at 40 kV and 40 mA. XRD data were collected from 3 to 50° (2θ) at a scanning rate of 2° min^{-1} , with a step size of 0.02° . To provide high-intensity, high-resolution measurements, parallel beam geometry and the D/teX Ultra, a high-speed one-dimensional X-ray detector, were used. In the LIG films on PDMS/PEG/Triton X-100, the diffraction peaks characteristic of graphene were positioned at slightly lower angles than reported in the literature, most likely due to interference with the substrate materials, as seen in Fig. 4. The diffraction peak at 26° , which corresponds to the graphene peak, overlaps with an amorphous halo at 22° , which is characteristic of PEG. The diffraction maximum at $\sim 35^\circ$ is associated with the (111) in-plane structure of the β -SiC crystallite, indicating laser modifications of PDMS in SiC.




Fig. 4. X-ray diffractogram.

3. Heartbeat Sensing

The graphene device is used as a sensor of heartbeat. To this end, the PDMS-PEG-TX-LIG composite device is attached to the wrist of a heartbeat

simulator arm [30], as depicted in Fig. 5a). The device is contacted by sputtering titanium contacts and connecting copper wires with silver paste to the contacts. The device is held fixed to the simulator with adhesive copper tape. The two contacts from the LIG device are routed to a Keithley 2450 SMU, which provides a current of $100 \mu\text{A}$ through the device, while measuring the voltage across the device. The sheet resistance acquired on these samples is on the order of $1 \text{ k}\Omega$. A computer is used to record the voltage signal from the device in real time, as in Fig. 5b). As the simulated heartbeat causes the wrist to pulse, the LIG device bends and straightens with the pulse, causing the voltage to change in real time accordingly. We use the peak prominence method, implemented in Python, to identify systolic and diastolic peaks of the heartbeat, as depicted in Fig. 6. It is evident that the signal from the LIG sensor is clear, tracking the generated pulse with high fidelity. The signal-to-noise ratio (SNR) is over 26 dB, indicating very low background noise. The number of rejected peaks, which are artefacts of the sensing method, is only 5 in this time segment, compared to 85 accepted peaks. The results are comparable to those obtained with LIG on polyurethane-based devices [31].




Fig. 5. a) Sensor attached to pulse simulator; b) pulse simulation measurement setup.

4. Conclusions

In conclusion, we have successfully made PDMS susceptible to laser induction of graphene by adding PEG and Triton X-100 to the PDMS. The composite structure is shown to be stretchable, with decent electrical sheet resistance, good hydrophobicity, and excellent sensor performance. Compared to our earlier work with PDMS/Triton X-100 only [27], the LIG described in this paper is of a higher quality, as demonstrated by the high SNR. Compared to our work with PDMS/PEG only, the present work improves on the adhesion of the LIG, hence enabling usability in devices. This result is an indication that biocompatible, easily moldable, widely accessible PDMS can be modified into a substrate for LIG-based wearable sensors.




Fig. 6. Response of sensor to simulated heartbeat.

Acknowledgements

This research was supported by the Science Fund of the Republic of Serbia, #4950, Polymer/graphene heterostructures for physiological sensors – Polygraph. The authors would like to thank the Ministry of Science, Technological Development, and Innovation of Republic of Serbia (Contract No: 451-03-33/2026-03/200026).

References

- [1]. H. Wang, Z. Zhao, P. Liu, X. Guo, A soft and stretchable electronics using laser-induced graphene on polyimide/PDMS composite substrate, *npj Flexible Electronics*, Vol. 6, Issue 1, 2022, 26.
- [2]. A. Dallinger, K. Keller, H. Fitzek, F. Greco, Stretchable and skin-conformable conductors based on polyurethane/laser-induced graphene, *ACS Applied Materials & Interfaces*, Vol. 12, Issue 17, 2020, pp. 19855-19865.
- [3]. T. Vičentić, et al., Laser-induced graphene for heartbeat monitoring with HeartPy analysis, *Sensors*, Vol. 22, Issue 17, 2022, 6326.
- [4]. M. M. Rodgers, V. M. Pai, R. S. Conroy, Recent advances in wearable sensors for health monitoring, *IEEE Sensors Journal*, Vol. 15, Issue 6, 2015, pp. 3119-3126.
- [5]. J. Lin, et al., Laser-induced porous graphene films from commercial polymers, *Nature Communications*, Vol. 5, Issue 1, 2014, 3142.
- [6]. A. F. Carvalho, A. J. Fernandes, R. Martins, E. Fortunato, et al., Laser-induced graphene piezoresistive sensors synthesized directly on cork insoles for gait analysis, *Advanced Materials Technologies*, Vol. 5, Issue 12, 2020, 2000630.
- [7]. E. R. Mamleyev, et al., Laser-induced hierarchical carbon patterns on polyimide substrates for flexible urea sensors, *npj Flexible Electronics*, Vol. 3, Issue 1, 2019, 2.
- [8]. Y. Chen, et al., UV laser-induced polyimide-to-graphene conversion: Modeling, fabrication, and application, *Small Methods*, Vol. 3, Issue 10, 2019, 1900208.
- [9]. X. Jiang, et al., In vitro and in vivo evaluation of a photosensitive polyimide thin-film microelectrode array suitable for epiretinal stimulation, *Journal of NeuroEngineering and Rehabilitation*, Vol. 10, Issue 1, 2013, 48.
- [10]. T. Stieglitz, H. R. Beutel, M. Schuettler, J. U. Meyer, Micromachined, polyimide-based devices for flexible neural interfaces, *Biomedical Microdevices*, Vol. 2, Issue 4, 2000, pp. 283-294.
- [11]. Y. Y. Chen, et al., Design and fabrication of a polyimide-based microelectrode array: Application in neural recording and repeatable electrolytic lesion in rat brain, *Journal of Neuroscience Methods*, Vol. 182, Issue 1, 2009, pp. 6-16.
- [12]. L. Groo, J. Nasser, D. J. Inman, H. A. Sodano, Transfer printed laser induced graphene strain gauges for embedded sensing in fiberglass composites, *Composites Part B: Engineering*, Vol. 219, 2021, 108932.
- [13]. A. Madevska Bogdanova, et al., Blood oxygen saturation estimation with laser-induced graphene respiration sensor, *Journal of Sensors*, Vol. 2024, 2024, 4696031.
- [14]. S. Hong, J. Kim, S. Jung, J. Lee, et al., Surface morphological growth characteristics of laser-induced graphene with UV pulsed laser and sensor applications, *ACS Materials Letters*, Vol. 5, Issue 5, 2023, pp. 1261-1270.
- [15]. S. P. Singh, Y. Li, J. Zhang, J. M. Tour, et al., Sulfur-doped laser-induced porous graphene derived from polysulfone-class polymers and membranes, *ACS Nano*, Vol. 12, Issue 1, 2018, pp. 289-297.
- [16]. A. Z. Yazdi, I. O. Navas, A. Abouelmagd, U. Sundararaj, Direct creation of highly conductive laser-induced graphene nanocomposites from polymer blends, *Macromolecular Rapid Communications*, Vol. 38, Issue 13, 2017, 1700176.
- [17]. W. Song, et al., Flexible, stretchable, and transparent planar microsupercapacitors based on 3D porous laser-induced graphene, *Small*, Vol. 14, Issue 1, 2018, 1702249.
- [18]. R. Ye, et al., Laser-induced graphene formation on wood, *Advanced Materials*, Vol. 29, Issue 36, 2017, 1702211.
- [19]. Y. Chyan, et al., Laser-induced graphene by multiple lasing: Toward electronics on cloth, paper, and food, *ACS Nano*, Vol. 12, Issue 3, 2018, pp. 2176-2183.
- [20]. F. Mahmood, H. Zhang, J. Lin, C. Wan, Laser-induced graphene derived from kraft lignin for flexible supercapacitors, *ACS Omega*, Vol. 5, Issue 24, 2020, pp. 14611-14618.
- [21]. T. Vičentić, et al., Laser-induced graphene on cross-linked sodium alginate, *Nanotechnology*, Vol. 35, Issue 11, 2024, 115103.
- [22]. R. Correia, et al., Biocompatible parylene-C laser-induced graphene electrodes for

- microsupercapacitor applications, *ACS Applied Materials & Interfaces*, Vol. 14, Issue 41, 2022, pp. 46427-46438.
- [23]. P. Zaccagnini, et al., Laser-induced graphenization of PDMS as flexible electrode for microsupercapacitors, *Advanced Materials Interfaces*, Vol. 8, Issue 24, 2021, 2101046.
- [24]. D. Pinho, B. N. Muñoz-Sánchez, C. F. Anes, E. J. Vega, et al., Flexible PDMS microparticles to mimic RBCs in blood particulate analogue fluids, *Mechanics Research Communications*, Vol. 100, 2019, 103399.
- [25]. J. C. Lötters, W. Olthuis, P. H. Veltink, P. Bergveld, The mechanical properties of the rubber elastic polymer polydimethylsiloxane for sensor applications, *Journal of Micromechanics and Microengineering*, Vol. 7, Issue 3, 1997, pp. 145-147.
- [26]. L. Tang, J. Zhou, D. Zhang, B. Sheng, Laser-induced graphene electrodes on poly(ether-ether-ketone)/PDMS composite films for flexible strain and humidity sensors, *ACS Applied Nano Materials*, Vol. 6, Issue 18, 2023, pp. 17802-17813.
- [27]. M. V. Pergal, et al., Laser-induced graphene on novel crosslinked poly(dimethylsiloxane)/Triton X-100 composites for improving mechanical, electrical and hydrophobic properties, *Polymers*, Vol. 16, Issue 21, 2024, 3157.
- [28]. M. Parmeggiani, et al., PDMS/polyimide composite as an elastomeric substrate for multifunctional laser-induced graphene electrodes, *ACS Applied Materials & Interfaces*, Vol. 11, Issue 36, 2019, pp. 33221-33230.
- [29]. A. Gavran, et al., Laser-induced graphene on biocompatible PDMS/PEG composites for limb motion sensing, *Sensors*, Vol. 25, Issue 17, 2025, 5238.
- [30]. S. D. Ilić, M. Spasenović, I. Pešić, M. Rašljic-Rafajilović, et al., Characterization of pulse waveform reproducibility in a commercial pulse simulator with a flexible novel MXene-based sensor, *Facta Universitatis, Series: Electronics and Energetics*, Vol. 39, Issue 2, 2026, pp. 309-331.
- [31]. V. Vojnović, et al., Pulse sensors based on laser-induced graphene transferred to biocompatible polyurethane networks: Fabrication, transfer methods, characterization, and application, *Chemosensors*, Vol. 13, Issue 4, 2025, 122.

Design and Validation of a Time-Frequency Data Analysis and Data Processing Platform

**Y. R. Xue^{1,2,3}, S. H. Zhao^{1,2,3}, L. D. Meng^{1,2,3}, Y. F. Sun^{1,3}, Y. W. Li^{1,3}, Y. Liu^{1,2,3}
and X. H. Li^{1,2,3}**

¹ National Time Service Center, Chinese Academy of Sciences, Xi'an 710600, China

² Key Laboratory of Time Reference and Applications, Chinese Academy of Sciences, Xi'an 710600, China

³ University of Chinese Academy of Sciences, School of Optoelectronics, Beijing 100049, China, China
Tel.: + 0086 029-83890447

E-mail: xueyanrong@ntsc.ac.cn, zsh@ntsc.ac.cn, menglingda@ntsc.ac.cn

Summary: Time-frequency data analysis serves as the cornerstone of modern precision metrology. However, with the explosive growth of data scale, traditional methods face severe challenges in computational efficiency, scalability, and software usability. This paper presents a comprehensive study on time-frequency data analysis methods, based on which the Time-Frequency Data Analysis and data processing platform (TFDA) is designed, encompassing algorithms for data preprocessing, drift analysis, and variance estimation, power-law noise simulation. The computational results are compared with those from classical time-frequency data analysis software, demonstrating complete consistency. These algorithms are integrated into the TFDA application, which features a modular layered architecture and native Chinese graphical interface. Experimental validation confirms that TFDA maintains consistency with classical software in analysis results, while exhibiting significant advantages in localization support and chart customization.

Keywords: Time-frequency, Data analysis, Frequency stability, Variance, TFDA platform.

1. Introduction

High-precision time-and-frequency technology constitutes a critical infrastructure of modern society, with a wide range of application areas spanning fundamental scientific research, as well as navigation and communications in daily life. For instance, high-energy-physics experiments require picosecond-level time synchronization to accurately record particle behaviors. In navigation and communications, high-precision time and frequency serve as the foundation for positioning, navigation and timing (PNT) of the Global Navigation Satellite System (GNSS). Nevertheless, generating high-precision time-frequency signals is only the primary step in system construction. In-depth analysis and evaluation of the quality of such signals are subsequently required, which directly determine the accuracy and reliability of the overall system.

Time-frequency data analysis refers to the technology for acquiring, processing, analyzing and evaluating time and frequency signals. Its core objective is to obtain key indicators such as drift characteristics, frequency stability and dominant noise types from noise-contaminated time-frequency data. These indicators are of great significance for evaluating the performance of atomic clocks or oscillators, optimizing time-synchronization accuracy, and detecting system faults [1, 2].

In terms of software tools, although existing classic time-frequency analysis software is feature-rich (e.g., Stable32, it is a classic software dedicated to the time-and-frequency field and is widely used worldwide. As a well-established time-frequency

data-analysis tool, Stable32 features comprehensive functions covering multiple time-domain and frequency-domain analysis methods, and can meet the requirements for frequency-stability analysis, noise analysis, drift analysis and other tasks), it still has certain shortcomings in architectural design and user experience. Taking Stable32 as an example, thanks to its comprehensive algorithm library and relatively authoritative outputs, it has long been regarded as a reference tool in the time-frequency analysis field. Nevertheless, given its early development and the lack of long-term updates, it comes with an outdated interface and rather cumbersome operation procedures, which fail to meet modern users' expectations for friendly, intuitive human-computer interaction [3-5].

In respect of computing architecture, traditional software Stable32 mostly employs a single-thread execution mode and provides no support for modern computing paradigms including parallel computing. Consequently, it responds slowly when processing large-scale data and can hardly satisfy practical analysis requirements of high concurrency and high throughput. Available software tools lack modular and extensible design concepts, so it is difficult to flexibly integrate custom algorithms or optimize computing workflows for specific tasks. Besides, traditional tools show inadequate performance in visualization and interaction. For example, garbled text occurs when domestic users plot Sigma-Tau diagrams using Stable32, which restricts its application in scientific-research and engineering scenarios.

So in the field of time-frequency measurement, although various variance estimation and noise

identification theories are quite mature, domestic researchers still face many shortcomings when using traditional time-frequency data analysis software, such as the lack of Chinese localization in existing software tools, lack of support for multi-channel comparative analysis, and the gap between algorithmic research and practical deployment has not been fully bridged. This paper addresses these challenges through the following methods: (1) systematic theoretical validation with Stable32; (2) development of the TFDA platform to connect algorithmic research with practical application deployment.

2. Time-Frequency Analysis Methods

This section presents the theoretical framework and validation results of the complete algorithms in the analysis platform. All algorithms were implemented in MATLAB and verified by comparison with Stable32 computational results.

2.1. Data Preprocessing and Drift Analysis

Phase-to-frequency deviation conversion through discrete differentiation is implemented, zero-mean normalization, data averaging (phase data decimation, frequency data group averaging), configurable threshold outlier detection based on median absolute deviation (MAD), and FFT-based frequency-domain filtering (low-pass, high-pass, band-pass, band-stop) and its inverse transformation [1, 2].for example, with the identical dataset and a given measurement interval, the comparison tables between the results obtained by implementing the above-mentioned procedure in MATLAB and those processed by Stable32 are presented in Table 1 and Table 2, respectively. The comparison demonstrates that under identical input conditions, the conversion results implemented in MATLAB in this paper are numerically identical to the outputs from Stable32 at multiple discrete data points (No. 150, 300, 450, 600, 750). This finding verifies the correctness of the phase-to-frequency conversion relationship adopted in this work, normalization results, band-stop filter results and other related results.

Table 1. Comparison of Phase-to-Frequency Conversion Results.

Data Set	Matlab	Stable32
150	-9.3590739e-3	-9.3590739e-3
300	2.67009884e-2	2.67009884e-2
450	-3.01265842e-2	-3.01265842e-2
600	2.75119331e-2	2.75119331e-2
750	1.89341977e-2	1.89341977e-2

Similarly, using the same dataset, the consistency between the normalization results obtained by the

MATLAB implementation and those by Stable32 is compared in Table 3. The results show that the normalization results produced by the MATLAB implementation are identical to those of Stable32.

Table 2. Comparison of Frequency-to-Phase Conversion Results.

Data Set	Matlab	Stable32
150	-3.0994962374960e+3	-3.0994962374960e+3
300	-4.1880121275803e+3	-4.1880121275803e+3
450	-5.9485882889692e+3	-5.9485882889692e+3
600	-4.8805002860460e+3	-4.8805002860460e+3
750	-7.2860225991854e+3	-7.2860225991854e+3

Table 3. Comparison of Normalization Results.

Data Set	Matlab	Stable32
150	-2.0334470675e+0	-2.0334470675e+0
300	4.774369726e-1	4.774369726e-1
450	-9.902538341e-1	-9.902538341e-1
600	-9.500221929e-1	-9.500221929e-1
750	-2.8935857535e+0	-2.8935857535e+0

Using the same data set with the measurement interval set as specified, the maximum and minimum Fourier frequencies obtained by implementing in MATLAB are 0.5 and 9.76563×10^{-4} , respectively, which are identical to the values calculated by Stable32 for band-stop filter and others.in Fig. 1.




Fig. 1. Band-stop filter result comparison

2.2. Frequency Stability Analysis

This paper implements comprehensive variance estimation: standard Allan variance, modified Allan variance (distinguishing white and flicker phase modulation noise), Hadamard variance (insensitive to linear drift), overlapping Allan and Hadamard variants, total variance (based on reflection extension), time variance, MTIE, and TIErms [3-13]. Using a set of phase data, all variance types achieve numerical

equivalence with Stable32 at every tested averaging time, with relative errors below 10^{-15} . Allan deviation result is Fig. 2. Comparison of ADEV results between TFDA and Stable32 is given in Table 4.




Fig. 2. Allan deviation result.

Table 4. Comparison Table of ADEV Results.

Tau/s	8.64e+4	1.73e+5	3.46e+5	6.91e+5	1.38e+6	2.76e+6	5.62e+6	1.11e+7
Matlab	3.38e-6	2.28e-6	1.74e-6	1.22e-6	7.73e-7	5.85e-7	3.07e-7	4.00e-7
Stable32	3.38e-6	2.28e-6	1.74e-6	1.22e-6	7.73e-7	5.85e-7	3.07e-7	4.00e-7




Fig. 3. Phase-modulated flicker -noise random data sequence.

3. Time-Frequency Data Analysis Platform Design and Implementation

To transform optimized algorithms into practical research tools, this paper develops the TFDA platform application for interactive analysis.

3.1. TFDA Platform Design

TFDA adopts a client-server architecture, with the backend based on Spring Boot providing high-performance RESTful API services, and the

2.3. Noise Simulation and Identification

Atomic clock noise is modeled as five independent power-law components $S_n(f) = h_a \cdot f^a$, where $a \in \{-2, -1, 0, +1, +2\}$ correspond to random-walk frequency modulation, flicker frequency modulation, white frequency modulation, flicker phase modulation, and white phase modulation noise, respectively [14]. This paper implements simulation algorithms for all five types: direct Gaussian generation for white phase modulation, cumulative summation for white and random-walk frequency modulation, and ARIMA models for flicker phase and flicker frequency modulation. The Lag-1 autocorrelation method provides consistent noise type identification. Validation on 10,000 synthetic datasets demonstrates that identification accuracy exceeds 97 % for all types, with white phase modulation, white frequency modulation, and random-walk frequency modulation achieving 100 % accurate identification [15]. Phase-modulated flicker -noise random data sequence is Fig. 3.

frontend based on Electron+Vue.js providing a native Chinese graphical interface. Time-Frequency data analysis platform is as follows in Fig. 4. Practically all data-analysis and processing tasks can be accomplished on this platform, which brings great convenience for users in China.




Fig. 4. Time-Frequency data analysis platform.

3.2. Validation and Results

This paper performs comprehensive validation using 1PPS measurement data (1-second sampling interval) from two time servers relative to UTC(NTSC). Table 5 confirms numerical consistency of Allan deviation between TFDA and Stable32. Table 6 summarizes functional advantages: TFDA provides equivalent analysis depth while adding

Chinese GUI and Chart Customization, complete Chinese localization, and flexible chart customization capabilities that Stable32 does not possess.

Table 5. Allan variance validation results (Time Server A).

Averaging Time (s)	TFDA	Stable32	Relative Error
8.64×10^4	4.00×10^{-7}	4.00×10^{-7}	$<10^{-15}$
1.73×10^5	3.38×10^{-6}	3.38×10^{-6}	$<10^{-15}$
3.46×10^5	3.07×10^{-7}	3.07×10^{-7}	$<10^{-15}$
6.91×10^5	1.22×10^{-6}	1.22×10^{-6}	$<10^{-15}$
1.38×10^6	7.73×10^{-7}	7.73×10^{-7}	$<10^{-15}$

Table 6. Functional comparison between TFDA and Stable32.



Feature	Stable32	TFDA
Variance Family (Allan, Hadamard, etc.)	Supported	Supported
Drift Analysis (8 methods)	Supported	Supported
Noise Identification	Supported	Supported
Chinese GUI	Not Supported	Supported
Chart Customization	Not Supported	Supported

4. Conclusions

This paper presents a comprehensive study on time-frequency data analysis methods and platform design. Through comparison with Stable32, on one hand, rigorous theoretical validation of the complete algorithmic system is achieved; on the other hand, the TFDA platform provides researchers with practical application tools for efficient processing and analysis of time-frequency data. Experimental validation confirms identical computational results with Stable32 and demonstrates significant functional advantages. This research establishes a complete technical chain from theoretical foundations through high-performance optimization to practical platform deployment, promoting the development of time-frequency metrology technology in the big data era.

References

- [1]. Y. Liu, X. Li, Z. Zhao, et al., Frequency Digital Measurement, Science Press, Beijing, 2025.
- [2]. H. Guo, S. Yang, H. He, Analysis of frequency drift characteristics for navigation satellite atomic clocks, GNSS World of China, Vol. 32, Issue 6, 2007, pp. 5-10 (in Chinese).
- [3]. D. W. Allan, Statistics of atomic frequency standards, *Proceedings of the IEEE*, Vol. 54, Issue 2, 1966, pp. 221-230.
- [4]. J. A. Barnes, The measurement of linear frequency drift in oscillators, in *Proceedings of the 15th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, 1983, pp. 551-582.
- [5]. D. W. Allan, M. A. Weiss, J. L. Jespersen, A frequency-domain view of time-domain characterization of clocks and time and frequency distribution systems, in *Proceedings of the 45th Annual Symposium on Frequency Control*, 1991, pp. 667-678.
- [6]. S. Bregni, Clock stability characterization and measurement in telecommunications, *IEEE Transactions on Instrumentation and Measurement*, Vol. 46, Issue 6, 1997, pp. 1284-1294.
- [7]. Y. Z. Tsytkin, K. Furuta, Frequency-domain criteria of robust stability – discrete-time and continuous-time systems, a unified approach, *International Journal of Robust and Nonlinear Control*, Vol. 5, Issue 3, 1995, pp. 207-222.
- [8]. W. J. Riley, Handbook of frequency stability analysis, NIST Special Publication 1065, National Institute of Standards and Technology, 2008.
- [9]. D. W. Allan, Should the classical variance be used as a basic measure in standards metrology? *IEEE Transactions on Instrumentation and Measurement*, Vol. 36, Issue 2, 1987, pp. 646-654.
- [10]. D. W. Allan, J. A. Barnes, A modified "Allan variance" with increased oscillator characterization ability, in *Proceedings of the 35th Annual Frequency Control Symposium (FCS)*, 1981, pp. 470-475.
- [11]. R. A. Baugh, Frequency modulation analysis with the Hadamard variance, in *Proceedings of the 25th Annual Symposium on Frequency Control*, 1971, pp. 222-225.
- [12]. S. Bregni, Measurement of maximum time interval error for telecommunications clock stability characterization, *IEEE Transactions on Instrumentation and Measurement*, Vol. 45, Issue 5, 1996, pp. 900-906.
- [13]. D. A. Howe, An extension of the Allan variance with increased confidence, in *Proceedings of the IEEE International Frequency Control Symposium (FCS)*, 1995, pp. 321-329.
- [14]. W. J. Riley, C. A. Greenhall, Power law noise identification using lag 1 autocorrelation, in *Proceedings of the 18th European Frequency and Time Forum (EFTF)*, 2004, pp. 439-444.
- [15]. C. Zhou, Q. Zhang, S. Yan, Power law noise identification using lag 1 autocorrelation by overlapping samples, in *Proceedings of the Joint Conference of the IEEE International Frequency Control Symposium and European Frequency and Time Forum (FCS/EFTF)*, 2011, pp. 1-4.



9788409900312