Proceedings of the Workshop on Innovation on Information and Communication Technologies

ITACA-WIICT 2025

Editors:
C. Fernandez-Llatas,
A. Martinez-Millana

Institute ITACA
Information and Communication Technologies

Program Committee

Organizing Committee

Antonio Martinez-Millana, Universitat Politècnica de València, Spain
Maria Guillem, Universitat Politècnica de València, Spain
David de Andres, Universitat Politècnica de València, Spain
Carlos Fernández-Llatas, Universitat Politècnica de València, Spain
Jose Mariano Dahoui, Universitat Politècnica de València, Spain

Program Committee

Carlos Fernandez-Llatas, Universitat Politècnica de València, Spain
Paulo de Carvalho, University of Coimbra, Portugal
Anna-Maria Bianchi, Politecnico di Milano, Italy
Jorge Munoz-Gama, Pontificia Universidad Catolica de Chile
Fernando Seoane, Karolinska Institutet
Jorge Henriques, University of Coimbra, Portugal
Cenk Demiroglu, Ozyegin University, Turkey
Johan Gustav Bellika, National Center of Telemedicine, Norway
Yigzaw Kassaye Yitbarek, University of Tromso, Norway
Raymundo Barrales, Universidad Autónoma Metropolitana de México
Onur Dogan, Istanbul Technical University,, Turkey
Frank Y. Li, University of Agder, Norway
Pilar Sala, MySphera., Spain
Alberto Bonastre, Universitat Politècnica de València, Spain
Jose Manuel Catala, Universitat Politècnica de València, Spain
Jose Carlos Campelo, Universitat Politècnica de València, Spain
Juan Vicente Capella, Universitat Politècnica de València, Spain
Antonio Mocholí, Universitat Politècnica de València, Spain
Sara Blanc, Universitat Politècnica de València, Spain
Juan-Carlos Baraza-Calvo, Universitat Politècnica de València, Spain
Juan Carlos Ruiz, Universitat Politècnica de València, Spain
Joaquin Gracia, Universitat Politècnica de València, Spain
David De Andres, Universitat Politècnica de València, Spain
Ricardo Mercado, Universitat Politècnica de València, Spain
Lenin Lemus, Universitat Politècnica de València, Spain
Vicente Traver, Universitat Politècnica de València, Spain
Antonio Martinez-Millana, Universitat Politècnica de València, Spain
Jose Luis Bayo-Monton, Universitat Politècnica de València, Spain

Juan Miguel García-Gomez, Universitat Politècnica de València, Spain
Elies Fuster, Universitat Politècnica de València, Spain
Carlos Saez, Universitat Politècnica de València, Spain
Ángel Perles, Universitat Politècnica de València, Spain
Luis José Saiz Adalid, Universitat Politècnica de València, Spain
Gema Ibañez, Universitat Politècnica de València, Spain
Pedro Yuste, Universitat Politècnica de València, Spain
Daniel Gil Tomàs, Universitat Politècnica de València, Spain
Sabina Asensio, Universitat Politècnica de València, Spain
Beatriz Garcia-Baños, Universitat Politècnica de València, Spain
Francisco Castells, Universitat Politècnica de València, Spain
Diogo Nunes, University of Coimbra, Portugal
Adriana Leal, University of Coimbra, Portugal
Zoe Valero Ramón, Universitat Politècnica de València, Spain
Yolanda Vives, Universitat Politècnica de València, Spain
Javier Urchueguia, Universitat Politècnica de València, Spain
Jose Vicente Oliver. Universitat Politècnica de València, Spain
Victoria Lerma, Universitat Politècnica de València, Spain
Edgar Lorenzo, Universitat Politècnica de València, Spain

Table of Contents

(POSTER) LALABY GLIO APP: Digital phenotyping of quality of life in glioblastoma patients Sabina Asensio-Cuesta, Elies Fuster-García, Juan M. García-Gómez, Jorge Soler, Inmaculada Maestu, Daniel Sánchez-García, Ángel Sánchez-García, Teresa Soria Comes, Maria De Julian Campayo and Carlos Sáez-Silvestre

Foreword to ITACA-WIICT 2025: Advancing Innovation Through Connection

Carlos Fernandez-Llatas Author[1] and Antonio Martinez-Millana[1]

ITACA, Universitat Politècnica de València. Valencia, Spain cfllatas@itaca.upv.es

The Workshop on Innovation on Information and Communication Technologies (ITACA-WIICT 2025) took place on July 10, 2025, at the Universitat Politècnica de València. The event brought together scientists, engineers, and professionals dedicated to research, development, and innovation in the field of information and communication technologies (ICT). Conceived as a meeting point for collaboration and knowledge exchange, the workshop aimed to strengthen connections between research groups, promote cooperation among professionals, and encourage technological transfer across disciplines.

The 2025 edition of ITACA-WIICT showcased the vitality and diversity of current ICT research, with contributions that spanned digital health, embedded systems, artificial intelligence, space technology, and process mining. Together, these works reflect ITACA's multidisciplinary approach and its commitment to bridging cutting-edge technological innovation with social and industrial impact.

Several papers explored how data-driven methodologies and interactive analytics can transform healthcare. One study presented a complete redesign of pain management pathways in hospital emergency departments using process mining tools and value-based healthcare principles, demonstrating measurable improvements in patient outcomes and efficiency. Another focused on a digital twin platform for personalized atrial fibrillation ablation planning, enabling doctors to simulate and compare treatment strategies in real time. Complementing these clinical advances, new mobile applications were showcased for continuous monitoring of patients with brain cancer, illustrating the growing potential of digital phenotyping to capture quality-of-life metrics unobtrusively.

Other works shifted the focus to the core of computing technologies. A study on neural networks implemented on Arduino systems analyzed the dependability of quantized models and error-correction mechanisms, proving that efficient artificial intelligence can be achieved even in constrained embedded environments. A complementary manifesto proposed the integration of process mining techniques within embedded and IoT systems, envisioning adaptive, context-aware devices capable of learning from their own operational data in real time.

The connection between technology and the environment also played a key role. A research project on wildfire-induced air pollution established links between satellite-based emission estimates and ground-level PM. concentrations, highlighting the need for predictive models to protect public health during extreme events. In parallel, another contribution expanded ICT toward space applications, demonstrating how microwave processing of lunar regolith could serve as a sustainable method for in-situ construction on the Moon, offering insights into future extraterrestrial habitat design.

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

Beyond these specific fields, the proceedings also include research on compensating for respiration-induced cardiac motion when solving the inverse electrocardiography problem. By proposing new algorithms to correct these physiological artifacts, this work contributes to more accurate non-invasive cardiac mapping — further evidence of how data analysis, modeling, and biomedical engineering intersect within ICT research.

Taken together, the works presented at ITACA-WIICT 2025 show the dynamism and creative capability of the research community at ITACA and its collaborators. They exemplify how advances in information and communication technologies are not confined to computation alone but extend to medicine, environmental monitoring, aerospace exploration, and beyond.

The workshop reaffirmed its role as a platform for interdisciplinary dialogue, where academic excellence meets practical innovation. Through initiatives such as these, ITACA continues to drive progress toward a more connected, intelligent, and sustainable future — one in which technology serves not only efficiency and precision but also well-being, resilience, and shared human value.

Revolutionizing Emergency Department Pain Management: A Value-Based Approach Through Interactive Process Mining and the DMAIC Cycle

Angeles Celda[1][0000−0003−2054−8236], Gema Ibanez-Sanchez[2][0000−0003−1824−281X], Zoe Valero-Ramon[2][0000−0002−4903−3896], and Carlos Fernandez-Llatas[2][0000−0002−2819−5597]

[1] Emergency Department, Hospital Universitario de Valencia, Av. de les Tres Creus, L'Olivereta, València, 46014-Valencia (Spain) angelescelda@gmail.com
[2] ITACA-SABIEN, Universitat Politècnica de València, Camino de Vera, sn, 46022-Valencia (Spain)
{geibsan,zoevara}@itaca.upv.es

Abstract. Healthcare systems globally face increasing pressures to deliver high-quality, patient-centric care while managing escalating costs and demands. This paper presents a comprehensive case study, PAINKEY, demonstrating how a structured methodology, the DMAIC cycle, combined with the analytical power of Interactive Process Mining, was successfully applied to redesign pain management pathways in an Emergency Department. Focusing on a common yet often problematic area of care, the project aimed to address high clinical burden, significant economic impact due to undertreated pain, and variability in care. By leveraging data-driven insights from PMApp, a specialized Interactive Process Mining Toolkit, the initiative led to tangible improvements in patient outcomes and operational efficiency, showcasing a transferable model for optimizing healthcare delivery within a Value-Based Healthcare framework.

Keywords: Process Mining · Value-Based Healthcare · Emergency Department · Pain Management · DMAIC Cycle.

1 Introduction: The Innovation in Healthcare Delivery

The contemporary healthcare landscape necessitates a fundamental shift from traditional, activity-centric approaches to more patient-centric, value-driven models. Healthcare systems worldwide grapple with the escalating costs, increasing patient demands, and the constant imperative to deliver superior quality care with finite resources [3, 12, 1]. This module explores a critical journey of transforming healthcare delivery models through a practical, real-world case study focused on the successful redesign of a pain management pathway within a hospital's Emergency Department (ED). The PAINKEY project serves as a living example of how data-driven methodologies and innovative thinking can yield significant improvements in patient outcomes and operational efficiency. The

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

primary drivers for exploring new care delivery models, particularly in this case, include:

- *High Clinical Burden and Patient Suffering:* Pain is remarkably prevalent in EDs, affecting over 60-80% of patients[8, 21]. Despite its ubiquity, it is frequently undertreated or managed with significant delays, leading to prolonged patient suffering, negative impacts on their experience, and potential worsening of clinical recovery [6, 10]. Patients anticipate rapid and effective pain relief, yet current systems often fall short, with average treatment times substantially exceeding patient expectations [22, 8, 21].
- *Significant Economic Impact:* Inefficient pain management is not merely a clinical concern but also an economic one. It contributes to extended lengths of stay (LOS) in the ED, increased resource consumption, and higher rates of return visits and readmissions. Untreated pain can also lead to chronic conditions, further escalating long-term healthcare costs[4]. Optimizing pain management can thus unlock substantial economic value for healthcare systems [15, 18].
- *Alignment with Value-Based Healthcare (VBHC) Principles:* The transition towards VBHC mandates that healthcare providers deliver the best possible health outcomes for patients per unit of cost [24, 11]. Pain management is an ideal candidate for VBHC initiatives due to its direct impact on patient quality of life [17], clear economic footprint, and significant margin for improvement. Enhancing pain care directly aligns with the core VBHC objective of maximizing patient-relevant outcomes with efficient resource utilization.
- *Existing Gaps and Variability in Care:* This case study highlights that despite the high prevalence of pain, there is often a lack of standardized assessment [13, 23] and considerable variability in how pain is managed across different patient groups and even within the same department. This variability leads to inefficiencies, suboptimal outcomes, and missed opportunities for early intervention.
- *Technological Enablement:* The advent of new technologies, particularly Interactive Process Mining (IPM), provides unprecedented opportunities to analyze complex healthcare processes in detail[9]. IPM allows for a shift beyond aggregated data to understand the real-world flow of patients, identify bottlenecks, and pinpoint the root causes of inefficiencies [9, 25]. This data-driven approach is a potent engine for change in modern healthcare [19, 16].

Through the lens of the PAINKEY project, we will explore how a structured methodology—the DMAIC cycle—combined with the analytical power of IPM and the PMApp toolkit, was utilized to address these critical motivators, leading to a new, more effective, and patient-centered model of pain care delivery.

2 Methodology: The DMAIC Cycle as a Framework for Process Improvement

The DMAIC (Define, Measure, Analyze, Improve, Control) cycle is a data-driven improvement cycle employed for enhancing, optimizing, and stabilizing business processes and designs. It serves as the core methodology of Lean Six Sigma [2, 20]. We will examine each phase through its practical application in the PAINKEY project, emphasizing the role of Interactive Process Mining (IPM) and its specialized toolkit, PMApp [14].

2.1 Define Phase: Understanding the Problem and Setting Goals

The "Define" phase is the foundational step where the problem is clearly articulated, the project scope is set, and the objectives are established. In the context of the PAINKEY project, this phase was crucial for understanding the specific challenges of pain management within the hospital's Emergency Department (ED). While scientific evidence broadly highlights issues in ED pain management, the first step was to confirm if these challenges were present in their specific ED. IPM was introduced early to provide a comprehensive, real-time view of the existing pain management process. Unlike traditional retrospective studies, IPM enabled the team to discover and analyze clinical pathways as they actually occurred, highlighting variability and bottlenecks. A key activity in this structured phase involved conducting "Data Rodeos," which are highly coupled, multidisciplinary interactive data analyses aimed at building process indicators that enable objective, comprehensive, and exploratory understanding, quantification, and qualification of processes and their changes. These sessions involved emergency nurses and Process Mining experts in co-creation, validating data quality and understanding the actual flow of events in the pain management process, as well as identifying patient flows that deviated from standard pathways. PMApp, specifically designed to initiate Data Rodeos, facilitates the application of the IPM methodology in real-world scenarios and promotes its acceptance among healthcare professionals in their daily practice. An illustrative example from this phase was the iterative refinement of initial questions; for instance, an early query like "What is the average time from triage to analgesia care?" initially yielded a finding of 9 minutes to triage plus 54 minutes to care. Further refinement, by filtering for triage levels II-III, revealed a refined average of 48 minutes to analgesia, with a concerning 18% of patients waiting over 60 minutes. This specific insight directly motivated the design of the PAINKEY pilot. During this stage, different "runners" were defined for IPM analysis, such as "episode" for calculating total length of stay (LOS) and "72-hour" or "30-day" for a more detailed analysis of return and readmission rates, which would be key metrics for success. The outcome of the Define phase was a clear understanding of the pain management problem in the ED, validated through real process data, and specific objectives for improvement, particularly regarding reducing time to analgesia and addressing variability.

2.2 Measure Phase: Quantifying the Problem

The "Measure" phase focuses on collecting data to quantify the extent of the problem identified in the Define phase. This involves selecting appropriate metrics, ensuring data accuracy, and establishing a baseline for future comparison. For the PAINKEY project, this crucial phase included the extraction of 121,981 patient episodes from January to December 2019 from the hospital's databases. Episodes not passing through triage (11,131) were excluded from the main analysis due to lack of initial complaint information, although a specific circuit was designed for them. An Interactive Process Indicator (IPI) was constructed by considering each key event in the patient journey with its timestamp (Figure 3). This allowed for a detailed, event-level reconstruction of the care process. PMApp plays a vital role in enabling the creation of logs from diverse data sources like CSV or SQL through its Ingestor Editor module, which also facilitates data quality management. The Ingestor Editor is organized into five tabs—Validators, Variables, Events, Trace Data, and Filters—each serving a specific purpose in refining the data. For instance, the Validators section features algorithms that accept or reject rows during data ingestion based on criteria set by healthcare professionals (HCPs). The Events tab is the core of the ingestion process, where blocks are employed to create events using existing variables and fields, defining nodes in the final process model. Specific Key Performance Indicators (KPIs) were chosen to provide concrete information on pain management and patient experience. These included: percentage of patients receiving first attention/analgesia within 30 minutes; percentage of patients with "intense" or "very intense" pain waiting more than 15 minutes; and percentage of patients with "moderate" pain waiting more than 30 minutes. Safety metrics such as return rates at 72 hours and 30 days, and readmission rates at 72 hours and 30 days, were also selected to detect ineffective treatments or complications. To estimate pain prevalence, the Manchester Triage System (MTS) chief complaint diagrams were utilized. This method was preferred over ICD coding due to its real-time availability at patient arrival and its ability to capture the initial reason for consultation, even if the final diagnosis differs. Diagrams were categorized into "With Pain" (e.g., abdominal pain, limb problems, headaches) and "Without Pain" groups based on the presence of pain, even in cases where the word "pain" did not explicitly appear in the definition. The culmination of the Measure phase was a robust dataset and a set of quantifiable KPIs that provided a clear baseline of the current state of pain management, highlighting significant delays and variability [10, 7].

2.3 Analyze Phase: Identifying Root Causes and Opportunities

In the "Analyze" phase, the collected data is meticulously examined to identify the root causes of the problem and pinpoint specific opportunities for improvement. This phase significantly leverages IPM's capabilities to uncover hidden patterns and process deviations. PMApp's Dashboard serves as the primary module

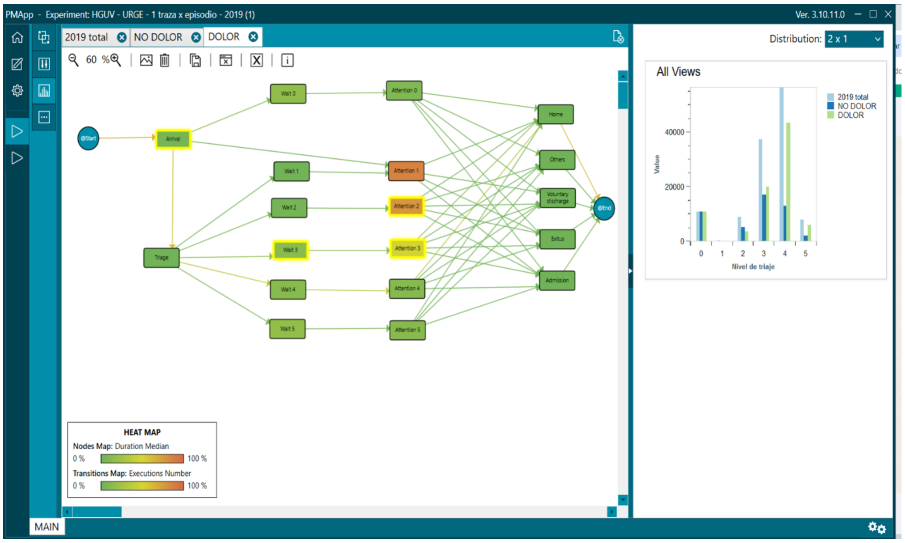

Fig. 1. PMApp

for visualizing and navigating the IPI. This module assembles personalized dashboards using configuration information from the Runner file and integrates new resources through installed plugins. The epidemiological profile analysis confirmed a high pain prevalence: 66% of patients passing through triage in 2019 presented with a pain-related complaint, aligning with existing literature. Comparative group analysis, facilitated by IPM, allowed for a detailed comparison between "With Pain" and "Without Pain" groups. It was found that "With Pain" patients generally had lower estimated complexity (predominantly triage levels III, IV, V; lower admission rates; shorter overall LOS), yet still experienced significant delays in care. The timeliness gap analysis revealed critical gaps: 66% of "With Pain" patients waited over 30 minutes for first attention, and a concerning 83% of patients with "intense" or "very intense" pain waited more than 15 minutes, clearly indicating a deviation from quality standards. Return/readmission pattern analysis showed that while initial return rates were low (4.1% at 72 hours), they tripled to 13.73% at 30 days, suggesting incomplete resolution of the pain problem for a significant proportion of patients. Abdominal pain specifically stood out with over 70% of returns/readmissions for the same reason, indicating a particular area for improvement in initial management. Granular analysis of the 8 most prevalent pain processes revealed considerable variability in triage levels, age distribution, admission rates, LOS, and return/readmission patterns. IPM was essential for visualizing and quantifying this process-specific variability. The Experiment Designer in PMApp, a module for creating formal schemes that define the Process Mining algorithm flow, allowed Process Miners to configure runners by dragging and dropping blocks that represent calls to func-

tions and algorithms to generate the IPI. This collaborative approach, involving discussions with HCPs during Data Rodeos, defines the sequence of operations to transform hospital data into a final IPI. Root cause identification pointed to delays in initial assessment and analgesia administration as primary areas for improvement, emphasizing that patients were waiting in pain. The outcome of the Analyze phase was a precise, data-backed understanding of the inefficiencies and root causes within the pain management process, providing clear targets for the improvement phase [8, 5].

2.4 Improve Phase: Implementing Solutions

The "Improve" phase focuses on developing and implementing solutions to address the root causes identified in the Analyze phase. These solutions are designed to be effective, efficient, and aligned with the project's objectives. Based on the thorough analysis, the PAINKEY project was conceived as a mobile web application (Web-app) for patient self-assessment upon arrival at the ED. The design was centered on Value-Based Healthcare (VBHC) principles, focusing on patient outcomes and overall process efficiency. PAINKEY integrated a comprehensive questionnaire that, while pain-focused, also included questions to identify other urgent or life-threatening conditions, ensuring equitable care. It collected comprehensive data on urgency, fragility, and complexity. The information from PAINKEY was then used to stratify patients and create differentiated care circuits within the SUH, aiming to expedite attention and place patients in the most appropriate care space. Patient responses generated alerts for triage staff, enabling proactive prioritization of calls and eliminating the "blind spot" between admission and triage. All data was automatically integrated into the patient's electronic health record. The redesign included implementing protocols for early analgesia initiation, directly addressing the identified delays. The development followed a rigorous approach, including a DELPHI study with pain and emergency experts to build consensus on key assessment items and prioritization criteria. PMApp's Experiment Designer, with its drag-and-drop block system, enhances the tool's customizability, empowering the Process Mining community to extend its functionality and reconfigure it to align with hospital information systems and legacy systems. This flexibility allows for the creation of custom dashboards aligned with the unique requirements of health scenarios for daily practice in the Production phase. PAINKEY was technically integrated with the hospital's information system, ensuring immediate data availability for staff and feeding data into the IPIs for continuous process analysis. The outcome of the Improve phase was the successful design and implementation of the PAINKEY Web-app and the associated redesign of care pathways, directly targeting the identified inefficiencies in pain management.

Figure 2 depicts a streamlined, technology-supported patient journey in the Emergency Department, focusing on pain management.

1. **Patient Arrival:** The process begins when the patient arrives at the Hospital Emergency Department.

2. **SMS Link to Webapp:** After admission, the patient receives an SMS containing a brief message and a link to the PAINKEY Webapp.
3. **Patient Self-Assessment:** The patient or a family member accesses the Webapp via the link and completes a valorization (assessment) questionnaire. In a second phase, a voice recording of the patient will be used to assess the intensity of the pain, suggesting future technological integration.
4. **Data Integration into Clinical History:** The data collected from the Webapp assessment is automatically incorporated into the patient's electronic clinical history.
5. **System Proposed Circuit:** Based on the assessed urgency, fragility, and complexity from the Webapp data, the system proposes a specific care circuit for the patient within the ED.
6. **Professional Confirmation:** A healthcare professional reviews and confirms this proposed circuit, acting as a guarantee of clinical safety and ensuring appropriate care.
7. **Communication to Patient:** Finally, the confirmed care pathway is communicated to the patient.

The workflow emphasizes early patient engagement through self-assessment, data-driven decision-making for care stratification, and a crucial human oversight for clinical safety, ultimately aiming to optimize patient flow and pain management.

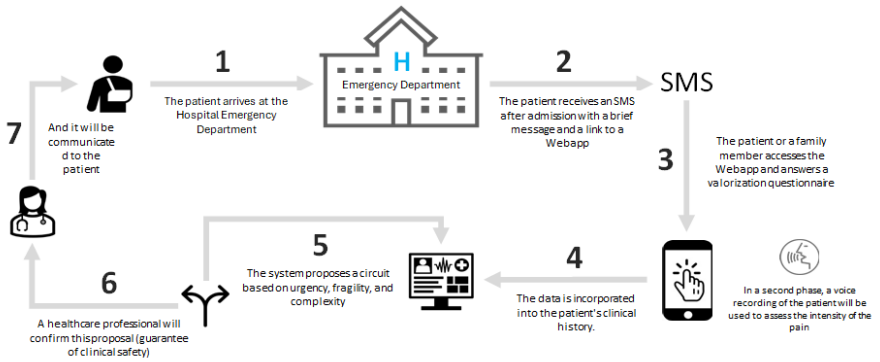

Fig. 2. PAINKEY flow

Figure 3 presents a screenshot of the application, where the patient is notified about the assigned circuit and a summary of the provided data.

2.5 Control Phase: Sustaining the Gains

The "Control" phase is crucial for ensuring that the improvements made are sustained over time and that the process continues to operate at the improved level.

Fig. 3. Screenshot of the data summary provided to the patient

This involves establishing robust monitoring systems and mechanisms for continuous feedback and adjustment. For the PAINKEY project, a specific control panel was configured in PMApp, incorporating the new data from PAINKEY self-assessments. This allowed for systematic analysis of the IPI by a multidisciplinary team. PMApp's Dashboard provides a platform for visualizing and navigating the IPI, acting as a container that assembles personalized dashboards using the configuration information from the Runner file and integrating new resources through installed plugins. Upon completion of the Dashboard customization, HCPs and Process Miners can initiate the exploration of the IPI, typically starting with the Main Perspective, which is the central element illustrating the model with all the data. The impact of PAINKEY was objectively quantified by measuring key value and quality indicators, such as time to first attention/analgesia, percentage of patients attended within optimal times, and potential reductions in return/readmission rates. Patient satisfaction was also evaluated. IPM allowed for detailed visualization and analysis of the new care circuits, helping to proactively identify any emerging bottlenecks or inefficiencies in the redesigned workflows, and also assessing the adequacy of defined care routes. The systematic analysis of data, including patient-reported outcomes (PROMs), revealed new underlying causes of problems or variability, such as characterizing chronic pain patients who frequently returned to the ED. Findings from IPM analysis systematically generated new research questions and guided proposals for specific, targeted improvement actions. This iterative and cyclical process ensured constant optimization of pain management, effectively closing the DMAIC cycle and reinforcing the VBHC approach. Furthermore, the control phase also recognized limitations, such as inclusion bias for smartphone users or data quality, and proposed mitigation strategies for future phases (e.g., touch terminals in examination rooms, real-time audits of records). The outcome of the Control phase was the establishment of a robust system for continuous monitoring and improvement, ensuring the sustainability of the benefits achieved by PAINKEY and providing a framework for ongoing optimization of pain management processes.

3 Conclusion: Lessons Learned and Future Directions

The PAINKEY case study provides a compelling example of how a structured, data-driven approach can significantly transform healthcare delivery models. By applying the DMAIC methodology in conjunction with Interactive Process Mining and the PMApp toolkit, the hospital was able to move beyond anecdotal evidence to precisely identify critical inefficiencies in pain management within its Emergency Department. Key takeaways from this case study include:

- *The Power of Data:* The ability to visualize and analyze real-world process data through IPM was fundamental. It allowed the team to quantify the extent of delays, understand the variability in care, and pinpoint specific areas for intervention, such as the alarming waiting times for patients with severe pain.

- *Patient-Centered Innovation:* PAINKEY exemplifies how technological solutions, when designed with the patient at the centre (validated through DELPHI studies), can directly address patient needs (e.g., rapid pain assessment) while also improving operational efficiency.
- *The Value of a Structured Approach:* The DMAIC framework provided a clear roadmap for the entire improvement journey, from problem definition to sustained control. This systematic approach ensures that interventions are evidence-based and their impact is measurable.
- *Continuous Improvement is Key:* Healthcare is dynamic. The Control phase, powered by IPM, highlights the necessity of ongoing monitoring and iterative adjustments to ensure that improvements are maintained and new opportunities for optimization are continuously identified.
- *Bridging the Gap between Theory and Practice:* This case demonstrates how theoretical concepts like Value-Based Healthcare can be translated into practical, implementable solutions that yield tangible benefits for both patients and the healthcare system.

While the PAINKEY project was a pilot in a specific setting, its methodology offers a transferable model for other healthcare challenges. The lessons learned underscore the importance of embracing data analytics, fostering multidisciplinary collaboration, and committing to continuous improvement in the pursuit of delivering high-value, patient-centered care. As future healthcare leaders, understanding and applying these principles will be crucial for navigating the complexities of modern healthcare and designing the care delivery models of tomorrow.

Acknowledgments. This work was supported by the "Reto Grünenthal Voz en Dolor Crónico", awarded to the PAINKEY team within the framework of the VI Hackathon de Salud. We extend our sincere gratitude to Grünenthal Pharma for launching this challenge.

References

1. Asplin, B.R., Magid, D.J., Rhodes, K.V., Solberg, L.I., Lurie, N., Camargo Jr, C.A.: A conceptual model of emergency department crowding. Annals of emergency medicine **42**(2), 173–180 (2003)
2. Barr, E., Brannan, G.D.: Quality improvement methods (lean, pdsa, six sigma). In: StatPearls [Internet]. StatPearls Publishing (2024)
3. Bernstein, S.L., Aronsky, D., Duseja, R., Epstein, S., Handel, D., Hwang, U., McCarthy, M., John McConnell, K., Pines, J.M., Rathlev, N., et al.: The effect of emergency department crowding on clinically oriented outcomes. Academic Emergency Medicine **16**(1), 1–10 (2009)
4. Blyth, F., MacMahon, M., Smith, B., Blyth, A., Brnabic, A.: The costs of chronic pain-long-term estimates. Pain **165**(4), e575–e587 (2024). https://doi.org/10.1097/j.pain.0000000000003058
5. Bond, M., Erwich-Nijhout, M., Phillips, D., Baggoley, C.: Urgency, disposition and age groups: a casemix model for emergency medicine. Emergency Medicine **10**(2), 103–110 (1998)

6. Chong, E., Prowse, M., Tang, J.: Problems and barriers of pain management in the emergency department: Are we ever going to get better? Pain Physician **12**(6), 945–956 (2009)
7. Constantí, V.A., de la Maza, V., Riera, M.P., Contreras, M.J., Cubells, C.L.: Diseño de indicadores de calidad: el proceso del dolor agudo en el servicio de urgencias. Acta Pediatrica Espanola **72**(11) (2014)
8. Cordell, W.H., Keene, K.K., Giles, B.K., Jones, J.B., Jones, J.H., Brizendine, E.J.: The high prevalence of pain in emergency medical care. The American journal of emergency medicine **20**(3), 165–169 (2002)
9. Fernandez-Llatas, C. (ed.): Interactive Process Mining in Healthcare. Springer, Cham (2021)
10. Hatherley, C., Jennings, N., Cross, R.: Time to analgesia and pain score documentation best practice standards for the emergency department - a literature review. Australasian Emergency Nursing Journal **19**(1), 26–36 (2016). https://doi.org/10.1016/j.aenj.2015.11.001
11. of Health Insurance, C.: Value-based health care (2023), https://vbhc.chi.gov.sa/assets/pdf/VBHC White Paper Version Final.pdf, accessed: 2025-06-07
12. Hoot, N.R., Aronsky, D.: Systematic review of emergency department crowding: causes, effects, and solutions. Annals of emergency medicine **52**(2), 126–136 (2008)
13. Hämäläinen, J., Kvist, T., Kankkunen, P.: Acute pain assessment inadequacy in the emergency department: Patients' perspective. Journal of Patient Experience **9**, 23743735211049677 (2022). https://doi.org/10.1177/23743735211049677
14. Ibanez-Sanchez, G., Fernandez-Llatas, C., Valero-Ramon, Z., Bayo-Monton, J.L.: Pmapp: an interactive process mining toolkit for building healthcare dashboards. In: International Workshop on Explainable Artificial Intelligence in Healthcare. pp. 75–86. Springer (2023)
15. IQVIA: No pain some commercial gain: Non-addictive prescription medicines (2024), https://www.iqvia.com/locations/emea/blogs/2024/10/no-pain-some-commercial-gain-non-addictive-prescription-medicines, accessed: 2025-06-07
16. Khan, M.A., Naqvi, S.A.M., Shuja, J., Akram, M.U., Shah, M.Z.: (pdf) the transformative impact of big data on modern healthcare: A technical analysis. ResearchGate (2024), https://www.researchgate.net/publication/389502287_The_Transformative_Impact_of_Big_Data_on_Modern accessed: 2025-06-07
17. Kwon, S., Cho, E.K., Kim, H.J., Cho, H.J., Lee, J.H., Jo, J.S.: Leg length discrepancy after total hip arthroplasty: A review of clinical assessments, imaging diagnostics, and medico-legal implications. Journal of Clinical Medicine **13**(12), 3580 (2024). https://doi.org/10.3390/jcm13123580
18. Liu, W., Lim, L., Lim, B., Gan, L., Sali, A., Lim, S., Lim, C., Chew, B., Wong, K., Tan, T.: Economic analysis of implementing virtual reality therapy for pain among hospitalized patients. JMIR Medical Informatics **6**(4), e11221 (2018). https://doi.org/10.2196/11221
19. MDClone: Guide to data-driven healthcare innovation. MDClone (2024), https://www.mdclone.com/wp-content/uploads/2024/02/Guide_to_Data_Driven_Healthcare_Innovation_by_MDClone.pdf, accessed: 2025-06-07
20. Monday, L.M.: Define, measure, analyze, improve, control (dmaic) methodology as a roadmap in quality improvement. Global journal on quality and safety in healthcare **5**(2), 44 (2022)

21. Mura, P., Serra, E., Marinangeli, F., Patti, S., Musu, M., Piras, I., Massidda, M.V., Pia, G., Evangelista, M., Finco, G.: Prospective study on prevalence, intensity, type, and therapy of acute pain in a second-level urban emergency department. Journal of Pain Research **10**, 2147 (2017). https://doi.org/10.2147/JPR.S137992
22. Nagpal, A.K., Gadkari, C., Singh, A., Pundkar, A.: Optimizing pain management in emergency departments: A comprehensive review of current analgesic practices. Journal of Pain Research **17**, 1–11 (2024)
23. Raja, A., Khan, S., Dahan, F., Huda, Z., Hassan, S., Khan, A.: Assessing pain in the ed including the use of pain scales (such as osbd, flacc, vrs, nrs, crs, and oucher). Current Emergency and Hospital Medicine Reports **4**(1), 1–9 (2016). https://doi.org/10.1007/s40138-016-0081-3
24. Saleh, F., Al Saeedi, S.F., Al-Maktoum, H.R., Al-Yaseen, M.R., Aldhuhoori, F.H., Al-Nuaimi, R.H., Al-Nuaimi, S.S., Al-Hammami, S.A., Al-Shamsi, S.B., Al-Mazrouei, S.J., Al-Blooshi, K.N., Al-Farsi, S.B., Al-Blooshi, M.K., Al-Blooshi, S.K.: Toward a clearer understanding of value-based healthcare: A concept analysis. International Journal of Environmental Research and Public Health **20**(10), 5882 (2023)
25. Virtana: The complete guide to infrastructure performance monitoring (ipm) (2025), https://www.virtana.com/guides/ipm-guide/, accessed: 2025-06-07

From Fire to Air: Connecting Wildfire $PM_{2.5}$ Emissions with Air Monitoring Station Concentrations

Helena van den Berg-Sesma[1], Edgar Lorenzo-Sáez[1], José-Vicente Oliver-Villanueva[1], and Victoria Lerma-Arce[1]

[1] Universitat Politècnica de València, ITACA Research Institute, Valencia, Spain

Abstract. Wildfire $PM_{2.5}$ emissions in eastern Spain exhibit high spatial heterogeneity, driven by complex interactions between fire severity, vegetation type, and atmospheric dispersion. This study estimates consumed biomass using satellite-derived fire severity metrics (ΔNBR) and severity-dependent biomass consumption rates. Pre-fire biomass values were assigned per vegetation type and adjusted at pixel level (30 m x 30 m resolution). Results reveal significant exceedances of WHO air quality guidelines, with localized $PM_{2.5}$ concentrations surpassing safe thresholds by over 400%. These findings underscore the need for refined, spatially explicit emission modeling to support air quality forecasting and public health responses. Future integration of AI-enhanced predictive models is recommended to anticipate emission hotspots and protect vulnerable populations amid a growing wildfire threat exacerbated by climate change.

1 Introduction

Wildfires are a recurrent phenomenon in Mediterranean-type ecosystems (MTEs) [1], characterized by hot, dry summers and mild, wet winters which occur in regions such as California [2], and the Mediterranean Basin [3]. But not only climate is a key factor in the occurrence of wildfires. In the Mediterranean basin, human activity has also intervened in the disruption of the natural patterns of wildfire occurrence. The abandonment of agricultural land has resulted in an accumulation of fuel and the continuous expansion of large unmanaged forest areas in the landscape [4], [5]. Also, climate change has been associated with an extended wildfire season [6] with an average one-week extension of the fire season in the Mediterranean region [7], and an increasing presence of fire in systems where it was not historically present [8].

One of the most important sources of $PM_{2.5}$ air pollution is from wildfire (bushfire) smoke exposure [9]. Smoke from wildfires disperses widely and affects large portions of the population away from the fire source [10]. $PM_{2.5}$ levels are significantly elevated during wildfire episodes [11], [12] and can exceed levels set by regulatory bodies (World Health Organization air quality guidelines for $PM_{2.5}$:5 μg/m^3 annual mean, 15 μg/m^3 daily mean) [13].

Exposure to Particulate matter from wildfire smoke has been associated with both acute and chronic health effects. Acute exposure can lead to respiratory issues

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

such as asthma exacerbations, bronchitis, and increased hospital admissions for respiratory diseases [14]. In addition, $PM_{2.5}$ has the largest body of evidence implicating adverse cardiovascular effects [15],[16],[17] and a growing literature links wildfire smoke to more severe outcomes such as cardiovascular mortality [18].

A common approach in public health research to assess population exposure to $PM_{2.5}$ from wildland fires is the use of ground-based air monitoring networks [19].

This study focuses on the spatial and temporal relationship between wildfire events and $PM_{2.5}$ concentrations recorded by regional air quality stations. Specifically, we analyze the daily $PM_{2.5}$ peaks observed during two wildfire events in June 2012 at stations located within a 50 km radius of the burned areas. By comparing these peaks to long-term reference values (annual medians, percentiles, and health-based thresholds), and correlating them with estimated wildfire emissions within each station's buffer zone, we aim to evaluate how fire activity translates into measurable short-term air quality impacts. This approach provides insight into the short-term air quality impacts of wildfires and evaluates whether the estimated wildfire emissions within each station's surrounding area can explain the magnitude of observed $PM_{2.5}$ peaks.

2 Materials and methods

A bottom-up approach was applied to quantify $PM_{2.5}$ emissions associated with the 2012 wildfires in the Valencian region. This method integrates spatial data on pre-existing biomass, land use, and fire severity. The resulting estimates of $PM_{2.5}$ emissions are expressed in tons per hectare, which are then used to assess the acute exposure of the local population during the wildfire period.

2.1 Study area

This study focuses on two major extreme wildfire events (EWEs) that occurred in late June 2012 in the interior of the Valencia region (eastern Spain): the Andilla and Cortes de Pallás wildfires (Figure 1). Both fires ignited under similar meteorological conditions—marked by unusually high temperatures and prolonged drought [20]—and occurred within a narrow time window, allowing for a combined assessment of their atmospheric impact.

Fig. 1. Location of Andilla and Cortes de Pallás wildfires

The Andilla wildfire burned approximately 20,439 hectares of predominantly pine forest and Mediterranean shrubland in the northwestern part of the Valencia province. The Cortes de Pallás wildfire, which broke out around the same time, affected 28,307 hectares in the southwestern part of the province, in a landscape similarly characterized by pine forest, shrubland and rural terrain.

Although the two wildfires were not contiguous, their spatial and temporal proximity, combined with prevailing wind directions, resulted in widespread smoke transport across the region. This led to measurable changes in air quality, even at considerable distances from the fire fronts. These events thus offer a valuable case study for analysing how wildfire emissions influence short-term $PM_{2.5}$ concentrations at regional air monitoring stations. We examine daily $PM_{2.5}$ peaks recorded during the fire period at stations located within a 50 km radius of the burned areas, with the aim of linking observed pollution episodes to fire-related emissions and spatial exposure patterns.

2.2 Estimation of pre-existing biomass

Pre-existing biomass values were calculated in tons per hectare ($t \cdot ha^{-1}$) using the methodology of Lerma-Arce et al. [21], tailored for Mediterranean-type ecosystems (MTEs). Biomass estimates were assigned to the forest and agricultural land use classes defined in the SIOSE (Spanish Land Use Information System) land cover database [22]. Tables 1 and 2 show the values of biomass per hectare for the three provinces of the region of Valencia and for the different forest and agricultural land uses, respectively.

Table 1. Average aerial biomass for each forestland use of SIOSE [t·ha^{-1}]

Province	Castellón	Valencia	Alicante
Pastureland	7.75	7.75	7.75
Conifers	121.51	90.66	96.69
Deciduous broadleaves	146.10	132.91	132.43
Perennial broadleaves	179.58	101.05	122.75
Shrubs	38.73	43.52	34.00

Table 2. Average aerial biomass for each agriculture land use of SIOSE [t·ha^{-1}]

System	Irrigation			Rainfed		
Province	Castellon	Valencia	Alicante	Castellon	Valencia	Alicante
Rice	7.47	7.47	7.47	**	**	**
Other herbaceous crops	12.58	12.20	11.47	6.85	6.81	7.05
Citric fruit trees	82.05	76.26	85.90	**	**	**
Non-citric fruit trees	77.01	73.87	70.55	31.51	31.40	32.00
Vineyards	22.86	22.86	22.86	22.86	22.86	22.86
Olive groves	33.72	33.72	33.72	33.72	33.72	33.72
Other woody crops	31.51	31.40	32.00	31.51	31.40	32.00
Meadows	7.74	7.74	7.74	7.74	7.74	7.74

Since the SIOSE dataset does not coincide exactly with the fire years, biomass was adjusted using annual biomass growth rates from the National Forest Inventory [23] for forest land uses (Table 3). No adjustment was made for agricultural areas under the assumption that annual growth is removed through pruning [21].

Table 3. Average annual yield of aerial biomass [t·ha-1·year-2]

Province	Castellón	Valencia	Alicante
Pastureland	0.83	0.83	0.83
Conifers	4.45	3.30	4.00
Deciduous broadleaves	5.86	6.23	5.34
Perennial broadleaves	6.64	4.52	5.96
Shrubs	1.61	1.85	1.43

2.3 Quantification of biomass consumption

Consumed biomass was estimated based on fire severity and vegetation type. Following De Santis et al. [24], this process involves calculating the proportion of pre-existing biomass that was burned. Although originally developed for California's Mediterranean climate, this methodology was applied to Eastern Spain by Chaves-Naharro [25]. Fire severity was determined by calculating the Normalized Burn Ratio (NBR) using Landsat satellite imagery taken before and after each fire. The NBR is a normalized difference between the near-infrared and short-wave infrared bands [26]. The difference between pre-fire and post-fire NBR (ΔNBR) was then computed to

assess the degree of fire severity, which was classified according to the EFFIS fire severity classification system [27], with minor modifications to class names (see Table 4).

Table 4. Relation between ΔNBR values and the fire severity class

ΔNBR	Fire severity
-500 – 100	Unburned
101 – 269	Low
270 – 439	Moderate
440 – 659	High
660 – 1300	Extreme

Each pixel in the satellite imagery (30 m × 30 m; 900 m²) was assigned a severity class, vegetation type, and corresponding biomass value. Biomass loss was then estimated using severity-dependent consumption percentages (Table 5), derived from the works of Chaves-Navarro [25] and De Santis et al. [24].

Table 5. Percentages of biomass consumed by severity class and vegetation type [%]

Fire severity	Shrubs	Conifers	Broadleaves
Low	0.71	0.25	0.25
Moderate	0.84	0.47	0.40
High	0.89	0.56	0.48
Extreme	0.95	0.65	0.56

These percentages were applied to the pre-fire biomass using the equation:

$$FC = B_i \times C_i \quad (1)$$

Where, FC is the fraction of consumed biomass ($t \cdot ha^{-1}$), B_i is the pre-existing biomass in the spatial unit, and C_i is the percentage of biomass consumed based on fire severity and vegetation type.

2.4 $PM_{2.5}$ emission factor

In this study, only particulate matter with a diameter less than 2.5 micrometers ($PM_{2.5}$) was considered. The emission factor used for $PM_{2.5}$ is 0.0193 tons of $PM_{2.5}$ per ton of dry biomass burned. This value was derived from empirical measurements taken during a series of extreme wildfire events that occurred in Portugal in the summer of 2009 [28].

2.5 Selection of Monitoring Stations within the 50 km Radius

To assess the short-term impact of wildfire emissions on ambient $PM_{2.5}$ concentrations, we selected ground-based air quality monitoring stations from the regional network located within a 50 km radius of the perimeters of the Andilla and Cortes de Pallás wildfires. This distance threshold was not only consistent with previous studies reporting measurable smoke transport from wildfires across tens of kilometers [29], [30] but also ensured the inclusion of key urban areas that were affected by the smoke plume. In particular, the 50 km buffer was defined to include the city of Valencia and its surrounding metropolitan area which suffered smoke intrusion during the fire period [31] (Figure 1).

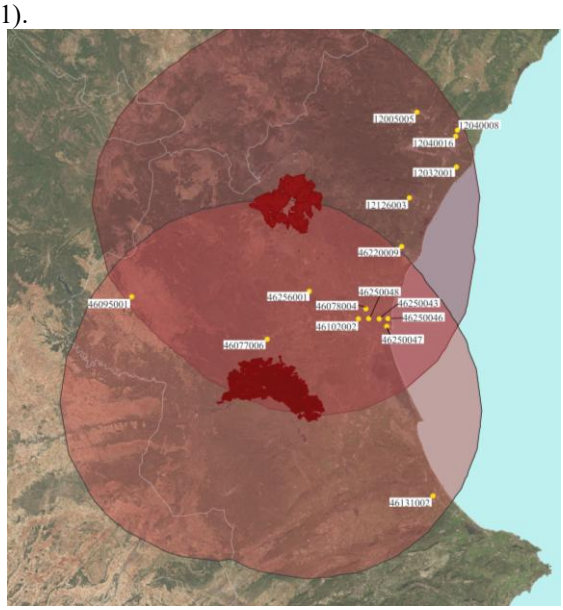

Fig. 2. Location of air quality monitoring stations within 50 km buffers of the Andilla and Cortes de Pallás wildfire perimeters.

A total of 16 air quality monitoring stations from the Valencian regional network were identified within these buffer zones. For each station, we extracted daily $PM_{2.5}$ data for the year 2012 and identified the maximum daily concentration during the fire period (June 28 – July 2, 2012). These peak values were then compared to:

- The station annual average for 2012,
- The 99th percentile of the annual $PM_{2.5}$ distribution, and
- The WHO 24-hour guideline value for $PM_{2.5}$ (15 $\mu g/m^3$) [13].

2.6 Reverse Spatial Approach: Estimating Emissions within Station Buffers

In a complementary analysis, we applied a reverse spatial approach to investigate whether observed pollution peaks could be predicted by the amount of wildfire emis-

sions surrounding each monitoring site. For this, we generated a 50 km buffer around each air quality station and calculated:

- The total area affected by wildfire within each buffer (in hectares),
- The total $PM_{2.5}$ emissions within that area (in tons), and
- The average emissions per hectare (tons/ha).

By intersecting the wildfire perimeters with the station buffers, we quantified the spatial overlap between burned area and each station's area of influence. We then used Pearson's correlation analysis to assess the strength and direction of the relationship between $PM_{2.5}$ exceedances (as a percentage) and the wildfire metrics (total burned area, total emissions, and emissions per hectare).

This two-part methodology enabled us to explore the spatial link between wildfire events and air quality impacts from two complementary directions: from the fires to the stations (direct influence) and from the stations to the fires (exposure estimation).

3 Results and discussion

The Andilla fire had higher pre-existing biomass (1,014,770 tons) and consumed biomass (665,145.07 tons) compared to the Cortes de Pallás event (820,648 and 519,134.71 tons, respectively). This resulted in greater $PM_{2.5}$ emissions for Andilla (12,837.30 tons) than for Cortes de Pallás (10,019.30 tons). The results illustrate the direct relationship applied during the methodology to convert the amount of biomass consumed into $PM_{2.5}$ emissions using a scale factor.

Table 6. Total pre-fire biomass, consumed biomass and $PM_{2.5}$ emissions in tones

Wildfire	Pre-existing biomass	Consumed biomass	$PM_{2.5.}$ emissions
Andilla	1,014,770	665,145.07	12,837.30
Cortes de Pallás	820,648	519,134.71	10,019.30

The spatial distribution maps (Figure 3) and box-and-whisker plots (Figure 4), offer a comprehensive view of the $PM_{2.5}$ emission patterns for the Andilla and Cortes de Pallas wildfires. The maps reveal notable contrasts in emission intensity and spatial heterogeneity. The Andilla wildfire shows extensive areas with high $PM_{2.5}$ emissions, that surpass 1.5 $t \cdot ha^{-1}$. In contrast, emissions in the Cortes de Pallás wildfire are predominantly below 1 $t \cdot ha^{-1}$ with no areas surpassing this threshold. These visual observations are supported by the statistical distribution shown in the boxplot. The Andilla fire displays a wider range of $PM_{2.5}$ values and a larger interquartile range, with several high outliers, reflecting localized emission hotspots. This indicates considerable variability in fire behavior and fuel characteristics across the burned landscape. Conversely, the Cortes de Pallás fire exhibits a more symmetric and constrained distribution, with a lower median and fewer extreme values, suggesting more homogeneous combustion conditions and less quantity of pre-existing biomass.

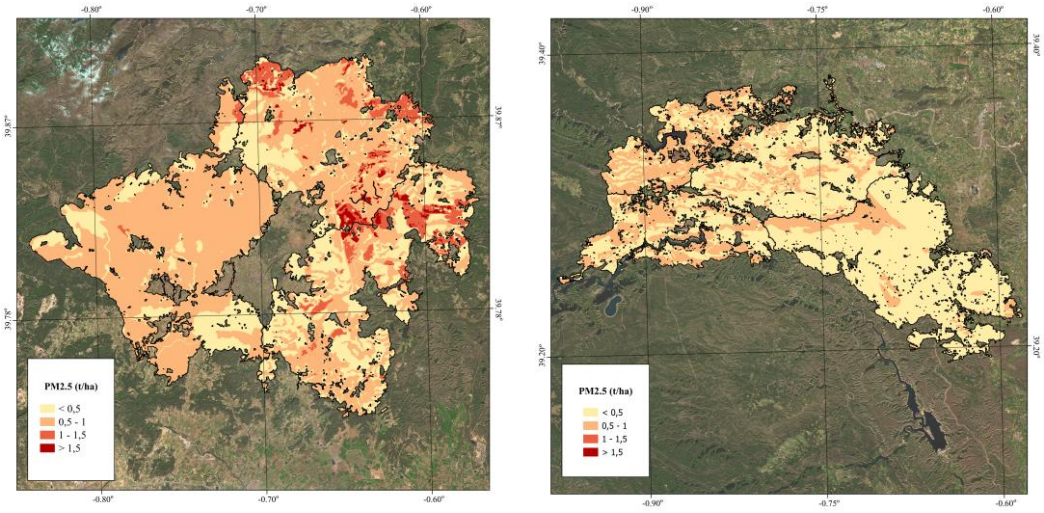

Andilla wildfire Cortes de Pallás wildfire

Fig. 3. Spatial distribution of $PM_{2.5}$ emissions ($t \cdot ha^{-1}$) from the Andilla (left) and Cortes de Pallás (right) wildfires. Emission intensity is categorized into four classes, highlighting areas with the highest particulate matter release. The maps illustrate differences in emission patterns, density, and spatial extent between the two fire events.

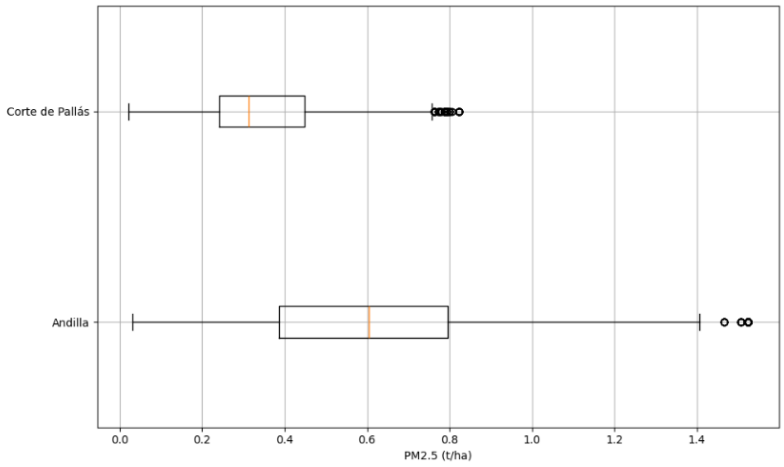

Fig. 4. Box-and-whisker plots of $PM_{2.5}$ emissions per hectare ($t \cdot ha^{-1}$) for the Andilla and Cortes de Pallás wildfires. The plots show the distribution, interquartile range, median, and outliers of emission values across the burned areas. Andilla exhibits greater variability and higher maximum values, indicating more heterogeneous fire behavior and fuel consumption compared to Cortes de Pallás.

To assess how wildfire emissions translated into air quality observations, we analyzed air monitoring stations located within a 50km radius of the Andilla and Cortes de Pallás fire perimeters. Table 7 provides baseline $PM_{2.5}$ concentrations from these stations in 2012, including average values and the 99th percentile, along with their distances from both wildfire sites.

The average background $PM_{2.5}$ levels range from 7.30 to 19.76 µg/m³, while the 99th percentile values reach up to 63.81 µg/m³, highlighting substantial variability in ambient air quality across the network. Several stations, such as 46256001 (located 18.60 km from Andilla and 22.39 km from Cortes de Pallás), lie well within the potential influence zone of wildfire smoke, making them critical for assessing short-range impacts. Stations that are farther from one wildfire are generally closer to the other, ensuring all selected stations are within 50 km of at least one fire perimeter. For example, station 46131002, although farther from Andilla, is located within 50 km of Cortes de Pallás, which is why it was included in this focused analysis.

These proximity data are essential for interpreting $PM_{2.5}$ anomalies observed during the fire periods and determining the extent to which local wildfire emissions influenced regional air quality. The high-end percentile values further serve as a reference for identifying exceedances potentially linked to smoke transport and accumulation during the fire events.

Table 7. Normal values for each air monitoring station for 2012

Station code	2012 average	Percentile 99	Distance to Andilla fire (km)	Distance to C. Pallás fire (km)
12005005	10,73	25,00	39,71	88,74
12032001	10,31	25,80	43,69	81,83
12040008	10,89	31,00	47,80	91,16
12040016	12,76	31,43	46,55	89,34
12126003	12,87	25,24	27,61	64,82
46077006	7,30	25,59	35,32	5,95
46078004	18,20	63,81	30,68	29,22
46095001	11,21	32,85	47,29	39,26
46102002	14,27	52,96	32,00	25,37
46131002	16,19	61,32	92,92	43,77
46220009	11,48	29,03	28,28	51,23
46250043	15,93	41,99	35,85	30,87
46250046	14,19	37,00	37,61	33,10
46250047	13,00	41,44	39,20	31,36
46250048	19,76	52,90	33,73	28,08
46256001	13,20	44,25	18,60	22,39

Building upon this baseline, Table 8 presents $PM_{2.5}$ concentrations recorded during the wildfire events. Wildfire period averages ranged from 9.90 to 29.74 µg/m³, with peak concentrations reaching as high as 75 µg/m³, clearly demonstrating the substantial impact of wildfire smoke on air quality. Stations closer to the fire perimeters, such

as 46256001 and 46250048, not only exhibited elevated wildfire-period averages but also some of the highest peaks, underscoring the influence of proximity on smoke exposure intensity. Additionally, the total $PM_{2.5}$ burden within the 50 km radius, expressed in tons, varies widely among stations—from a few hundred to over 22,000 tons—reflecting the spatial heterogeneity of wildfire emissions and their dispersion. When compared to the baseline conditions, these wildfire-period increases confirm the significant increase in particulate matter due to wildfire activity, emphasizing the importance of continuous monitoring to capture these episodic air quality deteriorations.

Table 8. Wildfire period values for each air monitoring station

Station code	Wildfire average	Peak	Total PM2.5 at 50km radius
12005005	9,90	22	6251,01
12032001	11,81	24	2691,53
12040008	13,62	28	395,86
12040016	14,25	26	1000,68
12126003	15,47	37	12752,64
46077006	11,57	33	20847,58
46078004	24,68	53	22369,31
46095001	16,19	45	4893,87
46102002	27,22	75	22245,18
46131002	27,80	62	912,85
46220009	13,50	32	12813,62
46250043	22,82	47	17181,50
46250046	17,76	38	14268,36
46250047	19,75	49	13354,62
46250048	29,74	71	20362,35
46256001	21,86	47	22855,25

To further quantify the severity of wildfire impacts on air quality, Table 9 presents the percentual exceedance of peak $PM_{2.5}$ concentrations observed during the wildfire period compared to three benchmarks: the 2012 annual average values, the World Health Organization (WHO) recommended a daily mean guideline of 15 $\mu g/m^3$, and the 99th percentile baseline concentrations from Table 2.

The data reveals that peak $PM_{2.5}$ levels at all stations substantially exceeded their respective 2012 averages, with exceedances ranging from approximately 105% up to over 425% (e.g., station 46102002). When compared to the WHO guideline, many stations showed exceedances well over 100%, with some surpassing 400%, underscoring the significant public health risks associated with wildfire smoke exposure.

Interestingly, exceedances relative to the 99th percentile baseline values were more variable. While several stations such as 12126003, 46077006, and 46102002 exhibited positive exceedances—indicating peak wildfire $PM_{2.5}$ surpassed typical high-end

concentrations—others showed negative values, reflecting that their peak wildfire concentrations did not always exceed historically high episodic levels. This variability likely reflects differences in local pollution sources and background variability alongside wildfire smoke influence. Overall, these exceedance metrics highlight the episodic yet intense nature of wildfire-driven air pollution events and emphasize the importance of monitoring and preparedness to mitigate health impacts during wildfire seasons.

Table 9. Percentual exceedance of the peak to normal values

Station code	Exc. 2012 average	Exc. OMS recommendation	Exc. 99th percentile
12005005	104,99	46,67	-12,00
12032001	132,78	60,00	-6,98
12040008	157,21	86,67	-9,68
12040016	103,70	73,33	-17,28
12126003	187,52	146,67	46,59
46077006	352,16	120,00	28,96
46078004	191,19	253,33	-16,94
46095001	301,35	200,00	36,99
46102002	425,68	400,00	41,62
46131002	282,86	313,33	1,11
46220009	178,67	113,33	10,23
46250043	194,99	213,33	11,93
46250046	167,74	153,33	2,70
46250047	277,00	226,67	18,24
46250048	259,27	373,33	34,22
46256001	256,10	213,33	6,21

Our analysis revealed a moderate positive correlation (Pearson's *r* = 0.52) between the exceedance of WHO's daily $PM_{2.5}$ guideline (daily mean of 15 μg/m³) and total wildfire-related $PM_{2.5}$ emissions, suggesting that monitoring stations experiencing more severe acute pollution peaks also tended to be spatially closer to a higher total of wildfire emissions. However, a univariate logarithmic regression model (Figure 5) explained only 32% of the variance in emissions (R^2 = 0.32), highlighting the limitations of relying solely on exceedance metrics for predictive purposes.

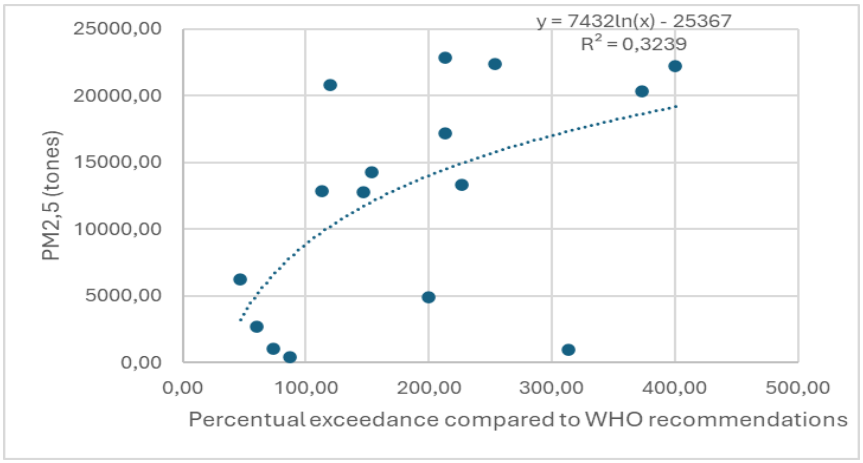

Fig. 5. $PM_{2.5}$ emissions versus WHO guideline exceedance (%) at monitoring stations. Logarithmic trend (R^2=0.32) shows moderate correlation, with variability indicating additional dispersion factors

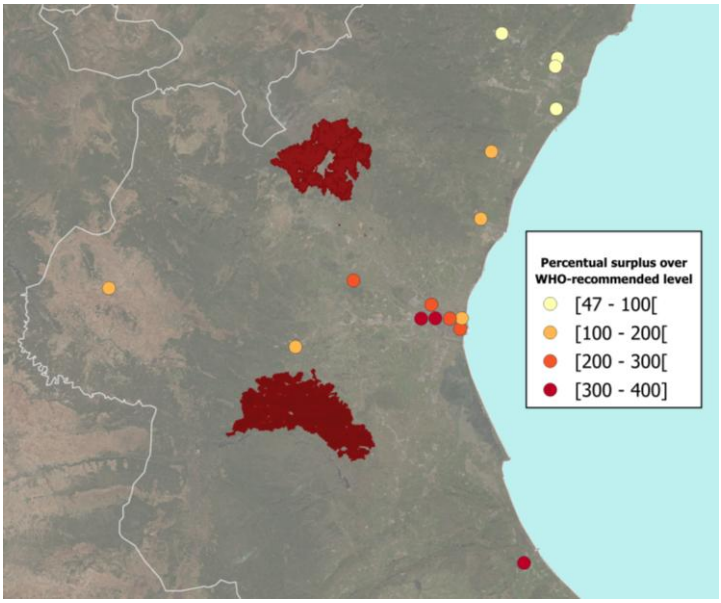

Fig. 6. Percent exceedance of WHO daily $PM_{2.5}$ guidelines (15 µg/m³) at monitoring stations.

Spatiotemporal patterns of $PM_{2.5}$ exceedances relative to WHO guidelines (Figure 6) reveal the limitations of distance-based assumptions. These anomalies align with our nonlinear regression model (Fig. 5, $R^2 = 0.32$), which confirms that acute exceedances explain only a fraction of total emissions variability.

Tools like HYSPLIT, CALPUFF, or WRF-Chem could elucidate how factors such as 1) Wind shear and boundary-layer dynamics alter plume advection, 2) Topographic channeling concentrates smoke in valleys and, 3) Fire behavior (flaming vs. smoldering phases) affects particle vertical mixing. Coupling these physical models with machine learning (e.g., random forests trained on dispersion outputs) could disentangle source contributions from background pollution, improving attribution of $PM_{2.5}$ anomalies to wildfire emissions versus local anthropogenic sources.

4 Conclusions

This study compared the Andilla and Cortes de Pallás wildfires, revealing that Andilla produced significantly higher $PM_{2.5}$ emissions due to greater biomass consumption and more variable fire behavior. Emission maps showed higher intensity and spatial heterogeneity for Andilla, while Cortes de Pallás displayed more uniform combustion. Air quality data within 50 km of both fires showed sharp increases in $PM_{2.5}$, with some stations exceeding WHO guidelines by over 400%. However, exceedance patterns did not consistently align with proximity, highlighting the influence of wind, topography, and combustion phase. A moderate correlation between emissions and $PM_{2.5}$ peaks suggests that exposure is driven by complex dispersion dynamics. To improve forecasting and attribution, future work should integrate physical models with machine learning approaches for better wildfire smoke impact assessment.

References

1. F. Moreira *et al.*, "Wildfire management in Mediterranean-type regions: paradigm change needed," *Environmental Research Letters*, vol. 15, p. 11001, 2020, doi: 10.1088/1748-9326/ab541e.
2. J. E. Keeley and A. D. Syphard, "Large California wildfires: 2020 fires in historical context," *Fire Ecology*, vol. 17, no. 1, p. 22, Dec. 2021, doi: 10.1186/s42408-021-00110-7.
3. J. G. Pausas, J. Llovet, A. Rodrigo, and R. Vallejo, "Are wildfires a disaster in the Mediterranean basin? - A review," *Int J Wildland Fire*, vol. 17, no. 6, p. 713, 2008, doi: 10.1071/WF07151.
4. F. Moreira *et al.*, "Landscape wildfire interactions in southern Europe: implications for landscape management," *J Environ Manag*, vol. 92, 2011, doi: 10.1016/j.jenvman.2011.06.028.
5. R. Delgado-Artés, V. Garófano-Gómez, J.-V. Oliver-Villanueva, and E. Rojas-Briales, "Land use/cover change analysis in the Mediterranean region: a regional case study of forest evolution in Castelló (Spain) over 50 years," *Land use policy*, vol. 114, p. 105967, Mar. 2022, doi: 10.1016/j.landusepol.2021.105967.
6. NOAA, Jul. 24, 2023. [Online]. Available: https://www.noaa.gov/noaa-wildfire/wildfire-climate-connection. [Accessed: Jun. 11, 2025].
7. S. El Garroussi, F. Di Giuseppe, C. Barnard, and F. Wetterhall, "Europe faces up to tenfold increase in extreme fires in a warming climate," *NPJ Clim Atmos Sci*, vol. 7, no. 1, p. 30, Jan. 2024, doi: 10.1038/s41612-024-00575-8.

8. M. M. Boer, V. R. De Dios, E. Z. Stefaniak, and R. A. Bradstock, “A hydroclimatic model for the distribution of fire on Earth,” *Environ Res Commun*, vol. 3, no. 3, p. 035001, Mar. 2021, doi: 10.1088/2515-7620/abec1f.
9. A. Haikerwal *et al.*, “Impact of Fine Particulate Matter ($PM_{2.5}$) Exposure During Wildfires on Cardiovascular Health Outcomes,” *J Am Heart Assoc*, vol. 4, no. 7, Jul. 2015, doi: 10.1161/JAHA.114.001653.
10. F. Johnston, I. Hanigan, S. Henderson, G. Morgan, and D. Bowman, “Extreme air pollution events from bushfires and dust storms and their association with mortality in Sydney, Australia 1994–2007,” *Environ Res*, vol. 111, no. 6, pp. 811–816, Aug. 2011, doi: 10.1016/j.envres.2011.05.007.
11. H. C. Phuleria, P. M. Fine, Y. Zhu, and C. Sioutas, “Air quality impacts of the October 2003 Southern California wildfires,” *Journal of Geophysical Research: Atmospheres*, vol. 110, no. D7, Apr. 2005, doi: 10.1029/2004JD004626.
12. F. Reisen, D. Hansen, and C. P. (Mick) Meyer, “Exposure to bushfire smoke during prescribed burns and wildfires: Firefighters’ exposure risks and options,” *Environ Int*, vol. 37, no. 2, pp. 314–321, Feb. 2011, doi: 10.1016/j.envint.2010.09.005.
13. World Health Organization, *WHO global air quality guidelines*. 2021.
14. C. E. Reid, M. Brauer, F. H. Johnston, M. Jerrett, J. R. Balmes, and C. T. Elliott, “Critical Review of Health Impacts of Wildfire Smoke Exposure,” *Environ Health Perspect*, vol. 124, no. 9, pp. 1334–1343, Sep. 2016, doi: 10.1289/ehp.1409277.
15. D. E. Newby *et al.*, “Expert position paper on air pollution and cardiovascular disease,” *Eur Heart J*, vol. 36, no. 2, pp. 83–93, Jan. 2015, doi: 10.1093/eurheartj/ehu458.
16. R. D. Brook *et al.*, “Particulate Matter Air Pollution and Cardiovascular Disease,” *Circulation*, vol. 121, no. 21, pp. 2331–2378, Jun. 2010, doi: 10.1161/CIR.0b013e3181dbece1.
17. S. Rajagopalan *et al.*, “Personal-Level Protective Actions Against Particulate Matter Air Pollution Exposure: A Scientific Statement From the American Heart Association,” *Circulation*, vol. 142, no. 23, Dec. 2020, doi: 10.1161/CIR.0000000000000931.
18. G. Chen *et al.*, “Mortality risk attributable to wildfire-related PM2·5 pollution: a global time series study in 749 locations,” *Lancet Planet Health*, vol. 5, no. 9, pp. e579–e587, Sep. 2021, doi: 10.1016/S2542-5196(21)00200-X.
19. C. E. Reid *et al.*, “Spatiotemporal Prediction of Fine Particulate Matter During the 2008 Northern California Wildfires Using Machine Learning,” *Environ Sci Technol*, vol. 49, no. 6, pp. 3887–3896, Mar. 2015, doi: 10.1021/es505846r.
20. AEMET, “Resumen anual climatológico,” 2012.
21. V. Lerma-Arce, H. Van den Berg, J. V. Oliver-Villanueva, and E. P. Coll-Aliaga, “Cartografía Territorial del Stock de Carbono en la Comunitat Valenciana,” 2020. [Online]. Available: https://politicaterritorial.gva.es/documents/20551069/174233262/Cartograf%C3%ADa+Territorial+del+Stock+de+Carbono+en+la+Comunitat+Valenciana.pdf/7e2501f8-2737-426e-80c5-617f33e98f36?t=1627301770832
22. SIOSE, “Descripción del modelo de datos SIOSE,” 2015.
23. MITECO, *Tercer Inventario Forestal Español*. 2006.
24. A. De Santis, G. P. Asner, P. J. Vaughan, and D. E. Knapp, “Mapping burn severity and burning efficiency in California using simulation models and Landsat imagery,” *Remote Sens Environ*, vol. 114, no. 7, pp. 1535–1545, Jul. 2010, doi: 10.1016/j.rse.2010.02.008.
25. J. Chaves-Naharro, “Incendios forestales y cambio climático. Metodología de cálculo de emisión de gases de efecto invernadero. Estimación de emisiones en los incendios de Andilla y Cortes de Pallás.,” 2014, *València*.
26. C. Key and N. Benson, “Landscape Assessment (LA) Sampling and Analysis Methods,” 2006, p. LA 1-51.

27. EFFIS, “Fire Severity.” Accessed: Jun. 12, 2025. [Online]. Available: https://forest-fire.emergency.copernicus.eu/about-effis/technical-background/fire-severity
28. C. Alves *et al.*, “Summer 2009 wildfires in Portugal: Emission of trace gases and aerosol composition,” *Atmos Environ*, vol. 45, no. 3, pp. 641–649, Jan. 2011, doi: 10.1016/j.atmosenv.2010.10.031.
29. A. Sapkota *et al.*, “Impact of the 2002 Canadian Forest Fires on Particulate Matter Air Quality in Baltimore City,” *Environ Sci Technol*, vol. 39, no. 1, pp. 24–32, Jan. 2005, doi: 10.1021/es035311z.
30. L. J. DeBell, R. W. Talbot, J. E. Dibb, J. W. Munger, E. V. Fischer, and S. E. Frolking, “A major regional air pollution event in the northeastern United States caused by extensive forest fires in Quebec, Canada,” *Journal of Geophysical Research: Atmospheres*, vol. 109, no. D19, Oct. 2004, doi: 10.1029/2004JD004840.
31. C. Vázquez and L. Ortega, “2000 personas combaten el fuego que devora el interior de Valencia,” *El Pais*, Valencia/Castellón, Jul. 02, 2012.

Development of an algorithm to compensate for the effect of heart motion from respiration in solving the inverse problem of electrocardiography

María Garrido[1], Ernesto Zacur[1,2], Ismael Hernández[1,2], Andreu Climent[1,2], and Maria Guillem[1,2]

[1] ITACA Institute, Universitat Politècnica de València, Spain
[2] Corify Care SL, Madrid, Spain
mgarper5@teleco.upv.es

Abstract. Respiratory motion induces displacements in the position of the heart that affect the electrical information that is measurable in the torso and should be considered to ensure the accuracy of the electrical maps obtained. In the present work, the effect of heart motion on electrocardiographic signals has been evaluated, and strategies have been developed to minimize the effects caused by heart motion during respiration within the framework of indirect estimation of cardiac activity from electrocardiographic measurements.

Keywords: Electrocardiography · Inverse Problem · Respiratory Motion · Heart Motion Compensation · Cardiac Mapping · Body Surface Potentials.

1 Introduction

We aim to evaluate the effects of respiration-induced cardiac motion on the electrical potential of the body surface and its effects on the estimation of electrical activity at the heart [1]. To assess this, we conducted a series of simulations and a series of measurements on healthy subjects, including their electrical activity using Acorys and their respiratory activity across different breathing patterns.

First, we will simulate sequences of synthetic heartbeats by mimicking the effects of respiration in a realistic way. We will make use of these synthetic signals to find a way to evaluate the motion of the heart and to choose a reliable metric that correlates heartbeats with the respiratory phase. The second part involves the acquisition and processing of the measurements obtained from healthy volunteers, on which the previously developed methods will be applied. The last part focuses on developing a method to take into account the respiratory motion in the estimation and visualization of cardiac electrical activity and to compensate its effect.

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

1.1 The Inverse Problem of Electrocardiography

The inverse problem in electrocardiography is used to estimate the electrical potentials of the heart from measurements recorded on the body surface [2]. This process is fundamental for obtaining non-invasive activation maps, such as those generated by the ACORYS® system. However, it is an intrinsically ill-posed and under-determined problem, where we have fewer metrics in the body surface (128) than in the epicardium nodes (≈ 4000), so its solution must be regularized. Therefore, small errors or noise in the measurements can lead to large deviations in the results. Adequately solving the inverse problem is essential for reconstructing accurate maps of cardiac activity and requires applying soft constraints to guide the solution towards realistic scenarios.

2 Artificial Heart Movement

This work aims to quantify and correct the effect of respiration-induced heart motion on the heart's electrical activity. To achieve this, synthetic heartbeats are simulated by applying translations and rotations to the heart's geometry along its Z axis (cranio-caudal direction), which allows for representing its displacement during respiration. These transformations are integrated into the calculation of the transfer matrices (A), generating ECGs that reflect this motion. It is observed in Fig. 1 that this motion modulates the amplitude of the signals, which can distort activation maps if respiratory phases are not considered and we average heartbeats that cannot be compared.

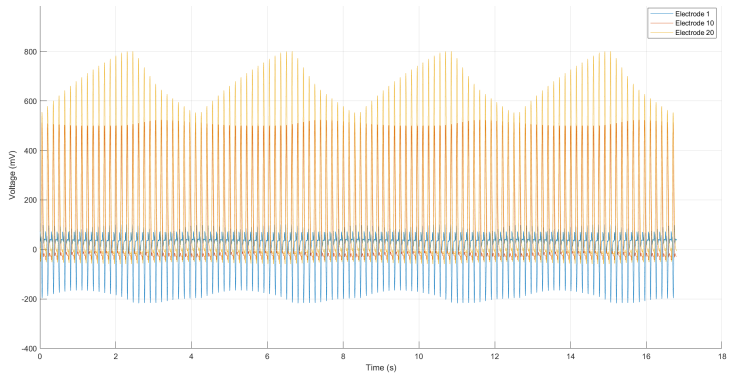

Fig. 1: Simulated electrical signals for 3 leads accounting for respiratory motion

2.1 Verification of the Heart Trajectory

To estimate the heart's trajectory during respiration, we will use the Opoint algorithm. This algorithm calculates 15 possible positions for the cardiac center

based on body surface potentials (BSPMs). These estimations vary over time when respiratory motion is simulated. Comparing the results with the static case, it is observed that motion alters the estimated positions. Finally, the distance between a fixed point on the heart and the curve estimated by Opoint is analyzed, observing that this distance remains almost constant only if the point moves along the same respiratory trajectory that we designed for computing the "breathed" EGMs, as we can see in Fig. 2 for the blue line.

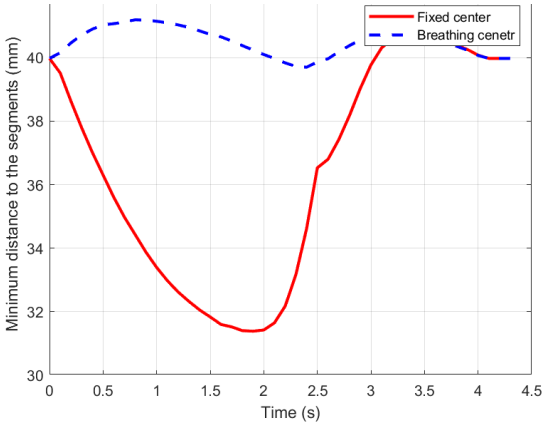

Fig. 2: Distance between an average point of the heart and the midpoint of the curve estimated by the Opoint for the case in which the heart is static (red) and when it's moving (blue)

2.2 Estimation of the Respiratory Phase

Until now, the calculation of electrograms (EGMs) and activation times (ATs) has been performed without considering the respiratory phase of each heartbeat. The next step is to detect this phase to apply more precise corrections. For this purpose, the Z-coordinate (cranio-caudal direction) of the midpoint of the 15 cardiac positions estimated by the Opoint algorithm is used as a metric, as it reflects displacement from head to toe of the heart during respiration. This value allows for identifying the respiratory cycle phase of each heartbeat, which facilitates appropriate triggering. Although this metric is smooth in ideal simulations, it is expected to also function on real signals, although it will be noisier.

3 Experimental Developments

Real electrical signals were acquired from two healthy volunteers using the ACORYS® system. This system includes a vest with 128 electrodes and a 3D recon-

struction of the subject's torso to accurately map surface potentials. Measurement conditions were carefully controlled to ensure signal stability. Tests were conducted under different respiratory patterns, with the best results observed in scenarios combining apnea and deep breathing.

3.1 Respiratory phase from Z metric

To estimate the respiratory phase of each heartbeat, the Z-coordinate of the cardiac center estimated by the Opoint algorithm was used, applied exclusively to the QRS segments to focus on ventricular activity. This approach allows for identifying whether the heartbeat occurred during inspiration or expiration, based on the heart's vertical position. In real signals, the average Z-coordinate was extracted for each heartbeat and compared with the ECG signal envelope, showing a clear correlation. During apnea, the estimated trajectory of Z-coordinates remained constant (≈ 5 mm), whereas in deep breathing, it presented oscillations of up to 25 mm. These results validate the use of this metric as a reference for respiratory phase-based triggering and show an undeniable correlation between both the position of the heart and the electric signals.

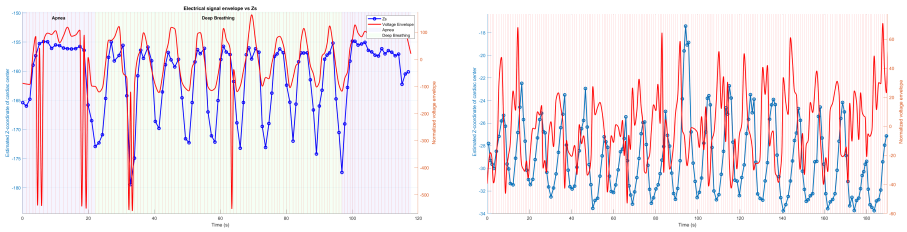

(a) Z trajectory of the cardiac center during apnea and deep breathing (blue) vs. envelope of the electrical signal (red).

(b) Z trajectory of the cardiac center during deep breathing (blue) vs. envelope of the electrical signal (red).

Fig. 3: Estimation of the cardiac center trajectory (Z coordinate) for different breathing patterns and electrical signal envelopes.

4 Improvements

4.1 Triggering

By estimating the cardiac center's position relative to the respiratory phase, a "triggering" approach averages only heartbeats within the same phase, aiming for more accurate electrograms (EGMs) and activation times (ATs). We identify the phase by inspecting the Zs metric, computed previously, specifically at its maximum or minimum points, which correspond to heartbeats within the same

respiratory phase. This phase-specific averaging yields clearer activation maps compared to random averaging, at least in the synthetic case, where we were able to reproduce a long respiratory session. This highlights the necessity of respiratory motion compensation to improve the precision of cardiac electrical activity reconstructions.

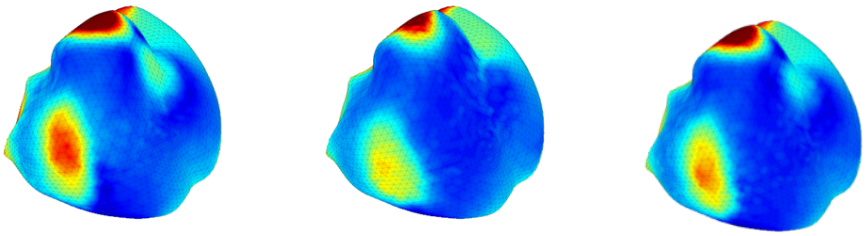

(a) AT map for Z peaks. (b) AT map for Z valleys. (c) AT map for any point.

Fig. 4: Activation time (AT) maps obtained by averaging beats at different respiratory phases in the synthetic case.

5 Results

By averaging heartbeats aligned to their respiratory phase (identified via cardiac Z-coordinate), increased cross-correlation between the selected heartbeats and SNR are achieved in synthetic and real data, alongside reduced variability and better QRS superposition. This phase-specific approach yields clearer activation maps, underscoring the need for motion compensation for accurate reconstruction of cardiac electrical activity.

References

1. Valderas, M., Vallverdú, M., Caminal, P., Schröder, R., Voss, A.: Extracción de la señal de respiración a partir del electrocardiograma. In: XXXVI Jornadas de Automática, Bilbao (2015).
2. Macfarlane, P.W., van Oosterom, A., Pahlm, O., Kligfield, P., Janse, M., Camm, J. (eds.): Comprehensive Electrocardiology. Springer, London, UK (2011).

Microwave Melting of Lunar Regolith: Toward Sustainable In-Situ Construction on the Moon

Beatriz García-Baños[1], Caroline O'Connell[2], Adrian Miró-Sanz[1], Sungwoo Lim[2], Jose M. Catalá-Civera[1], Andrea Lucca-Fabris[2]

[1] ITACA Institute, Universitat Politècnica de València, Camino de Vera s/n, 46022, Valencia, Spain
beagarba@upvnet.upv.es
[2] Surrey Space Centre, University of Surrey, Guildford, Surrey GU2 7XH, United Kingdom
c.o'connell@surrey.ac.uk

Abstract. Microwave processing of lunar regolith presents a promising approach for in-situ construction on the Moon, reducing the need for terrestrial material transport. This study investigates the dielectric behavior of a lunar regolith simulant as a function of temperature, with particular focus on its suitability for microwave-induced melting. Results show a significant increase in both dielectric constant and loss factor with temperature, especially beyond the glass transition (~650 °C) and melting point (~990 °C), where molecular and ionic mobility markedly enhance microwave absorption. These findings suggest a self-reinforcing heating effect that can improve energy efficiency and process speed. This knowledge has important implications for the design of microwave reactors capable of operating under lunar conditions, offering a potential solution for sustainable, on-site fabrication of structural materials for extraterrestrial habitats.

1 Introduction

Establishing a sustained human presence on the Moon is one of the key objectives in the current era of space exploration. One of the main challenges in achieving this goal is the need to develop construction methods that minimize dependence on Earth-supplied materials [1-3]. In-situ resource utilization (ISRU), which involves harnessing local resources such as lunar regolith, has emerged as a promising strategy for building infrastructure directly on the lunar surface. Lunar regolith, a fine and highly reactive layer of crushed rock and mineral dust covering the Moon's surface, is abundant and can be transformed into solid building materials using thermal processing techniques [4]. Among these, microwave energy has gained attention as a particularly attractive solution due to its ability to deliver volumetric, selective, and efficient heating in a vacuum environment without the need for physical contact [1].

Microwave melting of regolith offers a path toward fabricating structural elements such as bricks, landing pads, or radiation-shielding walls. However, the successful implementation of this technique depends on a deep understanding of the material's response to microwave energy, especially at high temperatures approaching or exceed-

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

ing the melting point. This includes knowledge of the dielectric properties of lunar simulants and how they evolve during heating [1]. Additionally, key questions remain regarding the efficiency, scalability, and reliability of microwave-based systems under lunar conditions. This study aims to address these challenges by characterizing the microwave melting behavior of regolith simulants, evaluating the physical and dielectric transformations involved to assist assessing the viability of this method as a sustainable construction approach for future lunar habitats.

2 Materials and Methods

The mare simulant JSC-1A was analysed in this work, being one of the most well-known and high quality simulants of lunar regoliths. Its main composition is shown in Table I. The JSC-1A simulant is a powder/gray sandlike material comprised of crushed basalt, with 1 mm particle size and lower.

Table 1. Composition of the JSC-1A regolith simulant

Major element composition	CAS #	% by Weight
Silicon Dioxide (SiO2)	14808-60-7	46-49
Titanium Dioxide (TiO2)	13463-67-7	1-2
Aluminum Dioxide (Al2O3)	1344-28-1	14.5-15.5
Ferric Oxide (Fe2O3)	1309-37-1	3-4
Iron Oxide (FeO)	1332-37-2	7-7.5
Magnesium Oxide (MgO)	1309-48-4	8.5-9.5
Calcium Oxide (CaO)	1305-78-8	10-11
Sodium Oxide (Na2O)	1313-59-3	2.5-3
Potassium Oxide (K2O)	12136-45-7	0.75-0.85
Manganese Oxide (MnO)	1344-43-0	0.15-0-20
Chromium Oxide (Cr2O3)	1308-38-9	0.02-0.06
Diphosphorus Pentoxide (P2O5)	1314-56-3	0.6-0.7

2.1 Dielectric properties measurement

The dielectric properties of the simulant have been measured in the equipment described in [5], and following the procedure of the ASTM Method B [6]. The method is based on the measurement of some microwave magnitudes as the resonant frequency and quality factor of a microwave resonant-cavity test fixture with the test specimen being heated with microwave energy. The presence of the specimen modifies the resonance properties of the microwave resonant cavity and these changes are employed to determine the dielectric properties. The measurements are typically carried out by a network analyzer. The resonant cavity is cylindrical and dielectric specimens are placed in a cylindrical quartz tube. The test has an accuracy of ±(1-2)% for the dielectric constant and ±(2-5%) for the loss factor.

Figure 1 shows the microwave cylindrical cavity employed for the dielectric measurements. The sample of material is placed inside a quartz holder and simultaneously heated with microwaves up to 1100ºC. The sample temperature is continuously measured with an infrared pyrometer (not shown). The sample was heated from ambient temperature to melting temperature at a rate of 0.25ºC/s under a nitrogen gas environment.

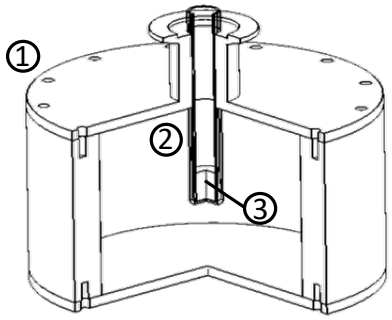

Fig. 1. Set up for dielectric characterization at microwave frequencies. 1) Microwave cylindrical cavity with insertion hole. 2) Quartz holder. 3) Regolith sample.

3 Results

Figure 3 presents the temperature-dependent dielectric properties of the regolith sample, specifically the real part of the permittivity (dielectric constant) and the dielectric loss factor. Both parameters exhibit a monotonic increase with temperature, indicating enhanced interaction of the material with the microwave field as thermal energy rises. A noticeable inflection point is observed around the glass transition temperature (~650 °C [Ref. 4, Caroline]), where the slope of increase becomes steeper. This behavior is attributed to the increased molecular mobility within the amorphous components of the regolith, which allows for more effective dipolar and interfacial polarization mechanisms to occur.

As the material approaches its melting temperature (~990 °C), it undergoes a transition from a rigid solid to a viscous or fully molten liquid phase. This structural breakdown significantly increases the freedom of movement of ions and polar species within the melt, further enhancing the dielectric loss and overall microwave absorption efficiency. The resulting increase in microwave energy coupling facilitates rapid and volumetric heating, which is advantageous for high-temperature in-situ processes such as sintering or melting for construction.

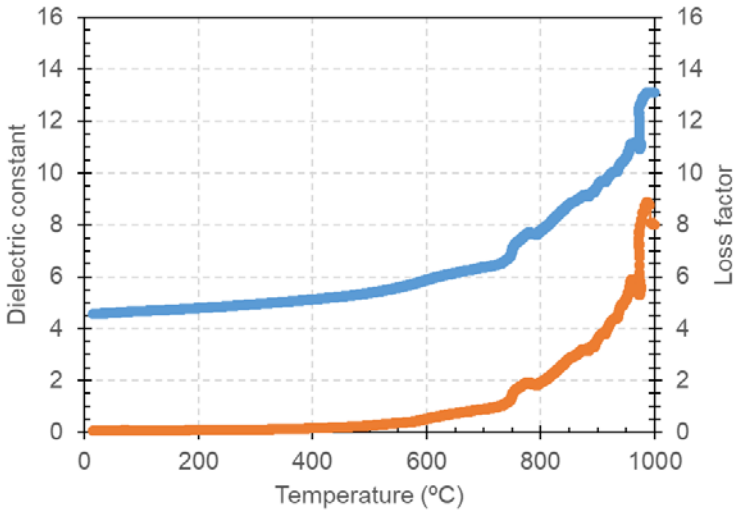

Fig. 2. Dielectric properties of regolith simulant sample as a function of temperature.

The observed increase in dielectric properties of the regolith with temperature, particularly above the glass transition and melting points, has significant implications for the design of microwave melting processes. The strong temperature dependence suggests that once initial heating is achieved, the material's ability to absorb microwave energy improves dramatically, leading to a self-reinforcing heating effect. This behavior enables the use of more energy-efficient and targeted heating strategies, potentially reducing overall power consumption and processing time. However, it also requires careful control of power input and temperature gradients to avoid thermal runaway or uneven melting, particularly due to the low thermal conductivity of this type of materials. The transition to a molten phase, with its markedly higher dielectric loss, indicates that process designs must accommodate changes in material state and ensure thermal and structural stability during operation.

4 Conclusion

This study has demonstrated the strong temperature dependence of the dielectric properties of lunar regolith simulants, highlighting their potential for efficient microwave processing in extraterrestrial construction applications. The observed increase in both the dielectric constant and loss factor with temperature—particularly above the glass transition and melting points—indicates that regolith becomes increasingly responsive to microwave energy as it heats. This behavior is advantageous for initiating and sustaining high-temperature processes such as sintering and melting, as it enables selective, volumetric heating with minimal energy losses.

These findings provide valuable design criteria for the development of microwave-based systems for in-situ resource utilization on the Moon. Understanding how dielec-

tric properties evolve during heating is essential for optimizing reactor configurations, energy delivery strategies, and process control protocols. The results support the feasibility of using microwave energy as a reliable and scalable method for converting lunar regolith into construction-grade materials in future lunar missions. Continued research into the real-time behavior of regolith under microwave fields will be key to advancing this technology and ensuring its successful application in space environments.

Acknowledgements

This work has been financed with the “Ayuda a Primeros Proyectos de Investigación (PAID-06-24)” of the Vicerrectorado de Investigación of the Universitat Politècnica de València.

References

1. S. Lim, J. Bowen, G. Degli-Alessandrini, M. Anand, A. Cowley, and V. Levin Prabhu, "Investigating the microwave heating behaviour of lunar soil simulant JSC-1A at different input powers," (in eng), Scientific reports, vol. 11, no. Jan (E-published), 01/22 2021, doi: 10.1038/s41598-021-81691-w.
2. Gerald B. Sanders and William E. Larson, “Progress Made in Lunar In Situ Resource Utilization under NASA’s Exploration Technology and Development Program”, J. Aerosp. Eng., 2013, 26(1): 5-17
3. Lukas Schlüter, Aidan Cowley, “Review of techniques for In-Situ oxygen extraction on the moon”, Planetary and Space Science 181 (2020) 104753
4. Nicholas James Bennett *, Damon Ellender , Andrew G. Dempster, “Commercial viability of lunar In-Situ Resource Utilization (ISRU)”, Planetary and Space Science 182 (2020) 104842.
5. J. Catala-Civera, A. Canos, P. Plaze-Gonzalez, J. Gutierrez, B. Garcia-Banos, and F. Penaranda-Foix, "Dynamic Measurement of Dielectric Properties of Materials at High Temperature During Microwave Heating in a Dual Mode Cylindrical Cavity," Microwave Theory and Techniques, IEEE Transactions on Microwave Theory and Techniques, 2015, doi: 10.1109/TMTT.2015.2453263.
6. Test methods for complex permittivity (Dielectric Constant) of solid electrical insulating materials at microwave frequencies and temperatures to 1650°” ASTM Standard D2520, American Society for Testing and Materials

DIGICOR
An Interactive Simulation Platform for Personalized Ablation Planning in Atrial Fibrillation

Elena Doallo[1], Clara Herrero[1], Rubén Molero[1,2], Andreu M. Climent[1,2], David Lundback[1,2], Raúl Moreno[1], and María S. Guillem[1,2]

[1] ITACA Institute, Universitat Politècnica de València, Spain
edoacer@alumni.upv.es
[2] Corify Care SL, Madrid, Spain

Abstract. Atrial fibrillation (AF) is the most common sustained cardiac arrhythmia, associated with significant clinical and economic burden. Although catheter ablation is an effective treatment, its success is limited in persistent cases due to the complexity of the underlying substrate. We present DIGICOR, a novel interactive platform that integrates electrocardiographic imaging (ECGI), computational heart models and digital twins to simulate AF and evaluate personalized ablation strategies. DIGICOR combines real-time simulation of atrial electrical dynamics using a GPU-accelerated cellular automaton with intuitive 3D visualization and functional map analysis. It allows users to interactively apply virtual lesions and assess their impact using metrics such as dominant frequency (DF), Stability map, and Reentrant Pattern Histogram. The system enables rapid comparison of scenarios and supports personalized treatment planning. DIGICOR is a step towards clinical integration of digital twins in electrophysiology workflows.

Keywords: Atrial fibrillation · Ablation · Functional Mapping · Digital Twin · Cellular Automaton

1 Introduction

Atrial fibrillation affects over 33 million people worldwide (Chugh et al., 2014) and is expected to double by 2060 (Krijthe et al., 2013). It results in ineffective atrial contraction, increasing stroke risk and reducing quality of life. Current ablation strategies are primarily guided by anatomical protocols, which often lack sufficient patient-specific adaptation (Boyle et al., 2019). This limitation is especially evident in persistent atrial fibrillation, where the underlying substrate exhibits complex structural and electrical remodeling. Although non-invasive tools such as electrocardiographic imaging (ECGI) and computational modeling offer considerable potential for enhancing treatment planning, their integration into clinical workflows remains a major challenge (Molero, 2023).

Digital twins have emerged as a powerful concept in personalized medicine, enabling simulation of disease dynamics and testing of interventions in a virtual environment. DIGICOR addresses key barriers to clinical translation by focusing on ease of use, fast simulation, and visual interpretation.

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

2 Objectives

The main goal of this work is to develop DIGICOR, an interactive simulation platform that facilitates the personalized planning of catheter ablation in atrial fibrillation. DIGICOR aims to bridge the gap between advanced computational modeling and clinical workflows by providing a user-friendly interface and real-time simulation capabilities. The specific objectives of the platform are as follows:

- To implement a computational framework that models atrial electrophysiology based on personalized 3D geometries and tunable simulation parameters.
- To allow the visualization of electrical propagation dynamics and the generation of functional maps such as Dominant Frequency, Stability Map, and Reentrant Pattern Histogram.
- To support the interactive application of ablation lesions and evaluation of their effects in near real-time, enabling rapid iteration of therapeutic strategies.
- To provide a modular architecture that integrates C++, Python, and MATLAB components, ensuring performance, flexibility, and extensibility.
- To include a scenario comparison system that facilitates the evaluation of multiple ablation strategies side by side, enhancing decision-making in a clinical research context.

3 Methods

DIGICOR integrates multiple technologies to provide an interactive simulation experience. As shown in Fig. 1, the platform's graphical interface is organized into several regions: the central panel displays the selected atrial geometry, while the bottom panel provides interactive controls to restart or play the simulation, navigate through frames, and edit ablation masks. On the right, the selection panel shows simulation status and available geometries.

The core simulation engine is implemented in C++ and CUDA, using a cellular automaton model specifically developed by our research group. This model simulates electrical propagation over 3D atrial geometries and supports configurable parameters such as diffusion coefficients, basal activation frequency, and region-specific heterogeneity. To ensure modularity and performance, the simulation engine has been compiled into a dynamic-link library (DLL) that integrates efficiently with the user interface.

The graphical user interface (GUI), developed in Python using the PyQt5 framework, manages the coordination between user inputs, simulation state, and rendering. Users can load geometries, adjust simulation parameters, apply virtual ablation lesions, and visualize results interactively.

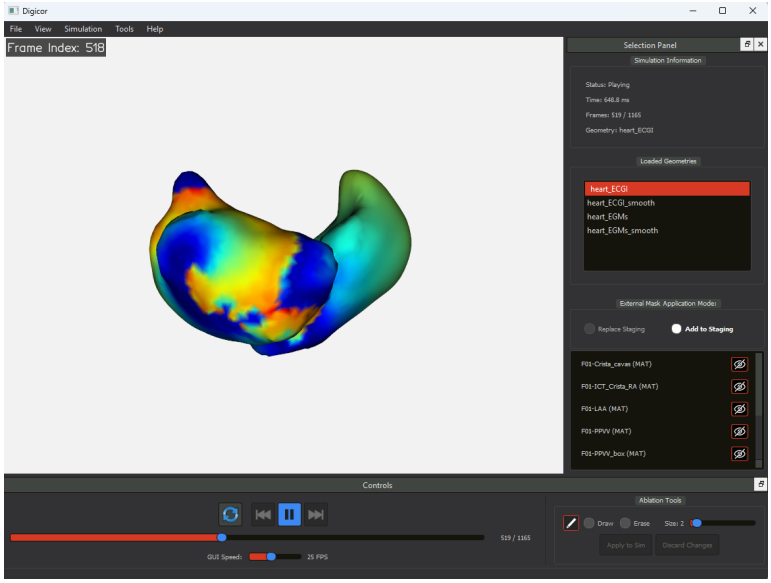

Fig. 1. Main DIGICOR interface during an atrial fibrillation simulation.

Signal preprocessing and functional map generation are performed using MATLAB scripts, also developed by our research group. These scripts analyze the activation signals and compute relevant metrics such as dominant frequency, frequency stability, and reentrant activity indicators. As shown in Fig. 2, these functional maps can be visualized simultaneously and provide insight into arrhythmic activity patterns.

Finally, real-time rendering is handled using the VTK and Vedo libraries, allowing users to visualize the evolution of activation on the atrial surface, rotate and zoom the geometry, and compare multiple ablation strategies side by side within synchronized viewer panels.

DIGICOR also incorporates a scenario management system that enables users to save, reload, and compare simulation states with different ablation configurations. This feature supports iterative experimentation and facilitates treatment strategy evaluation within a single session.

4 Results

DIGICOR enables the interactive exploration of atrial activation patterns under different physiological conditions and ablation strategies. Simulations are fast enough to provide feedback within seconds, supporting a near real-time workflow. The system supports both single-scenario mode and a comparison mode in which multiple strategies can be viewed.

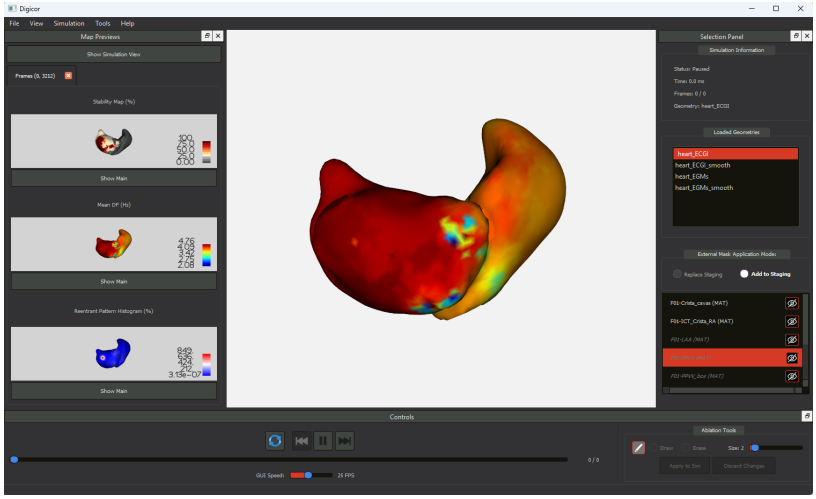

Fig. 2. Functional map visualization in DIGICOR. The interface shows the selected map in the central display, while the left panel provides quick access to all available maps. Users can toggle between functional maps and the main simulation view.

As shown in Fig. 2, DIGICOR visualizes functional maps in real time from simulated activation signals. Users can preview all maps in a side panel and display the selected one over the 3D atrial geometry. All maps are rendered using color-coded schemes imported from custom MATLAB colormaps and adjusted dynamically to the data range.

Functional maps generated by DIGICOR include:

- Stability Map: captures the variability and persistence of high-frequency components.
- Dominant Frequency: identifies the local frequency of activation across the atrial surface.
- Reentrant Pattern Histogram: highlights regions of circular conduction, indicative of potential rotors.

DIGICOR has been validated on synthetic geometries with known activation patterns, reproducing expected dynamics and allowing parameter adjustment. Ablation masks can be applied via a brush tool or imported externally to ensure reproducibility. To illustrate the effect of ablation strategies, Fig. 3 compares two simulations. Image (a) applies a partial left pulmonary vein isolation (PVI), where fibrillatory activity persists. Image (b) applies a more extensive ablation, bilateral PVI, posterior wall box, and LAA isolation (PVI+PW+LAA), successfully terminating the arrhythmia. This highlights the need for personalized treatments, a task simplified by DIGICOR's interactive tools.

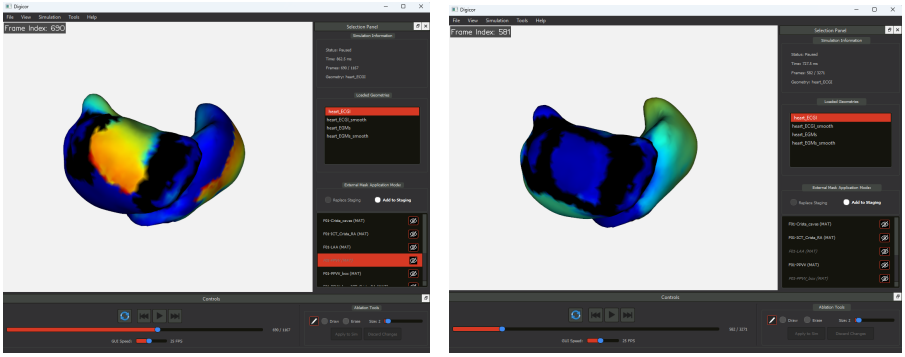

Fig. 3. Comparison of two ablation strategies: (a) PVI: fibrillatory activity persists. (b) Combined PVI+PW+LAA strategy: FA is successfully interrupted.

These results demonstrate how DIGICOR can assist clinicians in evaluating different lesion strategies and anticipating their effects prior to intervention. The ability to simulate and visualize outcomes in real time supports more informed and personalized ablation planning.

5 Conclusions

DIGICOR bridges the gap between advanced computational models and clinical decision-making. By enabling fast, visual and personalized simulation of atrial fibrillation and ablation planning, it supports a shift toward more tailored and effective treatments. In future developments, we aim to incorporate a wider range of atrial conditions beyond AF, including structural alterations such as fibrosis, in order to enhance the realism and clinical relevance of the simulations.

References

1. Chugh, S. S., Havmoeller, R., Narayanan, K., Singh, D., Rienstra, M., Benjamin, E. J., Gillum, R. F., Kim, Y., McAnulty, J. H., Zheng, Z., Forouzanfar, M. H., Naghavi, M., Mensah, G. A., Ezzati, M., Murray, C. J.: Worldwide epidemiology of atrial fibrillation. *Circulation* **129**(8), 837–847 (2013). doi:10.1161/circulationaha.113.005119
2. Krijthe, B. P., Kunst, A., Benjamin, E. J., Lip, G. Y. H., Franco, O. H., Hofman, A., Witteman, J. C. M., Stricker, B. H., Heeringa, J.: Projections on the number of individuals with atrial fibrillation in the European Union, from 2000 to 2060. *European Heart Journal* **34**(35), 2746–2751 (2013). doi:10.1093/eurheartj/eht280
3. Molero Alabau, R. (2023). Estimation of atrial electrical complexity during atrial fibrillation by solving the inverse problem of electrocardiography [Doctoral dissertation, Universitat Politècnica de València].

4. Boyle, P. M., Zghaib, T., Zahid, S., Ali, R. L., Deng, D., Franceschi, W. H., Hakim, J. B., Murphy, M. J., Prakosa, A., Zimmerman, S. L., Ashikaga, H., Marine, J. E., Kolandaivelu, A., Nazarian, S., Spragg, D. D., Calkins, H., Trayanova, N. A.: Computationally guided personalized targeted ablation of persistent atrial fibrillation. *Nature Biomedical Engineering* **3**(11), 870–879 (2019). doi:10.1038/s41551-019-0437-9

Embedded Systems meet Process Mining: A Manifesto

Zoe Valero-Ramon[1][0000−0002−4903−3896] and Carlos Fernandez-Llatas[1,2][0000−0002−2819−5597]

SABIEN group at ITACA Institute, Universitat Politècnica de València, Spain
{zoevara,cfllatas}@itaca.upv.es

Abstract. This paper explores the integration of process mining techniques within embedded systems, highlighting the potential benefits and challenges. Embedded systems are specialised computing systems designed to perform specific functions within larger mechanical or electronic systems. Process mining, an interdisciplinary field, extracts actionable insights from event logs to optimise business processes. The integration of process mining into embedded systems offers several advantages, including real-time optimisation, predictive maintenance, adaptive and autonomous systems, resource management, and anomaly detection. However, this integration faces significant challenges such as computational overhead, real-time constraints, data heterogeneity, and privacy and security concerns. Overcoming these barriers requires further research to develop efficient algorithms and integration methods that balance accuracy, latency, and privacy protection.

Keywords: Embedded Systems · Process Mining · IoT · Real-Time Processes · Adaptative Systems.

1 Introduction

Embedded systems are specialised computing systems designed to perform dedicated functions within larger mechanical or electronic systems [10]. Embedded systems in automotive, manufacturing, and related industrial sectors now favour architectures that combine high-performance computing and intelligent processing [1]. Unlike general-purpose computers, embedded systems are tailored for specific tasks, offering optimised performance, reliability, and efficiency in resource-constrained environments. They integrate hardware and software components to achieve seamless interaction with the physical world, often leveraging sensors, actuators, and communication interfaces [14].

These systems are ubiquitous across various domains, such as automotive control (e.g., engine management systems), healthcare (e.g., wearable medical devices), consumer electronics (e.g., smart appliances), and industrial automation (e.g., robotics and IoT devices). With the advent of advanced technologies like real-time operating systems (RTOS), machine learning at the edge, and

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

ultra-low-power architectures, embedded systems are evolving to support complex applications requiring higher computational capabilities, connectivity, and adaptability.

The increasing reliance on embedded systems in critical applications, coupled with advancements in hardware and software design, underscores the importance of developing efficient, reliable, and secure solutions. These systems face challenges such as balancing business requirements with security, managing connectivity, and ensuring cybersecurity [4]. Mixed-criticality systems, energy efficiency, reliability, and security are emerging design challenges that require novel approaches [13].

Process Mining (PM) [16] is an interdisciplinary field bridging data science and process management, focused on extracting actionable knowledge from event logs recorded by information systems. By analysing timestamped sequences of events, Process Mining techniques reconstruct, validate, and improve business processes, revealing bottlenecks, deviations, and inefficiencies invisible through traditional methods. Core methodologies like process discovery, conformance checking, and enhancement enable data-driven process optimisation across domains such as healthcare, manufacturing, and logistics [11].

Process mining has expanded into cyber-physical domains, enabling the extraction of process models from industrial control systems for anomaly detection and cyber-attack identification [19]. This integration with embedded systems allows for real-time monitoring and adaptive control in various applications, including emergency management and healthcare [8]. The SmartPM system demonstrates automated process adaptation in cyber-physical environments, addressing unanticipated exceptions and exogenous events through declarative task specifications, execution monitoring, and automated planning techniques [7]. The development of enterprise systems based on Cyber-Physical Systems (CPS) and Internet of Things (IoT) principles focuses on merging real and virtual objects, increasing complexity. Process mining in this context can generate business processes from sensor-acquired data and integrate this data with existing processes, enhancing the adaptability and flexibility of enterprise systems in unpredictable physical environments [12]. Other work discusses integrating process mining with CPS and IoT technologies to enable real-time monitoring and adaptive control in enterprise systems [18].

However, applying Process Mining to resource-constrained embedded environments introduces novel challenges, including real-time log processing, limited computational capabilities, or synchronisation of heterogeneous event streams, among others. Process mining, a technique for analysing and improving business processes using event logs, faces several challenges when applied to complex environments like healthcare and embedded systems. In healthcare, processes are often intricate and flexible, requiring adaptation of traditional process mining approaches [3]. In smart spaces and embedded systems, process mining must address differences between human activities and business processes, necessitating further research to overcome these obstacles [5]. A Delphi study identified 32 challenges for process mining in organisations, highlighting barriers to adoption

and areas for improvement [9]. These challenges include real-time log processing, limited computational capabilities, and synchronisation of heterogeneous event streams, particularly in resource-constrained environments like embedded systems. However, more research is needed in this regard.

While embedded systems and Process Mining have been regarded as separate topics in research and in practice, we believe that, on the one hand, Process Mining will greatly benefit from embedded systems (like Process Mining at the edge). On the other hand, Process Mining poses challenges that will require enhancements and extensions of the current state of the art in the embedded systems field. In this article, we question the extent to which these two paradigms can be combined and discuss emerging challenges. This paper investigates the viability of embedding Process Mining in edge devices and examines the emerging technical and methodological barriers to this convergence.

2 Background

In our increasingly interconnected world, a vast array of devices—embedded in electrical and electronic components—are equipped with sensors and actuators that enable them to sense, act, and react to their environments. These devices collect and exchange data through a variety of communication networks, including the Internet, forming the backbone of the Internet of Things (IoT). This network of embedded systems facilitates the continuous monitoring of physical phenomena via sensing technologies such as wearables, beacons, smartphones, and machine sensors.

Embedded systems are driving the evolution of smart environments, such as smart homes, connected cars, smart logistics, and Industry 4.0. By collecting and analysing data in real time, they enable actionable insights and feedback mechanisms to optimise processes and enhance system performance. As a result, embedded systems play a critical role in the big data ecosystem, serving as one of the three primary sources of information, alongside human-sourced and process-mediated data [15].

Process Mining, in particular, focuses on analysing timestamped tasks and activities across time and locations to uncover patterns, optimise workflows, and improve systems. When integrated with embedded systems, Process Mining can go beyond traditional analytics by leveraging real-time data streams from IoT devices for process execution, monitoring, and optimisation. This convergence of embedded systems and Process Mining offers a more comprehensive view of system performance, enabling organisations to identify inefficiencies, predict outcomes, and unlock untapped potential for optimisation. For example, the PROMOTE tool, developed for embedded system development, has demonstrated the ability to uncover bottlenecks in industrial projects [6]. As process mining combines process modelling with data science, it offers new ways to address compliance and performance issues in complex systems [17].

By embedding Process Mining capabilities into connected devices, systems can achieve unprecedented levels of autonomy and adaptability, turning raw data

into actionable insights. This integration not only enhances process analytics but also drives innovation in system development, enabling smarter, more efficient, and highly responsive systems in a wide range of domains.

2.1 How can Embedded Systems benefit from Process Mining?

Process mining can be applied to Internet of Things (IoT) environments and cyber-physical systems (CPS) by using stream processing and complex event processing to aggregate low-level IoT data into higher-level process events [2]. This enables both offline and online process mining, facilitating process discovery, conformance checking, and enhancement in smart factory settings. Embedded systems can significantly benefit from Process Mining when real-time data streams are utilised effectively. Above, there are included a non-exhaustive key ways Process Mining can enhance embedded systems and their real-time data.

Real-Time Process Insights and Optimisation. Embedded systems generate fine-grained event logs (e.g., sensor readings, state transitions) that Process Mining algorithms (e.g., conformance checking) can analyse to discover, monitor, optimise operational workflows or even detect deviations from expected workflows. This could have an impact with the immediate identification of hardware/software failures (e.g., an industrial robot arm skipping a step), or reduce downtime via predictive diagnostics (e.g., detecting wear patterns in motor controllers). For example, in industrial automation (e.g., Industry 4.0), real-time Process Mining can identify bottlenecks in machine operations, allowing embedded controllers to dynamically adjust parameters for improved efficiency and throughput.

Enhanced Predictive Maintenance. By analysing patterns in sensor data, Process Mining can predict equipment failures or maintenance needs before they occur. Or even Process Mining can discover inefficiencies in runtime processes (e.g., CPU scheduling, memory allocation). This could allow dynamic adjustment of resource usage (e.g., reallocating compute tasks based on process bottlenecks) and energy savings in IoT devices by pruning redundant operations. For example, embedded systems in manufacturing equipment can leverage Process Mining to analyse vibration, temperature, or usage patterns, ensuring timely maintenance and minimising downtime.

Adaptive and Autonomous Systems. Process Mining enables embedded systems to adapt workflows based on real-time insights, making systems more autonomous and responsive. For example, a smart home system can use Process Mining to analyse user behaviour (e.g., energy usage patterns) and autonomously adjust Heating, Ventilation, and Air Conditioning (HVAC) or lighting systems for efficiency and comfort.

Improved Resource Management. Process Mining helps embedded systems optimise the use of resources such as energy, bandwidth, or storage by identifying inefficiencies in real-time operations. For example, IoT devices in smart grids can use Process Mining to monitor and analyse power consumption patterns, enabling dynamic load balancing and energy savings.

Real-Time Monitoring and Anomaly Detection. Process Mining can detect anomalies or deviations from expected processes in real time, triggering immediate corrective actions by embedded systems. For example, in wearable medical devices, Process Mining can monitor patient vitals and detect unusual patterns (e.g., irregular heartbeats), allowing the system to alert healthcare providers instantly.

Process Conformance and Compliance. Process Mining ensures that embedded systems operate within predefined parameters and comply with regulatory or safety standards. For example, in automotive embedded systems, Process Mining can validate that processes like airbag deployment or engine calibration conform to safety protocols in real time.

Data-Driven System Development. Process Mining provides insights into how embedded systems are used in practice, enabling iterative design improvements. For instance, smart appliance manufacturers can use Process Mining to analyse user interactions and optimise device features or workflows for better usability and performance.

Context-Aware Decision Making. Embedded systems can make context-aware decisions based on real-time process data analysed by Process Mining. For example, in connected cars, Process Mining can integrate sensor data (e.g., traffic conditions, road hazards) to adapt navigation processes dynamically.

Synchronisation of Distributed Embedded Systems. Process Mining can help coordinate and optimise processes across multiple interconnected embedded systems. For example, in a smart logistics network, Process Mining can synchronise the operations of embedded systems in vehicles, warehouses, and inventory systems to optimise delivery routes and schedules.

Scalability and IoT Integration. Process Mining can handle the growing complexity of embedded systems in IoT networks, providing scalable solutions for analysing distributed event data. For example, in smart cities, Process Mining can analyse data from embedded systems in traffic lights, public transport, and utilities to optimise city-wide processes in real time.

Table 1 provides an overview of example use cases demonstrating the application of Process Mining techniques across various domains that utilise embedded systems. Each domain highlights specific challenges or objectives where Process Mining can deliver value by analysing event data to optimise processes, ensure safety, or predict failures. The table also includes the corresponding Process Mining techniques employed, showcasing the versatility of Process Mining in addressing domain-specific needs. These examples illustrate the practical synergy between Process Mining and embedded systems, emphasising their potential to enhance real-time decision-making and system performance.

Domain	Application	PM technique used
Industrial IoT	Optimising assembly line robotics	Conformance Checking
Healthcare devices	Validating infusion pump safety protocols	Deviation Mining
Automotive	Detecting ECU communication anomalies	Process Discovery
Smart grids	Predicting transformer failures	Predictive Process Monitoring

Table 1. Example of use cases.

2.2 How Process Mining Can Benefit from Embedded Systems

Process Mining, an interdisciplinary field bridging data science and process management, relies on event data to analyse, monitor, and optimise processes. Embedded systems, as the backbone of numerous IoT-enabled devices, serve as rich, real-time data sources that can significantly enhance the capabilities and applications of Process Mining. Below, we outline the scientific basis for how Process Mining benefits from embedded systems.

Real-Time Data Acquisition. Embedded systems generate high-resolution, timestamped event data from sensors and actuators in real time. This continuous stream of data provides Process Mining with detailed, accurate insights into system workflows and interactions as they occur. The integration of real-time data allows Process Mining techniques to move from retrospective analysis to real-time process monitoring and adaptation, making it possible to detect inefficiencies, bottlenecks, or anomalies on the fly. For example, in industrial automation, embedded systems in machinery can provide real-time feedback on production events, enabling Process Mining to instantly identify and address deviations from optimal workflows.

Fine-Grained Process Visibility. Embedded systems capture granular data from physical processes, which traditional enterprise systems may overlook. This enables Process Mining to gain a more comprehensive and detailed view of underlying processes. Fine-grained data enhances the accuracy of process discovery, conformance checking, and enhancement, particularly for complex, multi-step processes that involve physical interactions. For instance, in healthcare, wearable devices equipped with embedded systems can provide minute-by-minute data on patient activities, enabling precise monitoring of treatment adherence or recovery processes.

Distributed and Decentralised Process Mining. Embedded systems are often distributed across environments (e.g., IoT networks), enabling localised data collection and analysis at the edge. This decentralisation supports distributed Process Mining approaches. Leveraging edge devices for localised Process Mining reduces latency, minimises bandwidth usage, and allows for faster, more efficient process analysis without relying entirely on centralised systems. An example of this approach could be found in smart logistics, where embedded systems in vehicles collect geolocation and operational data, enabling Process Mining to optimise delivery routes and schedules in a decentralised manner.

Enabling Context-Aware Process Analysis. Embedded systems are equipped with sensors that capture contextual data (e.g., environmental conditions, user interactions). This context enriches Process Mining analysis by linking process events to their physical surroundings. Context-aware Process Mining allows for a deeper understanding of the factors influencing process performance, enabling more targeted and effective optimisation. This benefit could be seen in a smart factory, for example, where temperature and humidity data from embedded systems can be correlated with production quality metrics, allowing Process Mining to identify environmental conditions that impact process outcomes.

Integration with Actuation Mechanisms. Embedded systems not only collect data but also control actuators that can modify processes in real time. This bidirectional capability enables Process Mining to go beyond passive analysis and actively influence process execution. The integration of Process Mining with embedded systems enables closed-loop systems where insights derived from Process Mining directly trigger adjustments to optimise processes. For example, in connected homes, Process Mining can analyse energy consumption patterns and directly instruct embedded systems in HVAC units to adjust settings for improved energy efficiency.

Scalability in IoT-Enabled Environments. The proliferation of embedded systems in IoT networks generates a vast amount of event data. Process Mining can leverage this data to scale its analysis to increasingly complex, interconnected systems. By harnessing the scalability of IoT and embedded systems,

Process Mining can analyse large-scale, distributed processes in ways that were previously infeasible. For instance, in smart cities, embedded systems in traffic lights and public transportation vehicles provide data for Process Mining to optimise city-wide traffic flow and reduce congestion.

Real-Time Anomaly Detection and Predictive Analytics. Embedded systems provide continuous data streams that enable Process Mining to detect anomalies or predict failures in real time, enhancing system reliability. Predictive capabilities allow Process Mining to proactively address potential issues, reducing downtime and improving system availability. An example of this benefit is found in embedded systems in wind turbines that can monitor vibration and rotation patterns, enabling Process Mining to predict maintenance needs and prevent failures.

Bridging Cyber-Physical Systems. Embedded systems are key components of cyber-physical systems, where physical processes are tightly integrated with computational processes. Process Mining can utilise data from embedded systems to analyse, optimise, and align these interactions. Process Mining provides a framework for understanding and improving the interplay between physical and digital processes, enhancing overall system performance and adaptability. For example, in autonomous vehicles, embedded systems provide real-time data on sensor inputs and control actions, allowing Process Mining to optimise decision-making algorithms for safety and efficiency.

Enhanced Process Conformance in Critical Applications. Embedded systems are often used in safety-critical applications (e.g., healthcare, automotive). Process Mining can analyse data from these systems to ensure conformance to safety standards and regulatory requirements. By ensuring process conformance, Process Mining enhances the reliability and safety of embedded systems in critical domains. For instance, in medical devices, Process Mining can monitor and analyse processes to ensure compliance with health regulations and detect deviations that could compromise patient safety.

Enabling Process Mining at the Edge. With advancements in embedded system architectures, Process Mining can be deployed directly on resource-constrained devices, enabling real-time, edge-based analysis. Edge-based Process Mining reduces reliance on centralised infrastructure, enabling faster decision-making and supporting applications in remote or bandwidth-limited environments. For example, embedded systems in agricultural sensors can perform localised Process Mining to optimise irrigation schedules based on soil moisture and weather conditions.

Embedded systems provide Process Mining with real-time, granular, and context-rich event data, enabling improved process visibility, adaptability, and optimisation. By integrating Process Mining with embedded systems, researchers

and practitioners can unlock new possibilities for real-time monitoring, predictive analytics, and decentralised process management. This synergy is particularly critical in IoT-enabled and cyber-physical environments, where the seamless interaction between physical and computational processes is essential for efficiency and innovation.

3 Challenges

While the integration of Process Mining with embedded systems offers transformative potential, several technical and practical barriers must be addressed to fully realise this synergy. The resource-constrained nature of embedded systems and the demands of real-time data processing impose significant challenges. In this section, we examine key barriers hindering real-time Process Mining on embedded systems, focusing on computational, temporal, and data-related constraints, as well as privacy concerns. Figure 1 presents an overview of these barriers that are explained in more detail.

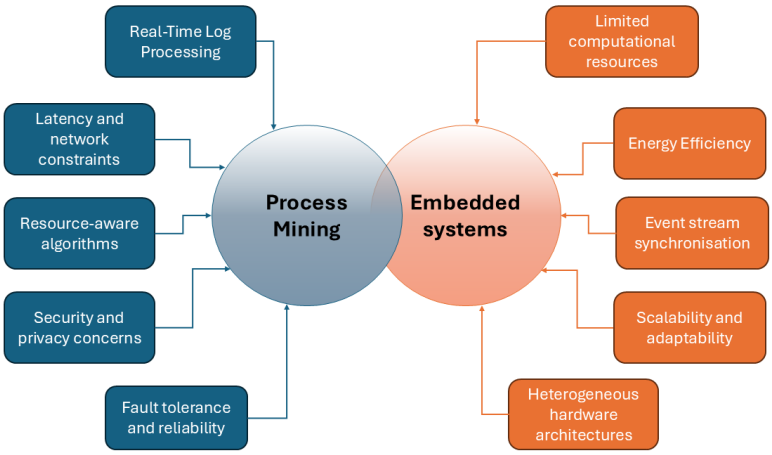

Fig. 1. Barriers hindering Process Mining on embedded systems.

Computational Overhead. Process Mining algorithms, such as the alpha-algorithm, often exhibit high computational complexity (e.g., $O(n^2)$ or higher for event log analysis). These algorithms are designed for enterprise-scale systems with substantial computational resources, which contrasts with the resource-constrained environments of embedded systems. Limited processing power, memory, and energy resources in embedded systems make it challenging to execute

computationally intensive Process Mining tasks, such as process discovery, conformance checking, or enhancement, without significant optimisations.

For instance, Process Mining applied to industrial embedded systems must operate efficiently within the constraints of microcontrollers or low-power processors, which are often optimised for specific tasks rather than general-purpose computation. Without lightweight, resource-aware Process Mining algorithms, the computational overhead may render real-time analysis infeasible.

Real-Time Constraints. Embedded systems often operate in domains where real-time responsiveness is critical, such as automotive systems, healthcare devices, or industrial control systems. In such domains, event log analysis must meet strict deadlines to ensure system safety and functionality. For example, in automotive applications, Process Mining-based anomaly detection may require response times under one millisecond to prevent accidents or ensure operational reliability.

Traditional Process Mining techniques are not inherently designed to meet such stringent latency requirements, as they often involve computationally expensive operations like log parsing, state-space exploration, or process model generation. Adapting Process Mining to meet real-time constraints necessitates the development of streamlined methods capable of on-the-fly analysis while adhering to hard deadlines.

Data Heterogeneity. Embedded systems generate data from diverse sources, including sensors, operating systems (OS), and communication networks. These data streams often differ in format, granularity, and semantics, making it difficult to fuse them into coherent event logs suitable for process analysis. For example, sensor data may provide high-frequency, low-level measurements, while OS logs may describe system-level events, and network logs capture communication patterns.

The challenge lies in standardising and integrating these heterogeneous data sources into unified process models without introducing errors or losing critical information. This is particularly complex in distributed or IoT-based embedded systems, where data flows across multiple devices and layers.

Privacy and Security Concerns. Embedded systems often operate in sensitive domains where data privacy and security are paramount, such as healthcare, finance, or critical infrastructure. Process Mining in such contexts must address the potential exposure of sensitive operational data, such as patient health information in medical devices or user behaviour patterns in smart homes.

Anonymising or encrypting event data to protect privacy adds further computational overhead and latency, which can conflict with the real-time requirements of embedded systems. Moreover, ensuring compliance with privacy regulations, such as the General Data Protection Regulation (GDPR) or the Health Insurance Portability and Accountability (HIPAA), complicates the implementation of Process Mining in embedded environments.

4 Conclusions

The integration of Process Mining techniques within embedded systems offers a transformative potential to enhance efficiency, adaptability, and reliability. By enabling real-time optimisation, predictive maintenance, and anomaly detection, embedded systems can achieve improved performance and resilience across a wide range of applications. However, realising this potential requires addressing significant challenges, including computational overhead, stringent real-time constraints, data heterogeneity, and privacy and security concerns.

To overcome these barriers, future research should prioritise the development of lightweight, resource-efficient Process Mining algorithms tailored specifically for the constrained environments of embedded systems. Designing edge-friendly Process Mining tools capable of delivering insights with minimal latency is essential for meeting real-time operational demands. Furthermore, addressing data heterogeneity will require standardised frameworks for cross-layer log integration to ensure coherent and accurate process analysis. Privacy concerns must also be addressed through robust anonymisation and encryption techniques, ensuring compliance with regulatory standards while maintaining system functionality.

An important direction for future work involves quantifying the trade-offs between competing requirements, such as achieving high accuracy versus minimising latency, or balancing computational efficiency with privacy protection. Domain-specific solutions, such as simplified process models for targeted applications or hybrid architectures that combine edge and cloud-based Process Mining, should be explored to address the unique needs of various industries.

Interdisciplinary collaboration between experts in Process Mining, embedded systems, and data security will be crucial for overcoming these challenges. As these issues are addressed, the fusion of Process Mining and embedded systems is poised to unlock new advancements in real-time, data-driven systems, enabling autonomous, self-optimising solutions across domains such as industrial automation, smart devices, healthcare, and beyond.

References

1. Bello, L.L., Mariani, R., Mubeen, S., Saponara, S.: Recent advances and trends in on-board embedded and networked automotive systems. IEEE Transactions on Industrial Informatics **15**(2), 1038–1051 (2018)
2. Cardoso, J., van der Aalst, W.M.: Path mining and process mining for workflow management systems. In: Encyclopedia of Data Warehousing and Mining, Second Edition, pp. 1489–1496. IGI Global (2009)
3. Chaydy, N., Madani, A.: An overview of process mining and its applicability to complex, real-life scenarios. In: 2019 International Conference on Systems of Collaboration Big Data, Internet of Things & Security (SysCoBIoTS). pp. 1–9. IEEE (2019)
4. Kala, N., Narasimhan, P.: Secure by design: Proactive approaches to embedded system security. International Journal of Scientific Research in Computer Science, Engineering and Information Technology **10**(6), 1009–1035 (2024)

5. Leotta, F., Mecella, M., Mendling, J.: Applying process mining to smart spaces: Perspectives and research challenges. In: Advanced Information Systems Engineering Workshops: CAiSE 2015 International Workshops, Stockholm, Sweden, June 8-9, 2015, Proceedings 27. pp. 298–304. Springer (2015)
6. Leppäkoski, A., Hämäläinen, T.D.: Promote: A process mining tool for embedded system development. In: International Conference on Product-Focused Software Process Improvement. pp. 529–538. Springer (2016)
7. Marrella, A., Mecella, M.: Adaptive process management in cyber-physical domains. Advances in Intelligent Process-Aware Information Systems: Concepts, Methods, and Technologies pp. 15–48 (2017)
8. Marrella, A., Mecella, M., Halapuu, P., Sardina, S.: Automated process adaptation in cyber-physical domains with the smartpm system (short paper). In: 2015 IEEE 8th International Conference on Service-Oriented Computing and Applications (SOCA). pp. 59–64. IEEE (2015)
9. Martin, N., Fischer, D.A., Kerpedzhiev, G.D., Goel, K., Leemans, S.J., Röglinger, M., van der Aalst, W.M., Dumas, M., La Rosa, M., Wynn, M.T.: Opportunities and challenges for process mining in organizations: results of a delphi study. Business & Information Systems Engineering **63**, 511–527 (2021)
10. Marwedel, P.: Embedded system design: embedded systems foundations of cyber-physical systems, and the internet of things. Springer Nature (2021)
11. Munoz-Gama, J., et al.: Conformance checking and diagnosis in process mining. Ph.D. thesis, Springer (2014)
12. Sacală, I.Ş., Moisescu, M.A.: The development of enterprise systems based on cyber-physical systems principles. Rom. Stat. Rev **4**, 29–39 (2014)
13. Sahoo, S.S., Kumar, A., Decky, M., Wong, S.C., Merrett, G.V., Zhao, Y., Wang, J., Wang, X., Singh, A.K.: Emergent design challenges for embedded systems and paths forward: mixed-criticality, energy, reliability and security perspectives. In: Proceedings of the 2021 International Conference on Hardware/Software Codesign and System Synthesis. pp. 1–10 (2021)
14. Stankovic, J., Lee, I., Mok, A., Rajkumar, R.: Opportunities and obligations for physical computing systems. Computer **38**(11), 23–31 (2005). https://doi.org/10.1109/MC.2005.386
15. Toorajipour, R., Oghazi, P., Palmié, M.: Data ecosystem business models: Value propositions and value capture with artificial intelligence of things. International Journal of Information Management **78**, 102804 (2024)
16. Van Der Aalst, W.: Process Mining.Data science in action. Springer (2016)
17. Van Der Aalst, W.M., Carmona, J.: Process mining handbook. Springer Nature (2022)
18. Vitale, F., Guarino, S., Flammini, F., Faramondi, L., Mazzocca, N., Setola, R.: Process mining for digital twin development of industrial cyber-physical systems. IEEE Transactions on Industrial Informatics (2024)
19. Xavier, M., Dubinin, V., Patil, S., Vyatkin, V.: Process mining in industrial control systems. In: 2022 IEEE 20th International Conference on Industrial Informatics (INDIN). pp. 1–6. IEEE (2022)

Dependability analysis of neural networks implemented in Arduino

Joaquín Gracia-Morán[1 [0000−0001−9715−8960]], Juan-Carlos Ruiz[1 [0000−0001−7678−3513]], David de Andrés[1 [0000−0002−4744−3795]], Luis-J. Saiz-Adalid[1 [0000−0002−4868−2050]], J. Carlos Baraza-Calvo[1 [0000−0001−7692−2309]], Daniel Gil-Tomás[1 [0000−0001−9225−1998]], and Pedro J. Gil-Vicente[1 [0000−0002−9364−7385]]

ITACA, Universitat Politècnica de València
Camino de Vera s/n, 46022 - Valencia (Spain)
{jgracia, jcruizg, ddandres, ljsaiz, jcbaraza, dgil, pgil}@itaca.upv.es
https://gstf.blogs.upv.es/

Abstract. The use of neural networks has expanded to environments as diverse as medical systems, industrial devices and space systems. In these cases, it is essential to balance performance, power consumption, and silicon area. Furthermore, in critical environments, it is necessary to ensure high fault tolerance.
Traditionally, the parameters of neural networks have been codified using 32-bit floating-point numbers, which entails high memory consumption and greater vulnerability to failures due to the aggressive scaling of CMOS technology. An effective strategy for optimizing these systems is to reduce parameter precision, using fewer bits and thus, reducing both the amount of memory required and the processing time.
However, several questions arise when implementing these types of networks in embedded systems: Do they maintain their reliability in critical environments, or do they require fault tolerance mechanisms? Are area and latency really reduced?
This work addresses these questions by reducing the precision of a neural network and implementing it in an Arduino-based system. In addition, Error Correction Codes have been incorporated and, using the fault injection technique, their reliability has been evaluated by comparing the same neural network with parameters encoded in 8, 16 and 32 bits.

Keywords: Error Correcting Codes · Dependability · Single Faults · Multiple Faults · Optimized Neural Networks · Embedded Systems

1 Introduction

Nowadays, neural networks have expanded to such an extent that they are now commonly found in fields as diverse as the automotive industry, aerospace applications, and consumer electronics [3]. Neural networks process a large amount of

C.Fernandez-Llatas and A. Martinez-Millana (Ed.):
Workshop on Innovation on
Information and Communication Technologies
(ITACA-WIICT 2025)

ISBN 978-84-09-77296-4

information through interconnected nodes (neurons) organized in multiple layers [39], basing their operation on a series of parameters (weights, biases, etc.).

The operation of a neural network is divided into two phases. In the first one, known as training phase, the aforementioned parameters are calculated. These values will determine the network's behaviour during the second phase, known as inference phase. In this second stage, predictions are made based on new input data. Thus, the integrity of the parameters calculated in the training phase is essential for the proper functioning of the neural network, as an erroneous value can seriously compromise the accuracy of the predictions.

During the initial training stage, the different values of the neural network (weights, biases, etc.) are represented in the IEEE 754 single-precision format [1], also known as FP32. To reduce the execution time during the inference stage, weights are usually stored in memory. However, the increasing integration scale of CMOS technology has increased the rate of multiple faults in memory devices [6], which can negatively affect the neural network inference process [21][24]. For example, corruption of just 13 bits in the CNN ResNet18, within a set of approximately 93 million bits, has been shown to cause a complete collapse of the network [29]. This vulnerability has caused a growing interest in researching the reliability of neural networks [34][28]. Results obtained in these works show that it is essential to incorporate fault tolerance mechanisms when implementing neural networks in critical applications [4][33][31].

To reduce memory footprint, a current trend is the compression of the neural network models, generating optimized models [22]. This reduction in memory usage decreases the likelihood of fault occurrences, thereby increasing the network's reliability [11].

While model compression is effective and optimized neural networks maintain an adequate level of accuracy for many applications, several questions arise when attempting to use these models in embedded systems operating in critical environments: Will the optimized network be sufficiently reliable, or will it be necessary to implement additional fault-tolerance mechanisms? What is the improvement in area consumption and latency of the optimized network?

In this paper, we have developed a simple neural network with parameters codified in FP32. Then, we have optimized it, by quantizing their parameters with different data word sizes. Next, we have implemented all neural network's models into an Arduino-based system, and we have studied the reliability of these neural networks, both unprotected and protected with different Error Correcting Codes, comparing their results.

This work is organized as follows. Section 2 briefly describes the neural network model compression processes and the application of error-correcting codes to them. Section 3 summarizes the embedded system used, the implemented neural network, and the optimization methods. Section 4 presents the experimental evaluation of the system, analyzing the impact of fault injection before and after applying different types of Error Correcting Codes. Finally, Section 5 presents the conclusions obtained from the results achieved.

2 Compression and Fault Tolerance in Neural Netowrks

2.1 Compression of neural networks

As discussed in the previous section, neural network parameters are usually represented in the IEEE 754 single-precision (FP32) format. However, using parameters in this format results in high memory usage and an increase in inference time. In this way, optimising both factors has become increasingly important. Different strategies can be used to compact neural networks [7][23]. At the hardware level, this optimization can be achieved through specialized devices, such as neural network accelerators, which optimize system architecture to improve parallelism and minimize memory accesses [25][20].

At software level, it is possible to reduce network complexity by designing smaller models with acceptable accuracy [37], or by decomposing the original parameters into multiple smaller matrices or tensors, focusing on the reduction of both, the memory footprint and the number of required operations [27].

Parameter pruning and optimization processes are widely used techniques for compressing and accelerating neural network models [5]. Both techniques try to eliminate redundant parameters that do not significantly affect network performance. They are typically applied in convolutional and fully connected layers. For example, the Brain floating-point format, also known as BF16 [19], is supported by various accelerators, such as NVIDIA Ampere GPUs [10] or Intel's Nervana accelerators [16]. These devices allow for inferences using BF16 arithmetic. Although other floating-point formats exist, their use is limited by a lack of hardware support, so they will not be included in this work.

Also, quantization allows weights, represented in FP32, to be converted into 8-bit integers, also known as INT8 [19]. This transformation reduces the size of parameters in memory and simplifies arithmetic operations by decreasing numerical precision.

While approximating FP32 values to formats with fewer bits generates rounding errors that, in some cases, can affect the accuracy of certain neural networks, those based on BF16 or INT8 offer the same accuracy as their FP32-based counterparts [8].

2.2 Protecting neural networks with error correcting codes

In general, the smaller the weights of a neural network, the greater its resilience [35][30]. Even so, fault tolerance mechanisms must be added if the neural network is to be used in a critical environment. In this section, we will focus on the protection of neural network parameters using Error Correcting Codes (ECC).

Many works have studied this type of neural network protection, among which we can highlight the following ones, as they propose ECCs similar to those used in this work. For example, [18] combines single error correction (SEC) and triple repetition codes. In this case, it is a specific solution for a specific CNN, rather than a generic methodology applicable to any network.

On the other hand, in [26], and depending on the bit values (30 : 28) of the neural network weights, various types of Hamming SEC are used to protect specific bits. The approach presented in [15] is applied to optimized neural networks. In this case, a training scheme is proposed that obtains at least seven non-significant bits within eight consecutive weights (8 bytes) to implement a SEC-DED (Double Error Detection) code.

A final approach is presented in [12], where a series of ECCs with different fault-tolerance properties have been implemented to protect the weights of several neural networks with parameters in BF16 format.

3 System Description

This section summarizes the main characteristics of both the embedded system and the neural network used in this work, as well as the optimization methods.

3.1 Embedded system and neural network

The embedded system has been used in a mobile robot based on an Arduino UNO R3 board. This system incorporates an ATmega328p microcontroller, 32 KB of FLASH memory, 2 KB of SRAM, 1 KB of EEPROM, and operates at 16 MHz. The robot has two DC motors for the rear wheels, three infrared proximity sensors (one E18-D80NK in the center and two KY-032 sensors on the sides), and a servomotor for steering, in addition to the necessary hardware to power the system, as shown in Fig. 1.

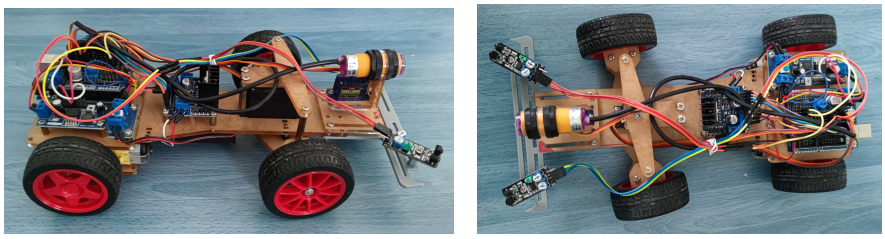

(a) Side view. (b) Top view.

Fig. 1. Arduino-based system.

To drive the mobile robot, we have generated and trained a small neural network with the following characteristics (see Fig. 2):

- 3-layer neural network: input, intermediate, and output.
- Input layer: receives readings from the three proximity sensors (left, right, and center).

– Intermediate layer with 5 neurons.
– Output layer: generates outputs corresponding to speed, direction of rotation, and forward direction.
– Two sets of weights, one between the input layer and the intermediate layer (hereafter referred to as Stage 1 weights), and another one between the intermediate layer and the output layer (hereafter referred to as Stage 2 weights).

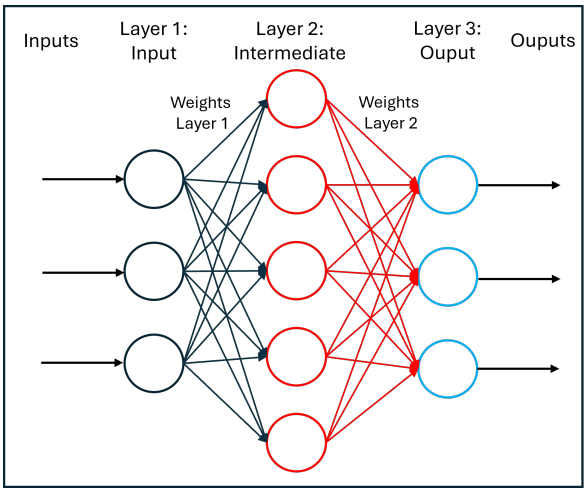

Fig. 2. Neural network scheme.

Once the weights were generated, inference was implemented in the mobile robot, such that the appropriate outputs were generated based on the proximity sensor values. Weights are originally encoded in the FP32 format. This format uses 32 bits distributed as follows (see Fig. 3): bit 31 stores the sign bit (S), bits 30 through 23 contain the exponent (E) in excess-127 format, and bits 22 through 0 represent the magnitude, with one implicit bit. Thus, the corresponding decimal value can be calculated according to Equation 1.

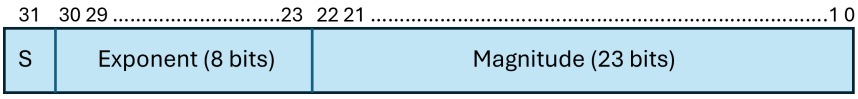

Fig. 3. FP32 format.

$$value = (-1)^S \times 2^{E-127} \times \left(1,0 + \sum_{i=1}^{23} b_{23-i} \times 2^{-i}\right) \quad (1)$$

3.2 Reduction of the neural network parameters' size

Despite the simplicity of the neural network used, its design allows us to evaluate the network's behavior in a real-world environment, providing a solid basis for analyzing the system's resilience to memory faults, and verifying both the feasibility of optimizing its parameters and its protection with ECC. In this way, we have carried out two types of parameter's size reduction, which we will describe next.

Summary of the BF16 format. We have converted the original FP32 parameters into the BF16 format just truncating the lowest 16 bits, as shown in Fig. 4. Thus, the optimized network weights are stored in memory in 16-bit variables. Inference is executed in real time, adjusting the robot outputs according to the sensor signals.

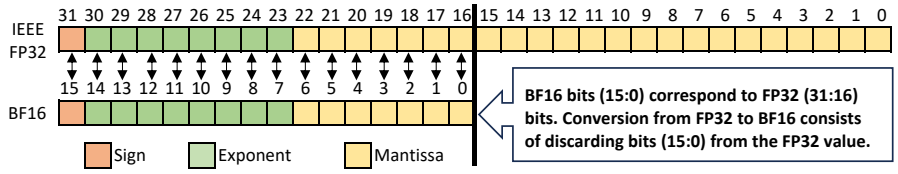

Fig. 4. FP32 vs BF16 format [32].

Although the weights in BF16 are stored in memory in 16-bit variables, the various operations related to the neural network are performed in 32-bit floating-point variables. To do this, we define a 32-bit floating-point variable, in which the weight in BF16 is stored in the 16 most significant bits, and the 16 least significant bits are set to '0'. This way, the same Arduino library can be used in both neural networks (FP32 and BF16) to perform the necessary floating-point operations, without having to implement a new library.

Summary of the quantization process: from FP32 to INT8. This section will detail the post-training quantization process carried out based on the general procedure presented in [17].

As already mentioned, quantization process reduces the computational and energy resources of a neural network. In the case of this work, this process consisted on converting the network's weights and biases, initially represented in 32-bit floating-point (FP32), into 8-bit integers (INT8).

To quantize the network, the first step is to quantize its weights and inputs. Each set of weights and the set of inputs are processed independently, since their quantized values depend only on their range of values, once the network has been trained, and on the range of values defined for quantization. Thus, each data set is traversed to locate its maximum and minimum values. Each data set is then iterated through again to adjust its values to the 8-bit integer range $[-128, 127]$ according to Equation 2.

$$\begin{aligned} scale &= \frac{(max - min)}{255} \\ zero &= -round(\frac{min}{scale}) - 128 \\ int8 &= round(\frac{float}{scale}) + zero \end{aligned} \tag{2}$$

scale allows to divide the range of real values into uniform partitions, while *zero* defines the integer value associated with the real value 0.0. Values greater than 127 or lower than -128 are truncated to those values.

Once this data have been quantized, it is necessary to determine how it affects the processing performed in each of the network's layers and what are the *scale* and *zero* values for these layers. This operation must be performed sequentially, following the propagation of information from the system's input to its output. Thus, in this case, the first stage to consider would be the intermediate layer in Fig. 2, which corresponds to a dense layer with a hyperbolic tangent as the activation function.

To do this, the dense layer is first applied to the input data set, but using the real values resulting from dequantizing the previously quantized inputs and weights, according to Equation 3.

$$float = scale \times (int8 - cero) \tag{3}$$

This allows us to monitor the operation performed by the dense layer, obtain the maximum and minimum values generated when processing all the input data, and thereby determine its *scale* and *zero*. From these values, using Equation 2, we obtain the quantized values provided by this layer when processing the quantized data.

Afterwards, it is necessary to determine the effect of quantization on the activation function. Since the hyperbolic tangent is a very expensive operation to perform, the simplest solution to fulfill it using integer arithmetic is to precalculate a lookup table for all possible values this function can have for the range of integers considered, $[-128, 127]$ in this case. As done before, it is necessary to consider the effect that the quantization of the dense layer from which the data processed by this activation function comes from. Therefore, the lookup table is calculated using Equation 4, from the dequantization of the output data of the dense layer (Equation 3).

$$
\begin{aligned}
scale &= \frac{(1.0 - (-1.0))}{255} = \frac{2.0}{255} \\
zero &= -round(\frac{-1.0}{scale}) - 128 = 0
\end{aligned}
\tag{4}
$$

$$
\begin{aligned}
&\forall\ value \in [-128, 127]\ : \\
&\quad float = scale_{Dense} \times (value - zero_{Dense}) \\
&\quad int8[value + 128] = round(\frac{tanh(float)}{scale}) + zero
\end{aligned}
$$

Thus, from this lookup table, the output values of the dense layer's activation function can be obtained after processing all the quantized values.

This same procedure must be applied successively to all the layers defined in the network. In our case, only one more dense layer with a hyperbolic tangent as the activation function remains.

Once the network is quantized, inference operations are performed using integer arithmetic, which simplifies calculations and significantly reduces latency. The convolutional and dense layers, which are the most computationally intensive, transform multiplication and accumulation operations into integer operations.

In fully connected layers, the output is calculated as the weighted sum of the inputs multiplied by the weights, plus a bias (translated as bias, shift, or threshold). When quantized, these values are scaled with a zero point, allowing floating-point operations to be transformed into integer operations with bit shifts and rounding. This transformation not only simplifies the required hardware but also allows operations to be parallelized, increasing efficiency. However, when implementing the quantized network in Arduino, we had an additional problem: to obtain maximum precision, the quantized network internally uses not only 8-bit integers, but also requires integers up to 64 bits. Specifically, there is a product of 32-bit integer values, which generates a 64-bit result. This is not supported by Arduino, since it only implements integers up to 32 bits and the arithmetic supported by them. Since the 64-bit result is subsequently shifted 32 places to the right and rounded, it was possible to solve this by using only 32-bit integers, alternating partial sums and one-bit shifts, instead of performing the entire shift at the end.

Integer arithmetic inference has the advantage of running faster and consuming less power, but it introduces small variations in the results due to the reduction in numerical precision. However, these variations are usually tolerable in many practical cases, especially when looking for lightweight and efficient systems.

4 Experimental evaluation

4.1 Results of fault injection experiments in the unprotected system

In order to evaluate the performance of the neural network when memory errors affect its weights, we have carried out different fault injection experiments [2] in the embedded system. To do this, single and multiple adjacent bit-flip faults were injected into all bits of all weights, analyzing whether these errors caused changes in one or more outputs. As can be seen in Table 1, even a single erroneous bit can cause changes in the outputs. Furthermore, the percentage of changed outputs increases as more faults are injected, an expected result.

Table 1. Percentage of outputs changed

Adjacent Faults	Layer 1	Layer 2	No Variation
FP32			
1	3,78	2,42	93,81
2	3,82	2,43	93,75
3	3,62	2,46	93,92
4	3,75	2,58	93,68
5	3,88	2,68	93,44
6	4,01	2,77	93,22
7	4,24	2,86	92,9
8	4,40	2,98	92,62
BF16			
1	6,41	3,81	89,78
2	6,51	3,83	89,67
3	5,77	3,85	90,38
4	5,80	3,87	90,33
5	5,84	3,9	90,26
6	5,89	3,93	90,18
7	6,15	3,97	89,88
8	6,25	4,02	89,73
INT8			
1	2,08	0,48	97,43
2	2,38	0,47	97,15
3	2,68	0,6	96,73
4	2,92	0,65	96,43
5	3,13	0,67	96,21
6	3,27	0,69	96,03
7	3,72	0,74	95,54
8	3,87	0,89	95,24

We can also see that the network with the INT8 neural network presents the lowest percentage of changed outputs, a common behavior in this type of

model [38]. As quantization limits the range of representable values, even in the presence of errors in the parameters, their value remains within the expected range, which limits the effect of the error. On the contrary, errors in the FP32 format can change the represented value by many orders of magnitude, with a devastating effect on the inference process.

On the other hand, BF16 network shows the highest percentage of changed outputs. This is because the BF16 format eliminates the superfluous bits that could be found in the FP32 format network, and whose modification does not affect the network's behavior. In other words, the BF16 format consists of fewer, but more significant, bits, so the variations are greater than in FP32 in proportion to the number of injections.

It can also be seen that the weights in Layer 1 are more likely to induce changes in the outputs than those in Layer 2. This is due to the complete connection between neurons in consecutive stages: an error in the intermediate stage propagates to all neurons in the output stage, while an error in the final stage only affects one neuron, which reduces the damage that can be caused.

Table 2 shows the percentage of failures that occurred with respect to the changed outputs, considering a variation in the output value greater than 5% as a failure. We can observe that the optimised neural networks (BF16 and INT8 networks) have a higher failure percentage than the FP32 network. As mentioned before, in the BF16 optimized network, all superfluous bits have been removed, leaving only the truly significant bits. On the other hand, although there is a lower percentage of changed outputs in the INT8 network (as seen in Table 1), these variations mainly provoke a change in outputs greater than 5%.

Table 2. Percentage of failures

	FP32		**BF16**		**INT8**	
Adjacent faults	**Layer 1**	**Layer 2**	**Layer 1**	**Layer 2**	**Layer 1**	**Layer 2**
1	12,56	3,91	15,79	5,21	75,00	66,07
2	11,56	5,33	15,03	7,41	71,43	70,69
3	4,93	4,89	7,14	7,14	72,22	74,07
4	5,21	4,85	8,12	7,69	73,47	67,31
5	4,93	4,86	8,33	8,33	71,43	68,89
6	4,95	4,57	9,09	8,59	78,79	70,59
7	6,49	4,07	12,90	8,33	84,00	69,57
8	6,76	3,52	14,88	8,02	92,31	76,92

It should be noted that, in more realistic environments, an optimized network is more robust, as the probability of errors decreases as fewer bits are required to store the different parameters. This can be verified theoretically with Equation 5, which shows the overall failure rate, which can be seen in Fig. 5 (considering λ equal to the failure rate -which in our case is 1 since we have performed an exhaustive fault injection in all bits of all weights-, n_cells the number of cells

occupied by the neural network weights relative to the overall number of cells, and $n\_failure$ the number of failures). In this figure, we can confirm that the optimized networks are more robust, with a higher failure rate of the FP32 network, followed by the BF16 failure rate.

$$Failure\_rate = \lambda \times n\_cells \times n\_failure \tag{5}$$

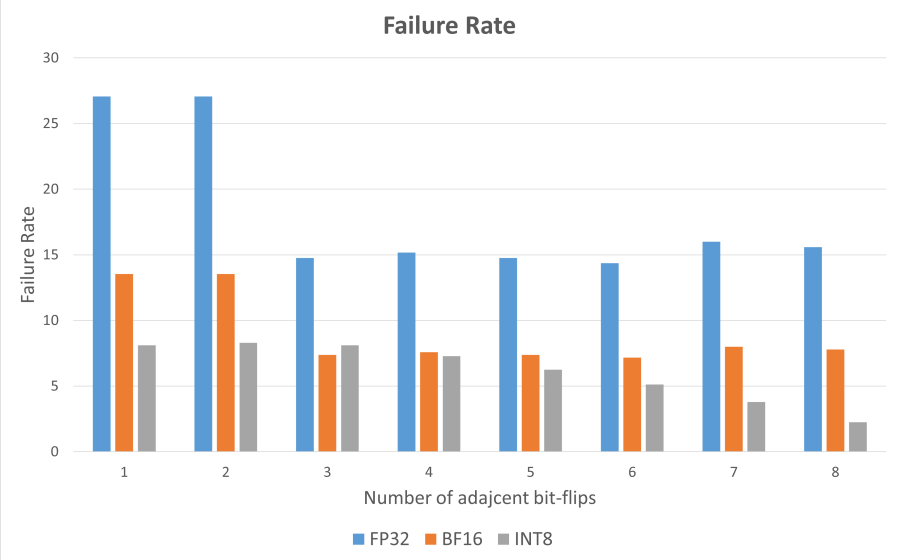

Fig. 5. FP32 vs BF16 format [32].

4.2 Error correcting codes used

As we have just checked, a variation in a single bit can cause changes in the outputs of the neural network. If we want to use it in a critical environment, the weights of the neural network must be protected. We have used a series of Error Correcting Codes (ECCs) to protect these weights.

- *Single Error Correction-Double Adjacent Error Correction (SEC-DAEC).* We have used different versions of this ECC, adapted to the word length of each data (FP32, BF16 and INT8). They are based on Hamming SEC codes and the interleaving technique [36].
- *Asymmetric ECC.* This type of ECC is able to protect some bits more than others. The idea behind applying Asymmetric ECC is to sacrifice some robustness of the neural network in exchange for improving both the size and

latency of the program, protecting some bits more than others [13]. As in the previous ECC, different versions of this ECC, adapted to the word length of each data (FP32, BF16 and INT8) have been used. Specifically, we have used:

  - UEP (40, 32) for the FP32 network. This ECC can correct single and double adjacent errors in the 16 most significant bits (the protected area); and single errors in both the protected area and the parity bits.
  - UEP (22, 16) for the BF16 network. This ECC protects the 10 highest-order bits against single or double adjacent errors. That is, if a single or double adjacent error occurs, regardless of the position, the 10 highest-order bits are guaranteed to be correct. If the error affects those bits, they are corrected. If the error affects unprotected bits, the decoding process is guaranteed to not affect the protected bits, ensuring their integrity.
  - UEP (12, 8) for the INT8 network. As in the previous case, this asymmetric ECC guarantees that the four highest-order bits will be protected against single or double adjacent errors regardless of the affected bits of the codeword.

– *FUEC-DAEC (23, 16) [14].* This ECC is capable of correcting single and double adjacent errors in 16-bit words with only 7 parity bits. Furthermore, it can also detect bursts of 3- or 4-bit errors. It has been used in the FP32 and BF16 networks.

4.3 Results of fault injection experiments in the protected system

As we have already seen, errors in the neural network weights cause failures in the outputs. To reduce these effects, we have added the ECCs just described to the different networks. Table 3 shows the failure rate of the different neural networks protected by the ECCs. With respect to the unprotected networks, the failure rate has decreased in all cases, an expected result. Specially significant are the results of single and double adjacent errors, as failure rate is 0 in almost all cases, according to the fault tolerance of the corresponding ECC. UEP ECC for the INT8 is the unique case where the failure rate is not 0 for single and double adjacent errors. This is the normal behaviour of the UEP ECC. In the cases of FP32 and BF16, the intrinsic redundancy of these networks causes a failure rate equal to 0 when using this type of ECC.

4.4 Analysis of the introduced overhead

The incorporation of fault-tolerance mechanisms in the system generates an overhead in both the software size and its latency. Table 4 presents the results obtained, where the values in parentheses indicate the percentage of memory used. It should be noted that, in the implementation of the different Error Correcting Codes, a literal translation of the design into hardware has not been made. Instead, software structures have been used that minimize the overhead introduced by the ECCs.

Table 3. Failure rate

FP32			
Adjacent Faults	**FUEC-DAEC**	**SEC-DAEC**	**UEP**
1	0,00	0,00	0,00
2	0,00	0,00	0,00
3	14,77	15,59	19,69
4	15,18	15,18	37,32
5	15,18	34,45	40,61
6	14,36	15,18	58,65
7	11,89	18,87	20,51
8	15,59	16,41	53,73
BF16			
Adjacent Faults	**FUEC-DAEC**	**SEC-DAEC**	**UEP**
1	0,00	0,00	0,00
2	0,00	0,00	0,00
3	6,56	6,56	6,56
4	14,15	14,15	14,15
5	10,87	10,87	10,87
6	7,79	7,79	7,79
7	9,02	9,02	9,02
8	8,82	8,82	8,82
INT8			
Adjacent Faults	**SEC-DAEC V1**	**SEC-DAEC V2**	**UEP**
1	0,00	0,00	2,67
2	0,00	0,00	2,46
3	8,20	7,69	8,82
4	7,49	7,49	9,02
5	6,15	6,36	6,87
6	4,61	5,13	5,02
7	2,87	2,87	3,79
8	2,26	1,13	2,26

Table 4. System overhead

FP32			
	Sketch	**Global variables**	**Execution time (us)**
Unprotected	5556 (17%)	265 (12%)	2576
FUEC-DAEC	8716 (27%)	341 (16%)	14122
SEC-DAEC	7870 (24%)	341 (16%)	13267
UEP	7908 (24%)	307 (14%)	14084
BF16			
	Sketch	**Global variables**	**Execution time (us)**
Unprotected	5452 (16%)	179 (8%)	2601
FUEC-DAEC	6592 (20%)	221 (10%)	6592
SEC-DAEC	6192 (19%)	221 (10%)	3717
UEP	6152 (19%)	221 (10%)	3596
INT8			
	Sketch	**Global variables**	**Execution time (us)**
Unprotected	5740 (17%)	658 (32%)	2916
SEC-DAEC V1	6236 (19%)	700 (34%)	3649
SEC-DAEC V2	6186 (196%)	700 (34%)	3666
UEP	6136 (19%)	700 (34%)	3552

Software size analysis. As can be seen in Table 4, memory consumption decreases significantly with the size of the parameters. If we also take into account that the ECCs used are also more compact, the final result is a clear decrease in the amount of memory used, even though the neural network is very small.

Regarding the ECCs themselves, the differences in memory consumption are minimal. One of the reasons is that all ECCs have similar error correction capabilities (they all correct single and double adjacent errors), which means that the formulas used for error correction are similar. It can be observed that the fewer the parity bits, the lower the memory consumption. In this sense, the UEP ECC requires the least memory, as it corrects the fewest errors (single and double adjacent errors in the protected area).

Regarding the FUEC-DAEC ECC, it occupies the largest memory space since it has the highest fault tolerance capabilities (it is capable of correcting single and double adjacent faults, and detecting bursts of 3 and 4 erroneous bits). It should be noted that in the experiments carried out with the FUEC-DAEC ECC, the 3- or 4-bit burst error detection offered by this ECC was not used, since we were initially only interested in error correction. However, considering the memory occupancy data and the time overhead associated with this ECC (analyzed in the next section), it could be beneficial to take advantage of this fault tolerance capability. For example, when detecting these types of errors, the calculations performed by the neural network could be repeated, that is, it is possible to implement temporal redundancy [9]. However, it is important to keep

in mind that implementing this additional functionality would imply an increase in both memory usage and program latency.

Runtime analysis. Regarding the execution time, given that the processor commands require approximately 750 ms to generate the desired effects on the mobile robot's mechanical components, there is sufficient time to perform the verification of the various ECCs without affecting the normal operation of the robot.

One notable finding is the decrease in execution time for the protected versions in the optimized neural networks. This significant reduction is due to the fact that the ECC error checking operations are performed with 8- or 16-bit data, stored as unsigned integers. This makes the operations performed on the Arduino more efficient, which directly affects execution time. Regarding the unprotected version, it can be seen that the execution time is similar in all versions of the neural network.

5 Conclusions and future work

In this work, we have analyzed the integration of different Error Correcting Codes (ECC) into a neural network implemented in Arduino with parameters in different formats (FP32, BF16 and INT8). Our main objective has been to protect the neural network weights from memory faults and to evaluate whether the generated overhead is manageable in terms of software size and execution time.

Although the neural network used is small, it is useful as a test case to study this type of protection against memory faults. To evaluate the system's behavior, we have exhaustively run several fault injection campaigns on the network's weights, both when protected and unprotected. In any case, we have observed that the inclusion of ECCs does not entail excessive overhead, making its use viable in this context. We have also seen a significant improvement in memory consumption and execution time when using reduced size parameters.

As future work, we plan to continue exploring new ECC and other fault tolerance mechanisms to evaluate their applicability in this type of embedded systems.

Acknowledgments. This work has been funded by the Spanish Government (DEFADAS project, Grant PID2020-120271RB-I00, MCIN/AEI/10.13039/501100011033) and by Universitat Politècnica de València (Convocatoria A+D, Proyectos de Innovación y Mejora Educativa, PIME/24-25/435).

References

1. IEEE standard for floating-point arithmetic. IEEE Std 754-2008 (2008)

2. Benso, A., Prinetto, P.: Fault Injection Techniques and Tools for Embedded Systems Reliability Evaluation. Springer Publishing Company, Incorporated, 1st edition (2010)
3. Brando, A., Serra, I., Mezzetti, E., Cazorla, F.J., Perez-Cerrolaza, J., Abella, J.: On neural networks redundancy and diversity for their use in safety-critical systems. Computer **56**(5), 41–50 (2023). https://doi.org/10.1109/MC.2023.3236523
4. Burel, S., Evans, A., Anghel, L.: Zero-overhead protection for cnn weights. In: 2021 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT). pp. 1–6 (2021). https://doi.org/10.1109/DFT52944.2021.9568363
5. Cheng, Y., Wang, D., Zhou, P., Zhang, T.: Model compression and acceleration for deep neural networks: The principles, progress, and challenges. IEEE Signal Processing Magazine **35**(1), 126–136 (2018). https://doi.org/10.1109/MSP.2017.2765695
6. Coelho, B.L., Dos Santos, F.F., Saveriano, M., Allen, G., Daniel, A., Guertin, S., Vartanian, S., Wyrwas, E., Frost, C., Rech, P.: Impact of radiation-induced effects on embedded gpus executing large machine learning models. IEEE Transactions on Nuclear Science, In Press pp. 1–1 (2025)
7. Deng, L., Li, G., Han, S., Shi, L., Xie, Y.: Model compression and hardware acceleration for neural networks: A comprehensive survey. Proceedings of the IEEE **108**(4), 485–532 (2020). https://doi.org/10.1109/JPROC.2020.2976475
8. Dhiraj Kalamkar, e.a.: A study of bfloat16 for deep learning training (2019), https://arxiv.org/abs/1905.12322, arXiv:1905.12322
9. Dubrova, E.: Fault-tolerant design. Springer-Verlag New York (2013)
10. Feng, B., Wang, Y., Geng, T., Li, A., Ding, Y.: Apnn-tc: accelerating arbitrary precision neural networks on ampere gpu tensor cores. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. SC '21, Association for Computing Machinery, New York, NY, USA (2021). https://doi.org/10.1145/3458817.3476157
11. Goldstein, B.F., Srinivasan, S., Das, D., Banerjee, K., Santiago, L., Ferreira, V.C., Nery, A.S., Kundu, S., França, F.M.G.: Reliability evaluation of compressed deep learning models. In: 2020 IEEE 11th Latin American Symposium on Circuits & Systems (LASCAS). pp. 1–5 (2020). https://doi.org/10.1109/LASCAS45839.2020.9069026
12. Gracia-Moran, J., Ruiz, J.C., de Andres, D., Saiz-Adalid, L.J.: Allocating ecc parity bits into bf16-encoded cnn parameters: A practical experience report. In: Proceedings of the 13th Latin-American Symposium on Dependable and Secure Computing. p. 75–80. LADC '24, Association for Computing Machinery, New York, NY, USA (2024), https://doi.org/10.1145/3697090.3697092
13. Gracia-Morán, J., Ruiz-García, J.C., Saiz-Adalid, L.J.: Uso de códigos de corrección de errores asimétricos en un sistema empotrado. In: VII Jornadas de Computación Empotrada y Reconfigurable (JCER2023), Jornadas SARTECO,. pp. 699–706 (2023)
14. Gracia-Morán, J., Saiz-Adalid, L.J., Gil-Tomás, D., Gil-Vicente, P.J.: Improving error correction codes for multiple-cell upsets in space applications. IEEE Transactions on Very Large Scale Integration (VLSI) Systems **26**(10), 2132–2142 (2018). https://doi.org/10.1109/TVLSI.2018.2837220
15. Guan, H., Ning, L., Lin, Z., Shen, X., Zhou, H., Lim, S.H.: In-place zero-space memory protection for CNN. Curran Associates Inc., Red Hook, NY, USA (2019)

16. Hickmann, B., Chen, J., Rotzin, M., Yang, A., Urbanski, M., Avancha, S.: Intel Nervana Neural Network Processor-T (NNP-T) Fused Floating Point Many-term Dot Product. In: 2020 IEEE 27th Symposium on Computer Arithmetic (ARITH). pp. 133–136 (2020). https://doi.org/10.1109/ARITH48897.2020.00029
17. Jacob, B., et al.: Quantization and training of neural networks for efficient integer-arithmetic-only inference. Tech. rep. (2017), https://arxiv.org/abs/1712.05877
18. Jang, M., Hong, J.: Mate: Memory-and retraining-free error correction for convolutional neural network weights. Journal of Information and Communication Convergence Engineering **19**(1), 22–28 (Mar 2021). https://doi.org/10.6109/jicce.2021.19.1.22
19. JunKyu Lee, e.a.: Resource-efficient convolutional networks: A survey on model-, arithmetic-, and implementation-level techniques. ACM Comput. Surv. **55**(13s) (Jul 2023). https://doi.org/10.1145/3587095
20. Lee, J., Kim, C., Kang, S., Shin, D., Kim, S., Yoo, H.J.: Unpu: An energy-efficient deep neural network accelerator with fully variable weight bit precision. IEEE Journal of Solid-State Circuits **54**(1), 173–185 (2019). https://doi.org/10.1109/JSSC.2018.2865489
21. Li, G., Hari, S.K.S., Sullivan, M., Tsai, T., Pattabiraman, K., Emer, J., Keckler, S.W.: Understanding error propagation in deep learning neural network (DNN) accelerators and applications. In: SC17: International Conference for High Performance Computing, Networking, Storage and Analysis. pp. 1–12 (2017)
22. Li, Z., Liu, F., Yang, W., Peng, S., Zhou, J.: A survey of convolutional neural networks: Analysis, applications, and prospects. IEEE Transactions on Neural Networks and Learning Systems **33**(12), 6999–7019 (2022). https://doi.org/10.1109/TNNLS.2021.3084827
23. Marinó, G.C., Petrini, A., Malchiodi, D., Frasca, M.: Deep neural networks compression: A comparative survey and choice recommendations. Neurocomputing **520**, 152–170 (2023). https://doi.org/https://doi.org/10.1016/j.neucom.2022.11.072, https://www.sciencedirect.com/science/article/pii/S0925231222014643
24. Mittal, S.: A survey on modeling and improving reliability of dnn algorithms and accelerators. Journal of Systems Architecture **104**, 101689 (2020). https://doi.org/https://doi.org/10.1016/j.sysarc.2019.101689, https://www.sciencedirect.com/science/article/pii/S1383762119304965
25. N. P. Jouppi, e.a.: In-datacenter performance analysis of a tensor processing unit. In: Proceedings of the 44th Annual International Symposium on Computer Architecture. p. 1–12. ISCA '17, Association for Computing Machinery, New York, NY, USA (2017). https://doi.org/10.1145/3079856.3080246
26. Nguyen, D.T., Ho, N.M., Chang, I.J.: St-DRC: Stretchable dram refresh controller with no parity-overhead error correction scheme for energy-efficient DNNs. In: 2019 56th ACM/IEEE Design Automation Conference (DAC). pp. 1–6 (2019)
27. Novikov, A., Podoprikhin, D., Osokin, A., Vetrov, D.: Tensorizing neural networks. In: Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1. p. 442–450. NIPS'15, MIT Press, Cambridge, MA, USA (2015)
28. Qutub, S., Geissler, F., Peng, Y., Gräfe, R., Paulitsch, M., Hinz, G., Knoll, A.: Hardware faults that matter: Understanding and estimating the safety impact of hardware faults on object detection dnns. In: Trapp, M., Saglietti, F., Spisländer, M., Bitsch, F. (eds.) Computer Safety, Reliability, and Security. pp. 298–318. Springer International Publishing, Cham (2022)

29. Rakin, A.S., He, Z., Fan, D.: Bit-flip attack: Crushing neural network with progressive bit search. In: IEEE/CVF International Conference on Computer Vision. pp. 1211–1220 (2019)
30. Reagen, B., Gupta, U., Pentecost, L., Whatmough, P., Lee, S.K., Mulholland, N., Brooks, D., Wei, G.Y.: Ares: A framework for quantifying the resilience of deep neural networks. In: 2018 55th ACM/ESDA/IEEE Design Automation Conference (DAC). pp. 1–6 (2018). https://doi.org/10.1109/DAC.2018.8465834
31. Ruiz, J.C., de Andrés, D., Saiz-Adalid, L.J., Gracia-Morán, J.: In-memory zero-space floating-point-based cnn protection using non-significant and invariant bits. In: Ceccarelli, A., Trapp, M., Bondavalli, A., Bitsch, F. (eds.) Computer Safety, Reliability, and Security. pp. 3–17. Springer Nature Switzerland, Cham (2024)
32. Ruiz, J.C., de Andrés, D., Saiz-Adalid, L.J., Gracia-Morán, J.: Tolerancia a fallos múltiples en redes convolucionales en coma flotante de 16 bits utilizando códigos correctores de errores. In: VIII Jornadas de Computación Empotrada y Reconfigurable (JCER2024), Jornadas SARTECO,. pp. 823–832 (2024)
33. Ruiz, J.C., de Andrés, D., Saiz-Adalid, L.J., Gracia-Morán, J.: Zero-space in-weight and in-bias protection for floating-point-based CNNs pp. 89–96 (2024). https://doi.org/10.1109/EDCC61798.2024.00028
34. Ruospo, A., Gavarini, G., de Sio, C., Guerrero, J., Sterpone, L., Reorda, M.S., Sanchez, E., Mariani, R., Aribido, J., Athavale, J.: Assessing convolutional neural networks reliability through statistical fault injections. In: 2023 Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 1–6 (2023). https://doi.org/10.23919/DATE56975.2023.10136998
35. Sabbagh, M., Gongye, C., Fei, Y., Wang, Y.: Evaluating fault resiliency of compressed deep neural networks. In: 2019 IEEE International Conference on Embedded Software and Systems (ICESS). pp. 1–7 (2019). https://doi.org/10.1109/ICESS.2019.8782505
36. Saiz-Adalid, L.J., Gil, P., Baraza-Calvo, J.C., Ruiz, J.C., Gil-Tomas, D., Gracia-Moran, J.: Modified Hamming Codes to Enhance Short Burst Error Detection in Semiconductor Memories (Short Paper) . In: 2014 Tenth European Dependable Computing Conference (EDCC). pp. 62–65. IEEE Computer Society, Los Alamitos, CA, USA (May 2014). https://doi.org/10.1109/EDCC.2014.25
37. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted Residuals and Linear Bottlenecks . In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4510–4520. IEEE Computer Society, Los Alamitos, CA, USA (Jun 2018). https://doi.org/10.1109/CVPR.2018.00474
38. Syed, R.T., Ulbricht, M., Piotrowski, K., Krstic, M.: Fault resilience analysis of quantized deep neural networks. In: 2021 IEEE 32nd International Conference on Microelectronics (MIEL). pp. 275–279 (2021). https://doi.org/10.1109/MIEL52794.2021.9569094
39. Y. Bengio, Y. Lecun, G.H.: Deep learning for AI. Communications of the ACM **64**(7), 58–65 (2021). https://doi.org/10.1145/3448250

LALABY GLIO APP: DIGITAL PHENOTYPING OF QUALITY OF LIFE IN GLIOBLASTOMA PATIENTS

Sabina Asensio–Cuesta[1], Elies Fuster-García[1], Juan M. García–Gómez[1], Jorge Soler[3], Inmaculada Maestu[3], Daniel Sánchez–García[1], Ángel Sánchez–García[1], Carlos Sáez-Silvestre[1], Teresa Soria Comes[3], Maria De Julian Campayo[3]

[1] Instituto de Tecnologías de la Información y Comunicaciones (ITACA), Universitat Politècnica de València, Camino de Vera s/n, 46022 Valencia, Spain;
[2] Consorcio Hospitalario Provincial de Castellón, Avda. Dr. Clará, 19, 12002 Castellón, Spain
[3] Hospital Universitario Dr. Peset de Valencia, Avenida de Gaspar Aguilar 90, 46017 Valencia, Spain

Abstract

This research presents Lalaby-Glio, an adaptation of the Lalaby app for monitoring quality of life (QoL) and capturing digital phenotypes in patients with glioblastoma. The interface was redesigned using nature-inspired imagery from the Valencian region to enhance user experience. A feasibility study with three male patients (mean age: 55.3) showed that two completed the 18-week protocol without issues; one withdrew after 7 weeks due to technological burden detected via the Lalaby dashboard.Patient 1 continued for 32 weeks, completing 203 daily symptom reports, 31 general QoL questionnaires (EORTC QLQ-C30), and 29 specific QoL questionnaires (EORTC BN-20), with an average app rating of 4.15/5. Patient 2 completed 18 weeks, submitting 101 daily reports, 17 general and 18 specific QoL questionnaires, with a rating of 3.1/5. Patient 3 contributed 7 weeks of data, completing 5 general and 7 specific QoL questionnaires, with a rating of 2.4/5.These findings support Lalaby-Glio's potential for digital phenotyping and non-invasive QoL monitoring in patients with glioblastoma. The app successfully integrated passive and active data collection with validated cancer-specific instruments. The nature-inspired redesign was positively received by the two long-term users. Future stages include a pre-pilot study with 15 participants to assess usability, engagement, and clinical relevance at scale.

Technical description of the result

- **What Is Lalaby-Glio?**
  A mobile app for real-time monitoring of quality of life in glioblastoma patients
- **Smartphone-based**
  Patients use their **own device**, with **no need for wearables**.
- **Multimodal Data Collection**
  Passive data from smartphone sensors
  Active input through patient-reported outcomes (PROMs)
- **Digital Phenotyping**
  The app generates a **digital phenotype** to detect changes in each patient's quality of life over time.

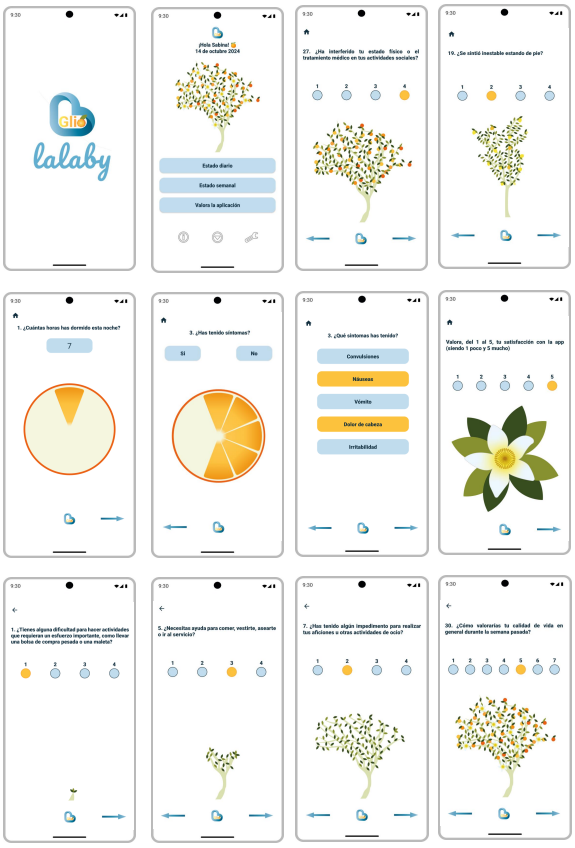

Figure 1. Screenshots of the Lalaby-Glio App Interface

Applications

- **Feasability study developed in** Hospital Universitario Doctor Peset de Valencia involving 3 patients.
- **Ethical committees**: H. Provincial de Castellón; UPV (P11_26-03-2025);
- **Multidisciplinary team:** UPV researchers; Oncology researchers from the Provincial Hospital of Castellón and Hospital Universitario Doctor Peset de Valencia (Fundación Fisabio)
- **Financing entity: GVA.** Proyecto CIAICO/2022/064

Tale 1. Patient Engagement and Data Collection Metrics in the Lalaby-Glio feasibility study

Patient	Weeks of Use	Daily Reports Completed	General QoL Questionnaires *(EORTC QLQ-C30)*	Specific QoL Questionnaires*(EORTC BN-20)*	App Rating (1–5 scale)	Completion Status
Patient 1	32	203	31	29	4.15	Completed (extended use)
Patient 2	18	101	17	18	3.10	Completed
Patient 3	7	45	5	7	2.40	Withdrew (technological burden)
Total	**57**	**349**	**53**	**54**	**Mean 3.2**	

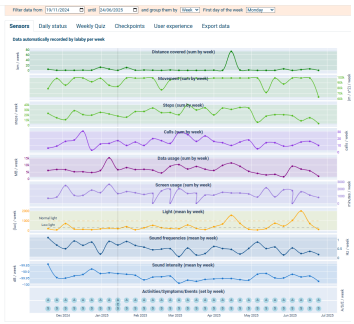

Figure 2. Lalaby-Glio Dashboard: Sensor Data Interactive Interface from a Patient Example

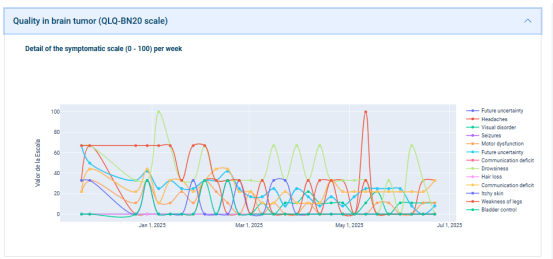

Figure 3. Lalaby-Glio Dashboard: Example of Quality of Life Scores from the QLQ-BN20 Scale for an Individual Patient

Additional comments

The Lalaby-Glio app demonstrates promising feasibility for real-time, non-invasive monitoring of quality of life (QoL) in patients with glioblastoma through digital phenotyping. The adaptation of the original Lalaby app—previously validated in lung cancer—has been successfully tailored to this new clinical context by integrating brain tumour–specific PROMs (EORTC QLQ-C30 and BN-20) and a nature-inspired visual redesign based on Valencian landscapes to enhance patient engagement.

Initial results from three male patients (mean age: 55.3) showed strong adherence in two participants who completed the 18-week protocol without technical issues. One patient extended usage to 32 weeks. The third participant discontinued use after 7 weeks, with app data indicating a perceived technological burden. Across users, the system collected over 350 daily reports and 119 QoL questionnaires, confirming the app's capability to integrate both passive and active data sources.

These findings support the technical and experiential viability of Lalaby-Glio, while also emphasising the importance of usability and personalisation for sustained patient participation. Upcoming phases will include a pre-pilot study involving 15 glioblastoma patients to further validate digital QoL markers and assess clinical utility at scale.

The ITACA-WIICT is a meeting forum for scientifics, technicians and other professionals who are dedicated to Information and communication technologies study and research. Its fundamental scope is to promote the contact among scientific and professionals, improving the cooperation as well as the technological transfer among professionals.

Institute ITACA
Information and Communication Technologies