JCC-BD & ET 2026

14th Conference on Cloud Computing, Big Data & Emerging Topics

(JCC-BD&ET 2026)

HTTPS://JCC.INFO.UNLP.EDU.AR

@CONF_CC_BD_ET

JCC@LIDI.INFO.UNLP.EDU.AR

POSTGRADO
FACULTAD DE INFORMÁTICA

III-LIDI
INSTITUTO DE INVESTIGACIÓN EN INFORMÁTICA - LIDI

14th Conference on Cloud Computing Conference, Big Data & Emerging Topics (JCC-BD&ET 2026)

La Plata, Buenos Aires, Argentina.

June 23−25, 2026

Instituto de Investigación en Informática LIDI (III-LIDI).
14th Conference on Cloud Computing, Big Data and Emerging Topics JCC-BD&ET 2026 ; Compilación de Marcelo Ricardo Naiouf ... [et al.]. - 1a ed. - La Plata : Universidad Nacional de La Plata. Facultad de Informática, 2026.
Libro digital, PDF

Archivo Digital: descarga y online
ISBN 978-950-34-2766-8

ISBN 978-950-34-2766-8

1. Actas de Congresos. 2. Computación. 3. Conferencias. I. Naiouf, Marcelo Ricardo, comp.

CDD 372.34

Preface

Welcome to the proceedings of the 14th Conference on Cloud Computing, Big Data & Emerging Topics (JCC-BD&ET 2026), held in a hybrid modality (both on-site and live online settings were allowed). JCC-BD&ET 2026 was organized by the III-LIDI and the Postgraduate Office, both from School of Computer Science of the National University of La Plata, Argentina.

Since 2013, this event has been an annual meeting where ideas, projects, scientific results and applications in cloud computing, big data and other related areas are exchanged and disseminated. The conference focuses on the topics that allow interaction between academia, industry, and other interested parties.

JCC-BD&ET 2026 covered the following topics: high-performance, edge and fog computing; internet of things; modelling and simulation; big and open data; machine and deep learning; smart cities; e-government; human-computer interaction; visualization; and special topics related to emerging technologies. In addition, special activities were also carried out, including 1 plenary lecture and 2 discussion panels.

Special thanks to all the people who contributed to the conference's success: program and organizing committees, authors, reviewers, speakers, and all conference attendees.

June 2026

Marcelo Naiouf
Franco Chichizola
Laura De Giusti
Leandro Libutti

Organization

Program Committee

María José Abásolo	Universidad Nacional de La Plata and CIC, Argentina
José Aguilar	Universidad de Los Andes, Venezuela
Jorge Ardenghi	Universidad Nacional del Sur, Argentina
Javier Balladini	Universidad Nacional del Comahue, Argentina
Oscar Bria	Universidad Nacional de La Plata, Argentina
Silvia Castro	Universidad Nacional del Sur, Argentina
Franco Chichizola	Universidad Nacional de La Plata, Argentina
Laura De Giusti	Universidad Nacional de La Plata and CIC, Argentina
Mónica Denham	Universidad Nacional de Río Negro and CONICET, Argentina
Javier Diaz	Universidad Nacional de La Plata, Argentina
Ramón Doallo	Universidade da Coruña, Spain
Marcelo Errecalde	Universidad Nacional de San Luis, Argentina
Elsa Estevez	Universidad Nacional del Sur and CONICET, Argentina
Pablo Ezzatti	Universidad de la República, Uruguay
Aurelio Fernandez Bariviera	Universitat Rovira i Virgili, Spain
Fernando Emmanuel Frati	Universidad Nacional de Chilecito, Argentina
Carlos Garcia Garino	Universidad Nacional de Cuyo, Argentina
Carlos García Sánchez	Universidad Complutense de Madrid, Spain
Adriana Angélica Gaudiani	Universidad Nacional de General Sarmiento, Argentina
Graciela Verónica Gil Costa	Universidad Nacional de San Luis and CONICET, Argentina
Roberto Guerrero	Universidad Nacional de San Luis, Argentina
Waldo Hasperué	Universidad Nacional de La Plata and CIC, Argentina
Francisco Daniel Igual Peña	Universidad Complutense de Madrid, Spain
Tomasz Janowski	Gdansk University of Technology, Poland
Laura Lanzarini	Universidad Nacional de La Plata, Argentina
Guillermo Leguizamón	Universidad Nacional de San Luis, Argentina
Leandro Libutti	Universidad Nacional de La Plata, Argentina

Edimara Luciano	Pontificia Universidade Católica do Rio Grande do Sul, Brazil
Emilio Luque Fadón	Universidad Autónoma de Barcelona, Spain
Mauricio Marín	Universidad de Santiago de Chile, Chile
Luis Marrone	Universidad Nacional de La Plata, Argentina
Marcelo Naiouf	Universidad Nacional de La Plata, Argentina
Katzalin Olcoz Herrero	Universidad Complutense de Madrid, Spain
José Angel Olivas Varela	Universidad de Castilla-La Mancha, Spain
Xoan Pardo	Universidade da Coruña, Spain
Patricia Pesado	Universidad Nacional de La Plata, Argentina
Mario Piattini	Universidad de Castilla-La Mancha, Spain
María Fabiana Piccoli	Universidad Nacional de San Luis, Argentina
Luis Piñuel	Universidad Complutense de Madrid, Spain
Adrian Pousa	Universidad Nacional de La Plata, Argentina
Marcela Printista	Universidad Nacional de San Luis, Argentina
Dolores Isabel Rexachs del Rosario	Universidad Autónoma de Barcelona, Spain
Nelson Rodríguez	Universidad Nacional de San Juan, Argentina
Enzo Rucci	Universidad Nacional de La Plata, Argentina
Juan Carlos Saez Alcaide	Universidad Complutense de Madrid, Spain
Aurora Sánchez	Universidad Católica del Norte, Chile
Victoria Sanz	Universidad Nacional de La Plata, Argentina
Remo Suppi	Universidad Autónoma de Barcelona, Spain
Francisco Tirado Fernández	Universidad Complutense de Madrid, Spain
Juan Touriño Dominguez	Universidade da Coruña, Spain
Gabriela Viale Pereira	Danube University Krems, Austria
Alvaro Wong	Universidad Autónoma de Barcelona, Spain
Gonzalo Zarza	Globant, Argentina

Table of Contents

Full papers 1

A Modular Framework for Insider Threat Detection and Analysis. *Mauricio Díeguez-Rebolledo (UFRO – Chile), Estéban Esparza (UFRO – Chile), Jorge Hochstetter-Diez (UFRO – Chile).* 2

A Scalable Architecture for Real-Time Public Health Vulnerability Assessment. *Luciano Perdomo (UNPSAJB / CENPAT-CONICET), Leo Ordinez (UNPSAJB / CENPAT-CONICET), Sergio Andrés Kaminker (UNPSAJB / CENPAT-CONICET).* 16

Building and Acceptance of Intelligent Immersive Assistants. *Pablo Pintado (UDA – Ecuador), Diego Amoroso (UDA – Ecuador), Amanda Roldán (UDA – Ecuador), Gabriela Duque (UDA – Ecuador), Daniel Iturralde (UDA – Ecuador), Esteban Mora (UDA – Ecuador), Daniel Capelo (UDA – Ecuador), José Fajardo (UDA – Ecuador).* 30

Feasibility Assessment for the Implementation of FinOps Practices in SMEs in the Province of Misiones: A Study Based on Local Infrastructure and Provincial Cloud Computing Services. *Miryan R. Puchini (UNaM), Nancy B. Ganz (UNaM), Yolanda E. Zado (UNaM).* 45

Lessons Learned from Developing a Blockchain-based Traceability Solution for Medicinal Cannabis Production in Argentina. *Facundo Paterno (UNCOMA), Alejandra Cechich (UNCOMA), Alan De Renzis (UNCOMA).* 58

Mobile Crypto Wallet User Experience Patterns Analysis. *Leandro Plaquin (UNLP), Francisco Gindre (UNLP), Matías Urbieta (UNLP).* 72

Multi-Agent Framework for Resilient Medical Triage in the Internet of Medical Things (IoMT). *Sergio Gramajo (UTN-FRRe), Reinaldo Scappini (UTN-FRRe), Salvador Nuñez (UTN-FRRe), Carlos Torres (UTN-FRRe), Raúl Montiel (UTN-FRRe), Jorge Roa (UTN-FRRe).* 88

Research Data Management in Stochastic Optimization: A Case Study on Heuristic Search and Retrospective Curation. *María Harita (UAB-España), Dolores Rexachs (UAB- España), Alvaro Wong (UAB- España), Emilio Luque (UAB- España).* 102

Short papers 118

A Blockchain-Based MFA Model for Decentralization, Privacy, and Asymmetric-Key Security. *Guillermo Daniel Romero Arregin (UCAECE), Miguel Méndez-Garabetti (UCAECE), Eduardo Piray (UCAECE / FOSSHLab).* 119

A Predictive Model on Percolation Theory and Machine Learning. *Eduardo Miguel Zarate (UNCA), Roman Bond (UNAJ), Diego Encinas (UNAJ / UNLP).* 125

Evaluating Large Language Models for Emergency Triage Prediction from Clinical Narratives: A Latin American Perspective. *José M. Manríquez-Troncoso (UCM – Chile), Xaviera A. López-Cortés (UCM – Chile), Felipe Tirado Marabolí (UCM – Chile).* 130

Full papers

A Modular Framework for Insider Threat Detection and Analysis

Mauricio Diéguez-Rebolledo[1,3,4][0000−0001−8728−7101], Estéban Esparza[2], and Jorge Hochstetter-Diez[1,3][0000−0002−8494−8100]

[1] Computer Sciences and Informatics department, Universidad de La Frontera, Temuco, Chile.
{mauricio.dieguez,jorge.hochstetter}@ufrontera.cl
[2] Master's Program in Informatic Engineering, Universidad de La Frontera, Temuco, Chile.
e.esparza02@ufromail.cl
[3] Centro de Investigación en Inteligencia Artificial Aplicada, Universidad de La Frontera, Temuco, Chile.
[4] Grupo de Ciberseguridad, Universidad de La Frontera, Temuco, Chile

Abstract. Insider threats remain a critical and frequently underestimated challenge in cybersecurity, particularly due to their origin from authorized users and the limitations of traditional detection approaches. Addressing this problem requires not only detection mechanisms, but also integrated approaches for security event analysis and decision support. This study presents the design and implementation of a modular framework for the detection and analysis of insider threats in organizational environments. The proposed system integrates a configurable event simulator, a data ingestion module, an analytical processing API, and an interactive dashboard, forming an end-to-end pipeline for security event analysis. This architecture enables the systematic processing and interpretation of security events, supporting monitoring and decision-making processes. The main contribution of this work lies in the development of an integrated and interpretable analytical framework, combining event classification, weighted risk scoring, and normalization to provide transparent and actionable insights for insider threat management. The system was evaluated in controlled environments using synthetic data derived from the CERT dataset, enabling systematic experimentation in the absence of real-world data. The results demonstrate the system's capability to detect anomalous behaviors, identify risk patterns, and operate efficiently under high-volume scenarios. These findings highlight the feasibility of the proposed approach as a practical foundation for modular insider threat analysis systems in organizational contexts.

Keywords: Insider Threat Detection · Cybersecurity · Risk Analysis

1 Introduction

Insider threats represent an increasingly significant challenge in the field of information security. Unlike external attacks, these threats originate from indi-

viduals with legitimate access to an organization's critical systems and data, making them inherently more difficult to detect and mitigate [11, 18]. Traditional cybersecurity solutions have been primarily designed to address external threats, leaving a gap in the ability to analyze internal user behavior and identify anomalous patterns among authorized users [1, 11, 26].

This challenge has become more pronounced in recent years. According to the Insider Threat Report 2024 [7], 83% of organizations reported experiencing at least one insider-related incident in the past year, with a significant increase in the frequency and scale of such events. Insider threats are also perceived as more difficult to manage than external attacks, with 53% of cybersecurity experts indicating greater difficulty in detecting and preventing internal threats [27]. This difficulty is largely attributed to the legitimate access of insiders, their familiarity with internal systems, and their ability to operate with discretion [27].

From an economic and operational perspective, the impact of insider threats is considerable. The Cost of Insider Risks Global Report [12] reports an average annual cost of USD 16.2 million per organization, with containment time playing a critical role in determining financial impact [21]. In addition to financial losses, insider incidents lead to significant consequences such as data loss, reputational damage, and disruption of critical processes [7]. These findings reinforce the need for effective tools capable of monitoring and analyzing internal behaviors in a timely manner.

In response to this problem, various technological solutions have been developed, including Security Information and Event Management (SIEM) systems [6, 9] and User and Entity Behavior Analytics (UEBA) platforms [8, 13]. While these tools enable the collection and analysis of large volumes of security data, they present important limitations, such as complex configuration requirements, high volumes of alerts, and limited support for intuitive and actionable decision-making [2, 19, 23, 28]. Consequently, despite these advances, a gap persists in the effective monitoring, interpretation, and response to insider threats.

This gap is not only technical but also operational, involving challenges in contextualizing user behavior and translating analytical results into actionable insights. At the same time, recent advances in cloud computing, real-time analytical APIs, and interactive visualization frameworks offer new opportunities for the development of more agile, modular, and integrated solutions capable of addressing these limitations.

In this context, this paper proposes the development and evaluation of a modular prototype system for insider threat detection, monitoring, and analysis. The proposed solution integrates event simulation, data ingestion, analytical processing, and visualization into an end-to-end pipeline that supports the identification of anomalous behaviors and facilitates decision-making through an interpretable risk scoring mechanism. The system has been evaluated under controlled conditions using synthetic data derived from the CERT dataset [14], a widely recognized benchmark in insider threat research.

Although the system has not yet been validated in real organizational environments, its modular and extensible design demonstrates strong potential for

integration into operational contexts. By focusing specifically on insider threats, the proposed approach enables more precise analysis of user behavior, contributing to bridging the gap between conceptual models and practical cybersecurity solutions.

To address these challenges, this work makes the following contributions:

- A modular end-to-end framework for insider threat management that integrates simulation, ingestion, analysis, and visualization components.
- A prototype implementation that enables controlled evaluation in the absence of real-world datasets.
- An interpretable risk scoring model based on event classification, weighting, and normalization.
- A decision-oriented visualization approach that supports the identification of anomalous behaviors and critical risk areas.
- A scalable foundation for future extensions toward real-time monitoring and predictive analytics.

The remainder of this paper is organized as follows. Section 2 presents the fundamental concepts and related work. Section 3 describes the system design and architecture. Section 4 presents the evaluation of the proposed solution. Finally, Section 5 concludes the paper and outlines future work.

2 Background

2.1 Fundamental Concepts

Insider threats have emerged as one of the most critical challenges in information security, primarily due to their origin from individuals with legitimate access to organizational systems and data [25]. Unlike external threats, which can be addressed through perimeter defenses, insiders can bypass traditional controls and operate without easily triggering security alerts [22]. This inherent advantage makes them particularly difficult to detect and mitigate.

According to CERT [3], an insider is defined as an individual with authorized access who uses that access, either maliciously or unintentionally, to compromise the confidentiality, integrity, or availability of organizational assets. This definition emphasizes three key aspects: legitimate access, variable intent, and potentially severe impact.

The difficulty in managing insider threats is further increased by the combination of privileged access and deep knowledge of internal systems, processes, and security mechanisms, enabling more effective evasion of traditional detection approaches [5, 17]. As a result, insider threats represent a multifaceted problem that requires more advanced and context-aware detection strategies.

Although insider threats have existed since the early days of computing, their impact has intensified due to digitalization and the increasing complexity of organizational environments. Notable cases, such as the Snowden disclosures [16], the Samsung intellectual property theft [24], and incidents affecting critical

infrastructure [20], demonstrate the wide range of financial, operational, and reputational consequences.

Recent reports confirm the growing relevance of this problem. The 2024 Insider Threat Report [7] indicates that 48% of organizations experienced an increase in insider attacks, while 51% reported multiple incidents within a year. This trend is associated with factors such as the increasing complexity of IT environments, the adoption of emerging technologies, and insufficient security controls. Additionally, a significant proportion of organizations report remediation costs exceeding one million dollars, highlighting the economic impact of these threats [7].

2.2 Main Approaches Identified in Industry and the Literature

To address insider threats, several technological approaches have been developed, aiming to overcome the limitations of traditional security systems in detecting malicious behavior from authorized users. The most widely adopted solutions include Security Information and Event Management (SIEM) systems [15], User and Entity Behavior Analytics (UEBA) [8], and Security Orchestration, Automation, and Response (SOAR) platforms [10].

SIEM systems enable real-time log collection and correlation to detect anomalous events; however, they often generate high volumes of false positives and require continuous tuning. UEBA approaches leverage machine learning and statistical analysis to identify abnormal behavioral patterns, but their effectiveness depends on the availability of large volumes of historical data. SOAR platforms focus on integrating data sources, automating workflows, and improving incident response, although their implementation can be complex and resource-intensive.

Beyond these technologies, academic research has explored insider threat management from multiple perspectives. A systematic mapping study by Esparza et al. [4], analyzing 47 proposals, highlights several key gaps in the literature. First, most approaches focus primarily on detection, with limited attention to predictive and preventive strategies. Second, there is a lack of practical and integrated tools suitable for real-world deployment, as many proposals remain at a conceptual or prototype level. Third, the literature emphasizes the need to integrate multiple approaches, combining behavioral analysis, machine learning, and continuous monitoring.

Additionally, predictive capabilities remain scarce despite their importance for anticipating potential threats, and many proposed solutions face significant implementation challenges, including high data requirements, integration difficulties, and operational costs. These limitations reveal a gap between existing approaches and the need for more practical, integrated, and proactive solutions for insider threat management.

3 Proposed Framework for Insider Threat Management

Traditional solutions such as SIEM, UEBA, and SOAR present limitations related to integration complexity, accuracy, and operational costs. In addition,

the systematic mapping by Esparza et al. [4] highlights key gaps, including the lack of practical and predictive solutions for insider threat management. These limitations motivate the need for new approaches that integrate detection and analysis capabilities in a more effective and adaptable manner.

In this context, this work proposes a modular prototype system designed to support insider threat detection in organizational environments. The current study focuses on the design and evaluation of the prototype under controlled conditions, aiming to validate its technical feasibility and identify its strengths and limitations.

The proposed system integrates complementary components to improve detection accuracy, support timely analysis, and facilitate integration with existing organizational infrastructures. Furthermore, it establishes a foundation for the future incorporation of predictive capabilities based on the analysis of anomalous behaviors identified by the system..

3.1 System Requirements

As an initial stage, system requirements were identified and prioritized, including both functional and non-functional aspects aligned with the project objectives. The prioritization considered their impact on system effectiveness, integration with existing infrastructures, and contribution to reducing incident response times. Table 1 summarizes these requirements.

Table 1. Requirements Prioritization

ID	Name	Priority
Functional Requirements		
RF1.1	Monitoring of Login Attempts	High
RF1.2	Monitoring of After-Hours Activity	High
RF1.3	Monitoring of External File Transfers	High
RF1.4	Access to Sensitive Files	High
RF1.5	Logging of Emails Sent Outside the Organization	High
RF2.1	Daily Periodic Analysis	High
RF2.2	Risk Classification	Medium
RF3.1	Interactive Dashboard for Administrators	High
RF3.2	Detailed Incident View	Medium
RF4.1	Dashboard Customization	Low
Non-Functional Requirements		
RNF-1	Ensure API response times below 2 seconds	High
RNF-2	Design a responsive interface for modern browsers and mobile devices	High
RNF-3	Handle up to 50,000 daily events without performance degradation	High
RNF-4	Provide an intuitive visual experience through interactive charts	High

Functional requirements focus on key capabilities for detecting insider threats, including monitoring login attempts (RF1.1), after-hours activity (RF1.2), ex-

ternal file transfers (RF1.3), access to sensitive files (RF1.4), and external email communications (RF1.5). In addition, the system incorporates periodic analysis (RF2.1) and risk classification (RF2.2), enabling continuous assessment and prioritization of incidents. Visualization capabilities, such as dashboards (RF3.1) and detailed incident views (RF3.2), support effective analysis and monitoring.

Non-functional requirements emphasize performance and usability aspects, including low response times (RNF-1), scalability to handle large volumes of data (RNF-3), and an interactive user experience (RNF-4), which are essential for deployment in dynamic organizational environments.

Overall, these requirements define a set of operational and analytical capabilities that support the detection of anomalous behaviors and timely response to security incidents.

3.2 System Design

The proposed solution is based on a modular architecture designed to address the limitations of traditional and academic approaches. The system is composed of four main components: Event Simulation, Data Ingestion, Analysis, and Visualization, as shown in Fig. 1.

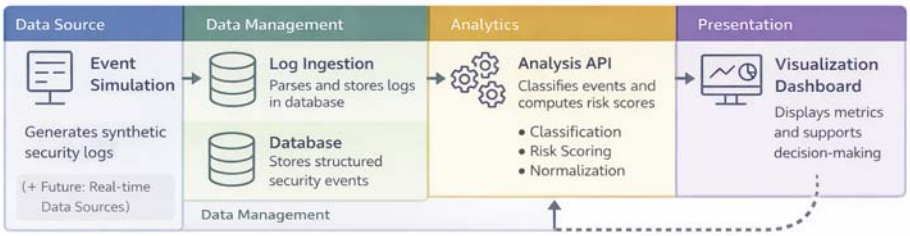

Fig. 1. Main System Components.

The Event Simulation component generates synthetic data that emulates realistic organizational behavior, including both legitimate and malicious activities, aligned with indicators from the CERT dataset [3]. This module enables controlled evaluation of the system and supports the validation of detection mechanisms. It is configurable in terms of user roles, activity patterns, and operational conditions. Although currently used for testing, it is designed to be replaced by real-time data acquisition in production environments.

The Data Ingestion component processes and structures logs generated by the simulator, ensuring their consistency and availability for analysis.

The Analysis component is implemented as an API that processes stored events to identify suspicious patterns and classify them into risk levels (e.g., low, medium, high), supporting decision-making processes.

The Visualization component provides an interactive interface that allows administrators to monitor key metrics, analyze trends, and explore detected incidents.

To illustrate the system structure, C4 diagrams are used. Figure 2 presents the context diagram, identifying two main actors: the security auditor, who interacts with the system, and organizational users, whose activities generate the analyzed data.

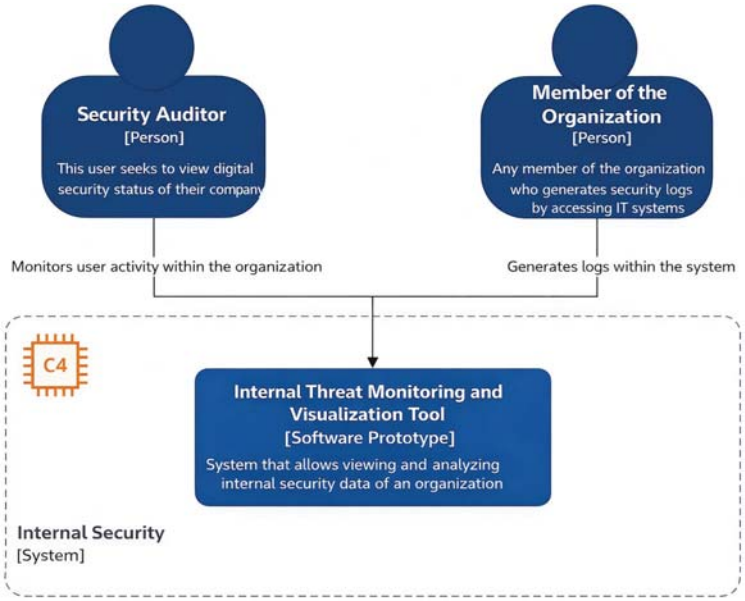

Fig. 2. C4 Context Diagram of the Proposed Framework.

Figure 3 shows the container diagram, including the Event Simulator, Log Ingestion module, Analytical API, Visualization Dashboard, and a relational database that stores processed data.

Event Simulation Component The simulation component generates synthetic data representing both legitimate and suspicious behaviors, aligned with indicators from the CERT dataset [3]. This enables controlled evaluation of the system by distinguishing between normal activities and potential insider threats. The simulator is configurable in terms of user roles, working hours, and event distributions, allowing adaptation to different organizational contexts.

The simulated behaviors include login attempts, access to sensitive files, external data transfers, email communications with attachments, and after-hours activity.

The component is structured into the following modules:

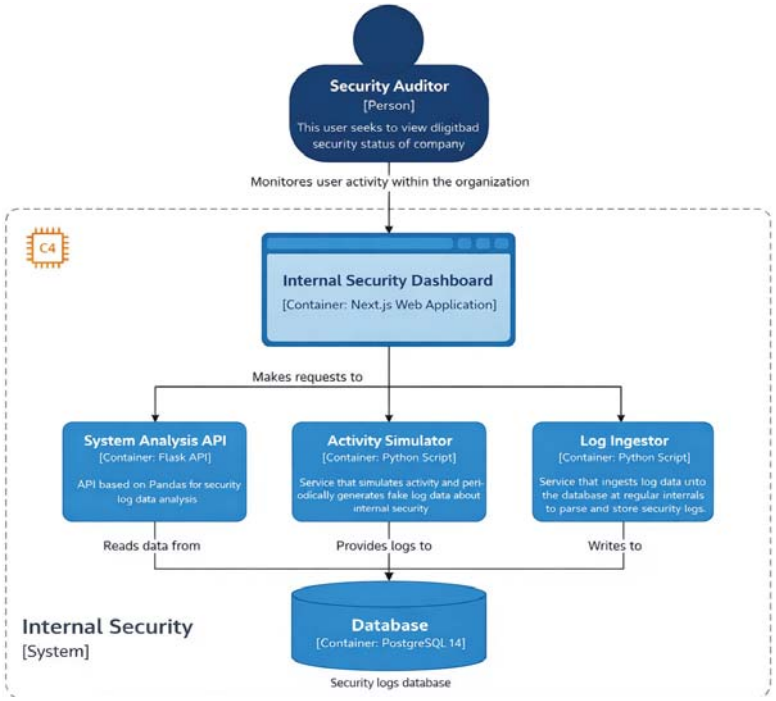

Fig. 3. C4 Container Diagram of the Proposed Framework.

- **Configuration:** Defines simulation parameters (e.g., users, schedules, event distributions) through a JSON file.
- **Event Generation:** Produces events dynamically based on configurable probabilities, determining whether they are legitimate or malicious.
- **Unauthorized Access Simulation:** Generates realistic unauthorized access attempts to sensitive resources.
- **Event Distribution:** Controls the frequency and temporal distribution of events.
- **Event Logging:** Stores generated events in structured logs for subsequent analysis.

Log Ingestion Component This component structures and centralizes data within a database, enabling its use for subsequent analysis and visualization. Although it is designed to be flexible and adaptable to different log formats, its initial configuration is tailored to the logs generated by the simulation component, as the database schema was designed to align with these records. This ensures seamless integration within the system pipeline and facilitates the analysis of data relevant to insider threats.

The main activities of this module are as follows:

- **Log Reading:** The module interprets and extracts relevant fields from each type of log generated by the simulation component.
- **Data Parsing:** Schema-defined patterns ensure that each record is correctly processed, extracting only the data required for storage.
- **Validation and Storage:** Each event is validated to prevent duplicates through uniqueness constraints defined in the database model.
- **Process Automation:** The module is implemented to perform ingestion periodically and automatically, ensuring that new logs are processed without manual intervention.

Event Analysis Component An analytical API was developed to process system-generated events and compute an organizational risk score. This component analyzes collected logs to identify suspicious behavioral patterns and provide interpretable insights to support decision-making. The approach combines classification, weighting, and normalization techniques, following practices derived from the CERT framework.

The main steps of the algorithm are:

Data Classification: Events stored in the database (e.g., timestamp, user, action, location) are categorized into risk classes such as unauthorized access, after-hours activity, external transfers, login failures, and external email communications.

Risk Scoring: Each event category is assigned a predefined weight according to its impact (e.g., unauthorized access $W_i = 5$, external transfers $W_i = 4$, after-hours activity $W_i = 3$, login failures $W_i = 2$). The raw risk score is computed as:

$$R_{\text{raw}} = \sum_{i=1}^{n} W_i \cdot f(E_i)$$

where $f(E_i)$ is the frequency of event type i.

Normalization: The score is normalized to a range of 0–100:

$$R_{\text{normalized}} = \min\left(100, \frac{R_{\text{raw}}}{N_{\text{total}}} \cdot 100\right)$$

where N_{total} is the total number of events in the analyzed period.

Risk Levels: The normalized score is mapped to qualitative levels (e.g., low, medium, high, critical), supporting prioritization of incidents.

Critical Pattern Identification: The analysis highlights key contributing factors, such as users with multiple suspicious events, abnormal temporal patterns, or repeated interactions with sensitive resources.

This approach provides a transparent and adaptable mechanism for risk assessment, enabling the extraction of actionable insights from existing data without requiring additional infrastructure. Its effectiveness depends on data quality and appropriate parameter configuration, while its modular design allows adaptation to different organizational contexts.

Visualization Dashboard A visualization dashboard was developed to support monitoring and analysis of internal security events. It integrates data generated by the simulation component and processed by the analytical API, enabling the identification of critical areas, behavioral patterns, and trends.

The dashboard follows a minimalist design and connects to the analytical API through REST-based HTTP requests, supporting both near real-time and historical data visualization.

The main functionalities include:

- Visualization of key metrics such as organizational risk score, event evolution, and critical areas.
- Interactive charts and tables for anomaly detection and trend analysis.
- Integration with the analytical API to ensure continuous data updates.
- Responsive interface optimized for desktop and mobile devices using `Tailwind CSS`.

Figures 4, 5, 6, and 7 illustrate representative views of the dashboard.

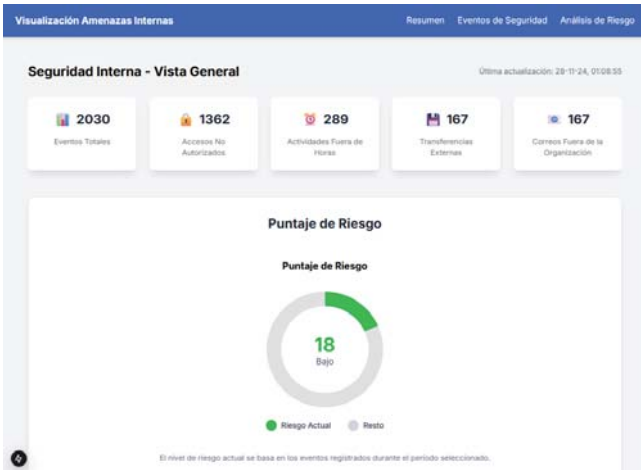

Fig. 4. Overview of the dashboard.

4 Performance Evaluation

All system components were integrated and evaluated to validate the end-to-end pipeline, from event simulation to visualization. Tests were conducted under controlled conditions to assess functionality and performance under varying workloads.

The evaluation included the generation of realistic scenarios with different user roles, event types (e.g., login attempts, file access, external transfers, and

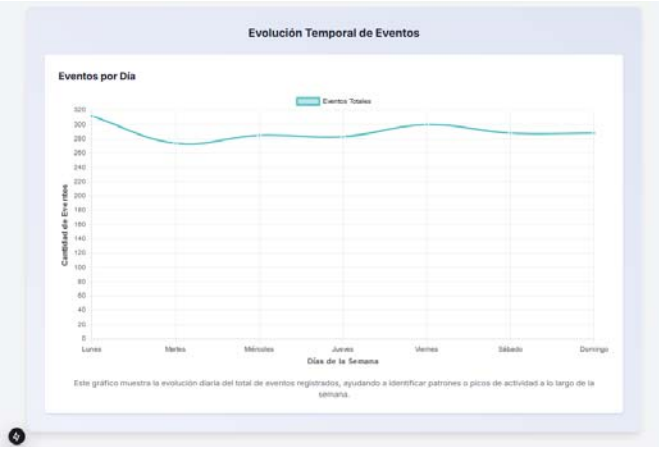

Fig. 5. Temporal evolution of security events.

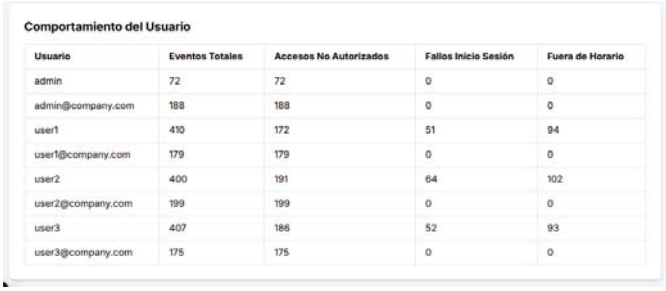

Comportamiento del Usuario

Usuario	Eventos Totales	Accesos No Autorizados	Fallos Inicio Sesión	Fuera de Horario
admin	72	72	0	0
admin@company.com	188	188	0	0
user1	410	172	51	94
user1@company.com	179	179	0	0
user2	400	191	64	102
user2@company.com	199	199	0	0
user3	407	186	52	93
user3@company.com	175	175	0	0

Fig. 6. Risk analysis and user behavior overview.

Fig. 7. Stacked bar chart of user activity.

emails), and temporal configurations. The system successfully reproduced these scenarios, enabling controlled validation of its behavior.

Integration tests confirmed correct communication between modules, while end-to-end execution verified the complete data flow across the system. Performance was evaluated through load testing under three scenarios: 10,000, 50,000, and 100,000 events over a 24-hour period.

The results show that the system accurately identifies anomalous behaviors and provides consistent visual reporting. In terms of performance, the system processed up to 50,000 events per day while maintaining response times below 2 seconds for both the API and the dashboard. Under higher load (100,000 events per day), response times increased to approximately 3 seconds, remaining within acceptable operational thresholds.

5 Conclusions

This work presents a modular prototype for insider threat management that integrates event simulation, analytical processing, and visualization to support the identification of anomalous behaviors and more proactive risk management. The results demonstrate the technical feasibility of the proposed approach and its effectiveness in modeling controlled scenarios using synthetic data, enabling the validation of its core components.

From a practical perspective, the system facilitates informed decision-making through clear visualization of risk metrics and trends, and shows potential for integration with existing security solutions such as SIEM platforms. Its modular design supports adaptability to different organizational contexts and evolving cybersecurity requirements.

However, the system has not yet been validated in real-world environments and relies on simulated data, which limits its ability to fully represent operational complexity. Addressing these limitations constitutes a key direction for future work.

Future developments include validation using real organizational data, incorporation of real-time monitoring capabilities, and the integration of machine learning techniques to enable predictive analysis of insider threats. Additionally, improving interoperability with existing security infrastructures will be essential to support scalability and adoption in practical settings.

Overall, the proposed solution establishes a foundation for the development of more advanced, proactive, and operationally integrated insider threat management systems.

Acknowledgments

Jorge Hochstetter-Diez is supported by Grant ANID, Chile, FONDECYT DE INICIACÓN EN INVESTIGACIÓN, Project Nº 11240115.

References

1. Aslan, , Aktuğ, S.S., Ozkan-Okay, M., Yilmaz, A.A., Akin, E.: A comprehensive review of cyber security vulnerabilities, threats, attacks, and solutions. Electronics **12**(6), 1333 (Mar 2023). https://doi.org/10.3390/electronics12061333, http://dx.doi.org/10.3390/electronics12061333
2. Ban, T., Takahashi, T., Ndichu, S., Inoue, D.: Breaking alert fatigue: Ai-assisted siem framework for effective incident response. Applied Sciences **13**(11), 6610 (May 2023). https://doi.org/10.3390/app13116610, http://dx.doi.org/10.3390/app13116610
3. Cappelli, D.M., Moore, A.P., Trzeciak, R.F.: The CERT Guide to Insider Threats: How to Prevent, Detect, and Respond to Information Technology Crimes. Addison-Wesley (2012)
4. Esparza, E., Diéguez, M., Hochstetter, J.: Prediction, Detection and Control of Insider Threats in Organizations, p. 378–395. Springer Nature Switzerland (2025). https://doi.org/10.1007/978-3-031-84078-4_26, http://dx.doi.org/10.1007/978-3-031-84078-4_26
5. Georgiadou, A., Mouzakitis, S., Askounis, D.: Detecting insider threat via a cybersecurity culture framework. Journal of Computer Information Systems **62**(4), 706–716 (2022)
6. González-Granadillo, G., González-Zarzosa, S., Diaz, R.: Security information and event management (siem): Analysis, trends, and usage in critical infrastructures. Sensors **21**(14), 4759 (2021)
7. Gurucul: 2024 insider threat report. White paper, Grucul - Cybersecurity Insiders (2024), https://gurucul.com/2024-insider-threat-report/, accessed 2026-03-15
8. IBM: What is user and entity behavior analytics (ueba)? (2022), https://www.ibm.com/think/topics/ueba, accessed: 2026-03-15
9. IBM: What is security information and event management (siem)? (2023), https://www.ibm.com/think/topics/siem, accessed: 2026-03-15
10. IBM: What is soar (security orchestration, automation and response)? (2023), https://www.ibm.com/think/topics/security-orchestration-automation-response, accessed: 2026-03-15
11. Inayat, U., Farzan, M., Mahmood, S., Zia, M.F., Hussain, S., Pallonetto, F.: Insider threat mitigation: Systematic literature review. Ain Shams Engineering Journal **15**(12), 103068 (Dec 2024). https://doi.org/10.1016/j.asej.2024.103068, http://dx.doi.org/10.1016/j.asej.2024.103068
12. Institute, P.: 2023 cost of insider risks global report. Industry report, Ponemon Institute (2023), https://www2.dtexsystems.com/2023ponemonreport, accessed 2026-03-15
13. Khan, M.Z.A., Khan, M.M., Arshad, J.: Anomaly detection and enterprise security using user and entity behavior analytics (ueba). In: 2022 3rd International Conference on Innovations in Computer Science & Software Engineering (ICONICS). pp. 1–9. IEEE (2022)
14. Lindauer, B.: Insider threat test dataset (2020). https://doi.org/10.1184/R1/12841247.v1, https://kilthub.cmu.edu/articles/dataset/Insider_Threat_Test_Dataset/12841247, accessed: 2026-03-15
15. López Velásquez, J.M., Martínez Monterrubio, S.M., Sánchez Crespo, L.E., Garcia Rosado, D.: Systematic review of siem technology: Siem-sc birth. International Journal of Information Security **22**(3), 691–711 (2023)

16. Meera, N.: Securing the expanding attack surface: Challenges and strategies for enterprise cybersecurity in the post-snowden era. International Journal of Trend in Scientific Research and Development **3**(4), 1954–1964 (2019)
17. Nassir, N.F.M., Rauf, U.F.A., Zainol, Z., Ghani, K.A.: Revealing the multi-perspective factors behind insider threats in cybersecurity. Journal of Media and Information Warfare **17**, 65–82 (2024)
18. Olajide O. Ogunbodede, Olumide S. Adewale, B.K.A.O.K.A.: Insider threat detection techniques: Review of user behavior analytics approach. International Journal of Research in Engineering and Science (IJRES) **12**(9), 109–117 (September 2024), https://www.ijres.org
19. Olaniyan, R., Rakshit, S., Vajjhala, N.R.: Application of User and Entity Behavioral Analytics (UEBA) in the Detection of Cyber Threats and Vulnerabilities Management, p. 419–426. Springer Nature Singapore (2023). https://doi.org/10.1007/978-981-19-8493-8_32, http://dx.doi.org/10.1007/978-981-19-8493-8_32
20. Ragusa, K.H., Chirayath, S.S.: Insider threat case studies from industry. In: Human Reliability Programs in Industries of National Importance for Safety and Security, pp. 61–72. Springer (2023)
21. Roessler, K.: $16.2m: The high cost of insider "risks" (2023), https://www.dtexsystems.com/blog/cost-of-insider-risks/, accessed 2026-03-15
22. Schoenherr, J.R.: Insider threats and individual differences: Intention and unintentional motivations. IEEE Transactions on Technology and Society **3**(3), 175–184 (Sep 2022). https://doi.org/10.1109/tts.2022.3192767, http://dx.doi.org/10.1109/tts.2022.3192767
23. Sheeraz, M., Paracha, M.A., Haque, M.U., Durad, M.H., Mohsin, S.M., Band, S.S., Mosavi, A.: Effective security monitoring using efficient siem architecture. Hum.-Centric Comput. Inf. Sci **13**, 1–18 (2023)
24. Six, S.: Year in review – insider risk in 2023. https://www.signpostsix.com/year-in-review-insider-risk-in-2023/ (2023), accessed: 2026-03-15
25. Subhani, A., Khan, I.A., Zubair, A.: Review of insider and insider threat detection in the organizations. Journal of Advanced Research in Social Sciences and Humanities **6**(4), 167–174 (Dec 2021). https://doi.org/10.26500/JARSSH-06-2021-0402, https://www.jarssh.com/ojs/index.php/jarssh/article/view/174
26. Vaideeswaran, N.: Insider threats explained. https://www.crowdstrike.com/en-us/cybersecurity-101/identity-protection/insider-threat/ (Jan 2025), https://www.crowdstrike.com/en-us/cybersecurity-101/identity-protection/insider-threat/, accessed: 2026-03-15
27. Whitelaw, F.: Understanding the shifting perceptions of insider threats over external cyber attacks (2024), https://www.securonix.com/blog/shifting-perceptions-of-insider-threats-vs-external-cyber-attacks/, accessed 2026-03-15
28. Yousef, R., Jazzar, M.: Measuring the effectiveness of user and entity behavior analytics for the prevention of insider threats. J. Xi'an Univ. Arch. Technol **13**, 175–181 (2021)

A Scalable Architecture for Real-Time Public Health Vulnerability Assessment

Perdomo Luciano[1,2], Ordinez Leo[1,2], and Sergio Andrés Kaminker[3]

[1] Laboratorio de Investigación en Informática (LINVI), Facultad de Ingeniería, Universidad Nacional de la Patagonia San Juan Bosco, Puerto Madryn, Argentina
[2] Instituto Patagónico de Ciencias Sociales y Humanas (IPCSH) CENPAT - CONICET, Argentina
lucianor.perdomo@gmail.com

Abstract. Public health surveillance in Argentina is constrained by both vertical and horizontal fragmentation of health information systems. Existing platforms, such as the National Health Surveillance System (SNVS) and the Integrated Health Record (Historia de Salud Integrada, HSI), operate as loosely coupled systems with limited interoperability, often relying on manual data exchange. This paper proposes a cloud-native architecture for real-time public health vulnerability assessment that integrates heterogeneous data sources, including clinical records, epidemiological data, and geospatial social determinants of health. The main contribution is a stream-centric modular monolith architecture that extends Hexagonal Architecture with an internal Kappa-style streaming layer, enabling continuous processing of domain events and computation of contextual indicators. These indicators are reintegrated into the domain model, establishing a controlled feedback loop between analytical processing and vulnerability assessment. The architecture combines Command Query Responsibility Segregation (CQRS) with the Transactional Outbox pattern to ensure consistent and reliable propagation of state changes throughout the system while preserving a unified deployment model. The proposed approach enables near-real-time vulnerability assessment at both individual and population levels, supporting context-aware decision-making in fragmented public health environments. The design is evaluated through a scenario-based architectural validation grounded in representative public health workflows.

Keywords: Software Architecture · Public Health Informatics · Stream Processing

1 Introduction

Public health surveillance systems in Argentina are characterized by significant fragmentation across institutional and technological boundaries. Systems such as the National Surveillance System (SNVS) [14] and the Integrated Health Record (HSI) [15] operate in a loosely coupled manner, resulting in limited interoperability and delayed access to integrated health data. In practice, data exchange is often performed through manual or batch-based processes, preventing timely population-level risk analysis.

This fragmentation severely limits the ability to perform real-time vulnerability assessment, as clinical and epidemiological data remain distributed across disconnected systems. Addressing this challenge requires not only data integration, but also an architectural approach capable of supporting continuous data ingestion, real-time processing, and controlled interaction between analytical and domain components.

Effective public-health vulnerability assessment therefore requires an architectural foundation capable of supporting continuous ingestion of heterogeneous data, near-real-time analytical processing, and incorporation of contextual information derived from external sources while preserving the autonomy of participating systems.

In this paper, we propose a cloud-native architectural model for assessing public health vulnerability in near-real-time that integrates heterogeneous data sources, including clinical records and geospatial social determinants of health. The architecture is designed to support continuous and context-aware risk evaluation at both individual and population levels. It also provides mechanisms for multi-source data consolidation, uncertainty-aware assessment, and policy traceability.

The proposed solution is based on a modular monolith structured using on Hexagonal Architecture (Ports and Adapters) [3] and Command Query Responsibility Segregation (CQRS) pattern [13]. A key design contribution is the integration of a Kappa-style streaming layer [11, 9] as a first-class architectural component. Unlike conventional approaches where streaming is treated as an external processing pipeline, the proposed design integrates the event stream as a first-class architectural component within the system boundary. This enables the generation of derived contextual indicators that may be selectively reintegrated into the domain model through a policy-driven mechanism under explicit constraints and contextual relevance criteria.

In addition, the architecture incorporates geospatial social determinants of health (SDOH) [18] as core inputs for vulnerability assessment. Using spatial processing capabilities (e.g., PostGIS [16]), the system associates individual clinical events with census-derived socio-environmental indicators [8], enabling contextual vulnerability scoring that reflect both clinical and environmental risk factors.

The main contributions of this paper are:

- A stream-centric adaptation of Hexagonal Architecture, where the event stream acts as both a communication backbone and a mechanism for generating derived contextual inputs.
- A unified architectural approach combining CQRS, the Transactional Outbox pattern, and Kappa-style streaming within a modular monolith.
- A vulnerability assessment model integrating policy-driven multi-source data consolidation, uncertainty-aware evaluation, and selective reintegration of domain-relevant analytical outputs.
- An architectural instantiation analyzed through representative public-health scenarios in the Province of Chubut.

By avoiding centralized data lakes and preserving data ownership within source systems, the approach reduces integration risks while enabling scalable, real-time public health intelligence for vulnerability assessment.

This work presents a reference architectural design that validated through scenario-based architectural analysis. The evaluation does not rely on empirical deployment or quantitative benchmarking, but on structured reasoning over representative operational scenarios derived from public health workflows.

The remainder of this paper is organized as follows: Section II reviews related work and identifies the research gap; Section III presents the proposed architecture; Section IV describes the prototype instantiation and scenario-based analysis; and Section V discusses limitations and future work.

2 Related Work

2.1 Data Fragmentation in Public Health Systems

Fragmentation of public health information systems is widely recognized as a major barrier to effective surveillance and decision-making [1, 2]. National-scale platforms, including electronic health

records and epidemiological reporting systems, are typically deployed as independent silos, limiting interoperability across institutional and administrative boundaries.

Existing integration approaches typically rely on batch-oriented ETL pipelines or manual reporting workflows [7], introducing significant latency and hindering timely analysis of population-level risks. Recent health informatics frameworks emphasize continuous, real-time data processing and interoperability through principles such as FAIR data management, enabling more scalable and explainable analytics pipelines [17]. These limitations are particularly critical in heterogeneous healthcare ecosystems, where clinical, administrative, and socio-environmental data are inherently distributed.

2.2 Event-Driven Architectures in Healthcare Systems

Event-driven architectures have been proposed as an alternative to batch-oriented integration in healthcare systems [10]. Streaming platforms such as Apache Kafka enable asynchronous event propagation and scalable integration across distributed infrastructures. Recent healthcare implementations using Apache Kafka demonstrate the effectiveness of event-driven architectures for integrating heterogeneous health data sources and improving system responsiveness [9].

However, most existing approaches treat streaming as an infrastructure-level concern. The stream is typically used for data transport and decoupling, rather than as an integral part of domain-driven processing. Comparative studies between event-driven and REST-based architectures indicate that while event-driven systems improve modularity and decoupling, they also introduce additional complexity in system design and reasoning [12]. As a result, the integration of heterogeneous health data—particularly when combining clinical records with geospatial social determinants—remains limited.

2.3 CQRS and Modular System Design in Complex Domains

CQRS has been widely adopted in domains requiring separation between transactional and analytical workloads, such as financial systems and enterprise platforms [13]. These approaches are grounded in Domain-Driven Design principles, which emphasize explicit modeling of complex domains and separation of concerns between domain logic and infrastructure [4]. When combined with modular monoliths or Hexagonal Architectures, CQRS improves maintainability, scalability, and system evolvability.

In healthcare systems, however, CQRS is often applied in isolated components rather than as part of a unified architecture supporting end-to-end data consolidation and analysis. Consequently, existing solutions frequently lack consistent traceability across the full data lifecycle.

2.4 Research Gap and Positioning

Despite significant advances in healthcare interoperability, event-driven integration, and domain-driven architectural design, existing approaches tend to address these concerns in isolation. Data integration frameworks focus on interoperability without enabling continuous analytical feedback, while event-driven healthcare platforms typically treat streaming infrastructure (e.g., Kafka) as external integration layers rather than first-class architectural components. Similarly, CQRS-based solutions focus on transactional scalability without incorporating contextual socio-environmental information into the domain model.

As a result, three key limitations remain insufficiently addressed in the literature:

– The absence of architectures treating the event stream as a first-class architectural component rather than an external processing pipeline.
– Limited integration of geospatial social determinants of health as first-class inputs in domain-driven models.
– Lack of feedback mechanisms that reintegrate derived analytical results into the domain model for continuous refinement and adaptation.

This work addresses these limitations by proposing a stream-centric architectural model treating the event stream as a first-class architectural component. The proposed approach enables selective, policy-driven reintegration of contextual analytical outputs into the domain model under explicit constraints.

Table 1. Comparison with representative approaches from the literature

Approach	EDA	CQRS	SDOH	Feedback	Domain Stream
EDA Approach [10]	Yes	No	No	No	No
Kafka Healthcare [9]	Yes	No	No	No	No
CQRS Approach [13]	No	Yes	No	No	No
Proposed Architecture	Yes	Yes	Yes	Yes	Yes

Table 1 highlights that existing approaches address individual concerns such as data integration, computational scalability, or asynchronous communication patterns in isolation. No existing work was identified that explicitly provides a unified architectural model integrating these capabilities within a consistent event-driven architecture while supporting selective reintegration of analytical outputs under explicit constraints.

3 Proposed Architecture

The proposed architecture defines a stream-centric, domain-driven system for real-time public health vulnerability assessment. It is implemented as a modular monolith structured using Hexagonal Architecture (Ports and Adapters) [3], ensuring strict separation between domain logic and external infrastructure while preserving deployment simplicity and system consistency.

The architecture adopts CQRS within the Hexagonal Architecture, separating a transactional write model from read-side projections to ensure consistent state transitions and scalable query processing. This separation enables independent evolution of transactional and analytical concerns while preserving a unified system boundary. The Transactional Outbox Pattern ensures reliable propagation of state changes from the write model to the event stream, enabling downstream processing and projection updates.

A central outcome of the design is the continuous refinement of vulnerability assessments using both raw clinical events and derived contextual information produced through stream-based processing. Rather than treating analytical results as external outputs, the system incorporates them as structured inputs for subsequent evaluations under explicit domain policies, enabling a controlled feedback loop beyond static scoring models. The overall architecture supports three fundamental capabilities: (i) integration of heterogeneous data sources with full provenance, (ii) continuous and context-aware vulnerability assessment at individual and population levels, and (iii) generation of auditable analytical projections for decision support in public health systems.

By preserving data ownership within source systems and avoiding centralized data lakes, the proposed architecture enables scalable interoperability across institutional boundaries while maintaining traceability and consistency of all domain-level decisions. The design addresses key limitations in the literature by (i) treating the event stream as a first-class architectural component, (ii) integrating geospatial social determinants of health as structured domain inputs, and (iii) enabling the selective reintegration of analytical outputs into subsequent domain evaluations.

3.1 System Overview

The system operates as an interoperability and real-time analytics layer over existing public health infrastructures. It supports (i) data consolidation with full traceability, (ii) continuous vulnerability assessment, and (iii) generation of auditable analytical outputs.

3.2 Architectural Design

Figure 1 illustrates the interaction between the domain core, streaming module, and external systems. The architecture follows the Hexagonal Architecture pattern, organizing the system into three main layers: (i) the domain core, (ii) application layer, and (iii) adapter layer. The **domain core** encapsulates business logic for data consolidation and vulnerability assessment, remaining independent from infrastructure concerns and operates exclusively on domain entities and value objects. The **application layer** orchestrates use cases, coordinating interactions between external inputs, domain services, and output generation. It exposes inbound ports that define the system's capabilities. The **adapter layer** implements both driving and driven adapters. Driving adapters include REST APIs, GraphQL interfaces, batch processes, and schedulers that trigger domain use cases. Driven adapters integrate with external systems such as PostgreSQL (write model), Kafka (event streaming), Elasticsearch and PostGIS (read models), Redis (caching), and notification services.

A central component of the architecture is the embedded streaming module, which is connected to the domain core through well-defined ports. This module consumes domain events produced via the Transactional Outbox pattern and performs asynchronous processing tasks, including data enrichment, projection updates, and contextual indicator computation. The results of this processing are not treated as external artifacts but may be reintroduced into the domain core as structured inputs for subsequent evaluations under explicit domain policies. This preserves bounded context isolation while enabling continuous analytical refinement. Domain state changes are persisted through the transactional write model and propagated to the streaming module via the transactional outbox pattern, enabling the generation of analytical projections and derived inputs for subsequent domain processing.

3.3 Stream-Centric Processing Model

The architecture integrates a Kappa-style streaming layer as an internal processing component connected to the domain core through well-defined ports and bounded contexts. The event stream is used for event propagation, contextual information derivation, and generation of analytical projections governed by domain-defined policies and event-triggered conditions. These projections are materialized and can be selectively incorporated into subsequent domain evaluations as structured inputs for vulnerability assessment. This design enables tight integration between analytical processing and domain-level decision logic while preserving bounded context separation and ensuring that the streaming layer remains within the system boundary rather than an external pipeline.

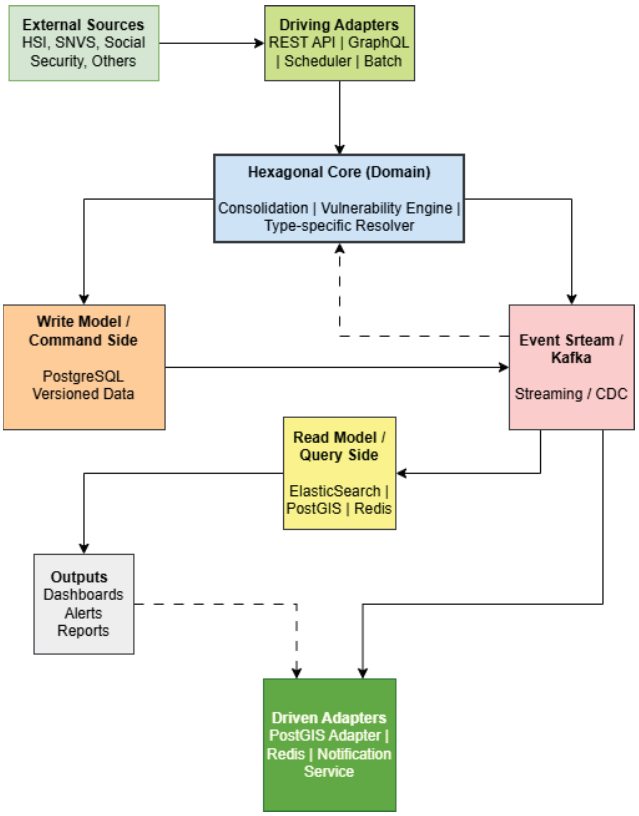

Fig. 1. The high-level architecture of the proposed system

3.4 Domain Model

The domain core is structured into two subdomains: (i) multi-source data consolidation and (ii) vulnerability assessment. The consolidation subdomain integrates heterogeneous inputs using resolution policies and type-specific resolvers, producing a consolidated representation of each entity. An uncertainty model quantifies the reliability of aggregated data based on completeness, source reliability, temporal decay, and inter-source conflict. When uncertainty exceeds predefined thresholds, escalation policies trigger human-in-the-loop conflict resolution to ensure semantic and clinical validity of the consolidated result. The assessment subdomain processes consolidated data through a Vulnerability Engine, integrating clinical, geospatial, and social determinants to compute a vulnerability score. A Risk Projector derives population-level insights from individual assessments. Core domain entities include *Person*, *Clinical Event*, *Vulnerability Assessment*, *Health Zone*, and *Multi-source Data*, ensuring traceability and reproducibility across evaluations.

3.5 Data Governance and Policies

The system defines explicit, versioned policies to manage heterogeneous and potentially conflicting data. Resolution, escalation, and versioning policies govern data consolidation, human-in-the-loop intervention, and reproducibility of assessments. This policy-driven approach enables transparent and auditable decision-making in environments with unreliable and fragmented data sources.

3.6 Traceability and Uncertainty Modeling

A fundamental requirement in public health systems is the ability to ensure full traceability and to quantify the reliability of analytical results. The proposed architecture addresses both concerns through explicit modeling of data provenance and uncertainty.

Traceability is achieved through the use of immutable *Source Value* objects, which associate each data element with its origin, timestamp, and metadata describing its acquisition context. This design enables complete auditability of the data lifecycle, allowing each vulnerability assessment to be traced back to the exact set of source observations and policies used in its computation. All domain events are persisted and propagated through the event stream, ensuring that both raw inputs and derived transformations remain observable and reproducible. This event-driven traceability supports retrospective analysis and recomputation through event replay under updated policies. Each domain event includes a persistent identifier linking it to its originating source value object, enabling deterministic replay of the system state under modified policies.

Uncertainty is modeled as a first-class domain concept. Each vulnerability assessment includes an uncertainty level computed as a composite metric based on four dimensions: (i) data completeness, (ii) source reliability, (iii) temporal decay, and (iv) inter-source conflict. These are aggregated into a single reliability metric. When uncertainty exceeds predefined thresholds, escalation policies trigger human-in-the-loop validation.

This design ensures that all analytical outputs are both explainable and auditable, which is critical in high-stakes public health decision contexts.

4 Methods

The proposed architecture is specified and analyzed through representative operational scenarios that illustrate its behavior under realistic public health conditions. A prototype implementation is currently under development to support future empirical validation in the Province of Chubut. The system is designed as a cloud-native, event-driven platform supporting near-real-time processing of heterogeneous health data through asynchronous communication and polyglot persistence.

4.1 Event-Driven Processing and Consistency

The system adopts the Transactional Outbox pattern to ensure reliable propagation of domain events. Each state transition within the Write Model is executed atomically, persisting both the domain state and its corresponding event within a single transaction. This approach aligns with event sourcing principles, where state changes are represented as a sequence of immutable events, enabling auditability and replay-based analysis [5].

A Change Data Capture (CDC) pipeline publishes events to Apache Kafka, which acts as the central event backbone. This provides at-least-once delivery and synchronization with downstream

read models. Derived contextual indicators are computed within the streaming layer and may be reintroduced into the system through controlled event-driven ingestion when relevant to subsequent assessments. This preserves the integrity of the CQRS model by ensuring that all domain decisions rely on validated and traceable data flows.

4.2 CQRS and Persistence Strategy

The system implements CQRS to separate transactional and analytical concerns. The Write Model, based on PostgreSQL, maintains normalized domain data and ensures strong consistency through transactional operations and event persistence. Each state change produces a versioned record. The Read Model is constructed asynchronously from the event stream using a polyglot persistence approach. Elasticsearch supports high-performance querying, PostGIS enables geospatial analysis, and Redis provides low-latency access to frequently requested data. This enables independent scaling of analytical workloads while preserving traceability.

4.3 Domain Subdivision

The domain layer is organized into two subdomainsaligned with the vulnerability assessment pipeline. The Multi-source Consolidation Subdomain constitutes the first step of the process. It addresses heterogeneous and potentially conflicting data sources through resolution policies, type-specific conflict resolution, and uncertainty estimation. A Multi-source Data structure represents an attribute with multiple source-specific values prior to consolidation. A Resolution Policy defines the strategy used to reconcile these conflicting values. Each applied policy is versioned, enabling auditability of automated decisions with respect to the clinical guidelines in effect at the time of processing. Type-specific Resolvers execute the corresponding Resolution Policies per data category. An Uncertainty Evaluator quantifies the reliability of consolidated outputs before propagation. After consolidation, the Vulnerability Assessment Subdomain is executed. This subdomain encapsulates the Vulnerability Engine, which integrates social, environmental, and clinical determinants to compute a vulnerability score. The resulting Vulnerability Assessment represents the versioned outcome of the scoring process, including score, uncertainty level, policy version, and evaluation timestamp. In addition, a Risk Projector derives higher-level population and territorial insights from individual assessments, supporting population-level public health analysis. This separation improves modularity, and enables independent evolution of resolution logic, uncertainty modeling, and vulnerability computation without introducing cross-domain coupling.

4.4 Operational Scenarios

To illustrate the behavior of the proposed architecture, we describe representative scenarios covering data ingestion, conflict resolution, human validation, and policy evolution.

Scenario A: External Data Integration (Social Security Input) When a Social Security System registers an event via API, the Driving Adapter invokes the corresponding Use Case. The Core validates the input and generates a *Domain Event*, which is persisted and published to the event log. The Streaming Module consumes this event to perform downstream processing tasks, including geospatial enrichment, aggregation of territorial metrics, and read model updates (Elasticsearch/PostGIS). When derived indicators are considered relevant for subsequent assessments, they

may be reintroduced into the Core through controlled event-driven ingestion. If the new data causes a vulnerability threshold to be exceeded, an *Alert Event* is emitted to the relevant stakeholders.

Scenario B: Multi-source Conflict Resolution When a Multi-source Data entity is updated, the Type-specific Resolver applies the active Resolution Policy to generate a Consolidated Data object. The Uncertainty Evaluator assesses the reliability of the consolidated result. If the resulting uncertainty exceeds predefined thresholds, a *ValidationRequired Event* is emitted. The Consolidated Data is then processed by the Vulnerability Engine, which recalculates the individual vulnerability score and produces an new Vulnerability Assessment. The resulting assessment is versioned and stored. Finally, the Streaming Module asynchronously updates geographic and executive projections to reflect the new risk profile.

Scenario C: Human-in-the-Loop Validation When a *ValidationRequired Event* is triggered, the system enters a supervised state. An epidemiologist interacts through the Surveillance Port to review the discordant source values (e.g., a critical diagnosis reported by the Social Security System but absent from the physician's record). After manual resolution, the specialist submits a consolidated value, generating a *ManualOverride Event*. This event allows the Core to override resolution policies, recalculate the individual's Vulnerability Assessment, and update the audit trail to reflect human intervention as the authoritative source.

Scenario D: Policy Evolution and Historical Recalculation When a Resolution Policy is updated (e.g., increasing the Score Impact of a specific respiratory determinant during a public health emergency), the Core persists a new policy version. Following the Streaming/Kappa architecture, the Streaming Module supports replay of historical events using the updated policy. This produces revised Risk Projections and Territorial Metrics without data re-ingestion, enabling retrospective analysis of population risk under the new criteria.

4.5 Deployment and Reliability

The architecture supports incremental adoption and containerized deployment. Reliability is supported through immutable event logs, idempotent consumers, and event replay capabilities, enabling recovery from partial failures. By preserving data ownership within source systems and enforcing controlled ingestion and governance policies, the system aligns with regulatory constraints typical of public health environments. Overall, these mechanisms illustrate how the proposed architecture enables consistent, traceable, and scalable processing of heterogeneous health data, enabling near-real-time vulnerability assessment in heterogeneous environments.

5 Evaluation

The proposed architecture is evaluated through scenario-based validation grounded in qualitative reasoning over public health workflows and architectural properties. Given the prototype stage, the evaluation focuses on structural correctness, consistency, and suitability for near-real-time heterogeneous data processing rather than quantitative benchmarks. Policies are treated as first-class versioned domain entities.

5.1 Scenario-Based Validation

The operational scenarios in Section 4.4 are used as architectural evaluation probes representing realistic workflows, including data ingestion, conflict resolution, human-in-the-loop validation, and

policy evolution. These scenarios illustrate the expected architectural behavior in terms of traceability, consistency, and responsiveness across distributed and heterogeneous data sources.

5.2 Architectural Properties Analysis

The architecture is analyzed in terms of the following properties:

Scalability: The separation between write and read models through CQRS enables independent scaling of analytical workloads without affecting transactional integrity. Read-side projections can be replicated and optimized according to query demand while maintaining a single authoritative write model.

Consistency: The use of the Transactional Outbox pattern ensures reliable event propagation, ensuring consistency between domain state and downstream projections. This design minimizes the risk of synchronization failures between transactional operations and event publication.

Traceability: Immutable event logs and source-level provenance tracking enable full reconstruction of system state and decision history. This supports auditing, policy review, and retrospective analysis of historical assessments.

Extensibility: The hexagonal architecture allows new data sources and analytical modules to be integrated without modifying the domain core. As a result, the platform can evolve incrementally as new public health systems become available.

Resilience: Event replay and idempotent consumers provide fault tolerance and recovery capabilities in case of partial system failures. Lost projections can be rebuilt from historical events without compromising domain consistency.

5.3 Comparative Discussion

In contrast to batch-oriented public health systems relying on ETL pipelines and centralized data lakes, the proposed architecture is designed to support continuous real-time vulnerability assessment. Compared to microservice-based event-driven systems, the proposed modular monolith is designed to reduce operational complexity while preserving separation of concerns and supporting scalability through well-defined internal architectural boundaries. Table 2 summarizes these differences across key architectural attributes.

Attribute	**Traditional Monolith**	**Pure Microservices**	**Proposed Architecture**
Complexity	Low	Very High	Medium
Data Consistency	Strong (ACID)	Eventual (Sagas/2PC)	Hybrid (Transactional Outbox)
Scalability	Vertical Only	Granular/Independent	Modular (via CQRS/Kafka)
Interoperability	Limited (Silos)	High (APIs)	High (Enriched Adapters)
Auditability	Difficult (State-based)	Complex (Distributed)	High (Immutable Event Log)

Table 2. Comparison of Architectural Approaches

Unlike conventional streaming systems that treat analytics as external pipelines, the proposed approach embeds the streaming layer within the system boundary, enabling the selective reintegration of domain-relevant analytical outputs into subsequent domain evaluations.

5.4 Limitations

The current design has limitations. First, the system has not been validated through large-scale deployment or benchmarking. Second, policy configuration may require domain expert involvement for tuning. Finally, integration depends on the availability and quality of external systems, which may vary across regions.

Despite these limitations, the proposed architecture provides a foundation for future development and evaluation of real-time, traceable, and context-aware public health analytics in fragmented environments.

6 Discussion

The proposed architecture is evaluated through a stream-centric, domain-driven design that addresses key limitations of fragmented public health information systems. By combining a modular monolith with Hexagonal Architecture, the system enforces separation between domain logic and infrastructure, while preserving traceability, reproducibility, and interpretability of all decisions.

The feedback mechanism between the streaming layer and the domain model enables selective reintegration of analytical outputs into domain-level decision processes.

The architecture incorporates an explicit uncertainty model intended to quantify the reliability of health data across multiple dimensions, including completeness, source reliability, temporal decay, and inter-source conflict. This design extends traditional scoring approaches by coupling each vulnerability assessment with a confidence measure, improving interpretability in high-stakes decision-making contexts.

From a systems perspective, the separation between transactional processing (write model) and analytical querying (read models) supports scalability without sacrificing consistency at the domain level. Event-driven propagation preserves traceability of state transitions, enabling auditability and retrospective analysis of system behavior.

Compared to batch-oriented surveillance systems relying on delayed ETL pipeline, the proposed architecture supports supports near-real-time interpretation of heterogeneous health data. In contrast to distributed microservice-based solutions, it emphasizes domain cohesion and architectural simplicity while preserving scalability through event-driven processing and internal modular decomposition.

Overall, the proposed design supports a shift from reactive surveillance systems toward proactive, context-aware public health intelligence platforms enabling continuous risk interpretation.

Conceptually, this approach shifts event streams from passive data transport toward active components supporting analytical processing and domain-level reasoning.

7 Limitations and Threats to Validity

While the proposed architecture provides a structured and scalable approach to real-time public health vulnerability assessment, several limitations and threats to validity remain.

7.1 Data Quality and Source Heterogeneity

The system depends on heterogeneous external data sources with varying completeness, accuracy, and temporal consistency. Although uncertainty is explicitly modeled, output quality depends on upstream data quality.

7.2 External System Dependence

The architecture operates as an interoperability layer rather than a replacement for existing public health systems. Consequently, its effectiveness depends on the availability, stability, and integration quality of external platforms such as electronic health record and epidemiological reporting systems. Failures or inconsistencies in upstream systems directly affect downstream analytical outputs.

7.3 Eventual Consistency in Analytical Views

Although the write model ensures strong consistency through transactional mechanisms, read models operate under eventual consistency due to asynchronous event propagation. This introduces temporary divergence between transactional state and analytical projections, which potentially affecting time-sensitive visualizations and dashboards.

7.4 Operational Complexity

The combination of event-driven processing, polyglot persistence, and geospatial analytics introduces significant operational complexity. While this design improves modularity and scalability, it increases deployment and maintenance complexity, particularly in resource-constrained environments.

7.5 Prototype Maturity

At the time of writing, the system is a prototype under active development. The architecture has been assessed through scenario-based analysis, while large-scale empirical evaluation and production deployment remain future work.

7.6 Generalizability

The proposed architecture is designed for fragmented public health ecosystems such as the Argentine healthcare context. While the underlying principles are generalizable, adaptation to other institutional or regulatory environments may require structural or policy-level adjustments.

8 Conclusion and Future Work

This paper presents a scalable stream-centric architecture for real-time public health vulnerability assessment designed to address the limitations of fragmented healthcare information systems. The proposed approach combines a modular monolith with Hexagonal Architecture and event-driven processing to enable consistent, traceable, and context-aware decision-making across heterogeneous data sources.

A key contribution of this work is the integration of multi-source clinical data with geospatial social determinants of health within a unified analytical framework. By incorporating uncertainty modeling, provenance tracking, and policy-driven data consolidation, the proposed architecture supports continuous vulnerability assessment while preserving the autonomy of existing health information systems.

The separation between a transactional write model and a polyglot analytical read model supports scalable analytical workloads without compromising the integrity of core clinical data. This design enables real-time interpretation of population health dynamics while maintaining auditability and consistency.

Overall, the proposed architecture outlines a feasible path from fragmented surveillance systems toward unified event-driven public health platforms capable of continuous, traceable, and uncertainty-aware vulnerability assessment.

8.1 Future Work

Interoperability can be further enhanced through the adoption of standardized healthcare data exchange protocols such as HL7 FHIR [6], enabling seamless integration across institutional and regional systems.

Predictive modeling remains a key direction for supporting individual and population-level risk forecasting. This includes longitudinal disease progression modeling and scenario-based evaluation of intervention strategies under varying clinical and socio-environmental conditions.

The platform can also be extended for real-time monitoring of healthcare system capacity, supporting early detection of system stress through hospital occupancy and critical resource tracking.

At the population level, advanced epidemiological forecasting techniques can further improve the evaluation of public health interventions over time.

Finally, the adoption of digital twin approaches represents a long-term direction, enabling simulation of individual trajectories and population dynamics under evolving policy and environmental conditions.

Acknowledgments. This work was supported by the Secretaría de Ciencia y Tecnología de Chubut and the National Council for Scientific and Technical Research (CONICET) through a doctoral scholarship.

Bibliography

[1] AbouZahr C, B.T.: Health information systems: The foundations of public health. Bulletin of the World Health Organization (2005)

[2] Casanova, R., V.G.F.A..B.B.J.W.: Architectural patterns for health information systems: A systematic review (2025). https://doi.org/10.3389/fdgth.2025.1694839, https://pmc.ncbi.nlm.nih.gov/articles/PMC12666560/, accessed: 2026-04-10

[3] Cockburn, A.: Ports and adapters pattern (hexagonal architecture) (2005), https://alistair.cockburn.us/hexagonal-architecture/

[4] Evans, E.: Domain-Driven Design: Tackling Complexity in the Heart of Software. Addison-Wesley (2003)

[5] Fowler, M.: Event sourcing (2005), https://martinfowler.com/eaaDev/EventSourcing.html

[6] HL7 International: Fhir: Fast healthcare interoperability resources. HL7 Standard (2019), https://www.hl7.org/fhir/

[7] Inmon, W.H.: Data warehousing and etl in healthcare. DM Review (2005)

[8] Instituto Nacional de Estadística y Censos (INDEC): Censo nacional de población, hogares y viviendas 2022: Resultados definitivos. https://www.indec.gob.ar/indec/web/Nivel4-Tema-2-41-165 (2023), accessed: 2026-04-10

[9] Katiyar, S.: Event driven architecture for a health sciences customer using kafka and java microservice. International Journal of Leading Research Publication (2023), https://www.ijlrp.com/research-paper.php?id=1910, accessed: 2026-04-10

[10] Kreps, J.: Kafka : a distributed messaging system for log processing (2011), https://api.semanticscholar.org/CorpusID:18534081

[11] Kreps, J.: Questioning the lambda architecture (2014), https://www.oreilly.com/radar/questioning-the-lambda-architecture/, accessed: 2026-04-10

[12] L. Lazzari, K.F.: Event-driven architecture and rest architectural style: An exploratory study on modularity. Journal of applied research and technology (2023), https://www.redalyc.org/journal/474/47480004002/, accessed: 2026-04-10

[13] Microsoft: Cqrs journey guide (2012), https://learn.microsoft.com/en-us/previous-versions/msp-n-p/jj554200(v=pandp.10)

[14] Ministerio de Salud de la Nación Argentina: Sistema nacional de vigilancia de la salud (snvs 2.0). https://sisa.msal.gov.ar/sisa/ (2020), accessed: 2026-04-10

[15] Ministerio de Salud de la Nación Argentina: Historia de salud integrada (hsi): Sistema de información clínica. https://www.argentina.gob.ar/salud/digital/hsi (2021), accessed: 2026-04-10

[16] PostGIS Project: Postgis: Spatial and geographic objects for postgresql (2023), https://postgis.net

[17] Siddiqi MH, Idris M, A.M.: Fair health informatics: A health informatics framework for verifiable and explainable data analysis. Healthcare (Basel, Switzerland) (2023). https://doi.org/10.3390/healthcare11121713, https://pmc.ncbi.nlm.nih.gov/articles/PMC10298118/, accessed: 2026-04-10

[18] World Health Organization: Social determinants of health (2021), https://www.who.int/health-topics/social-determinants-of-health

Building and Acceptance of Intelligent Immersive Assistants

Pablo Pintado [1[0000-0002-1356-0248]], Diego Amoroso [1[0009-0000-6944-9752]], Amanda Roldán [1[0009-0009-3548-2269]], Gabriela Duque [1[0000-0003-3783-1791]], Daniel Iturralde [1[0000-0003-2134-0453]], Esteban Mora [1[0000-0003-0313-3004]], Daniel Capelo [1[0009-0006-4739-5806]] and José Fajardo [1[0000-0002-5614-9643]]

[1] Universidad del Azuay, Cuenca, Ecuador
ppintado@uazuay.edu.ec, diegoandres.amoroso@es.uazuay.edu.ec, amy.3107@es.uazuay.edu.ec, gduque@uazuay.edu.ec, diturralde@uazuay.edu.ec, ejmora@uazuay.edu.ec, dcapelo@uazuay.edu.ec, jfajardo@uazuay.edu.ec

Abstract. Emerging technologies continue to evolve. As a result, institutions like Universidad del Azuay (UDA) are expanding support services beyond traditional omnichannel platforms. There is now a need to offer support through extended reality and voice-based environments, enabling immersive, natural user interactions. This study presents the building and evaluation of an Intelligent Immersive Assistant (IIA) that provides accurate information on multiple topics through immersive and voice interaction. The Design thinking methodology was used. A survey helped identify user needs and define the problem. An IIA prototype, UDAbot, was built using Unity 3D, XR Plugin Management, the Meta All-in-One SDK, and the Voice SDK with Wit.ai for voice interaction. Generative AI and Retrieval-Augmented Generation (RAG) enhanced response accuracy. Finally, acceptance of the IIA was evaluated via a survey that assessed 21 factors related to technology acceptance, usability, and immersive environments. The study identified 16 key factors influencing user adoption.

Keywords: intelligent immersive assistants, extended reality, retrieval augmented generation, generative ai, acceptance factors, udabot

1 Introduction

In line with the constant evolution of emerging technologies, institutions such as the UDA provide assistance across various topics through an omnichannel approach. Now there is a growing need to do so through extended reality and voice environments, with an acceptable immersive experience. This study identified previous work on the building and evaluation of IIA acceptance. For example, in [1], they present the immersive Math City platform to improve mathematics learning through the 5E teaching model and evaluate usability using the System Usability Scale (SUS). In [2], they present an AI-assisted Virtual Medical Simulator (VMS) to improve medical emergency training, created with Unity and integrated with OpenAI. In [3], a Virtual Reality (VR)

application is presented to help children with Autism Spectrum Disorder (ASD) understand social scenarios. They evaluated the application with the Quality of Experience Assessment of Extended Reality Meetings (ITU-T P.1320). In [4], they used the UTAUT model to analyze the acceptance of Augmented Reality (AR) applications in the context of learning. To our knowledge, there is a gap in the current state of the art that motivates the development and evaluation of an IIA to assist with various topics, providing accurate, real-time information, with acceptable immersive and voice interaction.

Design thinking [5] guided this research. A survey [6] defined the need and problem, leading to a tailored solution for the IIA context. The prototype was developed using Unity 3D, XR Plugin Management, the Meta All-in-one SDK [1], and Voice SDK with Wit.Ai [2]. Generative AI and RAG were integrated to ensure accurate responses and an immersive experience. Prototype testing and IIA UDAbot implementation followed, after which acceptance was evaluated with a survey [6] addressing 21 acceptance, usability, and immersion factors, ultimately identifying 16 relevant factors influencing user adoption.

The rest of the article is structured as follows. Section 2 presents a literature review of related works. Section 3 explains the research methodology. Section 4 describes the IIA's construction, followed by Section 5's acceptance evaluation of the IIA UDAbot. Section 6 presents research results, while Section 7 offers discussion and conclusions.

2 Background

Selective literature reviews on immersive assistants were conducted using SpringerLink (326), IEEE Xplore (90), ACM Digital Library (53), and Scopus (5056). After filtering by inclusion and exclusion criteria, 34 articles remained that included “Immersive,” “Assistants,” “Quest,” “Unity,” and “Acceptance” in their abstracts. These abstracts were reviewed for relevance to this research, yielding 15 primary studies. The reviews were conducted at three stages: construction of the IIA in October 2024; evaluation of the assistant’s acceptance in September 2025; and article development for publication in January 2026. Below is a summary of the selected articles.

2.1 Intelligent Immersive Assistants (IIA)

There are research studies on IIA. For example, [1] presents the Math City immersive platform for improving mathematics learning, built with Unity, and evaluates its usability using the SUS. [7] presents the VR Math Bridge application for teaching calculus in Quest 3's immersive virtual environments, which shows effectiveness in providing personalized, real-time responses to learners.

[1] https://developers.meta.com/horizon/downloads/package/meta-xr-sdk-all-in-one-upm/#documentation

[2] https://wit.ai/

In [8], the authors evaluate how immersive VR and AR experiences affect user enjoyment and cognitive and behavioral engagement. They compared several devices and concluded that Oculus VR delivers a better immersive experience. Furthermore, [9] proposes an intelligent system of STEM digital books based on the metaverse and Generative AI, developed in Unity, that provides personalized real-time responses.

Expanding on immersive technology, [10] proposes a 3D AR assistant developed in Unity and notes that environmental configuration and processing limitations affect the quality of the results. Similarly, [11] introduces a 3D AI agent for immersive interactive experiences and integrates a Large Language Model (LLM) with RAG. The authors demonstrate that the proposed model outperforms ChatGPT-4 in factual accuracy, reliability, and reasonableness.

The integration of VR in training is also discussed in [2], where an AI-assisted VR VMS with environment-sensitive natural language is presented to improve medical emergency training, built with Unity and integrated with OpenAI. In a similar vein, [12] introduces a VR application developed to train psychiatry students in diagnosis in a virtual environment. This virtual agent simulates a patient with depression, featuring facial expressions and vocal interactions.

In addition, [13] analyzes the convergence between Extended Reality (XR) and AI through a systematic review. The authors observe how the way of interacting between the virtual and the physical is changing, and highlight a synergy between various disciplines to maximize its potential.

2.2 Acceptance of immersive assistants

In [14], an immersive assistant for an ultrasound robotic system is proposed, and SUS is used to evaluate its usability. Results showed that interacting with the friendly, responsive avatar felt natural, boosting user confidence and comfort. Similarly, in [15], a virtual AR geriatric patient simulator was developed, with SUS and TAM employed to assess usability and acceptance. The authors conclude that AR enables the creation of realistic scenarios, though challenges like user comfort and cognitive load remain.

Extending this focus on immersive applications for specific populations, [3] presents a VR application developed in Unity to help children with ASD understand social scenarios. The application was evaluated with the Quality of Experience Assessment of Extended Reality Meetings (ITU-T P.1320).

Continuing with research on user acceptance, [16] extends TAM to integrate AR animation acceptance factors, finding that interactivity, visual appeal, perceived usefulness, ease of use, and attitude are key factors. Similarly, [4] applies UTAUT to analyze AR app acceptance in learning, showing that performance expectation and facilitating conditions influence usage intention and academic performance.

Expanding into broader digital services, [17] presents a case study that identifies 12 key factors influencing digital service acceptance and highlights differences by service activity and user demographics. [18] proposes an omnichannel Intelligent Virtual Assistant (IVA) to support users across digital channels, concluding that a flexible architecture enables future integration of new components, such as immersive environments.

The state of the art of IIAs for construction, using technologies such as Unity for immersive experiences, is oriented towards specific areas such as teaching and learning and culture, rather than assisting with a range of topics through immersive and voice interaction. Furthermore, research evaluating IIA acceptance focuses on specific topics, such as teaching and learning, psychology, and simulators. This research applies various models and factors, including SUS, UATUT, TAM, and ITU-T P.1320, and proposes both common factors for technology acceptance and use and specific factors for immersive environments.

We understand that there is a gap in the state of the art. Through this research, we propose the building and evaluation of the acceptance of an IIA for assistance in various topics, with immersive and voice interaction.

3 Methodology

For this research, the Design Thinking methodology [19] was used. First, the problem was defined with a focus on user needs. Next, a solution suitable for the diverse realities of an IIA was devised. A prototype was then developed, tested, and the Intelligent Immersive Assistant was implemented. Finally, the acceptance of the IIA UDAbot was evaluated as added value (see Fig. 1).

In August 2024, during the "empathy" stage, 61 UDA students surveyed [6] —both male and female, aged 17 and older—identified which topics they wanted IIA assistance with. Their primary interests included the academic calendar (67%), academic offerings (55%), student benefits (55%), study abroad (55%), and scholarships (54%).

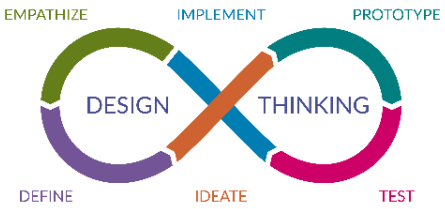

Fig. 1. Design thinking stages.

During the "definition" stage, the key problem was identified based on user perceptions and expert recommendations. The main user need is a tool that assists them with

topics of interest about the university, in an immersive environment for Meta Quest 3 devices, with intuitive XR interaction and voice commands. In addition, the goal is to evaluate user acceptance of the proposed IIA.

In the “ideation” stage, the problem was analyzed, and the solution was conceived. This solution extends the services IVA UDAbot through the immersive Meta Quest 3 channel to create the IIA. This IIA provides diverse information through intuitive XR interaction. It integrates interactive visual elements and can respond to users’ movements, hand and controller interactions, and voice commands. These features combine to deliver an immersive and interactive user experience.

In the “prototype” stage, the Unity 3D engine was used on Meta Quest 3 devices to build the IIA prototype. To enhance the environment, key features were added, including mixed reality using Passthrough from the Meta All-in-one SDK, interaction with virtual objects, and both graphical and voice interfaces. Furthermore, the prototype incorporates Generative AI, LLMs, and RAG to generate content and provide intelligent assistance on topics relevant to users. Additional construction details are provided in the “Building IIA UDAbot” section.

In December 2024, the “testing” stage was carried out, during which 15 UDA students evaluated the immersive experience and satisfaction with the IIA prototype on Meta Quest 3 devices. Results show that 60% of participants rated the experience as “very immersive”, 33% as “immersive”, and 7% as “moderately immersive”. For satisfaction, 53% rated the prototype as “very satisfied” and 47% as “satisfied”. Additionally, participants suggested adding a tutorial to provide preliminary information on available interactions with the IIA.

During the "implementation" stage, a tutorial was added at the start of the IIA interaction. On the technical side, performance and fluidity improved by reducing the number of scenes and 3D model elements, optimizing code, and adjusting configuration settings. This lowered resource consumption and improved fluid interaction with the IIA.

To evaluate acceptance of the IIA UDAbot, a 25-question survey [6] was developed to measure factors of acceptance, usability, and immersive experience. The survey achieved a Kuder-Richardson-20 (KR-20) reliability of 0,722, which is considered reliable [20]. In September 2025, 27 UDA students—both male and female from rural and urban areas—participated in the acceptance evaluation.

For the analysis of the IIA acceptance survey results, the “mean” was used as the measure of central tendency [21]. Of the 21 factors considered, 16 were identified as priorities for exceeding the mean. These are: 1) Perceived usefulness, 2) Attitude toward use, 3) Performance expectations, 4) Perceived ease of use, 5) Effort expectations, 6) Intention to use, 7) Hedonic motivation, 8) Perceived enjoyment, 9) Habit, 10) Usage behavior, 11) Usability, 12) Security, 13) Ergonomics, 14) Accessibility, 15) Experience, and 16) Visual aesthetics.

Additionally, differences in the relevant factors were identified by user gender and location. To synthesize these findings, the behavior of the priority factors was summarized in a Characterization [6] by defining patterns. Further details regarding this evaluation are provided in the sections “Acceptance IIA UDAbot” and “Results”.

4 Building IIA UDAbot

For the IIA UDAbot prototype, we set up Unity Hub, created a 3D project for extended reality and Meta Quest devices, and used the XR Plugin Management tool for virtual reality features. We also installed the Meta All-in-one SDK [3], which streamlines integrating LLMs like LLaMa and enables AI tasks such as inference. These tools made it easy to integrate the immersive assistant with Meta Quest 3, creating a seamless user experience.

Next, the prototype was developed. Several Meta All-in-one SDK functions were integrated, including the Mixed Reality (MR) “Passthrough” function. This feature allows the Quest 3's cameras to display a real-time view of the physical environment while combining it with virtual elements. This method created an immersive experience.

Hands were recognized in the extended environment. The “Hand Grab” feature enables users to interact with virtual objects by moving, rotating, or grabbing them. Fig. 2 shows the left hand being recognized.

Different ways to interact with virtual objects were added. “Poke Interaction” lets users touch and select items in a virtual space with their hands, as shown in Fig. 2 when grasping a virtual tablet. “Ray Interaction” allows users to point at and interact with distant objects using a virtual pointer. Fig. 2 also shows the pointer highlighting the information menu.

To make the environment feel realistic and engaging, gravity, collisions, and friction were applied to virtual objects so they respond naturally when you move or release them. For example, an object falls to the ground or bumps into other objects, just like in real life.

For voice interaction, Speech-to-Text and Text-to-Speech technologies were used via Meta's “Voice SDK” toolkit with the Wit.Ai service. This service offers natural language processing with user voice recognition. It converts spoken language into understandable text and produces responses in both text and voice formats. As shown in Fig. 3, the user interacts with the IIA via voice, asking who it is. The voice command is processed by the Wit.Ai service, which returns a response in both voice and text.

[3] https://developers.meta.com/horizon/downloads/package/meta-xr-sdk-all-in-one-upm/#documentation

Fig. 2. Hand Grab, Poke Interaction and Ray Interaction.

Fig. 3. IIA UDAbot integrated with voice SDK, Amazon Bedrock, and RAG.

Amazon Bedrock, an AWS service, was integrated to generate content and responses on various topics requested by the user. This platform enables the creation of generative AI applications and simplifies the use of advanced generative language models, including Meta's LLaMa, Anthropic's Claude, and Amazon's Titan. LlaMa 3.1, an LLM with 8, 70, and 405 trillion parameters, was used. This enables adaptation to various processing needs. LlaMa 3.1 is open-source and optimized for natural language processing tasks, including text generation and comprehension. The RAG technique was also implemented. It allows the IIA to generate accurate responses in real time by combining the LLM with external information. This approach improves response quality and user interaction.

5 Acceptance IIA UDAbot

To evaluate the acceptance of the IIA UDAbot, a descriptive survey research method [6] was used. This method collected standardized information from a specific sample and analyzed the acceptance and use of the IIA.

Factors influencing acceptance and use of digital services, identified in previous studies [17], [22] were used as input for this evaluation. These factors align with TAM, UATUT, ISO 9241-210, and ISO/IEC 25010. Key factors include 1) IT knowledge, 2) Perceived usefulness, 3) Attitude toward use, 4) Performance expectations, 5) Perceived ease of use, 6) Effort expectations, 7) Intention use, 8) Social influence, 9) Facilitating conditions, 10) Hedonic motivation, 11) Perceived Enjoyment, 12) Habit, 13) Usage behavior, 14) Usability, 15) Reliability, 16) Security, 17) Ergonomics, and 18) Accessibility. Since the IIA operates in an immersive environment, XR factors [16] — 19) Experience, 20) Interactivity, and 21) Visual aesthetics —were also considered. These variables define the evaluation context and support the questionnaire's structure to measure acceptance, usability, and immersive experience with IIA UDAbot [6].

The survey contains 25 questions, including three sociodemographic questions (Gray) and 22 yes/no questions related to factors: 17 about technology acceptance and use (Orange) and five about immersive experience (Green). Furthermore, consistency and reliability were assessed using the KR-20 coefficient (0,722; considered reliable) [20]. For more details on the relationships between questions and factors, see Table 1, where "FACTOR" lists the 21 factors, "DESCRIPTION" lists the 25 questions, and "ANCHOR" lists the references.

Table 1. Relate Factors to Questionnaire.

FACTOR	DESCRIPTION	ANCHOR
Age	¿How old are you?	
Gender	Indicate your gender.	0=F; 1=M
Location	Indicate your location.	0=Rural; 1=Urban
IT knowledge	¿Have you previously used devices or applications in immersive mixed reality environments?	0=No; 1=Yes
Perceived usefulness / Attitude toward use / Performance expectations	¿While using the prototype, did you find useful information about the university?	0=No; 1=Yes
	¿Do you think it would be useful to integrate this prototype into academic activities to achieve better learning outcomes?	0=No; 1=Yes
Perceived ease of use / Effort expectations	¿Was the prototype easy to use?	0=No; 1=Yes
Intention to use	¿Would you use this prototype if it were available in your courses?	0=No; 1=Yes
	¿Would you recommend the use of the prototype to other students?	0=No; 1=Yes
Social influence	¿Do people whose opinions you value use immersive environments?	0=No; 1=Yes

Facilitating Conditions	¿Does the university have the technological facilities to access the use of immersive environments?	0=No; 1=Yes
	¿If you had difficulties using the prototype, could you get help?	0=No; 1=Yes
Hedonic Motivation / Perceived Enjoyment	¿Did you find the prototype entertaining?	0=No; 1=Yes
Habit	¿Would you incorporate the prototype into your academic routine?	0=No; 1=Yes
Usage Behavior	¿Would you like to use the prototype more often?	0=No; 1=Yes
Usability	¿Was it easy to learn how to use the basic functions of the immersive environment?	0=No; 1=Yes
Reliability	¿Did the prototype have any glitches that forced you to restart or repeat a task?	0=No; 1=Yes
Security	¿Do you consider the prototype to be secure in terms of the information it handles?	0=No; 1=Yes
Ergonomics	¿Is the prototype geared toward the user's actual needs?	0=No; 1=Yes
Accessibility	¿Do you think the prototype can be operated by users with different abilities?	0=No; 1=Yes
Experience	¿Did you feel a clear sense of being inside the immersive environment?	0=No; 1=Yes
Interactivity	¿Were you able to accurately control your actions in the immersive environment?	0=No; 1=Yes
	¿Was the detection of controls and virtual objects adequate?	0=No; 1=Yes
	¿Was voice detection adequate?	0=No; 1=Yes
Visual aesthetics	¿Did you find the visual design of the prototype visually appealing and pleasant?	0=No; 1=Yes

In September 2025, the IIA was validated with 27 UDA students, aged 18-30, both male and female. The participants came from rural and urban areas and were studying Computer Science Engineering, Electronic Engineering, and Architecture. The survey results were analyzed using quantitative descriptive statistics. The survey considered 21 factors of acceptance, usability, and immersive experience. Of these, 16 were seen as relevant because they exceeded the overall average of 89,58% [21]. These factors are listed in Table 2: those marked in Orange are for acceptance and usability, while those marked in Green are for immersive experience.

Table 2. Factors of acceptance, usability, and immersive experience relevant to IIA.

FACTOR	PERCENTAGE
Perceived usefulness / Attitude toward use / Performance expectations	100%
Perceived ease of use / Effort expectations	100%
Intention to use	100%
Hedonic motivation / Perceived enjoyment	100%
Habit	96,30%

Usage behavior	96,30%
Usability	100%
Security	96,30%
Ergonomics	96,30%
Accessibility	92,59%
Experience	96,30%
Visual aesthetics	96,30%

The remaining five factors, such as IT knowledge (62,96%), Social influence (59,26%), Facilitating conditions (64,81%), Reliability (81,48%), and Interactivity (83,95%), are below average. Although not considered relevant, they should still be analyzed as part of actions to strengthen IIAs.

According to the gender analysis unit, there are differences in the relevant factors for males and females. For females, with a mean of 89,92%, 13 relevant factors were identified: 1) Perceived usefulness, 2) Attitude toward use, 3) Performance expectations, 4) Perceived ease of use, 5) Effort expectations, 6) Intention to use, 7) Hedonic motivation, 8) Perceived enjoyment, 9) Habit, 10) Usability, 11) Ergonomics, 12) Experience, and 13) Visual aesthetics. In contrast, the general approach gives higher priority to Usage behavior, Security, and Accessibility, which are not top priorities for females. For males, with an average of 89,46%, 16 priority factors closely align with those in the general approach, indicating similarities between males and the general group, while females show a distinct set of priorities.

According to the analysis by user location, users in rural areas have an average of 85,78% with 13 priority factors. Females also have 13 relevant factors; however, Usage behavior is a priority for rural users but not for females, whereas Visual aesthetics is a priority for women but not for rural users. Urban users have an average of 90,24% and share the same 16 priority factors as the general user group.

6 Results

The following findings are relevant results derived from the building and evaluation of the acceptance of the IIA UDAbot for Quest 3 devices:

The key finding is the adoption of the Unity 3D video game engine for IIA development. VR functionalities are enabled by the “XR Plugin Management” plugin, while the “Meta All-in-one SDK” integrates the immersive environment and assistant with Meta Quest 3. Main features include: 1) “Passthrough” for blending physical and virtual environments; 2) “Hand Grab” for recognizing and interacting with hands; 3) interaction with 3D objects using “Poke Interaction” (touching) and “Ray Interaction” (interaction with distant objects). (see Fig. 2).

The second key finding is the successful enablement of voice interaction using Meta's Voice SDK with Wit.Ai, supporting speech-to-text and text-to-speech for natural language understanding. In addition, Generative AI was integrated via Amazon Bedrock, enhancing content creation and enabling accurate, real-time responses. Furthermore, the RAG technique combined the LLM with external information. Collectively, these technologies led to the development of an IIA prototype that delivered an immersive, satisfactory experience (see Fig. 3).

The third finding comes from evaluating IIA UDAbot acceptance. 16 relevant factors related to acceptance, usability, and immersive experience were identified. Differences in priority factors appear for gender and user location units of analysis, with females and the rural sector each showing only 13 relevant factors. These differing priority factors are summarized in a **Characterization**, which defines behavior patterns for each unit of analysis (see Table 3). The 16 relevant factors are in the "FACTORS" column; factors in Orange relate to acceptance and usability, and those in Green relate to immersive experience. The "PATTERNS-DIMENSIONS" columns show the "General," "Age," "Gender," and "Location" approaches. Cells with "*" mark a factor as relevant, and "Gray" cells mean the factor is not relevant. The general approach, age, male gender, and urban sector present all 16 relevant factors. For the female gender, Usage behavior, Security, and Accessibility are not relevant. For the rural sector, Security, Accessibility, and Visual aesthetics are not relevant.

Table 3. Characterization of priority acceptance, usability, and immersive experience factors for IIA UDAbot

FACTORS	PATTERNS-DIMENSIONS					
			Gender		Location	
	General	Age	Female	Male	Rural	Urban
Perceived usefulness / Attitude toward use / Performance expectations	*	*	*	*	*	*
Perceived ease of use / Effort expectations	*	*	*	*	*	*
Intention to use	*	*	*	*	*	*
Hedonic motivation / Perceived enjoyment	*	*	*	*	*	*
Habit	*	*	*	*	*	*
Usage behavior	*	*		*	*	*
Usability	*	*	*	*	*	*
Security	*	*		*		*
Ergonomics	*	*	*	*	*	*
Accessibility	*	*		*		*
Experience	*	*	*	*	*	*
Visual aesthetics	*	*	*	*		*

7 Discussion and conclusion

The construction and evaluation of acceptance of the IIA UDAbot for Quest 3 devices provide significant insights into the technologies used and the factors of acceptance, usability, and immersive experience for IIA adoption.

Selective literature reviews on the construction and acceptance of immersive assistants were conducted in SpringerLink (326), IEEE Xplore (90), ACM Digital Library (53), and Scopus (5056). After applying inclusion and exclusion criteria and reviewing selected abstracts, 15 primary studies were obtained. The review reveals that prior studies concentrate on specific areas and topics.

The project followed the six-stage Design Thinking methodology [19]: 1) empathizing to detect needs, 2) defining the problem, 3) devising an IIA to assist on various topics in an immersive environment for Meta Quest 3 headsets, using XR intuitive interaction and voice control, 4) developing the prototype, 5) testing the prototype, and 6) implementing the IIA. The acceptance of the IIA UDAbot was also evaluated using a survey.

During prototype development, several contributions were made. The Unity 3D engine and XR Plugin Management enabled VR functionality. The Meta All-in-one SDK integrated the immersive environment and assistant with Meta Quest 3, incorporating features such as: 1) 'Passthrough' to blend physical and virtual environments; 2) "Hand Grab" for hand recognition and interaction with extended virtual objects; 3) "Poke Interaction" for touch-based control of 3D virtual objects; and 4) "Ray Interaction" for engaging with distant objects (see Fig. 2).

Meta's "Voice SDK" was enabled with the Wit.Ai service for Speech-to-Text and Text-to-Speech for natural language understanding. Voice recognition achieved 90% accuracy in quiet environments and 80% in noisy environments. To improve content generation and enable real-time responses, the Generative AI LLaMa 3.1 was integrated via Amazon Bedrock, using the RAG technique to combine the LLM with relevant external information. Together, these implementations enhanced the quality of responses and user interaction.

As an added value, the research evaluated IIA UDAbot acceptance. A 25-question survey covering acceptance, usability, and immersive experience was developed [6], showing a KR-20 reliability coefficient of 0,722 [20]. The survey included three sociodemographic questions, 17 on technological acceptance and usability, and five on immersive experience (see Table 1). Based on the evaluation, with a general approach, 16 relevant factors were identified [21] (see Table 2). The analysis by gender and location revealed varying priority factors: females and rural users showed 13 relevant factors. These differences were synthesized as a characterization of the IIA priority factors by units of analysis (see Table 3).

Combining acceptance evaluation results and the review of immersive environments produced several recommendations for future IIA implementations: 1) Since Wit.AI voice recognition was available only in English, Spanish pronunciation limited IIA response naturalness; updating to a Spanish-supporting version is advised. 2) Implement brief digital scaffolding to orient users to immersive interaction logic with discreet visual cues or messages, which impacts the IT knowledge factor [23]. 3) Mark actionable elements clearly and provide immediate feedback after each action to improve Reliability and Interactivity factors [24]. 4) Optimize startup fluidity and ensure automatic recovery from interruptions by saving progress and restoring last position or task, supporting Reliability and Ease of Use. 5) Enhance visual and sensory accessibility with high-contrast settings, font size adjustments, and alternative cues such as vibration for inclusive usability [16].

This research contributes to both the construction and evaluation of IIA acceptance in XR and voice environments. The study's findings show that integrating techniques and tools provided assistance across various topics, delivering accurate answers and intuitive, immersive interactions. Finally, 16 priority factors for IIA adoption were identified.

References

[1] M. Singh, D. Sun, and Y. Yang, “Design and implementation of Math City: An inquiry-based metaverse platform for mathematics education,” *Computers and Education: Artificial Intelligence*, vol. 9, Dec. 2025, doi: 10.1016/j.caeai.2025.100462.

[2] H. Sirandass, C. Jaiswal, and M. Sukumar, “Q-VMedSim: A Training Platform for Emergency Response Procedures through Virtual Reality and AI Assistance,” in *2025 IEEE World AI IoT Congress (AIIoT)*, 2025, pp. 285–291. doi: 10.1109/AIIoT65859.2025.11105232.

[3] Y. Zhang, C. Keighrey, and N. Murray, “A VR Intervention to Develop Social Skills for Children with ASD: An Expert Evaluation,” in *IMX 2024 - Proceedings of the 2024 ACM International Conference on Interactive Media Experiences*, Association for Computing Machinery, Inc, Jun. 2024, pp. 30–40. doi: 10.1145/3639701.3656322.

[4] S. J. Alotaibi, “UTAUT model for the use of Augmented Reality in education and training,” *Revista Romana de Informatica si Automatica*, vol. 33, no. 1, pp. 21–32, 2023, doi: 10.33436/v33i1y202302.

[5] L. Przybilla, K. Klinker, M. Lang, M. Schreieck, M. Wiesche, and H. Krcmar, “Design Thinking in Digital Innovation Projects - Exploring the Effects of Intangibility,” *IEEE Trans. Eng. Manag.*, vol. 69, no. 4, pp. 1635–1649, Aug. 2022, doi: 10.1109/TEM.2020.3036818.

[6] R. Hernández Sampieri, C. Fernández Collado, and P. Baptista Lucio, “Metodologia de la Investigación (Sexta edición),” *Mc Graw Hill*, 2014.

[7] H. Lai, A. Nassani, J. Blake, and J. Villegas, “VR Math Bridge: Bridging Interactivity in Online Education with AI and VR,” in *2025 IEEE Gaming, Entertainment, and Media Conference (GEM)*, Sep. 2025, pp. 1–6. doi: 10.1109/GEM66882.2025.11155841.

[8] I. Verhulst, A. Woods, L. Whittaker, J. Bennett, and P. Dalton, “Do VR and AR versions of an immersive cultural experience engender different user experiences?,” *Comput. Human Behav.*, vol. 125, Dec. 2021, doi: 10.1016/j.chb.2021.106951.

[9] Y. Luo, D. Ke, and J. Lin, “A Design Scheme of STEM Digital Textbook Based on Metaverse and Generative Artificial Intelligence,” in *ICETM 2024 - Proceedings of the 2024 7th International Conference on Educational Technology Management*, Association for Computing Machinery, Inc, Mar. 2025, pp. 299–308. doi: 10.1145/3711403.3711453.

[10] X. Huang, J. Twycross, and F. Wild, “A PROCESS FOR THE SEMI-AUTOMATED GENERATION OF LIFE-SIZED, INTERACTIVE 3D CHARACTER MODELS FOR HOLOGRAPHIC PROJECTION,” 2020.

[11] J. Lee, G. Ali, and J.-I. Hwang, “RAG based AI-Agent for Contextualized Analysis of High-Density Historical Records: Application to the Annals of the Joseon Dynasty,” in *2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*, 2025, pp. 893–894. doi: 10.1109/ISMAR-Adjunct68609.2025.00243.

[12] L. Graf, G. Gradl-Dietsch, and M. Masuch, “Depressed Virtual Agents: Development of a Playful VR Application for the Training of Child and Adolescent Psychiatry Students,” in *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents, IVA 2023*, Association for Computing Machinery, Inc, Sep. 2023. doi: 10.1145/3570945.3607306.

[13] A. Singh, S. Joshi, and M. Domb, “Embedded Conversational AI, Chatbots, and NLP to Improve Healthcare Administration and Practices,” in *2023 2nd International Conference on Automation, Computing and Renewable Systems (ICACRS)*, IEEE, Dec. 2023, pp. 38–45. doi: 10.1109/ICACRS58579.2023.10404985.

[14] T. Song, F. Pabst, U. Eck, and N. Navab, “Enhancing Patient Acceptance of Robotic Ultrasound through Conversational Virtual Agent and Immersive Visualizations,” *IEEE Trans. Vis. Comput. Graph.*, vol. 31, no. 5, pp. 2901–2911, 2025, doi: 10.1109/TVCG.2025.3549181.

[15] A. Ahmmed, E. Butts, K. Naeiji, L. Thiamwong, and S. Daher, “System Usability and Technology Acceptance of a Geriatric Embodied Virtual Human Simulation in Augmented Reality,” in *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2024, pp. 594–603. doi: 10.1109/ISMAR62088.2024.00074.

[16] X. Zhang, K. Zhang, L. Liang, J. Zhang, and J. Chen, “Research on the Factors and Pathways Affecting Users’ Willingness to Use AR Animations for Science Popularization: An Extension of the Technology Acceptance Model,” *IEEE Access*, vol. 13, pp. 60088–60107, Apr. 2025, doi: 10.1109/ACCESS.2025.3557006.

[17] P. Pintado, S. Wiesner, D. Prado, and E. Estevez, “Examining Factors of Acceptance and Use of Technology in Digital Services in the Context of Ecuador,” in *Communications in Computer and Information Science*, Springer Science and Business Media Deutschland GmbH, 2023, pp. 171–185. doi: 10.1007/978-3-031-40942-4_13.

[18] D. and D. G. and C. A. and I. D. and M. E. and F. J. Pintado Pablo and Mejia, “Building Omnichannel Intelligent Virtual Assistants,” in *Cloud Computing, Big Data and Emerging Topics*, L. and C. F. and L. L. Naiouf Marcelo and De Giusti, Ed., Cham: Springer Nature Switzerland, 2026, pp. 93–108.

[19] A. Albataineh and A. Qusef, “Why Companies in Jordan Should Adopt e-Business Solutions?,” in *2020 11th International Conference on Information and Communication Systems, ICICS 2020*, Institute of Electrical and Electronics Engineers Inc., Apr. 2020, pp. 272–276. doi: 10.1109/ICICS49469.2020.239498.

[20] G. F. Kuder and M. W. Richardson, “The theory of the estimation of test reliability,” *Psychometrika*, vol. 2, no. 3, 1937, doi: https://doi.org/10.1007/BF02288391.

[21] J. Sageder, A. Demleitner, O. Irlbacher, and R. Wimmer, “Applying voting methods in user research,” in *ACM International Conference Proceeding Series*, Association for Computing Machinery, Sep. 2019, pp. 571–575. doi: 10.1145/3340764.3344461.

[22] P. Pintado, I. Jaramillo, D. Prado, and E. Estevez, “The Taxonomy of Factors of Acceptance and Use of Technologies for Human Computer Interface in Digital Services,” 2023. doi: 10.24215/16666038.23.e14.

[23] J. Bacca-Acosta, J. Tejada, R. Fabregat, Kinshuk, and J. Guevara, “Scaffolding in immersive virtual reality environments for learning English: an eye tracking study,” *Educational technology research and development*, vol. 70, no. 1, pp. 339–362, 2022, doi: 10.1007/s11423-021-10068-7.

[24] T. Hirzle, F. Müller, F. Draxler, M. Schmitz, P. Knierim, and K. Hornbæk, “When XR and AI Meet - A Scoping Review on Extended Reality and Artificial Intelligence,” in *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, Apr. 2023. doi: 10.1145/3544548.3581072.

Feasibility Assessment for the Implementation of FinOps Practices in SMEs in the Province of Misiones: A Study Based on Local Infrastructure and Provincial Cloud Computing Services

Miryan R. Puchini[1, 2][0000-0003-0420-8688], Nancy B. Ganz[1, 2, 3][0000-0002-9891-3371] and Yolanda E. Zado[1][0009-0006-7396-4021]

[1] Universidad Nacional de Misiones, Facultad de Ingeniería, Obera, Misiones. Argentina.
[2] Universidad Nacional de Misiones, Facultad de Ciencias Exactas Químicas y Naturales, Posadas, Misiones. Argentina.
[3] Instituto de Materiales de Misiones (IMAM), Posadas, Misiones. Argentina.
miryanpuchini@fio.unam.edu.ar

Abstract. Cloud computing has established itself as a pillar of organizational digital transformation, offering flexible and scalable pay-per-use models. However, this paradigm introduces significant challenges in the financial management of technological resources, particularly in small and medium-sized enterprises (SMEs), where a lack of visibility and governance of spending can limit the expected benefits. In this context, FinOps (Cloud Financial Operations) practices are emerging as an organizational framework aimed at maximizing the economic value of the cloud through coordination between the areas of technology, finance, and business. This work evaluates the feasibility of implementing FinOps practices in SMEs in the province of Misiones, considering the characteristics of the regional productive fabric and the availability of cloud computing infrastructure and services provided by local actors, especially the state-owned company Marandú Comunicaciones S.E. The methodology combines a systematic review of recent literature on cloud computing, FinOps, and technology adoption in SMEs with semi-structured interviews with key actors in the provincial ecosystem. As a main contribution, it proposes a Matrix for Evaluating the Feasibility of Implementing FinOps Practices, structured in six dimensions that integrate technical, economic, organizational, and institutional aspects, under a triple helix approach that articulates business, university, and government. The results indicate that, although SMEs in Misiones have low to intermediate levels of digital maturity, there is a favorable environment for the progressive and contextualized adoption of FinOps, aimed at improving the efficiency, predictability, and sustainability of cloud use in regional economies.
Keywords: Cloud Computing, FinOps, Small and medium-sized enterprises (SMEs).

Resumen. La computación en la nube se ha consolidado como un pilar de la transformación digital organizacional, al ofrecer modelos flexibles y escalables de pago por uso. Sin embargo, este paradigma introduce desafíos significativos en la gestión financiera de los recursos tecnológicos, especialmente en las pequeñas y medianas empresas (PyMEs), donde la falta de visibilidad y gobernanza del gasto puede limitar los beneficios esperados. En este contexto, las prácticas FinOps (Operaciones Financieras en la Nube) emergen como un marco organizacional orientado a maximizar el valor económico de la nube mediante la coordinación entre las áreas de tecnología, finanzas y negocio. Este trabajo evalúa la factibilidad de implementar prácticas FinOps en PyMEs de la provincia de Misiones, considerando las características del entramado productivo regional y la disponibilidad de infraestructura y servicios de computación en la nube provistos por actores locales, en particular la empresa estatal Marandú Comunicaciones S.E. La metodología combina una revisión sistemática de literatura reciente sobre computación en la nube, FinOps y adopción tecnológica en PyMEs, con entrevistas semiestructuradas a actores clave del ecosistema provincial. Como principal aporte, se propone una Matriz de Evaluación de Factibilidad para la Implementación de Prácticas FinOps, estructurada en seis dimensiones que integran aspectos técnicos, económicos, organizacionales e institucionales, bajo un enfoque de triple hélice que articula empresa, universidad y gobierno. Los resultados indican que, si bien las PyMEs de Misiones presentan niveles de madurez digital bajos a intermedios, existe un entorno favorable para la adopción progresiva y contextualizada de FinOps, orientada a mejorar la eficiencia, previsibilidad y sostenibilidad en el uso de la nube en las economías regionales.
Palabras clave: Computación en la Nube, FinOps, Pequeñas y Medianas Empresas (PyMEs).

1 Introduction

Currently, digital transformation shapes the technological landscape of organizations and is fundamentally grounded in the cloud computing paradigm. This model has revolutionized the way computational resources, storage, and applications are delivered by replacing traditional technological infrastructure—conceived as an internal, rigid, and high-cost asset—with scalable, flexible, and on-demand services. In this sense, cloud adoption does not merely represent a technological upgrade but entails a profound economic, organizational, and operational shift in the management of Information Technology (IT) resources.

In the province of Misiones, most small and medium-sized enterprises (SMEs) are concentrated in traditional productive sectors such as tea, yerba mate, forestry, metal-mechanical, and sawmill industries. These organizations face significant challenges in advancing digital transformation processes and adopting Industry 4.0 models, mainly due to the limited incorporation of digital technologies, low maturity in cloud service management, and the absence of financial governance mechanisms aimed at controlling and optimizing costs in cloud environments.

Although cloud computing infrastructure and services are available from companies located in the Northeast region of Argentina (NEA), and even from local providers such as the state-owned company Marandú Comunicaciones S.E., there is insufficient empirical evidence to assess the feasibility and concrete benefits of implementing FinOps (Cloud Financial Operations) practices in the context of SMEs in Misiones. This situation creates uncertainty regarding these organizations' ability to efficiently manage cloud-related costs, optimize their technology investments, and generate sustainable competitive advantages.

The objective of this study is to identify the factors and barriers that hinder the adoption of cloud computing and FinOps practices among SMEs in the province of Misiones. To this end, a methodology is proposed based on the review of relevant academic literature and the conduction of interviews with representatives of cloud service providers, as well as members of executive and managerial teams of local SMEs. Understanding these obstacles is considered essential in order to propose strategies and interventions that facilitate an effective technological transition aligned with the actual capabilities of these organizations.

The knowledge generated from this study aims to contribute to informed decision-making by policymakers, researchers, and industry professionals, promoting the digital transformation of SMEs, which play a strategic role in developing countries as drivers of economic growth, job creation, and facilitators of economic self-sufficiency.

2 Theoretical Framework

2.1 Cloud computing

For the definition of the cloud computing concept, this study adopts the conceptual framework proposed by the National Institute of Standards and Technology (NIST), an internationally recognized authority in technological standardization. According to NIST, cloud computing is a model that enables ubiquitous, convenient, and on-demand network access to a shared pool of configurable computing resources—such as networks, servers, storage, applications, and services—that can be rapidly provisioned and released with minimal management effort and limited interaction with the service provider.

This model is based on five essential characteristics: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. In particular, the measured service characteristic introduces a substantial shift in the economic logic of IT by enabling pay-as-you-go schemes that transform technology costs into variables dependent on actual consumption [1].

2.2 Relationship between Cloud Computing and FinOps

The fundamental properties of cloud computing—elasticity, rapid provisioning, and consumption-based billing—enable financial models in which IT costs shift from fixed and predictable expenses to variable costs closely tied to the actual use of resources. While this approach offers significant advantages in terms of flexibility and scalability, it also introduces important challenges in financial management, particularly for small and medium-sized enterprises (SMEs).

The ease of resource provisioning, when not accompanied by adequate control and governance mechanisms, can lead to scenarios of overconsption, inefficiency, and costs that are difficult to anticipate. This issue highlights the need for specific financial management approaches tailored to cloud environments, capable of aligning technological decisions with the organization's economic and strategic objectives.

In this context, FinOps practices emerge as a collaborative framework that integrates technology, finance, and business teams with the aim of maximizing the economic value of cloud computing, improving cost visibility, and optimizing expenses associated with cloud service consumption. Thus, FinOps becomes an essential complement for the sustainable adoption of cloud computing, particularly for SMEs located in underdeveloped or developing countries.

2.3 FinOps (Cloud Financial Operations)

FinOps is defined as an organizational and cultural practice aimed at the efficient, predictable, and governed management of cloud computing expenditure, based on continuous collaboration among engineering, finance, and business teams. It emerges as a response to the challenges inherent in the pay-as-you-go model, in which the lack of visibility, control, and financial governance can lead to high and unpredictable costs.

Academic literature and industry guidelines converge in characterizing FinOps as a framework that promotes data-driven decision-making based on actual consumption metrics, emphasizing processes, metrics, and shared responsibilities rather than the use of specific tools or technologies. This approach is particularly relevant in large-scale cloud environments, such as those provided by Amazon Web Services, where the complexity of billing models and the elasticity of resources increase the risk of financial inefficiencies if appropriate control practices are not implemented.

From the perspective of cloud service providers, Google Cloud defines FinOps as an operational framework and a cultural shift that integrates technology, finance, and business with the aim of fostering financial accountability and accelerating business value realization from cloud transformation. This definition reinforces the cross-functional nature of FinOps and its explicit focus on value creation beyond mere cost control.

In this sense, FinOps enables organizations to enhance financial accountability in the use of cloud resources and maximize business value; it contributes to understanding the complexity and challenges of traditional IT financial management; and it facilitates the identification of key components and metrics necessary to achieve measurable and sustainable outcomes. Google Cloud also highlights the importance of defining clear success metrics to evaluate the impact of FinOps adoption, including accountability and enablement metrics, measurement and completion metrics, cost optimization metrics, planning and forecasting metrics, and metrics related to the use of tools and technological accelerators [2].

The FinOps framework is operationally structured into three main phases or pillars:

- Inform: focused on establishing financial visibility into cloud resource consumption through the consolidation and analysis of billing data, resource tagging, and the identification of usage patterns and anomalies.
- Optimize: aimed at improving technical and economic efficiency through proper resource sizing (rightsizing) and the strategic use of commitment-based pricing models, such as reserved instances, savings plans, and flexible capacity instances.
- Operate: oriented toward institutionalizing FinOps practices through continuous monitoring, automation, and the implementation of preventive controls to avoid inefficient configurations or those misaligned with financial objectives, supported by approaches such as Infrastructure as Code (IaC).

From this perspective, FinOps is distinguished by its provider-, technology-, and architecture-agnostic nature, which facilitates its progressive adoption in organizations of different sizes. For SMEs, particularly in regional contexts and developing economies, FinOps represents a key framework for maximizing the value of cloud computing and preventing operational and financial complexity from undermining the expected economic benefits [3] [4].

2.4 Case Study: Implementation of FinOps Practices in Fintech Organizations

The Fintech sector (short for Financial Technology) comprises companies and organizations that leverage digital technologies to deliver or enhance financial services. This sector represents one of the most advanced environments in the adoption of FinOps practices, due to its strong reliance on scalable cloud infrastructures, digital business models, and real-time transactional operations. Fintech companies operate in contexts characterized by high demand volatility—such as payment spikes, tagging strategies, and instant credit services—where cloud elasticity is strategic, but also entails financial risks if consumption is not rigorously managed.

Several case studies demonstrate that large-scale Fintech organizations have achieved significant cost reductions through the structured implementation of FinOps frameworks, combining:

- Intelligent workload distribution
- Automation and dynamic scaling (autoscaling)
- Resource rightsizing
- Multi-cloud strategies to optimize pricing and mitigate vendor lock-in
- Container-based architectures and serverless models.

Reported outcomes indicate cloud cost reductions ranging from 25% to 40%, while maintaining high standards of security and regulatory compliance. In particular, the adoption of serverless models has enabled neobanks and payment platforms to pay exclusively for actual execution time, eliminating costs associated with idle infrastructure. Additionally, the implementation of FinOps metrics, resource tagging, periodic audits, and centralized spending governance has allowed organizations to transform cloud consumption into a predictable

model aligned with business objectives. These practices have not only improved financial efficiency but also strengthened innovation capacity by freeing up resources for investment in new digital products.

From a strategic perspective, the Fintech case demonstrates that FinOps should not be understood solely as a cost-reduction tool, but rather as an enabler of financial sustainability, operational scalability, and competitive advantage in highly dynamic digital markets.

This background is particularly relevant to the present study, as it shows that maturity in cloud financial management constitutes a critical factor for maximizing the economic value of digital transformation, even in organizations operating under technology-intensive and highly complex models.[5]

2.5 Advanced FinOps: Cloud Consumption Forecasting and Artificial Intelligence-Based Techniques

Recent literature on FinOps-oriented solutions identifies resource usage forecasting as a key component for efficient financial management in cloud computing environments. In this context, two main approaches to consumption forecasting are distinguished: global forecasting models and local forecasting models. Global models enable the capture of general behavioral patterns and reduce operational complexity; however, they present limitations in adequately representing heterogeneous infrastructures with multiple workloads and differentiated consumption profiles. Conversely, local models offer greater accuracy when modeling individual resources, although their adoption is constrained by scalability issues and by the increase in computational and operational costs as the cloud environment grows.

As an intermediate alternative, several studies propose grouping similar time series for the application of shared predictive models. Nevertheless, these approaches often present limitations related to scalability, the lack of metrics explicitly aligned with FinOps, and limited integration of multivariate forecasting to support medium- and long-term strategic decision-making.

In parallel, the scale and complexity of billing schemes in hyperscale cloud environments have created favorable conditions for the application of Artificial Intelligence (AI) and Machine Learning (ML) techniques. These technologies represent a natural evolution of the FinOps framework, enabling a transition from rule-based automation to data-driven predictive optimization approaches capable of anticipating consumption patterns, reducing operational uncertainty, and improving the economic efficiency of cloud environments.

In this vein, the analyzed study proposes the MSFS system, a hybrid approach for multi-time-series forecasting aimed at optimizing cloud resources from a FinOps perspective. The system introduces the concept of Group-Specific Forecasting Models (GSFM), which cluster virtual machines with similar behavior through embedding and clustering techniques, allowing predictive models to be trained per group and achieving a balance between scalability and accuracy. It also integrates consensus-based anomaly detection mechanisms (HEADA), administrative rules, and metrics specifically aligned with FinOps—such as the Forecast Ratio—with the objective of improving financial decision-making.

Validated in a real-world telecommunications environment, the MSFS system demonstrates superior performance compared to traditional global and local approaches, achieving operational expenditure (OpEx) reductions of up to 44.7%. This highlights the potential of AI-based predictive techniques as advanced tools to support the mature implementation of FinOps practices.[4]

2.6 Small and Medium-Sized Enterprises (SMEs)

Small and medium-sized enterprises (SMEs) play a fundamental role in economies worldwide, serving as one of the primary drivers of economic and social development. Academic literature recognizes them as key contributors to growth, job creation, and economic self-sufficiency, particularly in developing countries and emerging economies.

Numerous studies agree that SMEs constitute the backbone of the productive structure in many regions of the world. Research conducted in countries across the Middle East and North Africa, for example, highlights that these enterprises make significant contributions to Gross Domestic Product (GDP) and are essential for economic stability and sustainability. These characteristics are shared by many regions with production structures based on traditional sectors, as is the case in large parts of Latin America.

One of the main concerns of SMEs lies in the need to reduce operational costs and ensure business sustainability over the medium and long term. In this regard, access to financing represents one of the most significant barriers to their development, particularly in contexts of lower levels of economic development, where financial institutions often impose stricter constraints on credit provision to this type of organization.

In this context, many authors agree that the adoption of cloud computing represents an effective alternative for SMEs to reduce costs, increase operational flexibility, and strengthen business sustainability. Cloud computing plays a significant role in fostering innovation and improving organizational productivity by enabling companies to dynamically adjust their IT resources according to actual business demands. Furthermore, access to computing

resources delivered over the Internet allows for reliable data storage, management, and protection without requiring substantial upfront investments in proprietary infrastructure. The adoption of these technologies contributes to the reduction of operational inefficiencies and the strengthening of SMEs' competitiveness, consolidating their role as a foundation for economic growth and job creation. Consequently, the adoption of cloud-based solutions, accompanied by appropriate management practices, emerges as a key factor in enhancing the efficiency, sustainability, and adaptability of small and medium-sized enterprises. [6].

2.7 Technological Maturity of SMEs in Argentina

Studies and reports conducted in recent years on digitalization and technological adoption in Argentina reveal a growing trend among SMEs toward investing in digital capabilities and migrating services to cloud computing environments [7]. Business surveys and market reports indicate that a significant proportion of these enterprises express intentions to invest in digitalization processes and the adoption of cloud solutions, primarily driven by the need to reduce costs, gain operational flexibility, and improve responsiveness to changes in the economic environment.

However, these advances coexist with persistent gaps in technological maturity and governance, which limit a homogeneous and sustainable adoption of cloud computing. These gaps are particularly evident in regions with lower technological density, such as the Northeast of Argentina (NEA) and the province of Misiones. In this regard, several studies agree that only a small number of companies are undertaking comprehensive digital transformation processes in the country; most exhibit partial and limited adoption of technologies, mainly concentrated in commercial and administrative areas.

Moreover, there is a marked disparity in levels of adoption and digital maturity depending on organizational size, with SMEs typically positioned in the early stages of digital transformation [8]. This situation increases the risk that cloud service adoption occurs without adequate practices of control, monitoring, and financial governance.

At the same time, the local cloud computing market is experiencing rapid growth, suggesting a favorable environment for greater adoption of these technologies, provided that barriers related to financing, availability of internal capabilities, and the implementation of appropriate management practices are effectively addressed. In this context, the consulting firm IMARC Group estimates that the cloud computing market in Argentina reached USD 4,620.46 million in 2024 and is projected to grow to USD 36,329.34 million by 2033, with a compound annual growth rate (CAGR) of 25.75% during the period 2025–2033 [9].

These national and regional insights constitute a key input for contextualizing the feasibility and requirements associated with the implementation of FinOps practices in SMEs in the province of Misiones, considering their level of technological maturity and the conditions of the local ecosystem [10].

2.8 Technology Service Providers in the NEA Region

While many SMEs in the NEA region rely on global cloud computing providers such as Amazon Web Services, Microsoft Azure, or Google Cloud Platform, there is also a set of local and regional providers that offer relevant advantages in the Argentine context, including geographical proximity, Spanish-language technical support, and local billing schemes.

Among these actors, companies such as Desarrollos NEA, based in the province of Corrientes, stand out for providing comprehensive consulting services, technological infrastructure, and software development with deployments in hybrid cloud environments. Likewise, Nextware, located in the province of Buenos Aires, specializes in cloud services, managed backups, and cybersecurity solutions specifically oriented toward Argentine SMEs [11].

At the national level, the state-owned company ARSAT offers, through its Data Center, public and private cloud services, S3-type storage, and hosting solutions for public institutions and companies in the NEA region, positioning itself as a key provider of technological infrastructure with nationwide reach [12].

Additionally, initiatives such as Silicon Misiones operate as technology hubs that facilitate access to cloud infrastructure for startups and local companies through agreements with major providers and support programs aimed at fostering the development of the technology sector [13].

2.9 Cloud Infrastructure and Services in the Province of Misiones: The Role of Marandú Comunicaciones S.E.

Marandú Comunicaciones S.E. is identified as the state-owned telecommunications company of the province of Misiones and constitutes a central actor in the development of the province's digital infrastructure. Its official website and public management reports disclose technical data related to connectivity, including the use of high-performance networking equipment, the evolution of data traffic—which recorded a 580% increase between 2021 and 2025—and the territorial expansion of its network nodes [14].

Various studies and institutional documents highlight Marandú's role in reducing the digital divide, particularly in rural and productive areas of the province. A report by the Instituto Nacional de Tecnología Agropecuaria

(INTA), published in 2022, analyzed connectivity in more than 300 rural locations across Argentina, using the province of Misiones as a pilot experience for the Rural Connectivity Plan. In that study, Marandú was identified as the implementer of fiber optic infrastructure in departments such as 25 de Mayo, underscoring its contribution to the development of so-called "digital agriculture" [15].

Technical reports from the Red Argentina de Instituciones de Ciencia y Tecnología also mention Marandú as an example of provincial public management that helps compensate for the lack of private investment in low-density population areas, strengthening digital inclusion and equitable access to technological services.

In terms of infrastructure, Marandú manages the "last mile" and the provincial fiber optic ring, which exceeds 1,450 km in length, and is integrated with the Federal Fiber Optic Network (REFEFO) operated by ARSAT. Its wholesale business model is characterized by not directly competing with small local Internet service providers, but rather offering them wholesale connectivity, thereby enabling the emergence of more than 60 technology SMEs in the province.

Specialized media outlets and knowledge economy portals, such as Economis and Canal 12, have documented Marandú's impact on strategic projects such as Silicon Misiones and the Posadas Industrial Park, highlighting the connectivity provided to more than 1,800 schools and hospitals throughout the province.

Among the most well-documented success cases is the implementation of a private cloud based on OpenStack technologies, developed in collaboration with Red Hat and Grupo Datco. This project enabled Marandú to reduce dependence on proprietary licenses and to consolidate a sovereign cloud that currently hosts more than 250 virtual private servers (VPS) used by public institutions and local technology companies, such as Silicon Misiones [16].

Research related to the knowledge economy and digital inequality also highlights Marandú's role as a provider of critical infrastructure—data center and connectivity—which has enabled the development of advanced technological initiatives such as FanIOT, as well as the centralization of data from rural schools and healthcare centers, contributing to the reduction of digital gaps in the province [17].

Currently, the company operates its own data center in the city of Posadas and an additional "shelter" in the Industrial Park. It is also integrated into the Capricornio Network, which allows it to have direct access to the Atlantic via Brazil, improving latency and the quality of its cloud services. Given its status as a state-owned enterprise, much of the available technical information on Marandú comes from public management reports, technological case studies, and sectoral analyses, rather than from traditional academic literature.

2.10 **Overview of SMEs in the NEA Region and the Province of Misiones**

The province of Misiones stands out within the NEA region for its business dynamism and sustained growth. According to data from the Provincial Institute of Statistics and Census (IPEC) and recent records, the business structure of Misiones is predominantly composed of microenterprises (62%), followed by small enterprises (22%), medium-sized enterprises (8%), and large companies (8%), the latter mainly linked to larger-scale industrial activities.

The provincial economy exhibits a strong orientation toward primary activities, although with an increasing process of industrialization at the source. The yerba mate and tea sectors position Misiones as a national leader in production, with 91 packaging companies and more than 9,000 producers by the end of 2025. The timber and sawmill sector is also highly relevant, considering that more than 98% of the country's wood-related companies are SMEs, and that Misiones concentrates one of the highest densities of sawmills and furniture factories in Argentina.

The metal-mechanical industry is mainly oriented toward supplying machinery for the forestry and agricultural sectors, while software development and IT services have experienced significant growth in recent years, driven by initiatives such as Silicon Misiones, the ICT Cluster, and the National University of Misiones. Currently, activities related to information technology represent one of the fastest-growing sectors in terms of new business creation in the province [18].

2.11 **Adoption and Impact of Cloud Computing in SMEs**

Despite the multiple benefits associated with cloud computing, various studies show that its level of adoption among small and medium-sized enterprises remains limited. In this context, it is essential to rely on specialized literature to identify and understand the main obstacles or barriers that hinder such adoption, as well as to develop strategies and interventions that facilitate an effective technological transition.

Analyzing these barriers contributes to evidence-based decision-making, providing relevant inputs for policymakers, researchers, and industry professionals interested in promoting the digital transformation of SMEs, particularly in regional contexts such as the province of Misiones.

The literature classifies barriers to cloud computing adoption into several interrelated dimensions:

- **Knowledge and awareness**: there is a lack of understanding of Industry 4.0 concepts and the transformative potential of cloud computing. The shortage of internal technical capabilities and specialized personnel hinders the transition to these technological environments.

- **Security and privacy**: one of the most significant barriers. Concerns about data breaches, denial-of-service (DoS) attacks, and the perceived loss of control over sensitive information generate distrust toward cloud services.
- **Economic and financial barriers**: include the perception of high costs associated with cloud services and difficulties in accessing financing or bank credit for technological projects.
- **Technological infrastructure**: issues related to connectivity quality, such as unstable connections, limited bandwidth, and high initial costs of communication equipment.
- **Organizational barriers**: lack of top management support and resistance to change. When executive teams do not perceive a clear return on investment or fail to understand the strategic benefits of cloud computing, budget allocation becomes constrained.
- **Legal and regulatory aspects**: regulatory gaps related to the use of digital technologies and data protection, particularly in developing countries.
- **Vendor lock-in**: concerns about dependency on a single cloud service provider, making future data and application portability more difficult.

Regarding strategies to overcome cloud adoption barriers, the literature consistently emphasizes that knowledge is a critical factor for the success of any technological adoption initiative. Successful cloud implementation requires IT staff to develop specific competencies to understand, operate, and effectively manage these environments.

The role of the State is fundamental from a systemic perspective, contributing through the implementation of educational and training programs aimed at reducing skill gaps in Information and Communication Technologies (ICT), as well as through policy interventions that support SMEs in their digitalization processes. These initiatives include training programs, consulting services, feasibility studies, and guidance throughout the technology adoption process.

It is also necessary to establish clear regulatory frameworks that support cloud computing adoption, providing legal certainty regarding data processing and protection.

At the organizational level, top management commitment is a determining factor. Executive teams must understand and support cloud adoption, promoting staff training (especially within IT departments) and highlighting the relative advantages of these technologies through practical examples and success cases. In this sense, the economic value of cloud computing must emerge from the balance between its technological potential and strategic managerial support.

Cloud service providers, in turn, can use the identified barriers as input to design strategies that improve technical support, accessibility, and security and privacy mechanisms. This involves offering service models tailored to different types of SMEs, effectively communicating the economic and operational benefits of cloud computing, and providing concrete evidence of the robustness of their security protocols in order to build trust and reduce perceived risk. The literature emphasizes that a solid adoption strategy should consider aspects related to infrastructure, legal and regulatory frameworks, the cloud economic ecosystem, human resources, and the financial implications of technological transformation [6].

2.12 Theoretical Frameworks for Analyzing Cloud Computing Adoption in SMEs

The adoption and impact of cloud computing in the context of SMEs can be examined through several well-established theoretical frameworks.

One of these is the Technology Acceptance Model (TAM), which provides a basis for understanding the factors that influence the intention of individuals and organizations to accept and use new technologies. The model is grounded in two core variables:

- Perceived usefulness: the degree to which an organization believes that using a technology will enhance its performance.
- Perceived ease of use: the degree to which the use of a technology is perceived as requiring minimal effort.

TAM, together with its extensions TAM2 (which incorporates social and workplace influences) and TAM3 (which adds aspects such as control and anxiety), is useful for explaining the motivations, perceptions, and attitudes of SMEs toward the adoption of cloud computing.

Another relevant framework is the Diffusion of Innovations (DOI) theory, which analyzes the social process through which innovations spread within a system. Unlike TAM, which focuses on individual perceptions, DOI emphasizes the characteristics of the innovation and their influence on the rate of adoption. According to this theory, adoption depends on five key attributes: relative advantage, compatibility, complexity, trialability, and observability. Applied to cloud computing, these attributes help explain why some SMEs adopt these technologies more rapidly than others. DOI also classifies adopters into five categories—innovators, early adopters, early majority, late majority, and laggards—which is particularly relevant for analyzing traditional productive contexts

where family-owned businesses and conservative organizational structures predominate. The innovation decision process unfolds in five stages: knowledge, persuasion, decision, implementation, and confirmation, providing a useful framework for analyzing the evolution of cloud adoption within SMEs.

Additionally, the Information Technology Business Benefits (ITB) framework offers a complementary perspective by enabling the evaluation of the impact of cloud computing beyond traditional return on investment (ROI). Under this approach, the value of IT should be analyzed in terms of its operational, tactical, and strategic contributions within the organization:

- Operational benefits: cost and time reductions in routine tasks, such as backup automation and the elimination of physical infrastructure maintenance.
- Tactical benefits: improvements in decision-making and medium-term resource management through access to real-time information and support for collaborative work.
- Strategic benefits: the creation of competitive advantages by enabling SMEs to access advanced technological capabilities and compete on equal footing with larger organizations.

From this perspective, cloud computing enhances SME competitiveness by increasing organizational agility, democratizing access to advanced technologies, and enabling human resources to focus on core business activities. Finally, the analysis of technological adoption barriers provides insight into challenges related to data security, resource availability, and system integration, offering a comprehensive framework for evaluating cloud adoption in SMEs [19].

3 Methodology and Methods

3.1 Literature Review

The literature review was conducted with the aim of identifying, systematizing, and analyzing the main theoretical and empirical contributions related to cloud computing, FinOps practices, and technological adoption processes in small and medium-sized enterprises (SMEs). A systematic review approach was adopted to ensure methodological rigor, traceability in source selection, and alignment with the research objectives.

In the first stage, search strings were defined by combining key terms associated with the conceptual axes of the study. The search strategy included Boolean expressions such as: (“FinOps” OR “Cloud Computing”) AND (“case study” OR “systematic review” OR “literature review” OR “rapid review” OR “mapping” OR “survey” OR “comparative analysis” OR “roadmap”) AND (“SMEs” OR “SMBs” OR “startup” OR “spin-off” OR “industry 4.0”). This formulation enabled the inclusion of conceptual studies, systematic reviews, and relevant comparative empirical analyses.

The bibliographic search was carried out using academic databases and specialized metasearch engines, such as Lens.org and Elicit, prioritizing peer-reviewed scientific publications. Inclusion criteria were defined as follows: (i) articles published between 2022 and 2025, to ensure topical relevance; (ii) studies focused on cloud computing, FinOps, or technological adoption in SMEs; and (iii) works presenting empirical evidence, structured reviews, or conceptual developments with clear methodological contributions. Duplicate or irrelevant articles, as well as those not meeting minimum academic quality standards, were excluded.

Subsequently, a systematic data extraction process was conducted, collecting information on study characteristics (type, context, sample), methodological approach, key findings, and conclusions. This stage enabled the identification of trends, convergences, and gaps in the recent literature, particularly regarding financial governance of cloud consumption and digital maturity in SMEs.

The review revealed that, although cloud computing has been widely studied from technical and strategic perspectives, the specific analysis of FinOps practices in SMEs—especially in regional or subnational contexts—remains incipient. The literature consistently indicates that cloud adoption without formal mechanisms for financial control and optimization can lead to inefficiencies, cost overruns, and reduced budget predictability.

Additionally, well-established theoretical frameworks for analyzing technological adoption in SMEs were identified. These include the Technology Acceptance Model (TAM) and its extensions (TAM2 and TAM3), which explain usage intention based on perceived usefulness and perceived ease of use; the Diffusion of Innovations (DOI) theory, which emphasizes attributes such as relative advantage, compatibility, and complexity; and the Information Technology Business Value approach, which expands evaluation beyond immediate financial return by incorporating operational, tactical, and strategic dimensions.

Complementarily, the literature highlights the importance of institutional and ecosystem factors in technological adoption processes, including the role of the State in establishing clear regulatory frameworks, training and funding policies, as well as the relevance of universities and technology providers in knowledge transfer and in reducing technical and cultural barriers.

Finally, the review process included a critical validation stage aimed at ensuring coherence between the identified findings and the research objectives. This analysis supported the development of the Feasibility Assessment Matrix proposed in this study, integrating technical, economic, organizational, and institutional dimensions under a systemic triple-helix approach.

3.2 Feasibility Assessment through Six Dimensions

In order to operationalize the developed theoretical framework and systematically address the research questions, a Feasibility Assessment Matrix for the Implementation of FinOps Practices in SMEs in the Province of Misiones is proposed (Table 1). This matrix constitutes an analytical tool that integrates technical, economic, organizational, and institutional dimensions, enabling a structured evaluation of FinOps adoption within the provincial productive and technological context.

The matrix is constructed based on the main contributions of the literature on cloud computing, FinOps, digital maturity, and technological adoption in SMEs, and is complemented by empirical evidence obtained through semi-structured interviews with key actors in the local ecosystem. These include executives and technical staff from SMEs, representatives from academia, policymakers, and stakeholders from the provincial provider of connectivity and cloud services, Marandú Comunicaciones S.E.

The adopted approach follows a triple-helix logic—industry, university, and government—recognizing that the feasibility of implementing FinOps practices does not depend solely on the internal capabilities of SMEs, but also on the availability of digital infrastructure, academic support, and the existence of favorable public policies and institutional environments. In this sense, the matrix enables a comparative analysis of the roles and contributions of each actor, as well as their interactions.

Furthermore, the matrix is articulated with the Provincial FinOps Maturity Model proposed in this study, allowing for the classification of the current level of cloud computing adoption and management in SMEs in Misiones, the identification of specific gaps and barriers, and the evaluation of potential evolution scenarios toward higher maturity levels, particularly within the framework of emerging technology parks and strategic productive clusters in the province.

Consequently, this tool not only fulfills a diagnostic function but also serves as an input for the formulation of strategic recommendations aimed at promoting a progressive, sustainable, and context-aware adoption of FinOps practices, aligned with the actual capabilities of the productive structure of Misiones and with regional development objectives.

3.3 Feasibility Assessment Matrix for the Implementation of FinOps Practices in SMEs in the Province of Misiones

Table 1 - Feasibility Assessment of FinOps Implementation in SMEs

	SMEs (Agro-industrial and Metal-mechanical)	**University**	**Government / Provincial Provider (Marandú)**
Technical Feasibility			
Digital maturity level	Predominantly low to intermediate maturity; traditional companies, some family-owned; incipient use of cloud solutions, mainly administrative and accounting SaaS.	High technical capacity; expertise in cloud computing, virtualization, and system architecture..	Mature infrastructure in terms of connectivity and data centers; experience in IaaS and private cloud..
Current cloud usage	Partial usage, without governance or cost control; low adoption of IaaS/PaaS.	Intensive use in teaching, research, and academic services.	Provision of IaaS and provincial cloud services with a focus on technological sovereignty.
Connectivity, latency, and availability requirements	Moderate needs; productive processes tolerant to moderate latency..	Variable requirements depending on research and simulation.	Compliance with adequate SLAs for regional SMEs; competitive latency through regional interconnection.
Most suitable cloud architectures	SaaS for management; basic IaaS for production systems; incipient PaaS.	IaaS and PaaS for research and technology transfer..	IaaS and managed services; potential expansion toward provincial PaaS.
Economic Feasibility			
Impact on operational costs	Potential reduction of OpEx through pay-as-you-go	IT budget optimization; improved use of shared	Economies of scale; optimization of public

	models and elimination of on-premise infrastructure.	resources; ICT staff can contribute to R&D.	spending on digital infrastructure.
ROI justification	Low current measurement level; absence of formal metrics.	Capability to model ROI and TCO; methodological support..	Ability to offer predictable and scalable cost models..
Economic barriers	Perceived costs, training requirements, financial uncertainty..	Lower barriers; institutional funding and R&D project generation.	High initial investment, amortized over the long term.
Provincial Infrastructure and Services			
Integration of cloud services with FinOps	Very low; lack of awareness of the FinOps framework.	High; potential role in methodological support; R&D projects to automate FinOps using AI and quantum computing.	High technical feasibility to integrate FinOps metrics into provincial services; highly skilled personnel in emerging technologies to implement optimized solutions.
Risk perception and trust	Greater trust in local/provincial providers than in hyperscale providers.	High technical confidence.	Key role as a sovereign and trusted provider.
Gaps between demand and supply	Simple needs; current supply is sufficient in most cases.	Advanced demand in certain projects.	Need to expand PaaS services and cloud financial management tools.
Organization, Culture, and FinOps Practices			
FinOps knowledge	Very low or non-existent.	High at both conceptual and technical levels; strong problem-solving capabilities.	Growing technical knowledge; FinOps as a strategic opportunity.
Organizational roles	No formal FinOps roles; informal functions.	Capacity to train FinOps and cloud management profiles.	Potential to incorporate cloud cost management roles.
Cultural barriers	Resistance to change; focus on immediate costs.	Innovation-oriented culture.	Culture oriented toward public service and sustainability.
Expected Impact and Outcomes			
Efficiency and scalability	Progressive improvement; reduction of operational inefficiencies.	Optimization of academic and research resources.	Efficient use of shared provincial infrastructure.
Indirect benefits	Increased competitiveness, production traceability, and export capacity.	Strengthening of the university's role in technology transfer.	Development of the provincial digital ecosystem.
Ecosystems and Public Policies			
Required public policies	Subsidies, training, and tax incentives.	Training and certification programs, along with strategic planning for the knowledge economy.	Regulatory framework, funding, and strategic planning for the knowledge economy.

Role of universities	Beneficiaries of technology transfer.	Central actor in coordination, education, and applied research.	Articulation with the educational and productive system.
Public-private partnerships	Necessary to reduce risks.	Facilitators of innovation and adoption.	Key to sustaining technology parks and cloud services.

3.4 Linkage with Technology Parks in Misiones

Technology parks under development or consolidation in Misiones function as accelerators of FinOps maturity, as they concentrate high-quality connectivity, access to provincial cloud infrastructure, university support, and public–private collaboration.

The specific role of technology parks focuses on:

- Reducing technical barriers: connectivity, latency, infrastructure
- Reducing economic barriers: shared costs and common services
- Reducing cultural barriers: fostering innovation through proximity
- Enabling collective FinOps: shared metrics and replicable models

Within the proposed framework, Marandú acts as a sovereign cloud platform, the National University of Misiones serves as a knowledge provider and training institution, and technology parks operate as controlled environments for advancing FinOps maturity.

3.5 Synthetic FinOps Maturity Model and Alignment with Technology Parks

Table 2 presents the FinOps maturity model and its alignment with technology parks. Specifically, the Infrastructure row describes where and how the technological resources used by SMEs are hosted and managed, as well as the degree of control, predictability, and financial governance enabled by each scenario. It does not refer solely to technology, but also to the organizational capacity to manage cloud-related costs.

In this sense, the Infrastructure dimension reflects a transition from rigid and isolated infrastructure to governed and scalable cloud infrastructure. It illustrates how technology parks and provincial cloud environments provide a realistic pathway for SMEs to adopt FinOps progressively, with lower risk and greater economic impact.

Table 2 – FinOps Maturity Model and Alignment with Technology Parks

Dimension	Level 0 – Initial	Level 1 – Reactive Cloud	Level 2 – Initial FinOps	Level 3 – Integrated FinOps
Cloud Usage	None	Isolated SaaS	Provincial / hybrid IaaS	SaaS + IaaS + PaaS
Cost Visibility	None	Low	Cost Visibility	High (FinOps metrics)
FinOps Practices	Non-existent	Informal	Partial	Integrated
Infraestructure	On-premise	External cloud	Provincial cloud + support	Scalable provincial cloud
Role of the University	None	Informal training	Technical assistance	Technology transfer
Role of the State	Regulator	Facilitator	Infrastructure provider	Ecosystem orchestrator
Alignment with Technology Parks	No	Low	High	Very high

The levels presented in Table 2 are described below:

- **Level 0 – Initial (On-premise):** The technological infrastructure is primarily based on proprietary or local servers, hosted either on company premises or through traditional hosting providers without elasticity. IT investments are rigid and involve high upfront costs (CapEx), with limited scalability and no pay-as-you-go models. From a FinOps perspective, there is no financial management of cloud resources, as there is neither measured consumption nor granular cost visibility. This scenario is typical of traditional SMEs operating outside technology parks.

- **Level 1 – Reactive Cloud (External cloud):** The organization begins to use cloud services from external providers (generally hyperscalers or commercial SaaS) in a reactive and unplanned manner. Usage is typically limited to specific needs such as email, billing, and backups, without a defined architecture or cost-control policies. Financial visibility is low, and costs are treated as general expenses, making the application of FinOps practices difficult. Vendor dependency and the risk of overconsumption are high.
- **Level 2 – Initial FinOps (Provincial cloud with support):** The infrastructure relies on provincial cloud services provided by actors such as Marandú Comunicaciones S.E., often combined with hybrid models. This environment offers improved cost predictability, lower regional latency, and local technical support, facilitating the initial adoption of FinOps practices. Institutional support, universities, and technology parks enable the introduction of basic consumption metrics, resource tagging, and financial reporting, reducing technical and economic barriers.
- **Level 3 – Integrated FinOps (Scalable provincial cloud):** This represents an advanced maturity scenario in which provincial cloud infrastructure is scalable, governed, and fully integrated into organizational processes. SMEs use combinations of SaaS, IaaS, and PaaS, supported by automation, cost-control policies, and advanced FinOps metrics. Technology parks function as controlled environments for experimentation and scaling, enabling replicable models and collective FinOps schemes. The State acts as an ecosystem orchestrator, ensuring sustainability and strategic alignment.

4 Conclusions

The analysis conducted leads to the conclusion that the implementation of FinOps practices in SMEs in the province of Misiones is viable from a technical, economic, and institutional perspective, provided that a progressive, context-aware approach aligned with the current levels of digital maturity of the regional productive structure is adopted. While cloud computing offers clear advantages in terms of flexibility and reduced upfront investment, its adoption without financial governance mechanisms limits value realization and exposes organizations to inefficiencies and cost unpredictability.

The findings indicate that the main barriers to FinOps adoption are not related to infrastructure availability, but rather to gaps in knowledge, internal capabilities, and organizational culture within SMEs. These limitations can be effectively addressed through training strategies, technical support, and knowledge transfer, particularly when actors from the scientific-technological ecosystem and the public sector are actively involved.

In this regard, the study highlights the strategic role of the provincial ecosystem, particularly the articulation among SMEs, universities, the State, and technology parks. The sovereign cloud infrastructure provided by Marandú Comunicaciones S.E., together with the educational capabilities of the university system and the collaborative environments of technology parks, creates a favorable scenario to accelerate FinOps maturity and reduce the risks associated with technological adoption.

The proposed Feasibility Assessment Matrix and FinOps Maturity Model constitute relevant methodological contributions, offering both diagnostic and prospective tools to guide strategic decision-making. Overall, these results reinforce the notion that FinOps should be understood not merely as a cost-control mechanism, but as a strategic instrument for sustainability, competitiveness, and digital transformation of SMEs, particularly in regional economies and developing country contexts.

References

[1] «NIST (National Institute of Standards and Technology)», Definición Cloud Computing. [En línea]. Disponible en: https://csrc.nist.gov/pubs/sp/800/145/final

[2]«Google Cloud», ¿Qué es Cloud FinOps? [En línea]. Disponible en: https://cloud.google.com/learn/what-is-finops

[3] D. E. J. Thorne, «A Systematic Review and Comparative Analysis of Amazon Web Services (AWS) Infrastructure: Architecture, Economic Implications, and Future Research Trajectories».

[4] M. Smendowski y P. Nawrocki, «Optimizing multi-time series forecasting for enhanced cloud resource utilization based on machine learning», *Knowl.-Based Syst.*, vol. 304, p. 112489, nov. 2024, doi: 10.1016/j.knosys.2024.112489.

[5] H. Kumar, «Optimizing cloud cost management for large-scale fintech operations», *Int. J. Multidiscip. Res. Growth Eval.*, vol. 6, n.º 1, pp. 2131-2133, 2025, doi: 10.54660/.IJMRGE.2025.5.1.2131-2133.

[6] Z. Alrababah, «BARRIERS TO CLOUD COMPUTING ADOPTION AMONG SMEs IN THE MIDDLE EAST: A SYSTEMATIC REVIEW», . *Vol.*, n.º 17, 2023.

[7] «Inversor LATAM», El 88% de las PyMEs argentinas invertirá en su digitalización este 2024. [En línea]. Disponible en: https://inversorlatam.com/?s=El+88%25+de+las+PyMEs+argentinas+invertir%C3%A1+en+su+digitalizaci%C3%B3n+este+2024

[8] F. Ribas, M. B. Laberta Lazcano, y D. Sánchez Abrego, «Transformación digital en empresas argentinas: una revisión de alcance», *Costos Gest.*, sep. 2024, doi: 10.56563/costosygestion.107.e2.

[9] «imarc - Transforming Ideas into Impact», Argentina Cloud Computing Market Size, Share, Trends and Forecast by Service, Workload, Deployment Mode, Organization Size, Vertical, and Region, 2025-2033. [En línea]. Disponible en: https://www.imarcgroup.com/

[10] «ARGENCON», Estudios económicos sobre la industria del conocimiento. [En línea]. Disponible en: https://www.argencon.org/argenconomics-2025/

[11] «Empresa servicios -NEXTWARE». [En línea]. Disponible en: https://nextware.com.ar
[12] «ARSAT», Servicios Ofrecidos. [En línea]. Disponible en: https://www.arsat.com.ar/nube-y-datacenter/
[13] «Silicon Misiones - hub de innovación tecnológico», Eonomia del Conocimiento - Misiones.
[14] «Marandú Comunicaciones SE», Conectividad. [En línea]. Disponible en: https://marandu.com.ar/conectividad/#:~:text=Actualmente%20MCSE%20opera%20y%20mantiene%20una%20infraestructura,Virtual%20Private%20Server%20(VPS)%20o%20servidores%20virtuales
[15] «ECONOMIS», Misiones, prueba piloto del plan de conectividad rural. [En línea]. Disponible en: https://economis.com.ar/misiones-prueba-piloto-del-plan-de-conectividad-rural/
[16] «Grupo DATCO», Infraestructura para la implementación de OpenStack en Marandu Comunicaciones. [En línea]. Disponible en: https://www.grupodatco.com/infraestructura-para-implementacion-openstack-en-marandu/
[17] «IPEC_Misiones_Economia-del-conocimiento-2014-2024».
[18] «IPEC- Instituto Provinical de Estadisticas y Censos». [En línea]. Disponible en: https://www.ipec.misiones.gov.ar/
[19] R. Oriza y Maulidar, «Adoption and Impact of Cloud Computing in Small and Medium Enterprises A Systematic Review», *J. Inform. Educ. Manag. JIEM*, vol. 6, n.º 2, pp. 8-15, jun. 2024, doi: 10.61992/jiem.v6i2.79.

Lessons Learned from Developing a Blockchain-based Traceability Solution for Medicinal Cannabis Production in Argentina*

Facundo Paterno, Alejandra Cechich[0000−0003−4804−6270], and Alan De Renzis

GIISCO Research Group
Departamento de Ingeniería de Sistemas - Facultad de Informática
Universidad Nacional del Comahue
Neuquen, Argentina
facundo.paterno1@gmail.com,
{alejandra.cechich,alanderenzis}@fi.uncoma.edu.ar

Abstract. Cannabis cultivation is a process that, although it follows a well-defined life cycle, always encounters unique details in each crop. Throughout the different stages, from germination to harvest, it is essential to adjust the light, water, and nutrient conditions to suit the specific needs of each plant. Establishing secure, transparent, and structured access to operational and medicinal information about the final product promotes not only consumer confidence but also regulatory compliance and the optimization of production processes. Currently, support systems exist at the national and local levels with a variety of services; however, very few cover all phases of cultivation. Besides, the most comprehensive systems reflect practices adopted from large-scale companies, making it difficult to implement traceability in small and medium-sized organizations. Our system attempts to address some of these shortcomings by providing complete but easy-to-adopt coverage of the phases. This paper discusses lessons learned during its development and implementation, including design decisions that might be considered when blockchain technology is deployed in practice.

Keywords: Blockchain · Traceability · Medicinal Cannabis.

1 Introduction

Cannabis is a medicinal plant native to Asia and is currently cultivated and used worldwide. It has become popular today due to its therapeutic properties and psychoactive substances. Medicinal cannabis production processes must control the quality of the production system from the initial to the final stage, following government regulations. To do so, the World Health Organization

* Supported by Universidad Nacional del Comahue, Project 04/F019-Modelado de Variedad en Sistemas Big Data

(WHO) has recommended that the herbal industry comply with Good Manufacturing Practices (GMP)[1]. The standard ensures that the production system achieves pharmaceutical quality and minimizes production risks; for example, cross-contamination with pathogens, chemicals, pesticides, and physical contaminants. Furthermore, it adds *traceability*, according to which the company must guarantee the process by tracing the material or product at least one step backward and one step forward. Then, a traceability system is one of the best solutions to guarantee the quality of cannabis medicines and their proper use (not abuse or misuse) [1].

Laws and regulations regarding cannabis differ from country to country; and most of them still lack clear tracking data and control plans. In Argentina, Law 27.669[2] regulates the production, national marketing, and export of the cannabis plant, its seeds, and its derivative products for human, veterinary, nutritional, cosmetic, industrial, plant health and fertility use, as well as uses arising from scientific research and technological and industrial development. This law extends and regulates the National Law 27.350, Medical and Scientific Research on the Medicinal Use of the Cannabis Plant and its Derivatives[3]. Furthermore, the Regulatory Agency for the Hemp and Medicinal Cannabis Industry (ARICCAME)[4] is the competent body for regulating, controlling, and granting administrative authorizations for the use of cannabis seeds, cannabis itself, and its derivative products.

On the other hand, blockchain technology has exthensively supported traceability in several domains. Blockchain platforms facilitate real-time product tracking, guarantee authenticity, and ensure strict compliance with regulatory requirements for management purposes from seed to sale [2, 8, 9]. For example, several systems offer traceability of cannabis production for various industrial uses, including medicinal [7]. However, many are tailored to the regulations of their target market or are too broad to properly support small and medium-sized enterprises (SMEs), which are the most frequent type in Argentina; for example, Cannabud.ai[5], a cannabis cultivation management software for the European Union, or BioTrack[6] in the US market. In the Argentinian market, some companies have begun offering traceability services; however, in some cases the offerings do not cover all phases of cultivation, and in other cases they are geared towards large producers, leaving small and medium-sized enterprises without accessible support. Therefore, the contributions of this paper are: (1) the introduction of *Araucann* - a blockchain-based system for traceability during cannabis production in SMEs; and (2) lessons learned during its development and implemen-

[1] https://www.who.int/teams/health-product-policy-and-standards/standards-and-specifications/norms-and-standards/gmp
[2] https://www.argentina.gob.ar/justicia/derechofacil/leysimple/salud/desarrollo-de-la-industria-del-cannabis-medicinal-y-el-canamo
[3] https://www.argentina.gob.ar/normativa/nacional/ley-27350-273801
[4] https://www.argentina.gob.ar/ariccame
[5] https://cannabud.ai/
[6] https://biotrack.com/

tation, so they can be considered when developing and/or selecting supporting software in this domain.

This article is organized as follows. The next section briefly discusses background on cultivating cannabis and software platforms commonly used in Argentina. Then, *Araucann* is introduced as a summary of challenges, design decisions and lessons learned, which are discussed in terms of the development and transferring stages. Conclusions and future work are addressed at the end.

2 Background and related tools

2.1 Cannabis Cultivation Process: From Seed to Final Product

Cannabis cultivation is a process that, although it follows a well-defined life cycle, always presents unique details in each grow [3, 5, 6]. The more thorough the production process is in terms of the details of the life cycle, the better the results will be [4]. The cultivation stages are:

- *Germination [2 to 10 days].* Germination marks the beginning of the cannabis life cycle, when the seed is activated and begins to develop its primary root. To germinate successfully, the seed needs constant moisture, temperatures between 21°C and 32°C, and darkness. Finally, the cotyledons appear and, at this point, the seedling is ready to be transplanted to its final growing medium.
- *Seedling Stage [2 to 3 weeks].*
  The seedling is the young stage of the cannabis plant. In this stage, the roots begin to spread in the soil while the aerial part of the plant grows slowly. The first true leaves appear with their characteristic serrated edges. As more leaves develop and the root system expands, the plant strengthens and is ready to enter the vegetative stage.
- *Vegetative Phase [3 to 15 weeks].* In this phase, the cannabis plant experiences accelerated growth, developing strong stems, broad leaves, and an extensive root system. During this period, the plant does not produce flowers but focuses on accumulating energy for flowering. A strong root system allows for efficient nutrient absorption, which directly impacts the size and future yield of buds[7].
- *Pre-flowering Phase* It occurs when the plant detects a change in the photoperiod, going from 18 hours of light to 12 hours in indoor grows or when the days begin to shorten outdoors. At this stage, the plant experiences considerable stretching, potentially doubling or even tripling its height.
- *Flowering Phase [7 to 14 weeks].* It is the stage in which the cannabis plant produces buds and maximizes the production of cannabinoids and resin. During flowering the demand for phosphorus and potassium increases, while

[7] A bud is the part of the cannabis plant that is harvested, dried, and cured for consumption. Simply put, it is the "flower" or smokable part of the marijuana plant.

nitrogen decreases. The buds begin to develop, and the flowers produce trichomes[8]. A controlled environment is required, where humidity should be kept low to prevent mold from growing on the buds. The optimal harvest time is determined by observing the trichomes with a magnifying glass or digital microscope.

– *Harvesting, Drying, and Curing.* Once the buds are ready, the harvesting process begins. The plants are cut and hung upside down in a dark, well-ventilated place to dry (Figure 1 (a)), a process that takes between 7 and 14 days. After drying, the curing process begins (Figure 1 (b)). The buds are stored in airtight jars and opened periodically to release moisture. This process can last from two weeks to several months, improving the flavor, aroma, and effects of the cannabis. The figure shows the curing of cannabis with a dome to control humidity and a sensor to validate the correct parameters.

Fig. 1. Drying room (a) - Curing (b)

2.2 Software platforms commonly used in Argentina

To analyze digital traceability of platforms used by the argentinean cannabis industry, we took into consideration technical information published by each provider, its official descriptions, and current regulations, such as the Argentine Law 27.669 and international standards such as GACP (Good Agricultural and Collection Practices)[9], GMP (Good Manufacturing Practices), and GS1[10]. Each platform is then summarized and the reasons for highlighting certain functions are explained, based on its design, technology, and main objective. A comparative summary can be seen in Table 1, where the assignment of functionalities was based on official sources and technical documentation provided by each vendor.

[8] In cannabis, trichomes are tiny resinous glands that primarily cover the flowers and, to a lesser extent, nearby leaves.
[9] https://www.who.int/publications/i/item/9241546271
[10] https://www.gs1.org/standards

Table 1. Cannabis platforms commonly used in Argentina vs. Araucann

Feature	**Trazacann**	**Cannect**	**GrowerIQ**	**Araucann**
Origin	Argentina	Argentina	Canada	Argentina
Main focus	Regulatory compliance	Agronomic optimization	Comprehensive business management	Comprehensive traceability & regulatory compliance
Traceability	Yes (GS1 + blockchain)	Yes (crop)	Yes (seed to sale)	Yes (end-to-end, blockchain, hash anchoring)
Crop Management	Parcial	Yes	Yes	Yes
Manufacturing Management	No	No	Yes	Yes
Quality Management	No	Partially	Yes	Partially
Commercial Management	Yes	Partially	Yes	Yes
Regulation Support	High	High	High	High
Key Technology	Blockchain, GS1	IoT, IA	ERP, traceability, compliance	ERP, blockchain

Trazacann [11]. It is an Argentinian digital platform designed to register, validate, and document each stage of the cannabis production process. It was developed by KYAS, INTA Patagonia Norte, Ciencia Sativa, and GS1 Argentina, and is intended for medicinal cannabis and industrial hemp projects. Its structure is based on GS1 standards and blockchain technology, ensuring reliable, transparent, and immutable traceability from planting to sale. While it does not offer a complete ERP system, Trazacann includes basic business functionalities such as sales reports and inventory control. It also features a mobile app with offline support, event geolocation, task alerts, and compliance with local regulations. Its primary focus is on regulatory compliance, prioritizing data standardization and legality.

Cannect[12]. Developed by IT Patagonia, it offers an IoT-based platform for monitoring environmental, water, and soil conditions in real time. It is specifically designed for the digitalization of agricultural environments, integrating sensors with artificial intelligence (AI) analysis. This combination allows for optimized agronomic decisions, disease prevention, and the assurance of stable production conditions. Its functionality is focused on the cultivation stage, offering traceability during this phase, but without integration with ERP, manufacturing, or quality modules. It lacks point-of-sale (POS) functionalities and advanced business tools, although its differentiating value lies in the accuracy of its environmental monitoring and its support for agronomic decisions.

[11] Source: Trazacann Official Website https://www.trazacann.com.ar/
[12] Source: Cannect Official Website https://www.cannect.com.ar

GrowerIQ[13]. It is a Canadian-based platform offering a comprehensive ERP system for the cannabis industry, used in several countries, including Argentina. It includes integrated modules for cultivation, manufacturing, quality (compliant with GMP standards), sales, CRM, and inventory. Its traceability covers the entire chain from seed to final sale. The tool complies with international standards such as GACP, GMP, and GPP, and is designed to facilitate audits, export, and regulatory documentation.

As we can see, the diversity of existing solutions reflects the different profiles of operators within the cannabis industry. Some tools, such as GrowerIQ, are designed to meet high international standards throughout the production chain, while others, such as Cannect or Trazacann, specialize in agronomic traceability and compliance with local regulations. The last column in Table 1 shows our own proposal, which overcomes some of the current platform's limitations through several decisions discussed in the following section.

3 Araucann: A blockchain development and its lessons learned

3.1 The Business Model

Our model for the overall process for cannabis production and traceability (Figure 2)[14] involves three main services requested by different roles in the production chain: (1) the *Seed Supply* service, which serves the *Supply Manager* and is performed by the *Purchase Seeds and Cuttings* process; (2) the *Biomass Cultivation* service, which serves the *Cultivation Manager* and is performed by the *Grow Biomass* process; and (3) the *Product Generation* service, which serves the *Production Manager* and is performed by the *Generate Products* process.

Within the general production process, one of the first fundamental stages is the acquisition of seeds and cuttings. The model reflects how this activity involves receiving the purchased batches (*Arrival of Seed/Cutter Batches* event) along with the purchase information, represented by the business object *Purchased Batch*. As a result, the availability of seeds or cuttings in stock is formalized (business object *Seeds and Cuttings* in Figure 2).

Then, when a market need arises for the production of cannabis derivatives (*Market need* event), the *Grow Biomass* process gathers information on the availability of seeds and cuttings, carries out the cultivation process, and records the harvesting information (*Harvested biomass* business object). This information is then fed into the *Generate Products* process, which, upon detecting a production requirement (*Pending biomass* event), generates the derived and packaged products (recorded in the *Packaged products* business object).

⇒ *Challenge*: Modeling traceable inputs from seed to sale.

[13] Source: GrowerIQ Official Website https://groweriq.ca/

[14] In ArchiMate notation: https://www.opengroup.org/archimate-forum/archimate-overview

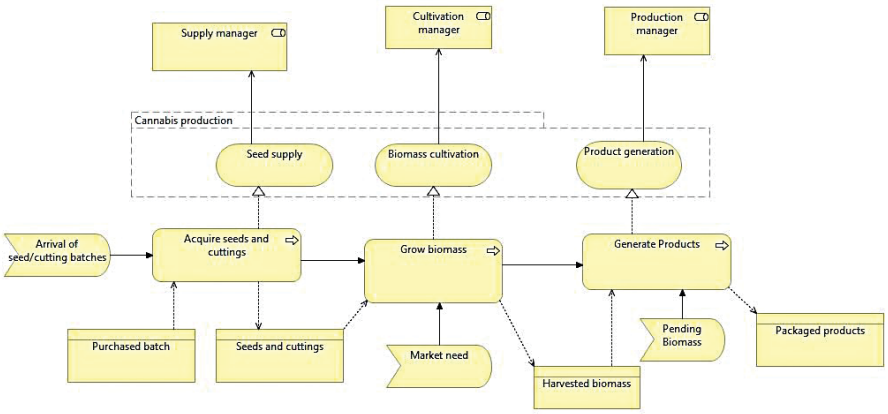

Fig. 2. General Cannabis Production Process

⇒ *Design decisions*: One of the first decisions was to ensure that the process begins with traceable inputs. This means that every seed or cutting that enters the system must be registered, validated, and has all its documentation in order. In this way, it is guaranteed that everything cultivated has a controlled origin, which is key to complying with regulations and maintaining transparency.

When a market need arises, it translates into a cultivation project, which establishes what will be cultivated, how much, and with what objectives. This begins an ordered cycle that includes the creation of plant batches and the monitoring of the cultivation throughout its different stages: germination, vegetative phase, flowering, and harvesting. Another key decision was to separate the biomass processing stage as an intermediate step. Before becoming final products, the biomass is transformed into raw material (such as dried flowers or hemp grain), and only then it is processed according to the desired product type: oils, edibles, fibers, and others.

⇒ *Lessons learned*: Throughout the entire process, data is captured manually and automatically using sensors. This information, from various sources, must be analyzed differently to ensure that the history of each batch is complete and consistent for future audits. Separating the processing stages facilitated the identification and recording of information for tracking. In this way, the final packaged products can ensure quality and regulatory compliance at every step.

3.2 The Application Model Exemplified

Due to space limitations, this section only describes the components of the application layer (light blue in Figure 3) of the *Grow Biomass* process; where software modules support the corresponding busines process (yellow).

The process begins when the cultivation manager defines a new project. Through the *Create Project* component, key data such as the crop objective,

quantity to be produced, estimated cost, and expected harvest date are entered. All of this is done through the *Crop Project Interface*, and an application object — the *Crop Project* — is automatically created, ready to begin fieldwork. This object constitutes the starting point for production traceability, as it consolidates the technical and economic planning of the cycle. Every subsequent action, such as creating batches, recording tasks, or harvesting, is hierarchically linked to this project, allowing for complete backward traceability.

With this project as a foundation, the next step is to create the plant batches. This step is essential for systematically tracking the growth and care of each group of plants. From the *Manage Batches* component, batches with unique identifiers can be created using the *Batches* interface. In this way, the *Plant Batch object* is built, which will accompany each group throughout its productive life. The *Plant Batch object* groups the biological units under a single identifier, facilitating operational and technical monitoring. It represents the basic unit of control in the field, on which the tasks and measurements that support daily traceability are recorded.

Then comes the time for daily work, where the application transforms into an active monitoring tool. Every watering, pruning, fertilization, or other action can be easily recorded from the *Manage Growth* component using the *Record Crop Data* function. Furthermore, the system will integrate with sensors to automatically record variables such as humidity, temperature, and light levels through the *Record IoT Environment Data* component. Both manual and automatic data sources primarily feed the *Operational Logs* object, which consolidates the technical and temporal information of all tasks performed. The *Environmental Logs* object, while complementary, plays a supporting role in documenting IoT conditions and reinforcing contextual traceability.

When the plants reach their optimal maturity point, the system enables the *Biomass harvesting* service. From the *Manage Harvest* component, the team can upload the harvest results using the *Input harvest data* and *Harvesting recording* functions, generating *Crop Reports*.

Then, from the *Process Biomass* component, the collected material (flowers, leaves, stems, etc.) is classified into individual units, formally recorded as a *Harvested Biomass* data object, which is one of the system's traceable pillars: it stores the final information of the plants from each batch, including its weight, type, and exact origin. Processing records and future raw materials are derived from it. In addition, every processing task is stored as a *Biomass processing log*.

The key aspect of this workflow is that each object created —whether a batch, a record, a biomass unit, or a project — is linked within the system through the *Blockchain access* interaction connected to the *Manage Blockchain* component. Thus, any audit, quality control, or subsequent analysis can be performed with precision, knowing at all times what was cultivated, how, when, and under what environmental conditions.

⇒ *Challenge:* Ensure consistency between the defined traceability model and its implementation.

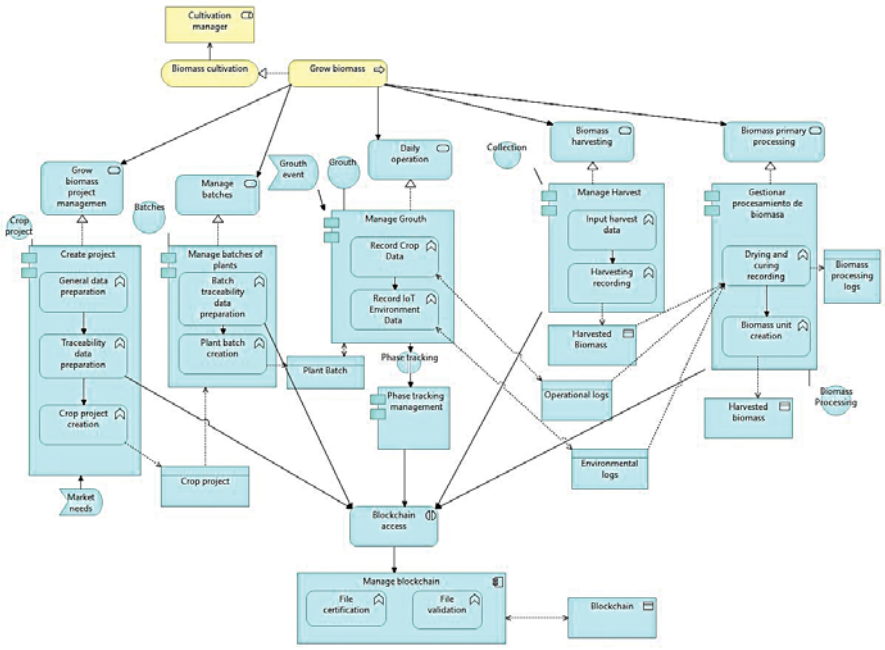

Fig. 3. Application layer that supports the process *Grow Biomass*

⇒ *Design Decisions:* Traceability at Araucann is built upon the articulation of three modeling levels: the business layer, the application layer, and the data model. These levels are complementary, allowing every action performed in the field or on the production floor to be reflected in the system as traceable and queryable data.

⇒ *Lessons Learned:* The value of the model lies not only in each layer individually, but also in how they are linked. The business layer establishes what needs to be done; the application layer defines how each step is digitally implemented; and the data model ensures that the information is consistently recorded and can be consulted or audited. Thus, a packaged bottle can be traced back to the biomass of origin and, in turn, to the plant and seed from which it came. This cross-functional integration is what makes the system a robust traceability tool, aligned with regulatory requirements and offering the added value of trust and transparency throughout the entire production chain.

However, a complete correspondence between system levels should also be carried out considering the infrastructure, as discussed in the following section.

4 Infrastructure and development

Figure 4 shows the technologies selected after careful consideration of implementation in initial or small-scale projects, which may be suitable for the needs

of medium-sized companies. Following, we discuss one challenge that guide the entire development process.

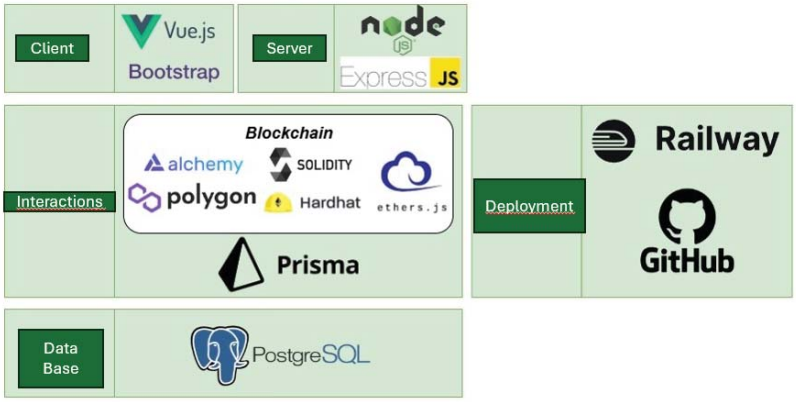

Fig. 4. Selected technologies

⇒ *Challenge:* Select a technological architecture that supports the execution of the defined application services, ensuring that the actions performed in the business layer (from acquiring seeds to packaging final products) are backed by robust data management tools, efficient user interfaces, and advanced traceability mechanisms based on public blockchain.

4.1 Selecting the blockchain

⇒ *Design Decision:* In the first stage of the project, we evaluated integrating the Argentine Federal Blockchain (BFA)[15] into the traceability system. The decision made perfect sense from both an academic and institutional perspective: it is a national network that seeks to offer a framework of trust for various projects. However, as the project progressed, technical and infrastructure limitations emerged that, in practice, made its use unfeasible in the context of an inicial project. For example, to participate in BFA, whether with a transactional or sealing node, it is mandatory to set up your own node. Unlike other chains, there are no public RPCs (Remote Procedure Call) available; therefore the only option was to manage a node independently.

The BFA was an excellent academic opportunity and had significant added value locally, but it was neither viable nor scalable in our situation. It required the organization to set an infrastructure administrator and assume fixed costs, when the necessary use at the biggining was limited to anchoring traceability records to the blockchain.

[15] https://bfa.ar/

For these reasons, it was decided to migrate to Polygon[16] , a public blockchain compatible with EVM (Ethereum Virtual Machine)[17]. The main advantage is that Polygon doesn't require setting up your own node by connecting through free, public RPC providers. It also offers significantly faster confirmation times and extremely low transaction costs.

⇒ *Lesson Learned:* Choosing the blockchain on which to base the system is not an easy task. Several factors are involved, from its nature to the associated costs of its use and maintenance. For instance, a public option, such as the BFA, proved unfeasible after analyzing the costs of operating and maintaining the system on that network. Carefully balancing attributes like performance versus cost is part of any development; however, starting with blockchain technology has the valuable characteristic of allowing you to begin with a version that isn't fully functional, with easy migration and/or expansion later (even with cloud-based offerings). Trade-offs must be carefully considered in each case.

4.2 Implementing the blockchain

⇒ *Design Decision:* After analyzing various technological integration alternatives, a minimalist anchoring model on blockchain was ultimately chosen.

⇒ *Lesson Learned:* This approach does not limit the technology's potential but rather constitutes an initial adoption strategy: instead of uploading complete crop information to the network, which would be costly and inefficient, blockchain was used to guarantee the immutability and verifiability of critical events in the production process. This provides a reliable external traceability mechanism while maintaining the privacy of internal data and the possibility of scaling to more complex models in the future.

⇒ *Design Decision:* For implementation, a single, generic smart contract was developed in Solidity[18]. Instead of having multiple separate contracts (one for each production stage), a single contract was designed that emits an event called Anchored. This event includes the following basic fields: the generated hash, the process stage (germination, vegetative, flowering, harvest, etc.), the objectId that references the system's internal registry, the address of the account that signed the transaction, and the block timestamp. This ensures simplicity in maintenance and future scalability. It was also decided that not all system events would be anchored to the blockchain, but only those considered critical to ensuring external traceability and independent verification. In this regard, representative milestones were established at different stages of the process.

⇒ *Lesson Learned:* It is necessary to carefully analyze and define what the milestones will be in traceability, as they are associated with the key stages supported by the blockchain. A greater number of milestones than necessary can reflect a lack of understanding of business processes and, consequently, a model

[16] https://polygon.technology/
[17] https://ethereum.org/what-is-ethereum/
[18] https://soliditylang.org/

that does not meet the needs of regulatory or legal frameworks. In this way, each key stage is supported by an immutable anchor that acts as an external certification, while the rest of the operational data remains in the internal database for detailed queries and analysis.

5 Transfer to small and medium-sized industry

A first prototype, which includes all design decisions previously discussed, is now available at https://www.araucann.com/bienvenida.

Based on our experience after a year of releasing the tool, we could summarize challenges that stem from our design decisions and extend beyond general industry issues (see Table 2). For example, during industry transfer we found the following.

⇒ *Challenge:* Demonstrate the usefulness of the tool in terms of its usability (ease of understanding based on the services it provides and the organization of its business processes); acceptance (in terms of performance and capabilities); and value (based on attributes such as traceability, security, integrity, privacy, etc.).

⇒ *Transferring decisions:* A gradual rollout to user groups new to the industry, with personalized support, allowed for a closer assessment of the tool's acceptance and value. Simultaneously, it enabled adjustments to the tool's user guides, identifying both general and specific needs.

On the other hand, maintaining a history of events and usage allowed for recording the differences between the services offered and those requested in each pilot case. This would lead to potential improvements in future versions of the tool, as well as a more user-centric approach.

⇒ *Lessons learned:* User guides tailored to the diverse needs facilitate the consolidation of services around defined uses, with potential for customization and reusability. A phased transfer, identifying potential gaps along the way, requires implementing deployment organizational capabilities, with personnel possessing strong communication skills, business acumen, and expertise with the tool. Such professionals are not easy to find, whether due to a lack of widespread knowledge in blockchain implementations or other factors, such as general traceability knowledge and a specific understanding of the cannabis industry. Maintaining a record of events and usage for future analysis also requires data analytics capabilities. All these considerations mean that the transfer to the industry demands a diverse skill set, encompassing expertise in the implemented tool, analytical capabilities, and customer relationship management adjusted to this particular domain.

6 Conclusion and Future Work

In this paper, we have introduced a blockchain-based system for traceability of cannabis production in the argentinean market; and we have focused on challenges, design decisions and lessons learned in the road, as summarized in Table

Table 2. Summary of challenges and lessons

Challenge	**Lessons**
Modeling traceable inputs from seed to sale	– Separating the processing stages facilitated the identification and recording of information for tracking. In this way, the final packaged products can ensure quality and regulatory compliance at every step.
Ensure consistency between the defined traceability model and its implementation	– The value of the model lies not only in each layer individually, but also in how they are linked. This cross-functional integration is what makes the system a robust traceability tool, aligned with regulatory requirements and offering the added value of trust and transparency throughout the entire production chain.
Select a technological architecture that supports the execution of the defined application services, ensuring that the actions performed in the business layer are backed by robust data management tools, efficient user interfaces, and advanced traceability mechanisms based on public blockchain.	– Choosing the blockchain on which to base the system is not an easy task. Several factors are involved, from its nature to the associated costs of its use and maintenance. Trade-offs must be carefully considered in each case. – Using a minimalistic approach does not limit the technology's potential but rather constitutes an initial adoption strategy: instead of uploading complete crop information to the network, which would be costly and inefficient, blockchain was used to guarantee the immutability and verifiability of critical events in the production process. – It is necessary to carefully analyze and define what the milestones will be in traceability, as they are associated with the key stages supported by the blockchain. A greater number of milestones than necessary can reflect a lack of understanding of business processes and, consequently, a model that does not meet the needs of regulatory or legal frameworks.
Demonstrate the usefulness of the tool (and blockchain) in terms of its usability, acceptance, and value.	– User guides tailored to the diverse needs facilitate the consolidation of services around defined uses, with potential for customization and reusability. – A phased transfer, identifying potential gaps along the way, requires implementing deployment organizational capabilities, with personnel possessing strong communication skills, business acumen, and expertise with the tool. Such professionals are not easy to find, whether due to a lack of widespread knowledge in blockchain implementations or other factors, such as general traceability knowledge and a specific understanding of the cannabis industry. – Maintaining a record of events and usage for future analysis also requires data analytics capabilities. All these considerations mean that the transfer to the industry demands a diverse skill set, encompassing expertise in the implemented tool, analytical capabilities, and customer relationship management.

2. In addition, some pendings arised from its implementation. For instance, the creation of more specific or interconnected smart contracts capable of executing their own logic within the production flow remained pending. As another extension, we should consider that the system uses its own unique coding scheme to identify each element (batches, plants, biomass, and products). However, a natural evolution would be to incorporate the international GS1 standard, used in Argentina for product identification and marketing. Its implementation would require a code management system, but would provide compatibility with official traceability and dispensing systems.

References

1. Ahmed, W.A., MacCarthy, B.L.: Blockchain-enabled supply chain traceability – how wide? how deep? International Journal of Production Economics **263**, 108963 (2023). https://doi.org/https://doi.org/10.1016/j.ijpe.2023.108963
2. Azevedo, P., Gomes, J., Romão, M.: Supply chain traceability using blockchain. Operations Management Research **16**, 1359–1381 (2023). https://doi.org/10.1007/s12063-023-00359-y
3. Duque Schumacher, A.G., Pequito, S., Pazour, J.: Industrial hemp fiber: A sustainable and economical alternative to cotton. Journal of Cleaner Production **268**, 122–180 (2020). https://doi.org/https://doi.org/10.1016/j.jclepro.2020.122180
4. Kaur, G., Kander, R.: The sustainability of industrial hemp: A literature review of its economic, environmental, and social sustainability. Sustainability **15**(8) (2023). https://doi.org/10.3390/su15086457, https://www.mdpi.com/2071-1050/15/8/6457
5. Malabadi, R.B., Kolkar, K.P., Chalannavar, R.K.: Cannabis sativa: Industrial Hemp (fiber-type)- An emerging opportunity for India. International Journal of Research and Scientific Innovation **10**(3), 01–09 (March 2023), https://ideas.repec.org/a/bjc/journl/v10y2023i3p01-09.html
6. Manaia, J.P., Manaia, A.T., Rodriges, L.: Industrial hemp fibers: An overview. Fibers **7**(12) (2019). https://doi.org/10.3390/fib7120106, https://www.mdpi.com/2079-6439/7/12/106
7. Moore, B.: Using blockchain and data analytics for improved medical cannabis efficacy. ONdrugDelivery) (128) (2021), https://www.ondrugdelivery.com/wp-content/uploads/2021/12/128_2021_Dec_Connecting_Drug_Delivery_GCAC.pdf
8. Onu, P., Mbohwa, C., Pradhan, A.: Blockchain-powered traceability solutions: Pioneering transparency to eradicate counterfeit products and revolutionize supply chain integrity. Procedia Computer Science **232**, 1420–1427 (2024). https://doi.org/https://doi.org/10.1016/j.procs.2024.01.140, 5th International Conference on Industry 4.0 and Smart Manufacturing (ISM 2023)
9. Xiao, L., Sun, W., Chang, S., Lu, C., Jiang, R.: Research on the construction of a blockchain-based industrial product full life cycle information traceability system. Applied Sciences **14**(11) (2024). https://doi.org/10.3390/app14114569, https://www.mdpi.com/2076-3417/14/11/4569

Mobile Crypto Wallet User Experience Patterns Analysis

Leandro Plaquin[1], Francisco Gindre[2], and Matías Urbieta[2]

[1] Facultad de Informática UNLP, La Plata, Buenos Aires, Argentina
leaplaquin@gmail.com
[2] LIFIA – CONICET, La Plata, Buenos Aires, Argentina
{fgindre,murbieta}@lifia.info.unlp.edu.ar

Abstract. This paper examines UX design patterns in non-custodial cryptocurrency wallets and their implications for usability, safety, and adoption. A hands-on evaluation of widely used mobile wallets (e.g., MetaMask, Trust Wallet, and Exodus) identifies recurring solutions across critical flows: Initial Setup and Seed-phrase Generation, Seed Import, Portfolio Visualization, Recipient Address, and Transaction Execution. We identify five UX Patterns which are described using a uniform pattern template that explains each one of them in terms of *what is it*, *when to use it*, and *why it works*. The resulting catalog provides actionable guidance for designers and developers of multi-chain, mobile-first wallets and, suggests directions for standardization.

Keywords: User Experience · User Interface · Design Patterns · Mobile · Cryptocurrency · wallet

1 Introduction

The rapid rise of blockchain technology has reshaped how we interact with digital finance. Cryptocurrencies, which initially were conceived as an alternative to traditional finance, have evolved into a core component of the global economy. Wallet applications play a crucial role in the widespread adoption of cryptocurrencies. Wallet applications let users send and receive crypto-assets. Moreover, self-custodial wallets let the user store and hold custody of their funds without intermediaries. Design is a fundamental tool not only to deliver an intuitive, accessible User Experience (UX) but also to avoid user mistakes that often lead to loss of funds [5]. The analysis of UX patterns applied to self-custody cryptocurrency wallets is an important tool to document and communicate common elements of user experience design across the ecosystem. According to CoinLaw's study on cryptocurrency adoption trends [3], 820 million unique cryptocurrency wallets are active globally as of 2025, where Hot wallets account for 78% of all crypto wallets in that same year. Binance's 2024 report cites a 47% year-over-year increase to 250 million registered users [1], further illustrating widespread engagement. A market study on *wallet* adoption trends [4] indicates that MetaMask [7] and Trust Wallet [15] are the most popular wallets with an estimate of 158 million users.

Documenting UX design patterns is an established practice in other application domains as can be seen on the UI Patterns catalog website[14]. Each pattern is described in detail with examples. For instance, the *Pricing Table* pattern[14] visually compares plans or services by features and price using a clear tabular layout, helping users make informed decisions. Likewise, a *Dashboard* [13] that presents the overall portfolio status contributes to visual clarity: a central panel shows real-time summaries organized into cards or charts that present balance, assets, and recent transactions in an intuitive way.

In this work, we present five cryptocurrency UX design patterns found in the most popular non-custodial wallets. Their aim is to provide design guidelines for existing and new solutions in order to improve the overall user experience.

2 Related Work

Tidwell [12] asserts that patterns have three key traits: they are proven solutions, they are context-dependent rather than absolute rules, and they function as a common vocabulary among designers. In cryptocurrency *wallets* this idea is especially relevant for domain-specific problems such as designing secure transaction flows or presenting complex *blockchain* information. Prior work, including Punchoojit et al. [11], underscores that patterns are reusable solutions to recurring human–computer interaction problems. We considered key resources such as *UI-Patterns.com* [14], Theresa Neil's book "Mobile Design Pattern Gallery: UI Patterns for Smartphone Apps" [8], and "Designing Search: UX Strategies for eCommerce Success" [9]. *UI-Patterns.com* is an online repository of UX and UI design patterns with practical examples and best practices. It is useful for analyzing recurring patterns in cryptocurrency *wallets*, including *onboarding*, forms, and visual *feedback*. Neil's book offers a structured guide to mobile-specific patterns, organized by navigation, forms, and search, which helps evaluate how *wallets* implement these patterns in a mobile context. Although its primary focus is e-commerce, the "Designing Search: UX Strategies for eCommerce Success" contributes lessons on user-centered search and navigation. Concepts such as clear visual *feedback* and efficient information organization apply to cryptocurrency *wallet* design, especially for transaction search and status presentation. Together, these sources provide a solid theoretical foundation and concrete examples for identifying and analyzing usability patterns in cryptocurrency *wallets*.

Despite the existence of these and other relevant works, none of them address in depth design patterns applied specifically to cryptocurrency *wallets*, where unique challenges include private-key management, *blockchain* network selection, irreversible transaction validation, and the absence of any intermediary to recover lost funds. This work aims to fill that gap with a detailed, domain-specific analysis of the patterns that directly bear on adoption, security, and usability of self-custodial mobile *wallets*. In sum, design patterns are fundamental tools for addressing usability challenges in a consistent and validated way. They are reusable elements of design to recurring problems and also serve as a shared vocabulary and terminology within different application domains.

3 Research Methodology

To identify and evaluate design and usability patterns present in multiple digital *wallets*, we carried out a three-stage process detailed in the following subsections.

3.1 Pattern-mining process

Preparation and wallet selection. We selected digital *wallets* based on market popularity, coverage of different crypto-assets (Bitcoin, Ethereum and stablecoins), and diversity of features (new wallet creation, seed-phrase import, and asset exchange). Each *wallet* was installed on a Samsung Galaxy S22 running Android 12 to simulate real use and access all features and user flows. We configured each *wallet* using standard account-creation and backup procedures, enabling a clean, from-scratch user experience.

Systematic flow exploration. With each app running, we exhaustively traversed key user flows: creating a new wallet, importing via seed phrase, selecting a network, and transferring assets. To document the interface, interactive elements, and action sequences, we captured screens when possible and photographed the device when security settings blocked screenshots.

Qualitative assessment. Each identified pattern was evaluated against usability and clarity criteria:

- *Error minimization:* designs that prevent invalid or confusing input (e.g., "Max" buttons, currency toggles, seed phrase validation).
- *Visibility:* how easily users can locate and recognize the component.
- *Consistency:* uniformity with other elements in the same app and with industry patterns.
- *Feedback:* the ability to inform users about operation status (errors, success, and wait times).

In sum, this methodology enabled a rigorous identification and evaluation of design patterns in cryptocurrency *wallets* through real-use simulation, comprehensive analysis of key flows, and a qualitative appraisal grounded in usability criteria.

3.2 Pattern Template

To describe UX design patterns in cryptocurrency *wallets*, we draw on general elements and guidance from the structure used in the thesis "Using Design Patterns in User Interface Design"[2], which organizes each pattern into the following attributes:
What is it: A high-level description of the pattern, including a brief explanation of its characteristics.

Table 1: Patterns instantiation in Wallets.

Wallets	**Portfolio Visualization**	**Recipient Address**	**Initial Setup and Seed Phrase Generation**	**Import Seed Phrase**	**Transaction Creation**
Metamask	×	●	●	●	●
Trust	●	●	●	●	●
Exodus	●	■	■	●	■
Ywallet	×	●	×	■	×
Phantom	■	●	●	■	●
Edge	●	×	×	×	●
Cake Wallet	×	●	●	■	●
Electrum	×	×	■	×	×
Monerujo	×	×	×	×	×

When to use it: The specific context in which the pattern is appropriate and the problems it helps solve in cryptocurrency *wallet* design.
Why it works: How the pattern addresses the identified problem and the benefits it brings to the user experience in a cryptocurrency *wallet*.
Examples: A walk-through of an implementation of the pattern on well-known *wallets* that will illustrate how the pattern is applied in real products.

3.3 Pattern Catalog

Table 1 summarizes the five design patterns analyzed in this work applied to nine reference *wallets*: MetaMask, Trust Wallet, Exodus, Phantom, Edge, Cake Wallet, Electrum, YWallet and Monerujo. These apps were selected for their global popularity. We included Privacy Coin wallets to also cover blockchain protocols that do not have a fully transparent ledger in our pattern analysis.

Each cell in the table is coded with a symbol: a ● indicates the pattern is present and clearly implemented; a ■ indicates the pattern is present but needs more detail or clarity; and × indicates the pattern is missing or insufficiently implemented.

4 Pattern: Portfolio Visualization

What is it? The portfolio visualization pattern is found in multi-asset wallets. It presents a clear, organized visual summary of a user's digital assets, combining data visualization principles, visual hierarchy, and contextual access to details.

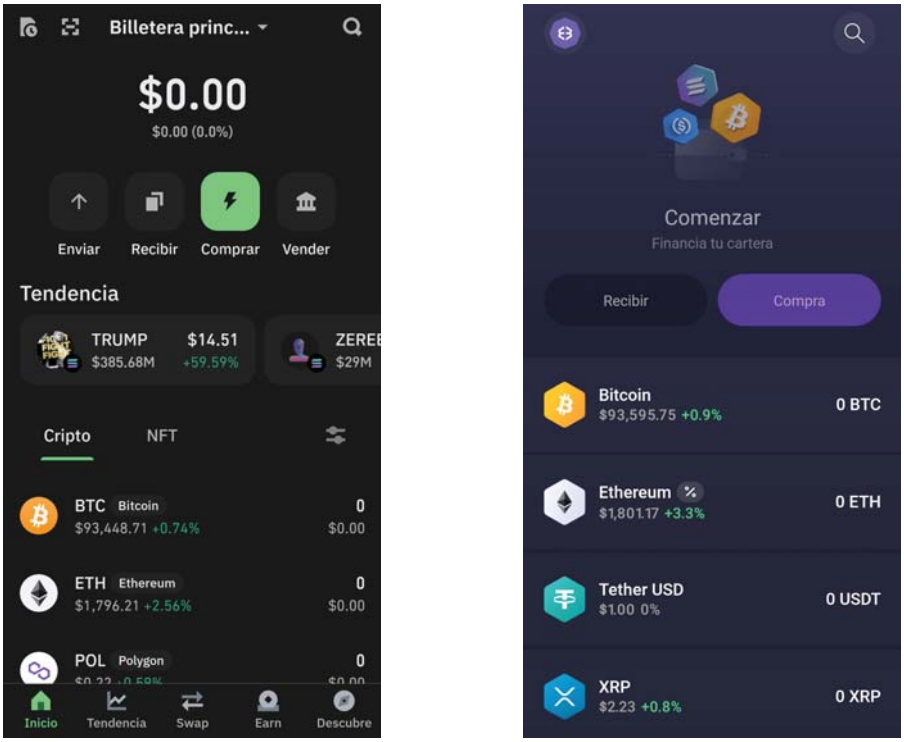

(a) Trust Wallet portfolio.

(b) Exodus Wallet portfolio.

Fig. 1: Portfolio views: Trust Wallet (left) and Exodus Wallet (right).

It helps users answer, at a glance: "How much do I have?", "What am I holding?" and, "How has my portfolio evolved over time?"

When is it used? When users need to monitor their cryptocurrency holdings continuously. It is important for frequent activity, quick decisions, or simple checks of total net worth. It also serves as a starting point for drilling into movements, transaction history, or performance charts.

Why does it work? It lowers cognitive load by presenting financial information visually, with structure and hierarchy. It provides immediate feedback on asset status, allows comparison across coins or *tokens*, and builds confidence by giving a clear view of available funds. In the interface shown in Figure 1a, the total portfolio amount appears prominently at the top left, expressed in US dollars as the *fiat* reference. Below it is a list of all cryptocurrencies in the user's portfolio. Each entry shows the asset name and current market price with an adjacent percentage indicator of recent price change—allowing users to see immediately whether an asset is up or down. The list also displays the user's total units for each asset and the corresponding dollar value, providing an up-to-date view of

both individual and overall portfolio value. In Figure 1b, Exodus wallet shows the same pattern with a design focused on visual simplicity and clear information.

5 Pattern: Recipient Address

What is it? This pattern covers how a receiving address is presented and managed in a *wallet*. In the case where the wallet operates across different cryptocurrency blockchains and their respective networks, the pattern also includes both the address display and the options for selecting the correct network.
When is it used? It is used when sharing a destination address from a *wallet* in order to construct a transaction. It typically appears under "Receive" or "Deposit". After selecting a cryptocurrency, the app shows the corresponding address along with a network selector when the asset is available on multiple blockchains. The address is usually presented as an alphanumeric string and a QR code, with buttons to copy, share, or view full-screen. Addresses presented as strings respect the encoding defined in the respective protocols.
Why does it work? It reduces mistakes by presenting explicit, visual, and validated information about the receiving address. String-Encoded addresses are shown in the encoded format defined in their protocol specification. These strings are either displayed in full length or abbreviated. This last form has a minimum length that ensures that users validate an address that is computationally expensive to perform an *address poisoning* attack [6]. Legends remind users to send assets only on the matching network, and the combination of QR code plus action buttons like copy, share and set amount make the task straightforward. The pattern aligns with principles of error prevention, system visibility, and cognitive load reduction.
Example: generating a QR code or sharing an address. Figure 2a shows the detail view of choosing Bitcoin recipient address. At the top, the asset's identifying icon is shown, followed by a large, centered QR code representing the receiving address. Below the QR, the full alphanumeric address can be copied via a dedicated button. An explanatory message reinforces that only *BTC* on the Bitcoin network should be sent. Although it was found that other wallets emphasize this with visible warnings to make the user aware of potential loss-of-funds if the wrong network is used (especially in ERC20-like tokens). Finally, an action button allows sharing the address, making it simple to send via other applications.
Example: sharing an address in multi-currency wallets. As shown in Figure 2b, the screen provides a search bar at the top and a drop-down filter labeled "All networks" allowing users to filter or directly select the *blockchain* network they want to use. Below, a list of popular cryptocurrencies appears with each item's name, identifying logo, corresponding network and a shortened version of its address. Next to each address, there are two action buttons: one to copy the address to the clipboard, and another to display the QR code, enabling accurate input through QR code reading for peer-to-peer interactions. This design lets users manage multiple addresses across networks and attempts to reduce the

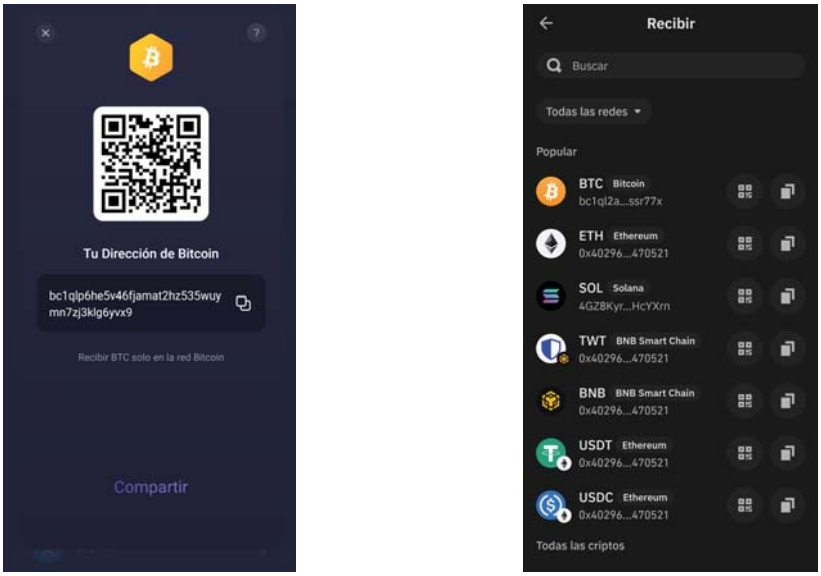

(a) An address in Exodus Wallet. (b) Addresses in Trust Wallet.

Fig. 2: Bitcoin receiving addresses.

chance of making common mistakes when receiving assets that exist on several blockchains, which improves operational safety and the overall user experience.

6 Pattern: Initial Setup and Seed Phrase Generation

What is it? The goal of this pattern is to guide the user through the startup and creation process of a *wallet*, during which a BIP-39 [10] seed phrase is generated and presented to the user. The seed phrase is essential for accessing, recovering, or restoring *wallet* funds. The pattern is intended to ensure that users retain full ownership and control of their assets, consistent with decentralization principles.
When is it used? It is used when the user chooses "Create new wallet". From that point, the system generates a unique seed phrase that the user must store securely. Users are typically asked to confirm they have recorded the phrase correctly—e.g., by verifying word order, selecting specific words, or retyping them.
Why does it work? It ensures that control of the *wallet* resides exclusively with the user. By presenting the seed phrase clearly and with security warnings, the pattern reinforces the concept of self-custody. Effective interfaces also use visual emphasis to capture attention, alert users to the risks of losing the phrase, and include confirmation steps that verify the phrase has been recorded correctly.
Example: MetaMask . Figure 3a shows the initial setup screen of the MetaMask mobile app. The button hierarchy guides inexperienced users towards cre-

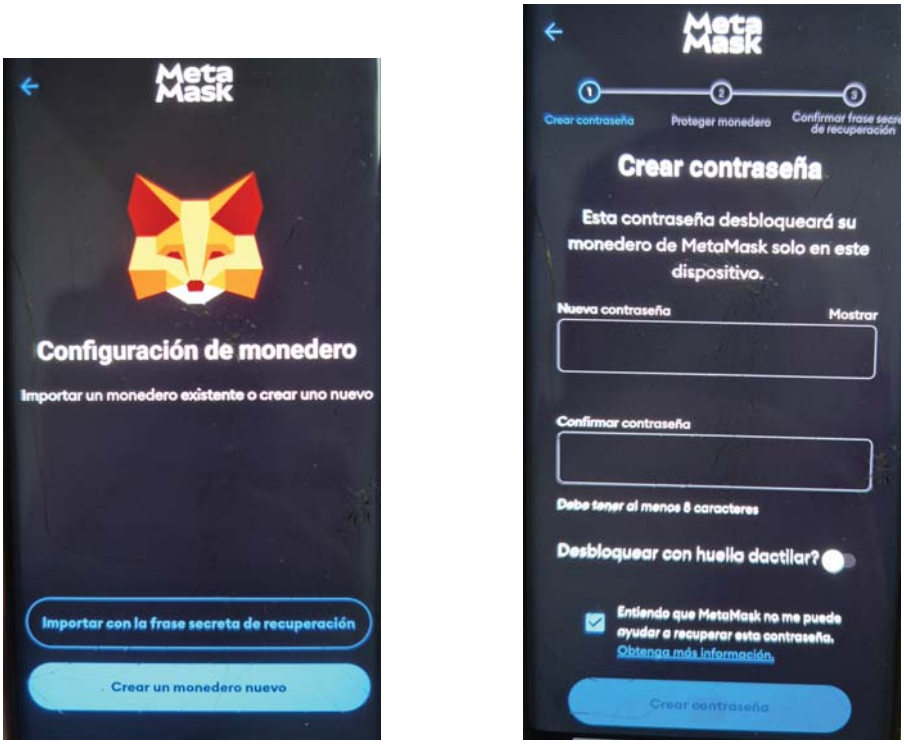

(a) App start screen. (b) Password creation.

Fig. 3: MetaMask start and password-creation screens

ating a new *wallet* without confusing them with advanced recovery flows, while also offering an alternative path for users who need to import an existing *wallet* via seed phrase—addressing both user profiles from the first interaction. Figure 3b shows the first step of the setup process: creating a local password to encrypt the key material and lock/unlock the *wallet* on the device. The UI uses a well-known pattern for password creation. It also offers an option to enable fingerprint unlock, improving accessibility without compromising security. Figure 4a shows the second step in MetaMask's setup: displaying the 12-word randomly generated seed phrase. The interface guides the user to write the phrase on paper and store it securely, because it will be requested again in the next step for validation. This design reflects security and error-prevention principles by encouraging a manual, offline record rather than digital storage, reducing leakage risk. Figure 4b shows the final step of *wallet* setup in MetaMask: confirming the secret recovery phrase. The goal is to ensure the user has recorded the seed phrase correctly and understands its importance. To proceed, the user needs to complete the puzzle to be allowed to continue. This mechanism has a

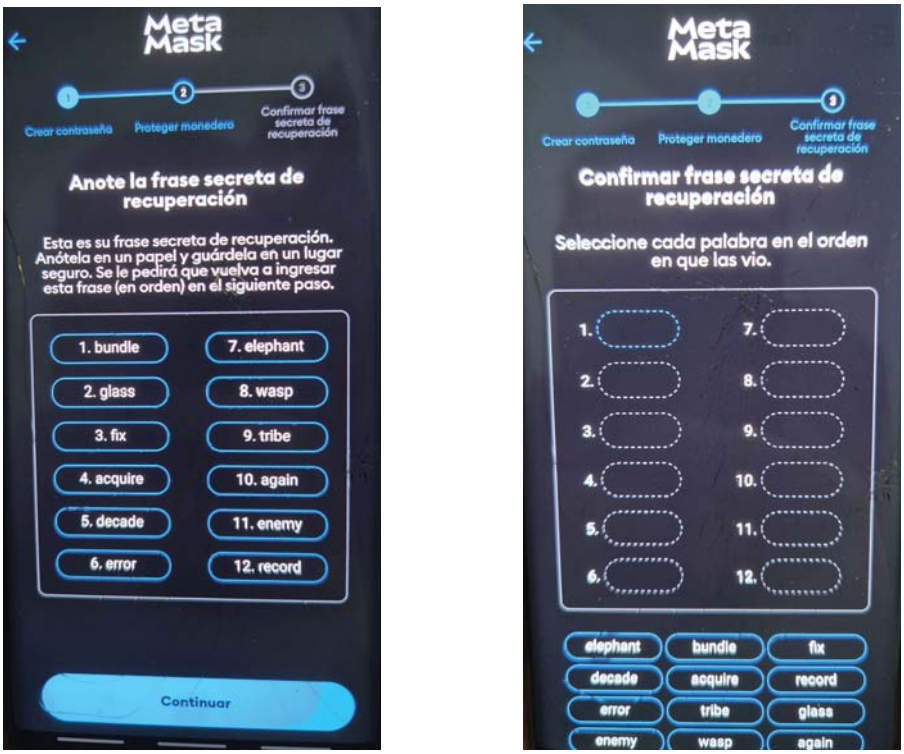

(a) Seed phrase in MetaMask. (b) Seed-phrase confirmation.

Fig. 4: Seed-phrase display and confirmation in MetaMask.

dual purpose: it verifies that the phrase was properly recorded and reinforces personal responsibility over *wallet* access.

Example: Trust Wallet As shown in Figure 5, the app opens with buttons for "Create new wallet" and "Add existing wallet", covering both user profiles from the first screen. The backup-method screen (center) highlights the secret-phrase option in green with a "Recommended" badge, while keeping a *Swift* (beta) alternative accessible. Before the phrase is shown, the user must tap through a disclaimer (right) confirming that Trust Wallet holds no copy of it. As shown in Figure 6, the 12-word BIP-39 [10] phrase (left) is displayed in a numbered grid; the subsequent spot-check (right) asks the user to select the correct word at randomly chosen positions, reducing friction compared with MetaMask's full drag-and-arrange puzzle while still verifying the backup.

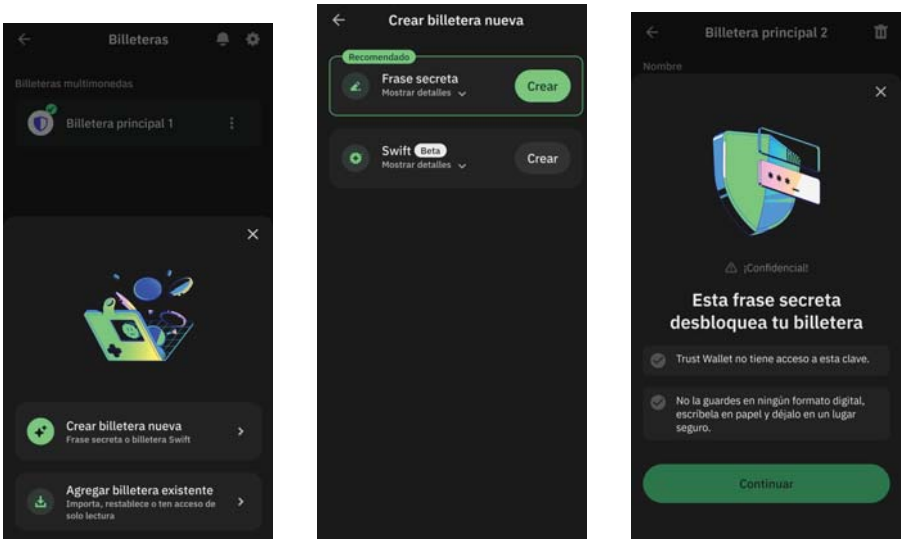

(a) Create or add wallet. (b) Backup method selection. (c) Self-custody disclaimer.

Fig. 5: Trust Wallet: initial screen, backup-method selection, and self-custody disclaimer.

7 Pattern: Import Seed Phrase

What is it? The seed-phrase import pattern lets users restore an existing cryptocurrency *wallet*—created earlier on another device or app—using a 12 or 24-word secret recovery phrase.
When is it used? When the user re-installs the app on the same device or moves to a new one; during migration between different *wallets* that support the BIP-39 specification[10]; to recover access after accidental deletion; and in multi-device setups where the user needs to operate the same *wallet* across environments.
Why does it work? It provides a mechanism for the user to exercise sovereign custody over their funds. The design guides users step by step and validates the phrase before granting access, reducing errors. Prominent security warnings build early awareness of the phrase's importance. Consistent visuals and copy help non-technical users complete the process independently. By adhering to standards such as BIP-39[10], the pattern also enables interoperability across applications.
Example: MetaMask. Figure 7a combines all elements needed to restore a wallet: a large field where the 12-word seed phrase can be pasted or scanned; two fields to create and confirm a local password with a strength indicator and a show/hide option; a switch to enable fingerprint unlock; and a prominent "IMPORT" button accompanied by a legal notice confirming user consent. The

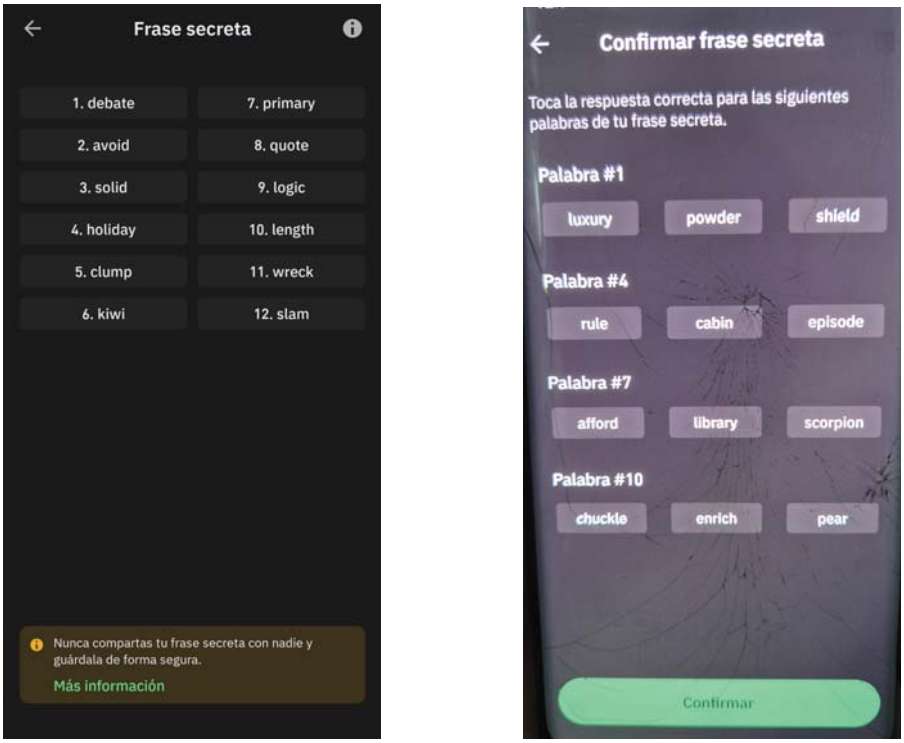

(a) 12-word seed phrase display. (b) Spot-check confirmation.

Fig. 6: Trust Wallet: seed-phrase display and spot-check confirmation.

design works thanks to clear visual hierarchy, icon and color cues that guide attention, and the ability to toggle visibility for both the seed phrase and the password—improving security and usability. This flow enables even non-technical users to recover keys and access assets safely.

Anti-Pattern: YWallet. While Metamask and Ywallet (Figure 7b) screens seem at a glance similar to each other, the Ywallet screen allows users to import different kinds of key material on the same screen and also shows options considered advanced for the average crypto user. Users have several opportunities to make mistakes while importing their keys, which could lead to loss-of-funds, or at least the perception that funds have been lost.

8 Pattern: Transaction Creation

What is it? Defines the user-interface and experience flow for sending digital assets from one *wallet* to another. It organizes the critical steps: select the asset, enter the destination address, specify the amount, choose the network (if applicable), review network fees, and confirm.

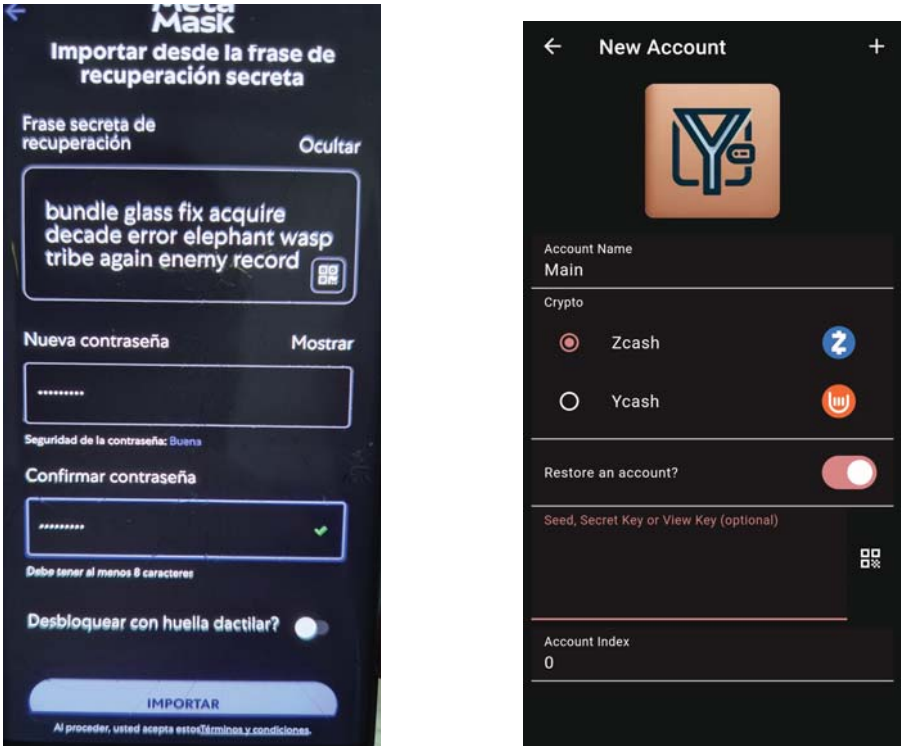

(a) MetaMask. (b) YWallet

Fig. 7: Seed-phrase import screens

When is it used? When a user wants to transfer a cryptocurrency from their *wallet* to another address.

Why does it work? It breaks the transaction creation process into clear, sequential steps, reducing the chance of errors. The interface provides validated fields (e.g., address format checks), options to select compatible *blockchain* networks, and a review screen summarizing the transaction information before its confirmation. It also shows the cost of the transaction (commonly transaction fee or transaction gas), which promotes transparency and user control.

Example: Trust Wallet . Figure 8a shows Trust Wallet's "Send" screen right after tapping the send button. A search bar and a list of spendable funds and coins appear. In this example the wallet only holds Cardano tokens (ADA). This design makes it clear that before entering the recipient or amount, users should confirm the correct network—preventing compatibility errors. Figure 8b shows the choose amount and recipient step. The recipient input field offers Paste, Contacts, and QR Scanner buttons, reducing transcription errors and making it fast and safe to enter a destination. Below it, the "Amount" field

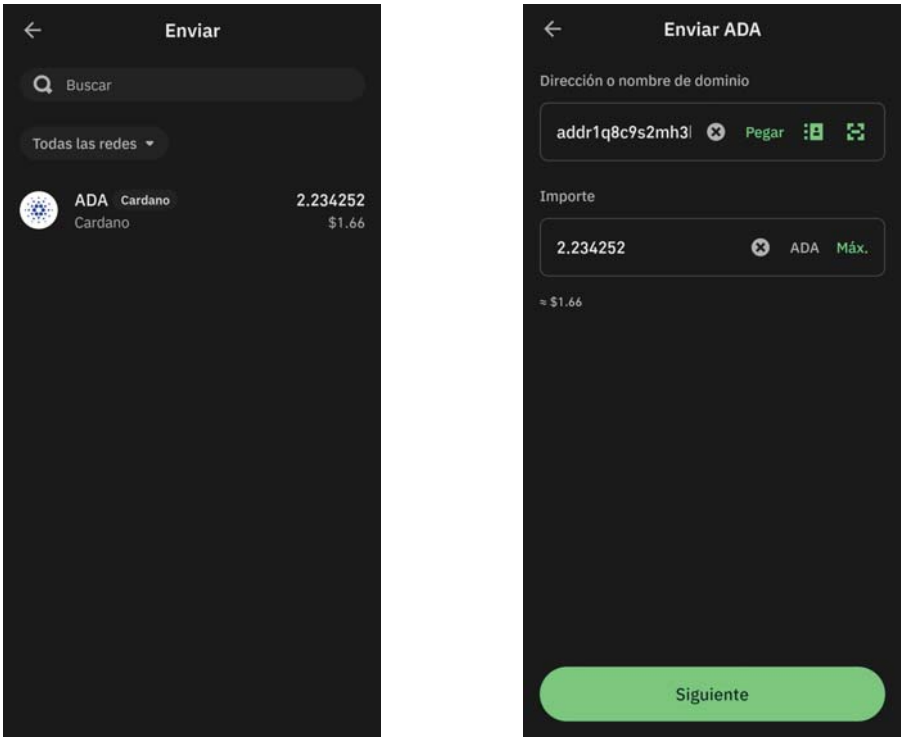

(a) Network selection for the transfer.

(b) Choosing amount and destination address.

Fig. 8: Transfer flow in Trust Wallet.

displays the quantity to send (2.234252 in this case), followed by the ADA label and a "Max" button that the user can tap to indicate the intent of sending the maximum amount possible minus fees. Users can switch between fiat and ADA denomination by tapping the ADA accessory button in the input field. In Figure 9a, relevant details are presented clearly before execution, with the information laid out hierarchically by relevance. A green "Confirm" button at the bottom highlights the final action. The following action is the completion step depicted in Figure 9b, which resembles the confirmation screen and adds any information that can be known at the time, such as the number of confirmations (2) and a "More details" button for additional information. The "Completed" label and share icon in the header confirm success unambiguously.

Example: MetaMask Figure 10 shows MetaMask's send flow. The opening screen (left) displays the active network, the source account with its balance, and a destination field that accepts typed addresses, QR scan, or a quick pick from the user's own accounts. The amount step (not shown) lets the user toggle between native token and fiat denomination and fill the full sendable balance with a "Use

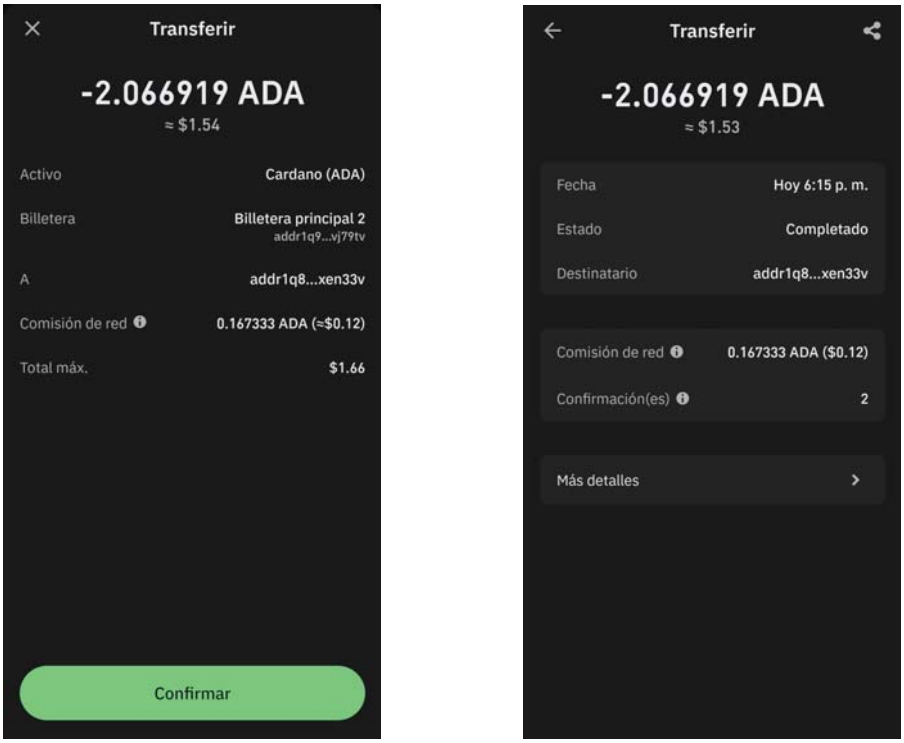

(a) Transfer confirmation. (b) Transfer completed.

Fig. 9: Confirmation and completion steps (Trust Wallet)

Max" button. The confirmation screen (right) is split into two scrollable panels: the first summarises accounts and the exact transfer amount with an "Estimated change" block; the second itemises the network fee with both estimated and maximum values, a processing-time hint, and the total cost. This layout ensures all critical data, including fees, is reviewed before the irreversible action is taken.

9 Conclusion

We presented five patterns: "Portfolio Visualization", "Recipient Address", "Initial Setup and Seed Phrase Generation", "Import Seed Phrase" and "Transaction Creation". These findings show that the crypto wallet industry makes use of UX Design Patterns as elements of reusable design intended to improve usability, accessibility, security and user trust in a context where responsibility is largely individual and certain actions cannot be undone. Standardizing validated fields, state indicators and clear typographic hierarchies provides a foundation on which new *wallets* can be built, accelerating adoption and reducing the inherent friction

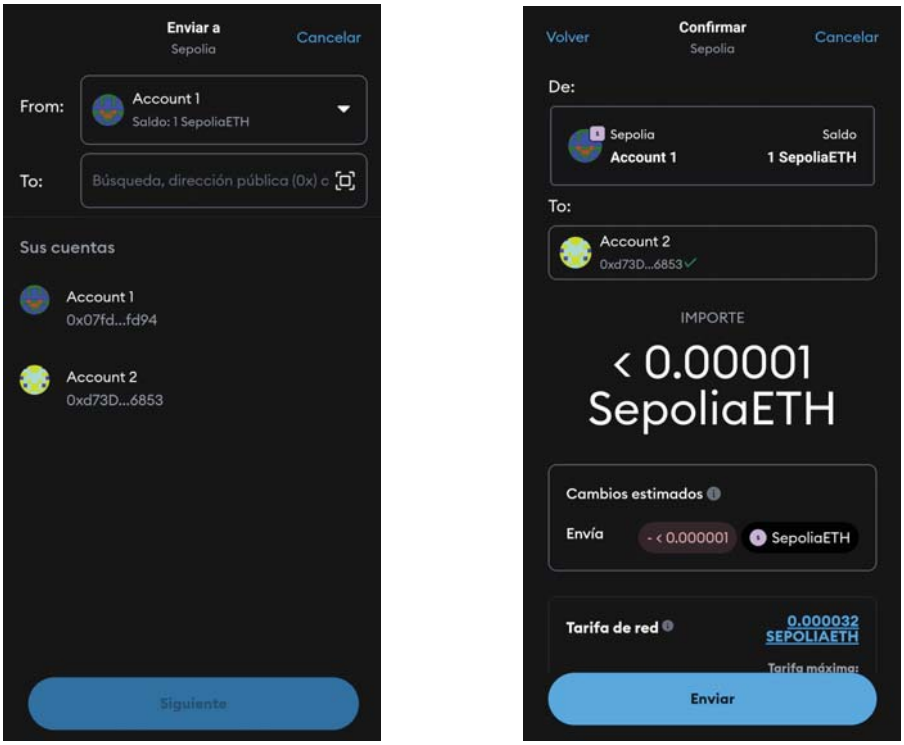

(a) Destination entry.

(b) Confirmation: accounts, amount and fee.

Fig. 10: MetaMask transaction: destination entry and confirmation.

of these flows by establishing a common language users and product owners can share. This work lays a foundation for understanding how UX Design Patterns contribute to the effectiveness of cryptocurrency *wallets* and opens avenues for future research aimed at refining the experience for an increasingly diverse and demanding user base in the blockchain space.

References

1. Binance: Binance year-end report 2024: Quarter billion users in transformational final quarter. `https://www.binance.com/en/square/post/19542070539505` (2025), [Accessed 27-02-2025]
2. Chase, C.: Using design patterns in user interface design. Division of Graduate Studies and Research of the University of Cincinnati pp. 1–92 (2012)
3. Elad, B.: Cryptocurrency Wallet Adoption Statistics 2025: E-Commerce, Payments & Regional Growth — coinlaw.io. `https://coinlaw.io/cryptocurrency-wallet-adoption-statistics/`, [Accessed 10-10-2025]

4. Elad, B.: Cryptocurrency Hot Wallet Statistics 2025: Growth & Trends • CoinLaw — coinlaw.io. https://coinlaw.io/cryptocurrency-hot-wallet-statistics/ (2025), [Accessed 11-10-2025]
5. Fröhlich, M., Gutjahr, F., Alt, F.: Don't lose your coin! investigating security practices of cryptocurrency users. In: Proceedings of the 2020 ACM Designing Interactive Systems Conference. p. 1751–1763. DIS '20, Association for Computing Machinery, New York, NY, USA (2020). https://doi.org/10.1145/3357236.3395535, https://doi.org/10.1145/3357236.3395535
6. Ledger: What Are Address Poisoning Attacks in Crypto and How to Avoid Them? | Ledger — ledger.com. https://www.ledger.com/academy/topics/security/what-are-address-poisoning-attacks-in-crypto-and-how-to-avoid-them (2025), [Accessed 13-10-2025]
7. Metamask: MetaMask - Crypto Wallet - Apps en Google Play — play.google.com. https://play.google.com/store/apps/details?id=io.metamask&hl=es_AR (2025), [Accessed 02-06-2025]
8. Neil, T.: Mobile Design Pattern Gallery: UI Patterns for Smartphone Apps. O'Reilly Media, Inc (05 2014)
9. Nudelman, G.: Designing Search: UX Strategies for eCommerce Success. Wiley Publishing, Inc (2011)
10. Palatinus, M., Rusnak, P., Bowe, S., Voisine, A.: Mnemonic code for generating deterministic keys. Bitcoin Bips (09 2013), https://github.com/bitcoin/bips/blob/master/bip-0039.mediawiki
11. Punchoojit, L., Hongwarittorrn, N.: Usability studies on mobile user interface design patterns: A systematic literature review. Advances in Human-Computer Interaction **2017**, 1–22 (11 2017). https://doi.org/10.1155/2017/6787504
12. Tidwell, J.: Designing Interfaces Patterns for Effective Interaction Design. O'Reilly Media, Inc. (2020)
13. Toxboe, A.: Dashboard design pattern — ui-patterns.com. https://ui-patterns.com/patterns/dashboard, [Accessed 30-05-2025]
14. Toxboe, A.: UI-Patterns.com — ui-patterns.com. https://ui-patterns.com/, [Accessed 30-05-2025]
15. wallet, T.: Trust: Crypto & Bitcoin Wallet - Apps en Google Play — play.google.com. https://play.google.com/store/apps/details?id=com.wallet.crypto.trustapp&hl=es_AR (2025), [Accessed 02-06-2025]

Multi-Agent Framework for Resilient Medical Triage in the Internet of Medical Things (IoMT)

Sergio Gramajo, Reinaldo Scappini, Salvador Nuñez, Carlos Torres, Raul Montiel, and Jorge Roa

CInApTIC (Centro de Investigación Aplicada en TICs), Universidad Tecnológica Nacional, Resistencia, Argentina
{sergiogramajo, rscappini, salvador.nunez, carlos.torres, raulmontiel, roajorge}@gfe.frre.utn.edu.ar

Abstract. Remote monitoring has evolved from reactive systems to complex Internet of Medical Things (IoMT) ecosystems. However, centralized architectures face critical vulnerabilities: high latency, absolute dependence on connectivity and its state, and a high rate of false positives in detections due to the system's own conditions, such as fall detection. This paper presents a distributed multi-agent framework designed to be executed at the *edge* using ESP32/Single Board Computer (SBC) hardware and MicroPython. The architecture replaces the monolithic logic of a system with a society of autonomous agents (ECG, Falls, Environmental, etc.) that collaborate through the *Micro Agent Communication Protocol* (μACP). The novelty lies in cooperative reasoning; the system emulates clinical judgment by correlating biological and environmental dimensions through Bayesian fusion. The results project the validation of an accuracy close to 100% in the discrimination of activities of daily living (ADL) and a latency reduction compared to cloud-based models.

Keywords: IoMT · Multi-Agent Systems · Medical Triage · μACP · Edge Computing

1 Introduction

The evolution of healthcare is at a turning point. Population aging and the rise of chronic diseases demand systems that transcend traditional hospital infrastructure. The Internet of Medical Things (IoMT), defined as the interconnection of medical devices and applications, enables the generation of massive biological data streams in real time [1]. This transformation, accelerated by the COVID-19 pandemic, has demonstrated that remote monitoring is vital for the continuity of medical care and the optimization of limited resources [2, 4, 6].

The traditional medical triage process has historically functioned as a critical mechanism in emergency departments to stratify patients based on their physiological stability and predict the amount of hospital resources necessary for their timely care. Standardized five-level clinical tools, such as the Emergency Severity Index (ESI), focus on rapidly identifying high-risk situations and patients requiring immediate life-saving interventions compared to those who are

stable [1]. However, recent global events, such as the pandemic and demographic aging, have evidenced the urgent need to extend this concept of care beyond physical hospitals, driving home hospitalization models and continuous remote monitoring systems supported by the Internet of Medical Things (IoMT) [2].

To ensure that remote clinical decision-making achieves the same efficacy as in-person triage, the underlying architecture requires minimal latency levels in data transmission [3]. Traditional infrastructures reliant on the centralized cloud often suffer from delays that are unacceptable for emergency situations. Therefore, the adoption of an Artificial Intelligence (AI)-driven hybrid *Fog-Edge* computing architecture optimizes data processing near patients and alleviates bandwidth overhead [4–6]. In urban and rural scenarios where network connectivity is intermittent, it is essential to adopt mesh network topologies based on LoRa technologies without the need for central gateways (gateway-free) to provide decentralized multi-hop communication with intrinsic self-healing and fault-tolerance capabilities [7]. Furthermore, the energy and operational efficiency of these long-range wireless transmissions is maximized by integrating low-power hybrid detection algorithms into wearable devices [8]. Complementarily, software design must account for offline operations in which a local data model temporarily stores information and, through hierarchical conflict resolution strategies, automatically synchronizes with the cloud when the connection is restored [9].

Equipping these sensors and devices with intelligence at the edge presents massive computational challenges, as they are severely constrained in terms of memory, energy, and processing capacity [10]. To overcome this barrier, the literature has adopted algorithmic compression strategies, such as quantization and model pruning. Particularly in computer vision-based medical monitoring systems, models detecting falls through video sequences are mathematically reduced using methodologies such as the Sparse Low-Rank (SLR) method applied to dense layers; this preserves diagnostic performance while reducing the parametric weight of the convolutional network by up to 90%, allowing it to execute effectively on resource-constrained devices [11].

At the clinical level, a priority application where this technology converges is the detection and prediction of falls in older adults, whose prompt attention drastically reduces severe complications [12, 13]. AI has evolved toward the implementation of powerful multimodal deep learning frameworks, such as architectures fusing convolutional neural networks and LSTMs (CNN-LSTM) to process both inertial acceleration dynamics and vital signs (heart rate, SpO_2) in a single, high-precision inference [14]. In turn, modern systems implement interactive validations through home IoT devices, such as smart speakers, to confirm the patient's status after an alert [15]. Additionally, cooperative AI-based predictive schemes integrating fuzzy logic and Deep Belief Networks (DBN) continuously evaluate instrumental activities of daily living and vital signs, forecasting long-term fall risk and enabling preventive action [16]. To achieve this perception of human behavior and mitigate high false alarm rates, clinical planning employs Adaptive Probabilistic Fusion Networks (APFN); these dynamically weight each

sensor signal using Gaussian mixture models to autonomously suppress noisy information [17]. Such continuous vital sign detection and modeling are also paramount for effectively coordinating the medical evaluation and evacuation of patients with special needs in disaster scenarios [18].

Data manipulation across this entire biomedical spectrum demands extreme resilience against systemic failures or cyber manipulation. Comprehensive resilience is orchestrated by combining the advantages of Software-Defined Networking (SDN) for global monitoring and routing, Machine Learning (ML) for diagnostics, and Microservices Architecture (MSA) to support disruption-tolerant access controls [19]. Given that in the clinical monitoring domain a vital sign alteration can be due to real deterioration, a sensor defect, or a cyberattack, frameworks such as *Intelligent Fusion* combine supervised XGBoost ensembles to address network intrusions (NIDS) with unsupervised LSTM autoencoders, tasked with capturing technical deviations, yielding a fused alert that prevents unnecessary escalations and safeguards patient safety [20]. Along the same lines of confidentiality, cutting-edge infrastructures like SecureEdge-MedChain combine Federated Learning at the edge with post-quantum cryptography-based distributed ledgers, ensuring that identities, diagnoses, and clinical predictions remain inviolable against future computational threats [21].

Finally, orchestrating the interoperability of all this algorithmic and telemetric complexity requires a collaborative Multi-Agent System (MAS) capable of discerning and debating contradictory information, an attribute proven effective in reducing asymmetries or misinformation in health data analysis through multi-agent debate [22–24]. Semantic coordination through traditional standards, such as FIPA-ACL, provides enormous expressiveness but introduces prohibitive levels of computational load for constrained environments [25–27]. In response, the technological ecosystem has matured toward the creation of modern distributed interoperability protocols such as MCP, ACP, A2A, and ANP [28]. For the most restrictive base of these devices, the nascent Micro Agent Communication Protocol (μACP) establishes a new standard through rigorous formal calculus; empirically demonstrating that a minimal base of verbs (PING, TELL, ASK, OBSERVE) is semantically complete to guarantee medical consensus under strict memory, energy, and bandwidth limitations [29–31]. It is the coherent summation of all these interconnected domains that comprehensively shapes a complete paradigm and gives rise to a true multi-agent framework for resilient medical triage within the next-generation Internet of Medical Things. In the triage domain, explainability and clinical precision are improved by deploying multi-agent systems like ED-Triage-Agent (ETA) to decompose tasks into specialized agents [32, 33], as well as by integrating machine learning (ML) techniques [34, 35]

Despite these advances, current systems suffer from "alarm fatigue." Sensors often operate in isolation, triggering critical alerts for sudden movements that do not represent a clinical risk. Furthermore, the centralized cloud-based model presents a single point of failure, where the loss of domestic connectivity disables triage at the most critical moment [36]. This article proposes a multi-

agent framework where intelligence resides in the patient's device, specifically by deploying autonomous agents on a ESP32/SBC [37] that functions as a dedicated node. By shifting triage to the edge, a resilient system is achieved, capable of self-managing the veracity of alerts through debate between biological and contextual agents.

To reflect this work, the article is organized as follows: Section 2 presents the different IoMT Layered Architectures; Section 3 details the modeling of the Proposed Multi-Agent Framework; Section 4 outlines the Resilient Triage methodology employed; and finally, the results discussion and conclusions are presented.

2 IoMT Layered Architecture and Hardware Fundamentals

The centralized model, depicted on the left side of Figure 1, illustrates the traditional architecture upon which most current remote monitoring solutions are based [2]. This paradigm is characterized by a monolithic and reactive structure,

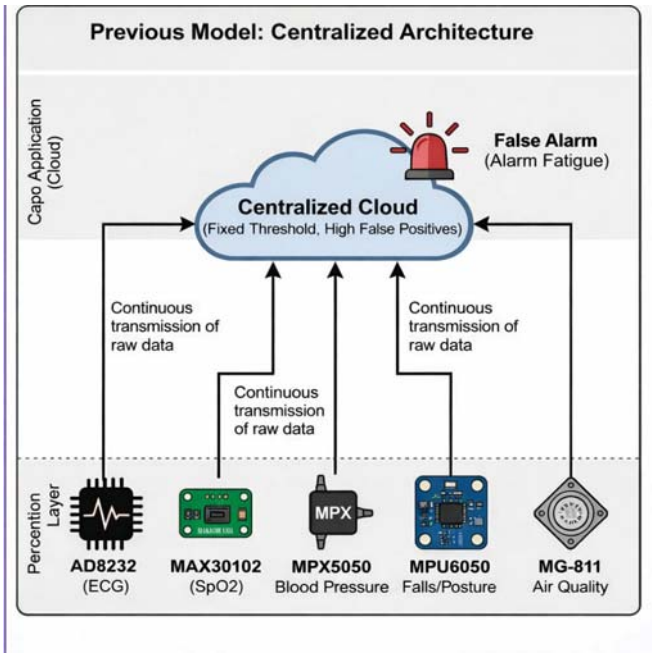

Fig. 1. Centralized Layered Architecture

with the following key points:

– Raw Data Continuous Stream: Sensors located in the perception layer function as simple collectors sending data through separate or uncoordinated

channels directly to the cloud. There is no intelligent preprocessing or local decision-making on the patient's device.
- Monolithic Cloud Intelligence: In this approach, the analytical capacity or intelligent "reasoning" resides predominantly in remote servers. The cloud receives the metrics and applies "static thresholds" configurable by professionals. If a particular sensor reading exceeds a predetermined value (e.g., a sudden movement in the accelerometer), an alarm is triggered.
- Isolated Processing and False Positives: The main design flaw of this model is that it processes sensor signals in isolation. As the figure 1 illustrates, this results in a high rate of false positives; a sudden peak in an accelerometer could be erroneously interpreted as a fall if it is not correlated with pulse stability, blood pressure, or even environmental variables.
- Alarm Fatigue: The excess of false alarms generated by this isolated analysis creates "alarm fatigue" in medical staff, a documented problem that dangerously reduces the effectiveness of response services.

The proposed system (see Fig. 2) is structured into hierarchical layers optimized for operational resilience and communication efficiency. These are detailed in the following subsections.

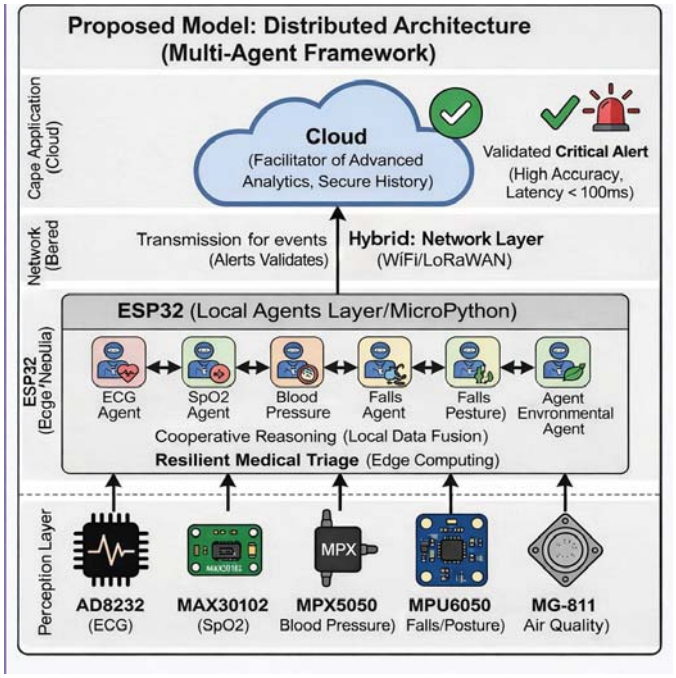

Fig. 2. Proposed Architecture

2.1 Perception Layer

This layer corresponds to the physical sensors (AD8232 for ECG, MAX30102 for SpO2, MPX5050 for pressure, MPU6050 for falls, and MG-811 for the environment) that constitute the physical base of the system and collect raw data sensed from individuals or patients under remote health care. The sensor details are as follows:

- **AD8232 (ECG Module):** Captures analog biopotentials to analyze cardiac stability and detect arrhythmias.
- **MAX30102 (Pulse Oximeter):** Uses photoplethysmography (PPG) technology to measure blood oxygenation (SpO2) and tissue perfusion.
- **MPX5050 (Pressure Transducer):** Enables non-invasive blood pressure (NIBP) monitoring using the oscillometric method, essential for evaluating shock states or hypertensive crises.
- **MPU6050 (Accelerometer/Gyroscope):** Analyzes patient kinematics to detect impacts and loss of verticality (falls/posture).
- **MG-811 (Environmental Sensor):** Evaluates air quality to correlate the patient's environment with possible episodes of respiratory distress.

2.2 Edge/Fog Layer

The physical base is the dual-core ESP32 microcontroller (160/240 MHz) or SBC. Its architecture allows it to host concurrent processes: while one core manages signal capture, the other executes the agent logic.

The ESP32/SBC block represents the key paradigm shift. It visually illustrates how intelligence is delegated to the network edge through a society of local agents (ECG, SpO2, Pressure, Falls, Environmental) that perform "Cooperative Reasoning" to fuse data in situ.

- **ECG Module (AD8232):** Captures cardiac biopotentials (Lead II) for arrhythmia analysis.
- **Motion Module (MPU6050):** 6-axis accelerometry for impact and posture detection.
- **SpO2 Module (MAX30102):** Photoplethysmography for oxygen saturation.
- **Additional Modules:** Different sensors that can enter the agent cooperation ecosystems.

2.3 Network Layer: Hybrid Connectivity Strategy

A resilient medical triage cannot depend on a single communications channel. A dual strategy is implemented:

- **WiFi or 4G/5G mobile network (Primary Channel):** Used for high-frequency streaming (e.g., ECG) when infrastructure is available.
- **LoRaWAN (Backup Channel):** Uses *Chirp Spread Spectrum* (CSS) modulation to send critical alerts kilometers away, even with low signal penetration or electrical failures.

2.4 Application Layer (Cloud)

It reflects the transformation of the data's final destination. On the left, we see the traditional centralized cloud processing signals in isolation, resulting in high false positive rates and the problem of "Alarm Fatigue". The cloud has become an "Advanced Analytics Facilitator" and secure historical record, receiving critical alerts already confirmed with high precision.

In the proposed framework, the functioning of the application layer undergoes a transformation; it now becomes an advanced analytics facilitator and a visualization interface for healthcare professionals. The logic for immediate and urgent triage has shifted to the lower layers (fog and edge), ensuring that the system can respond in milliseconds to critical situations. As a result, the cloud is freed from real-time processing to focus on maintaining a holistic view of the patient's state, stored centrally to facilitate diagnoses.

To achieve this level of specialization, the application layer manages mass storage services and secure databases, utilizing tools like PostgreSQL or InfluxDB for optimized time-series management.

Atop this infrastructure, other deep learning algorithms are deployed and shown in visual interfaces, such as Grafana dashboards, allowing medical staff to observe the long-term outlook and make informed treatment decisions.

The operation of this layer is also governed by a design focused on data privacy and anonymization for training, enabling audits.

3 Modeling of the Proposed Multi-Agent Framework

We define the system as a tuple $\mathcal{S} = \langle \mathcal{A}, \mathcal{O}, \mathcal{C}, \mathcal{R} \rangle$, where:

- $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ is the set of autonomous agents.
- $\mathcal{O}$ is the observation space of the environment and the patient.
- $\mathcal{C}$ is the communication protocol.
- $\mathcal{R}$ is the reasoning engine based on triage rules.

3.1 Triage Agent (A_{TRJ})

The A_{TRJ} is the coordinating agent whose function is to synthesize the overall state of the patient. Its behavior is governed by a Belief-Desire-Intention (BDI) model.

- **Beliefs (B):** B_1: Current state of the sensors; B_2: Personalized thresholds for the patient; B_3: Current triage level.
- **Desires (D):** D_1: Monitor that the patient is in the "Green" (stable) triage level; D_2: Validate threshold exceedance; D_3: Change triage level; D_4: Send an alarm signal.
- **Intentions (I):** I_1: Monitor patient; I_2: Validate critical event; I_3: Change triage level; I_4: Execute emergency protocols if a critical event is validated.

3.2 Cooperative Reasoning for False Positive Reduction

The core innovation of our proposal in this work is the cross-validation algorithm. Consider the fall event E_{CAI}. In traditional systems:

$$Alarma(t) = 1 \iff |\boldsymbol{a}(t)| > \tau_{impacto} \tag{1}$$

Where $\boldsymbol{a}$ is the acceleration and $\tau_{impacto}$ the threshold. This model fails if the sensor is accidentally dropped. In our model, the A_{TRJ} calculates a probability of a real event $P(E_{real})$:

$$P(E_{real}|E_{CAI} = 1) = \frac{P(E_{CAI} = 1|E_{real}) \cdot P(E_{real})}{P(E_{CAI} = 1)} \tag{2}$$

To refine this probability, the agent consults the physiological agents (A_{ECG}, A_{SpO2}). If the hemodynamic state is stable post-impact, $P(E_{real})$ decreases, classifying the event as a technical false positive.

4 Resilient Triage Methodology

The triage process is divided into three severity levels, managed autonomously by agents at the edge.

Red Alert: Confirmed Emergency ($CONFIRMED$) The **Red** state is reached when the function $Validar(Sospecha, Respuestas)$ returns a high truth value. This occurs when there is a **positive physiological correlation** between the detected event and the patient's vital signs. **Activation Criteria:**

- **Multiple Evidence:** The emitting agent A_i detects an anomaly and at least one reference agent A_j confirms a reactive alteration.
- **Clinical Example:** The Falls Agent detects an impact ($A_{CAI} \rightarrow Sospecha$) and the ECG Agent reports immediate tachycardia (> 120 BPM) or severe post-impact bradycardia.
- **System Action:** Immediate activation of external emergency protocols (ambulance, coronary unit) and critical notification to the monitoring center.

Yellow Alert: Uncertainty and Precaution ($UNCERTAIN$) The **Yellow** state is the major innovation in terms of resilience. It is activated when the system detects an anomaly but cannot confirm its physiological origin due to contradictory or incomplete data.

Activation Criteria:

- **Data Inconsistency:** The emitting agent reports a severe anomaly, but support agents report stable signs.
- **Partial Network/Sensor Failure:** If a reference agent does not respond within the window $\Delta t = 2000ms$, the system does not discard the alert due to "silence", but escalates it to Yellow as a precaution.

– **Clinical Example:** The SpO2 sensor detects a sudden desaturation ($<$ 85%), but the patient exhibits vigorous movement according to the accelerometer (possible movement artifact or detached sensor).
– **System Action:** Notification to local caregiver or family member for visual verification. The system increases the sampling frequency of all sensors for the next 5 minutes.

Decision Logic Summary
Table 1 summarizes the state transition based on cooperative validation:

Table 1. Triage Agent Decision Matrix

Suspicion (A_i)	**Support (A_j)**	**Final State**	**Clinical Meaning**
Impact (Fall)	Tachycardia	**RED**	Confirmed trauma with stress.
Hypotension	Dizziness/Environment	**RED**	Detected syncope.
Impact (Fall)	Stable rhythm	**YELLOW**	Possible device drop.
Tachycardia	High Temp.	**YELLOW**	Possible thermal stress.
Desaturation	No response A_j	**YELLOW**	Sensor failure / Technical doubt.
High pulse	High activity	**GREEN**	Normal patient exercise.

The **Distributed Validation Algorithm** is based on a quorum query among agents. In the proposed multi-agent framework, the Triage Agent (A_{TRJ}) categorizes alerts based on evidence convergence. This section details the clinical and technical logic behind the **Red** and **Yellow** states. The algorithm is broken

Algorithm 1 Cooperative Validation of Critical Alerts

Emitter Agent (A_i): Detects anomaly $X > \tau_i$
$A_i \rightarrow A_{TRJ}$: INFORM(Sospecha(X))
Triage Agent (A_{TRJ}):
$List_Agentes_Ref \leftarrow \{A_{ECG}, A_{AMB}, A_{SpO2}\} \setminus \{A_i\}$
for all $A_j \in List_Agentes_Ref$ **do**
    $A_{TRJ} \rightarrow A_j$: QUERY-IF(Estado_Signos_Vitales)
    end for all
    $Respuestas \leftarrow$ Wait for synchronization ($\Delta t = 2000ms$)
    if $Validar(Sospecha, Respuestas) == CONFIRMED$ **then**
        $Triaje \leftarrow ROJO$ (Activate External Emergency)
    else if $Validar(Sospecha, Respuestas) == UNCERTAIN$ **then**
        $Triaje \leftarrow AMARILLO$ (Notify local caregiver)
    else
        $Triaje \leftarrow VERDE$ (Discard as False Positive)
    end if

down into four critical stages that guarantee system resilience and accuracy in remote diagnosis:

- **Initial phase (Lines 1-2):** Execution begins asynchronously when a sensor agent (A_i) detects that a variable X has exceeded the critical threshold τ_i. Instead of escalating the alert directly to medical services (which would generate a false positive if the sensor is misplaced), the agent informs its "suspicion" to the Triage Agent (A_{TRJ}).
- **Dynamic Orchestration (Lines 4-7):** The A_{TRJ} assumes the role of coordinator. It identifies available reference agents in the local IoMT network. It is crucial to note that the emitter is explicitly excluded to avoid circular redundancy. Through *QUERY-IF* messages, the current state of vital signs and environmental variables is requested from the peripheral agents.
- **Temporal Synchronization (Line 8):** This step is vital for resilience. The system defines a time window $\Delta t = 2000ms$. This controlled delay compensates for possible network latencies at the edge (*Edge*) and allows secondary physiological responses (such as reactive tachycardia following an impact) to be captured by the other agents.
- **Triage Deduction (Lines 9-15):** The internal $Validar$ function processes the collected evidence:
  - **CONFIRMED:** The suspicion is supported by concurrent data (e.g., sudden accelerometry + heart rate increase).
  - **UNCERTAIN:** Occurs when there is packet loss or contradictory data. In this case, the system opts for a conservative stance, notifying a local caregiver for visual inspection.
  - **GREEN (False Positive):** If vital signs are stable and there is no contextual correlation, the event is discarded, optimizing emergency system resources.

The proposed logic ensures that no decision errors exist due to a lack of data, as the system is "aware" of an agent's absence. If the A_{ECG} fails or disconnects, the A_{TRJ} adjusts the confidence weights in the decision, increasing the sensitivity of the A_{SpO2} as compensation. This constitutes **structural resilience**.

Regarding **Security and Privacy: Zero Trust** Given the sensitive nature of medical data, the framework implements a **Zero Trust** architecture [**?**]. Each μACP interaction requires mutual authentication and AES-256 encryption supported by the ESP32/SBC cryptographic engine. Furthermore, **Federated Learning** is utilized, enabling agents to train personalized models locally and only share parameter updates, ensuring that raw biological data never leaves the edge.

5 Discussion of Expected Results

Due to the nature of the proposed model, the original centralized system was compared with the proposed multi-agent framework and the theoretical improvements achieved.

The results demonstrate that collaboration among agents not only reduces noise in alerts but also optimizes latency. By processing validation at the local

Gateway (Edge), the response to a real fall is almost instantaneous, eliminating dependence on the round-trip time (RTT) to a cloud server.

From a medical perspective, the inclusion of environmental agents (A_{AMB}) explains physiological variations that would otherwise be alarming. For example, mild tachycardia (110 bpm) in a room with a temperature of $38°C$ (common in the NEA region's climate) may be a normal thermoregulatory response rather than an impending heart failure.

6 Conclusions and Future Work

The development of a multi-agent framework for medical triage in the IoMT represents a fundamental advance toward resilient digital health. The transition from centralized and fragile intelligence to a society of autonomous agents resolves the problem of alarm fatigue and guarantees constant surveillance even in digital isolation. The use of the μACP protocol and MicroPython democratizes access to hospital-grade technology at home, providing precise, explainable, and human triage. The proposed multi-agent framework establishes a solid foundation for a resilient IoMT. The transition from passive systems to collaborative agents enables:

- Threat Mitigation of network failure impacts in critical health applications.
- Provision of a more human and contextualized triage.
- Facilitation of the adoption of home hospitalization.

As future work, the implementation of Federated Learning among agents from different homes is planned to improve anomaly detection algorithms without compromising patient data privacy.

Acknowledgments. The authors thank CInApTIC and UTN Facultad Regional Resistencia for their technical support.

Disclosure of Interests. The authors declare that they have no conflicts of interest related to this research.

References

1. Gilboy, N., et al.: Emergency Severity Index (ESI) Handbook Fifth Edition. EMSC (2020). `https://media.emscimprovement.center/documents/Emergency_Severity_Index_Handbook.pdf`. Last accessed 2026/04/02.
2. Gramajo, S., Torres, C., Nuñez, S., Scappini, R., Roa, J., Montiel, R. (2025). Integrated IoT System for Remote Health Monitoring in Home Hospitalization. In Actas del XXII Congreso Latinoamericano de Software Libre y Tecnologías Abiertas, (pp. 598-602). Porto Alegre: SBC. https://doi.org/10.5753/latinoware.2025.16539
3. Mounika, K., Akram, S.V.: Wireless IoT-based continuous ECG monitoring using ESP32 and thingspeak based performance analysis across age groups. Sensors and Actuators A: Physical **399**, 117392 (2026). https://doi.org/10.1016/j.sna.2025.117392

4. Islam, U., Alatawi, M.N., Alqazzaz, A., Alamro, S., Shah, B., Moreira, F.: A hybrid fog-edge computing architecture for real-time health monitoring in IoMT systems with optimized latency and threat resilience. Sci. Rep. **15**(1), 25655 (2025). https://doi.org/10.1038/s41598-025-09696-3
5. Yıldırım, E., Cicioğlu, M., Çalhan, A.: Fog-cloud architecture-driven Internet of Medical Things framework for healthcare monitoring. Med. Biol. Eng. Comput. **61**(5), 1133–1147 (2023). https://doi.org/10.1007/s11517-023-02776-4
6. Shah, B., Junaid, M., Rustam, H., Habib, M., Anwar, S.: AI-Driven Fog-Edge Computing for IoMT Systems: Architecture and Use Cases. Proceedings of the AAAI Symposium Series **6**(1), 42–48 (2025). https://doi.org/10.1609/aaaiss.v6i1.36023
7. Arregui Almeida, D., Chafla Altamirano, J., Román Cañizares, M., Játiva, P.P., Guaña-Moya, J., Sánchez, I.: Gateway-Free LoRa Mesh on ESP32: Design, Self-Healing Mechanisms, and Empirical Performance. Sensors **25**(19), 6036 (2025). https://doi.org/10.3390/s25196036
8. Chen, L., et al.: Energy-Efficient Fall-Detection System Using LoRa and Hybrid Algorithms. MDPI (2022). `https://www.mdpi.com/2313-7673/10/5/313`. Last accessed 2026/04/02.
9. IoT For All: From 'Always Connected' to 'Fully Off-Grid': A Case Study in Offline IoT Solutions (2024). `https://www.iotforall.com/offline-iot-system-engineering`. Last accessed 2026/04/02.
10. Emergent Mind: Resource-Constrained IoT Devices. Knowledge Base (2024). `https://www.emergentmind.com/topics/resource-constrained-iot-devices`. Last accessed 2026/04/06.
11. Garmendia-Orbegozo, A., Anton, M.A., Nuñez-Gonzalez, J.D.: Reduction of Vision-Based Models for Fall Detection. Sensors **24**(22), 7256 (2024). https://doi.org/10.3390/s24227256
12. PapersFlow: Fall Detection Algorithms: Research Guide & Papers — Context-Aware Activity Recognition Systems. PapersFlow (2024). `https://papersflow.ai/research/topics/context-aware-activity-recognition-systems/fall-detection-algorithms`. Last accessed 2026/04/10.
13. Habib, M.A., et al.: Fall Prediction and Prevention Systems: Recent Trends, Challenges, and Future Research Directions. PMC (2017). `https://pmc.ncbi.nlm.nih.gov/articles/PMC5713074/`. Last accessed 2026/04/02.
14. Zhou, L, Wang, L.: A Multi-Modal CNN-LSTM Framework with Multi-Head Attention and Focal Loss for Real-Time Elderly Fall Detection. arXiv:2603.22313v1 (2026). `https://arxiv.org/html/2603.22313v1`. Last accessed 2026/04/02.
15. Lin, H.-C., Chen, M.-J., Lee, C.-H., Kung, L.-C., Huang, J.-T.: Fall Recognition Based on an IMU Wearable Device and Fall Verification through a Smart Speaker and the IoT. Sensors **23**(12), 5472 (2023). https://doi.org/10.3390/s23125472
16. Mohan, D., Chong, P.H.J., Gutierrez, J.: A Novel Cooperative AI-Based Fall Risk Prediction Model for Older Adults. Sensors **25**(13), 3991 (2025). https://doi.org/10.3390/s25133991
17. Chen, J., Qiu, Y., Xu, Y., Li, L., Wu, M., Wang, L.: Multi-sensor fusion for AI-driven behavior planning in medical applications. Front. Phys. **13**, 1588715 (2025). https://doi.org/10.3389/fphy.2025.1588715
18. Konopko, K., Janczak, D., Walendziuk, W.: Implementation of Vital Signs Detection Algorithm for Supervising the Evacuation of Individuals with Special Needs. Sensors **25**(23), 7391 (2025). https://doi.org/10.3390/s25237391
19. Tomer, V., Sharma, S., Davis, M.: Resilience in the Internet of Medical Things: A Review and Case Study. Future Internet **16**(11), 430 (2024). https://doi.org/10.3390/fi16110430

20. Pastore, F., Anwar, R., Jabeur, N., Ali, S.: Intelligent Fusion: A Resilient Anomaly Detection Framework for IoMT Health Devices. Information **17**, 117 (2026). https://doi.org/10.3390/info17020117
21. Ravisankar, S., Maheswar, R.: SecureEdge-MedChain: A Post-Quantum Blockchain and Federated Learning Framework for Real-Time Predictive Diagnostics in IoMT. Sensors **25**(19), 5988 (2025). https://doi.org/10.3390/s25195988
22. Chen, C.H., Tsai, C.H., Peng, Y.S.: Exploring Health Misinformation Detection with Multi-Agent Debate. In: Accomazzi, A., Ghosal, T., Grezes, F., Lockhart, K. (eds.) Proceedings of the Third Workshop for Artificial Intelligence for Scientific Publications. pp. 16–21. Association for Computational Linguistics (2025). https://doi.org/10.18653/v1/2025.wasp-main.3
23. De La Iglesia, D.H., Villarrubia González, G., Sales Mendes, A., Jiménez-Bravo, D.M., Barriuso, A.L.: Architecture to Embed Software Agents in Resource Constrained Internet of Things Devices. Sensors **19**(1), 100 (2018). https://doi.org/10.3390/s19010100
24. Li, C., Wu, J., Du, Q., Yu, S., Zou, R., Yu, K., et al.: ACPs: Agent collaboration protocols for the internet of agents. In: 2025 9th IEEE International Conference on Network Intelligence and Digital Content (IC-NIDC). pp. 342–346. IEEE (2025).
25. SmythOS: An Introduction to FIPA Agent Communication Language: Standards for Interoperable Multi-Agent Systems (2024). `https://smythos.com/developers/agent-development/fipa-agent-communication-language/`. Last accessed 2026/04/02.
26. Bellifemine, F., Poggi, A., Rimassa, G.: Developing multi-agent systems with a FIPA-compliant agent framework. Software: Practice and Experience **31**(2), 103–128 (2001).
27. Wang, J., Chen, M., Zhou, J., Li, P.: Data communication mechanism for greenhouse environment monitoring and control: An agent-based IoT system. Inf. Process. Agric. 7(3), 444–455 (2020). https://doi.org/10.1016/j.inpa.2019.11.002
28. Ehtesham, A., Singh, A., Gupta, G.K., Kumar, S.: A survey of agent interoperability protocols: Model Context Protocol (MCP), Agent Communication Protocol (ACP), Agent-to-Agent Protocol (A2A), and Agent Network Protocol (ANP). arXiv preprint arXiv:2505.02279 (2025).
29. Mallick, S., Chebolu, I : μACP: A Formal Calculus for Expressive, Resource-Constrained Agent Communication (v2). arXiv:2601.00219v2 (2026). `https://arxiv.org/html/2601.00219v2`. Last accessed 2026/04/10.
30. Mallick, A., Chebolu, I.: The Micro Agent Communication Protocol (μACP). Internet-Draft draft-mallick-muacp-02, Internet Engineering Task Force (2026). `https://datatracker.ietf.org/doc/draft-mallick-muacp/02/`
31. Mallick, S., Chebolu, I.: [2601.00219] μACP: A Formal Calculus for Expressive, Resource-Constrained Agent Communication. arXiv (2026). `https://arxiv.org/abs/2601.00219`. Last accessed 2026/04/02.
32. Jasim, A.A., Ata, O., Salman, O.H.: Multisource data framework for prehospital emergency triage in real-time IoMT-based telemedicine systems. International Journal of Medical Informatics **192**, 105608 (2024).
33. Ganjali, R., Golmakani, R., Ebrahimi, M., Eslami, S., Bolvardi, E.: Accuracy of the Emergency Department Triage System using the Emergency Severity Index for Predicting Patient Outcome; A Single Center Experience. Bull. Emerg. Trauma **8**(2), 115–120 (2020). https://doi.org/10.30476/BEAT.2020.46452
34. Mlak, S., Yahyaoui, A., Abdellatif, T.: IoT emergency healthcare system using ML-based triage. International Journal of Sensor Networks **45** (2024). https://doi.org/10.1504/IJSNET.2024.10064029

35. Sharma, K., Sivadas, H., Reddy, S.: ED-Triage-Agent: A Framework for Human-AI Collaborative Emergency Triage. medRxiv (2026).
36. Mallick, S., Chebolu, I.: [2601.00219] μACP: A Formal Calculus for Expressive, Resource-Constrained Agent Communication. arXiv (2026). https://arxiv.org/abs/2601.00219. Last accessed 2026/04/02.
37. Bhattacharya, P., Mukherjee, A., Bhushan, B., et al.: A secured remote patient monitoring framework for IoMT ecosystems. Scientific Reports **15**, 22882 (2025). https://doi.org/10.1038/s41598-025-04774-y

Research Data Management in Stochastic Optimization: A Case Study on Heuristic Search and Retrospective Curation

Maria Harita[1 [0000−0002−3732−4717]], Dolores Rexachs[1 [0000−0001−5500−850X]], Alvaro Wong[1 [0000−0002−8394−9478]], and Emilio Luque[1 [0000−0002−2884−3232]]

Universitat Autònoma de Barcelona, Barcelona, Spain, 08193
Computer Architecture and Operating System Department
corresponding author: mariadelosangeles.harita@uab.cat
http://https://webs.uab.cat/hpc4eas/

Abstract. The complexity of stochastic optimization and large-scale benchmarking has elevated Research Data Management (RDM) from a peripheral task to a core methodological requirement. While the FAIR principles, that ensures data is Findable, Accessible, Interoperable, and Reusable, provide high-level guidance, their practical implementation in doctoral-level computational research remains insufficiently documented. In fact, existing guidelines are often difficult to apply as they do not immediately adapt to the diverse and specific scenarios of high-dimensional research. This paper presents a case study on the RDM lifecycle of a heuristic framework, documenting the transition from non-standardized experimental logs to FAIR-compliant datasets in the CORA.RDR repository. We identify critical challenges such as data heterogeneity, volume, and the technical debt generated by stochastic variability across thousands of independent runs. In response, we propose a lightweight, practical framework based on hierarchical structures, self-documenting naming conventions, and multi-level metadata schemas. The contribution of this work lies in bridging the gap between theoretical RDM and the non-linear reality of doctoral research. By providing actionable lessons on retrospective curation and interpretable transparency, we offer a roadmap for researchers to ensure the long-term reusability and integrity of complex computational artifacts.

Keywords: Research Data Management · Stochastic Optimization · Open Science · FAIR principles · Reproducibility

1 Introduction and Context

The increasing reliance on computational methods across scientific disciplines has led to substantial growth in the volume, complexity, and heterogeneity of research data. In fields such as optimization, artificial intelligence, and simulation-based decision support, experimental workflows often involve large-scale benchmarking, stochastic algorithms, and multi-dimensional problem spaces. Despite

this complexity, data management is frequently treated as a secondary administrative concern rather than a core component of the research methodology [13]. Recent initiatives, particularly the FAIR (Findable, Accessible, Interoperable, and Reusable) guiding principles [1], advocate for these four core pillars in data management. Similarly, the growing emphasis on reproducible research highlights the importance of transparent data and code sharing practices [2]. However, while these frameworks provide high-level guidelines, the existing documentation and RDM guidelines often prove difficult to apply in practice, as they do not immediately adapt to the diverse and specific scenarios encountered in computational research. This mismatch limits their practical utility for researchers dealing with high-dimensional and stochastic workflows.

In optimization research, this gap is particularly evident. Benchmark-driven evaluation requires multiple independent executions to account for stochastic variability, resulting in thousands of individual runs. As illustrated in Figure 1, the initial experimental phase follows a linearized execution model. The algorithm undergoes multiple independent runs, where each execution generates a complete set of outputs, including performance metrics, stochastic seeds, and granular logs. Without a dedicated management layer, this process leads to a linear accumulation of storage requirements, where the total footprint grows as $O(N \cdot C_{run})$, being N the number of configurations and C_{run} the data cost per run. This unchecked growth rapidly results in data entropy, where traceability is compromised as the link between specific seeds and their corresponding logs becomes fragmented [10].

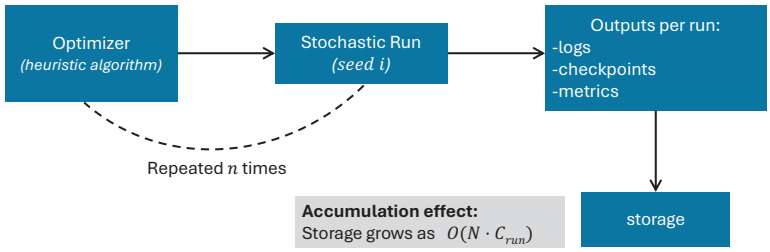

Fig. 1. Initial workflow of stochastic optimization without a management layer: repeated experimental runs generate independent sets of raw logs and metrics. This results in a storage overhead that scales linearly with the number of executions, complicating the traceability of stochastic seeds and configuration parameters.

The paradigm of Open Science has shifted the focus from isolated discovery to a transparent ecosystem, largely driven by the recognition of a reproducibility crisis [9]. Trust in research findings must now be earned through demonstrable trustworthiness and radical transparency [3]. To mitigate these issues, proponents of Open Science emphasize a tripartite focus: responsible research practices for rigorous conduct, transparency for complete reporting, and open science

for dissemination [6]. Central to this shift are Data Management Plans (DMPs) [4] and institutional infrastructures like the *Consorci de Serveis Universitaris de Catalunya [CSUC] - Repositori de Dades de Recerca (CORA.RDR)*, which ensure long-term stewardship [5].

However, a significant gap remains between these idealized guidelines and the messy reality of research. In practice, researchers grapple with non-linear workflows and a technical debt accumulated during a doctoral thesis. Real-world data management is frequently an iterative and retrospective effort, as data availability tends to decline rapidly if not managed from the outset.

This paper reflects on the data management journey of a doctoral thesis focused on a heuristic search algorithm [17]. The research involved the progressive generation of complex data that required significant post-hoc curation before its release on CORA.RDR [19] [18]. By describing this transition, we move beyond theory to provide an honest assessment of the limitations encountered.

In practice, operationalizing guidelines involves specific technical hurdles: designing scalable directory structures, defining consistent naming conventions, and preserving traceability in non-deterministic environments. These decisions, often shaped by individual workflows rather than universal standards, introduce critical methodological considerations. To address this, we propose a conceptual framework (Figure 2) where the experimental executions and their outputs are not isolated events but are governed by a transversal data management layer. This layer ensures that every algorithmic output is indexed and tied to its specific configuration and metadata, bridging the gap between raw execution and FAIR publication.

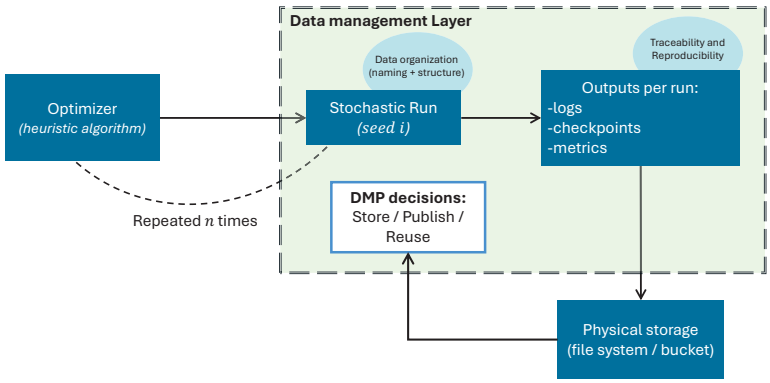

Fig. 2. Conceptual framework for data management in stochastic optimization. The experimental pipeline generates outputs, while the transversal data management layer governs organization, traceability, and curation decisions.

The main contribution of this work is twofold. First, it provides an in-depth analysis of data management challenges in stochastic optimization grounded in a

real research workflow. Second, it introduces a practical and reproducible framework that can be adapted by researchers facing similar experimental complexities, effectively bridging the gap between high-level principles and computational research practice.

The remainder of the paper is structured as follows: Section 2 reviews relevant concepts in research data management, reproducibility, and the engineering challenges of production-level. Section 3 introduces the experimental case study, detailing the stochastic nature of the heuristic algorithms used and the resulting data swamp. Section 4 analyzes the practical challenges encountered during the doctoral research, focusing on technical liability and scalability. Section 5 presents the proposed lightweight data management framework, including its structural and metadata components. Section 6 details the retrospective curation process and the final dataset publication in CORA.RDR. Finally, Sections 7 and 8 discuss the lessons learned and provide concluding remarks on the future of adaptive Open Science policies.

2 Related Work and Practical Frameworks

The transition from conceptual Open Science to technical execution is documented as a path filled with engineering challenges. While institutional guidelines provide the normative what, recent literature focusing on production-level AI and Big Data highlights the how through a more pragmatic lens.

2.1 Data Management as Engineering: from AI and Big Data

Managing data for complex computational models introduces a layer of complexity that standard RDM guidelines often overlook. According to Munappy et al. [12], the data pipeline is as critical as the algorithm itself. In fields such as Deep Learning or stochastic optimization, they identify three pillars that resonate with our heuristic research:

- **Data Versioning and Lineage:** Small changes in hyperparameters or seeds lead to divergent results. Maintaining strict lineage is essential for reproducibility in production-grade environments.
- **Continuous Curation:** Unlike the collect-once, publish-once model, high-stakes computational research requires continuous monitoring to avoid data drift or loss of context.
- **The Pipeline Approach:** There is a necessary transition from manual handling to automated workflows where data management is integrated into the software engineering process.

This engineering-centric view is further supported by Theodorakopoulos et al. [11], who argue that data cleaning and quality assurance represent over 80% of the engineering effort. They specifically highlight the issue of **technical debt**, which occurs when rapid prototyping precedes structured management, a scenario which is nearly universal in doctoral research.

2.2 The Human Factor: Research Integrity and Documentation

Despite these technical frameworks, the human element remains a significant vulnerability. Bouter and Haven et al. [8] provide empirical evidence that the reproducibility crisis is often fueled by poor documentation. Survey findings indicate that 51.3% of researchers admit to Questionable Research Practices (QRPs), such as keeping inadequate notes.

In the context of a doctoral thesis, this suggests that technical pipelines (like those proposed by Munappy et al.) are often bypassed under research pressure. This pragmatic gap leads to the rapid decay of data availability (estimated at 17% annually, according to [7]), unless institutional repositories are utilized to anchor the data in a permanent, FAIR-compliant infrastructure.

The existing literature excels at describing production-level AI or systemic integrity, but there is a lack of documentation on how a lone researcher can mitigate the gap between messy experimental logs and FAIR datasets. This work builds upon these engineering and integrity principles to propose a lightweight, post-hoc curation framework tailored for the iterative nature of heuristic search.

3 Case Study: Data Generation in Heuristic Optimization

This section details the experimental context of a doctoral thesis focused on the development of a heuristic framework for high-dimensional optimization. In practice, this research faced the typical doctoral dilemma: during the intensive phases of algorithm design and publication, data management was often relegated to a secondary status, treated as a pending notebook of notes and local logs. The transition from this ad-hoc storage to a FAIR-compliant repository (CORA.RDR) revealed the contrast between prescriptive theory and experimental reality.

As illustrated in Figure 3, the idealized DMP assumes a linear, planned path. However, the heuristic research reality is a cyclical process where intensive stochastic generation leads to mounting technical debt.

3.1 Experimental Context and Methodology

The research centers on a methodology combining Montecarlo-based exploration with clustering mechanisms to solve problems in high-dimensional search spaces. This approach was evaluated across two divergent domains:

- **Synthetic Benchmarking:** Using standardized test suites (e.g., CEC benchmarks) to evaluate global optimization performance [16] [14].
- **Applied Scenarios:** Simulation-based decision support for operational configurations, such as resource allocation in emergency departments [15].

This dual strategy significantly increased the heterogeneity of the data, as each domain produced different artifacts, ranging from simple objective values to complex simulation traces.

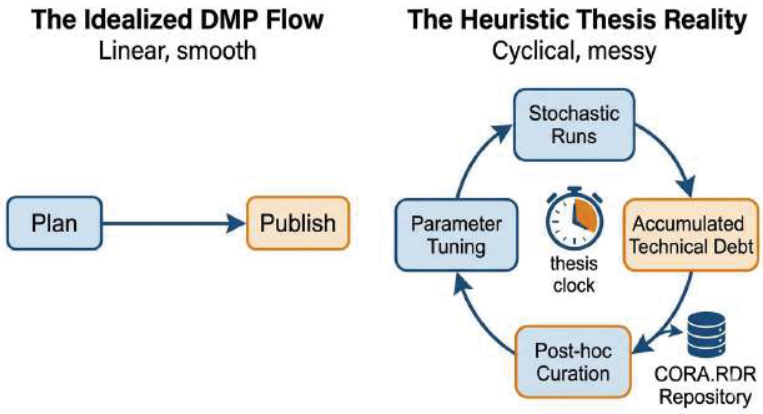

Fig. 3. The Heuristic Thesis Reality vs. the Idealized DMP. Intensive stochastic runs generate fragmented, undocumented logs that require significant post-hoc curation to resolve technical debt before repository deposit.

3.2 Execution Protocol and Data Volatility

To account for the stochastic nature of the heuristic, the protocol required 200 independent runs per configuration. Each run generated a unique set of outputs:

- **Raw Logs:** Step-by-step records of algorithm behavior and intermediate states.
- **Stochastic Context:** Critical but often forgotten metadata, such as random seeds and hardware environment.
- **Aggregated Results:** Statistical summaries (mean, best-so-far, convergence curves).

The scale of this process is non-trivial. With dozens of benchmark functions and hundreds of runs, the system rapidly produced thousands of individual data outputs. In practice, managing this volume without an automated pipeline led to a data swamp where linking specific results to their original configuration became increasingly difficult over time.

3.3 The Turning Point

The complexity described above made ad-hoc organization insufficient for reproducibility. As a doctoral researcher, the realization of this issue occurred at the moment of publication: the need to answer where the data should go, should it be deleted, kept in local drives, or published?

Initially, there was a lack of clarity regarding what data management actually entailed. The requirements for scalable storage, consistent naming conventions, and traceability between raw and derived data were not met by design, but rather identified through the necessity of post-hoc curation. The challenges encountered

during this retrospective organization process (technical, organizational, and human), are detailed in the next discussion on implementation hurdles.

4 Practical Challenges: From Experimental Chaos to FAIR Data

The experimental workflow of a doctoral thesis in heuristics reveals a breach between high-level RDM guidelines and the reality of large-scale stochastic optimization. This section analyzes the difficulties encountered when attempting to transition from raw execution logs to a structured, reusable dataset for CORA.RDR.

4.1 Scalability and the Data Swamp

The most immediate challenge is the exponential growth of artifacts. In our case study, executing 200 independent runs per configuration for dozens of benchmark functions generated thousands of individual files.

- **Volume vs. Structure:** Naive storage strategies (flat directories) quickly become navigable nightmares, hindering retrieval and post-processing.
- **Heterogeneity:** As shown in the curation pipeline (Figure 4), we deal with a noisy mix of sequential raw logs, aggregated performance metrics, and compact configuration files. Without a standardized mapping, integrating these diverse formats is an engineering bottleneck.

4.2 The Reproducibility-Traceability Dilemma

In stochastic environments, reproducibility is not just about sharing code; it is about preserving the exact execution context.

- **Stochastic Context:** Ensuring that random seeds and hardware/software versions are consistently linked to each of the 200 runs is critical. A single missing seed renders an entire experimental branch non-reproducible.
- **Traceability:** Maintaining the lineage between a specific version of the heuristic, its parameters, and the resulting convergence trajectory becomes increasingly difficult as the thesis evolves. Ambiguous naming conventions often lead to orphan data, where results exist but their originating conditions are lost.

4.3 Operational Overhead and Decision Making

Adequate documentation often feels like a trade-off against research speed, which is a form of technical debt that researchers accumulate to meet publication deadlines.

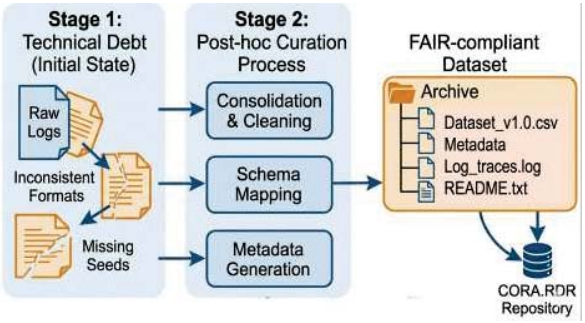

Fig. 4. The post-hoc curation pipeline required to resolve the Data Swamp. It illustrates the engineering effort needed to move from technical debt (left) to a FAIR-compliant archive (right).

- **Metadata Fatigue:** Determining the right granularity for metadata is non-trivial. Too much detail hinders usability; too little compromises integrity. In practice, metadata is often treated as a retrospective task, leading to inconsistencies.
- **Curation Decisions:** A key challenge was deciding what to publish. Should we upload the gigabytes of raw logs or only the processed summaries? Balancing completeness with repository constraints (like those of CORA.RDR) requires a selective strategy that is rarely addressed in standard RDM guides.

While FAIR principles provide a vision, they lack the "low-level" instructions for structuring execution-level data or managing stochastic variability. This work intends to move beyond these abstractions, proposing a framework that operationalizes these needs into a concrete, reproducible workflow.

5 A Lightweight Framework for Stochastic Data Management

To narrow the difference between idealized DMP guidelines and the experimental reality of heuristic research, we propose a structured, lightweight framework. This framework is not merely a set of folders, but a transversal data management layer designed to govern the lifecycle of stochastic artifacts, ensuring they move from raw logs to FAIR datasets systematically.

The framework operationalizes RDM through three main pillars: structural organization, standardized lineage, and multi-level metadata.

5.1 Structural Organization and Scalability

The first pillar addresses the Data Swamp by enforcing a hierarchical directory structure. Unlike flat or ad-hoc organizations, this design separates the immutable experimental evidence (runs/logs) from the volatile analysis results.

```
root/
benchmarks/ (Problem definitions & instances)
runs/ (Execution-level data)
<benchmark_ID>/
run_001/ (Individual stochastic output)
results/ (Aggregated & processed data)
metadata/ (Machine-readable CSV/TXT schemas)
```

This separation ensures that raw data remains untouched, facilitating independent validation, while the `results/` directory can be regenerated or updated as analysis evolves.

5.2 Standardized Lineage: Naming and Seeds

To ensure traceability without the need for complex database lookups, the framework implements a Self-Documenting Naming Convention. Each file acts as a primary key for its own context.

A typical filename follows the pattern:
`<benchmark>_<dimension>_<runID>_<seed>.ext` (e.g., `F12_30D_run045_seed1234.csv`).

By embedding the stochastic seed and the runID directly into the filename, we eliminate the risk of orphan data and ensure that any single file can be traced back to its specific execution branch in the stochastic search space.

5.3 The Metadata Layer: From Logs to Knowledge

Reproducibility in heuristics depends on capturing context that code alone cannot provide. Our framework adopts a layered metadata approach, visualized in Figure 5.

- **Global level:** Hardware environment, software versions, and core heuristic operators.
- **Run level:** Timestamps, specific operator parameters, and performance metrics (runtime, fitness).

By using machine-readable formats (TXT/LOG), this metadata layer allows for automated ingestion by repository harvesters or future meta-heuristic analysis tools.

5.4 Implementation and Repository Deposit

The final stage of the framework governs the selection of data for public release. Recognizing storage constraints in repositories like CORA.RDR, the framework prioritizes:

1. **High-Value Raw Data:** A representative subset of full logs for reproducibility.

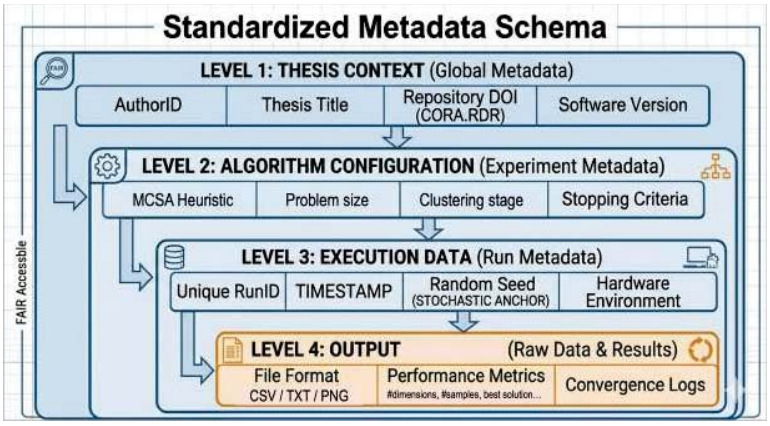

Fig. 5. Hierarchical Metadata Schema for Heuristic Research. The proposed framework organizes data via a multi-level inheritance structure. In practice, this Standardized Schema (the map of what to save) acts as the bridge to resolve the problem identified in Figure 4. Level 1 provides global context (Thesis, AuthorID), which is inherited by Level 2 (Algorithm Configuration, such as problem size). Level 3 links this configuration to a Unique RunID and its critical Random Seed, ensuring that Level 4 (Output CSVs, logs or Figures) remains fully traceable and interoperable with CORA.RDR.

2. **Processed Aggregates:** The complete set of results used for publication.
3. **Integrity Documents:** The code, the metadata schema, and a detailed README.

This selective strategy ensures that the resulting deposit is not just open, but usable and interpretable by the community, fulfilling the transition from an internal thesis reality to a professional scientific contribution.

6 Dataset Publication and Reproducibility

In contrast to idealized data management pipelines, the datasets in this work were retrospectively structured from heterogeneous experimental outputs.

This section describes the practical process followed for organizing, curating, and publishing the experimental datasets generated throughout the research life-cycle. Unlike idealized data management pipelines, the datasets presented in this work were not initially designed with publication in mind. Instead, they were retrospectively structured from heterogeneous experimental outputs accumulated over multiple years of research activity.

6.1 Temporal Evolution and Data Accumulation

The experimental data were generated over an extended period of approximately one year per publication cycle, beginning early in the doctoral research process.

Initial experiments were conducted incrementally and stored in local directories and institutional cloud storage (e.g., OneDrive), without an explicit long-term publication structure in mind.

As research outputs were progressively prepared for conference submissions, previously generated results were reused and reorganized to construct experimental tables and performance comparisons. This iterative process led to a non-linear accumulation of datasets, where early benchmarks (e.g., classical optimization functions with limited runs) coexisted with more recent large-scale experiments, including high-dimensional benchmark suites and applied simulation scenarios.

6.2 Retrospective Curation: From Local Logs to CORA.RDR

The transition from active research to data publication required a significant retrospective effort. This section describes the process of transforming a heterogeneous collection of local files into the structured datasets currently hosted in CORA.RDR.

6.3 Initial State: Technical Debt and Fragmentation

During the early stages of the doctoral thesis, data was stored intuitively in dedicated directories to preserve experimental separation. However, the lack of a global naming convention led to the coexistence of multiple schemas:

- **Benchmark runs:** `parallel_benchmarks_cec2017`
- **Knapsack optimization:** `EX01_instance`, `EX02_instance`, n `instance...`
- **Heuristic/Simulation outputs:** `HeuristEDStaffs*`, `sensor-processed-[num]*`

These naming patterns reflected the evolution of the research rather than a unified design, resulting in significant technical liability. The diversity of formats (CSV, TXT, LOG, and even exploratory PNG visualizations) further complicated the integration task.

6.4 Consolidation and Cleaning Pipeline

The need for systematic organization emerged during the preparation of conference and journal contributions. We implemented a two-stage consolidation process:

1. **Domain Categorization:** Grouping datasets by research line (e.g., stochastic benchmarks vs. applied simulation).
2. **Alignment with Publication Units:** Mapping raw artifacts to specific scientific papers to ensure that every result reported in the literature had a corresponding traceable data source.

Cleaning operations focused on resolving the reality of the thesis: removing duplicate files from storage transfers, excluding incomplete runs, and filtering out proof-of-concept executions. This process was not about data reduction, but about ensuring that every published file was part of a valid, traceable execution chain.

6.5 Final Dataset Structure and Reproducibility Strategy

The resulting datasets were organized into two major collections. The **Benchmarking collection** contains thousands of files spanning various dimensionalities (4D to 10D) with 200 independent runs each. The **Simulation collection** includes Montecarlo configurations and post-simulation indicators (LoS, occupancy, etc.) derived from a validated ED model (EDSimulator [20] [21]).

A key concern during this phase was the reproducibility of experiments involving external dependencies. Since the heuristic search algorithm and the simulation environment are subject to external ownership or academic restrictions, we adopted an Output-Driven Reproducibility strategy. While the full software environment cannot be redistributed, we provide:

- Detailed README files describing the experimental methodology and file interpretation.
- Structured tables of parameters and configurations to compensate for any missing raw context.
- Fully interpretable outputs that allow researchers to validate the methodology's findings against the published results, even without direct access to the source code.

6.6 The Impact of Curation on Data Sustainability

The retrospective effort described here is the difference between data loss and scientific transparency. As shown in Figure 6, research data without active management is subject to rapid decay, the Data Swamp or Data Dumpster effect, where contextual trust is lost over time.

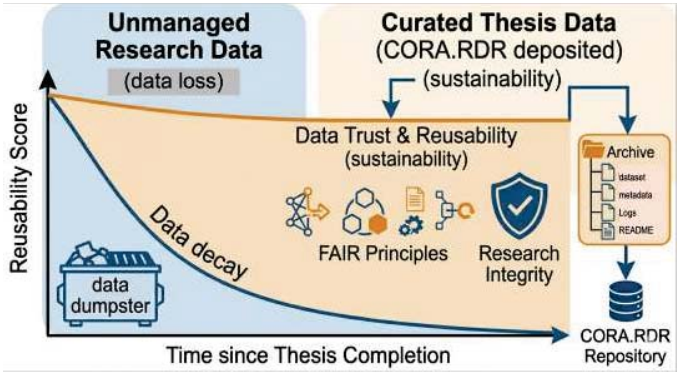

Fig. 6. Impact of RDM on data reuse and viability over time. Unmanaged data suffers from rapid decay and loss of contextual trust. In contrast, the post-hoc curation and deposit in CORA.RDR promotes sustained reusability, directly supporting research integrity.

By investing in curation, we ensure that the datasets remain viable beyond the duration of the doctoral thesis, fulfilling the commitment to Open Science despite the initial non-standardized start.

7 Lessons Learned and Practical Recommendations

This section synthesizes the insights gained from the transition between exploratory research and FAIR-compliant publication. Rather than presenting an idealized workflow, these lessons reflect the messy reality of a doctoral journey and provide a pragmatic roadmap for researchers in similar computational fields.

7.1 Metodological Lessons: Proactive vs. Retrospective Effort

The most significant realization was the high cost of retrospective curation. In the early stages, data management was implicit and driven by immediate needs. Transitioning this to an explicit, public-facing format months later was cognitively demanding and time-consuming.

- **Lesson:** Data management is not a final task, but an integral part of the research methodology.
- **Recommendation:** Treat data as a primary research output from day one. Even in exploratory phases, establishing a minimal Metadata Schema (as proposed in Section 5) prevents the accumulation of missing documentation.

7.2 Technical Lessons: Documentation over Perfection

In long-term research, heterogeneity is inevitable. Attempting to fully homogenize every dataset across different domains (benchmarks vs. simulation) proved impractical. However, we found that comprehensive documentation can mitigate structural inconsistencies.

- **Lesson:** A well-written README is more valuable than a perfect but poorly explained directory structure.
- **Recommendation:** Prioritize human-readable documentation. In our case, detailed descriptions of variable meanings and file-naming logic (e.g., embedding seeds in filenames) compensated for the lack of initial standardization.

7.3 Managing Constraints and Partial Reproducibility

A common fear among doctoral researchers is that partial data (due to proprietary algorithms or restricted simulation environments) is not worth publishing. Our experience shows that transparency about limitations strengthens trustworthiness.

- **Lesson:** Full reproducibility is the ideal, but interpretable transparency is a valid and valuable scientific contribution.

– **Recommendation:** Even if the algorithm is private, publish the outputs, the configuration parameters, and the statistical summaries. This allows the community to validate the results without requiring access to restricted environments.

7.4 Implications for Open Science Policy

This case study highlights a gap between high-level policy (FAIR, DMPs) and the iterative nature of a thesis. Based on this, we suggest:

1. **DMPs as Living Documents:** Policies should recognize that Data Management Plans must evolve. They are not administrative contracts but flexible roadmaps.
2. **Repositories as tools:** Platforms like **CORA.RDR** are more than storage; their metadata requirements act as a quality control mechanism that forces researchers to improve their own internal logic.
3. **Support for Post-hoc Curation:** Recognizing that much of the curation will be retrospective, institutions should provide lightweight tools and training for cleaning and restructuring legacy doctoral data.

7.5 Final Reflection: What would be done differently?

If this research were restarted, the focus would shift from managing files to designing a pipeline. A consistent naming convention, systematic logging of parameters, and early adoption of a metadata schema would have reduced the final curation effort by an estimated 70%.

In conclusion, Open Science in computational research is not about achieving formal perfection, but about documenting the research journey with enough honesty and structure that others can build upon it. The framework proposed here demonstrates that even with an imperfect start, any researcher can achieve a high standard of data integrity through disciplined, lightweight practices.

8 Conclusions

This paper has provided a pragmatic perspective on research data management (RDM) by documenting the transition of a doctoral project from initial experimental to FAIR-compliant publication. Our analysis demonstrates that while high-level frameworks like the FAIR principles provide a necessary foundation, their practical utility is limited by the complexity of adapting their general requirements to the engineering and operational realities of stochastic optimization.

The core contribution of this work is the demonstration that meaningful alignment with Open Science is achievable retrospectively. By applying a lightweight framework focused on structural organization, standardized lineage (seeds and run IDs), and layered metadata, we successfully transformed thousands of fragmented artifacts into reusable datasets. This process highlights that Openness

should not be measured by formal perfection from day one, but by the eventual interpretability and trustworthiness of the data deposited in repositories like CORA.RDR.

Our findings suggest that the scientific community must move toward more adaptive RDM models. For the doctoral researcher, data management should be recognized not as an administrative burden, but as a core engineering component of the thesis lifecycle. We argue that prioritizing Interpretable Transparency (providing well-documented outputs even when code or environments are restricted) is a more realistic and effective path for many applied research domains than the binary open vs. closed approach.

Ultimately, this case study serves as both a roadmap and a call to action for early-career researchers. The preservation of algorithmic transparency depends on the discipline of today's experimental workflows. By embracing RDM as a foundational research practice, we ensure that the results of today's heuristic explorations remain a stable foundation for tomorrow's scientific advancements.

Acknowledgments. This research has been supported by the Agencia Estatal de Investigación (AEI), Spain and the Fondo Europeo de Desarrollo Regional (FEDER) UE under contractPID2023-147955NB-I00, and partially funded by the *Fundación Escuelas Universitarias Gimbernat* (EUG).

References

1. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L.B., Bourne, P.E., et al.: The FAIR guiding principles for scientific data management and stewardship. Scientific Data **3**, 160018 (2016). https://doi.org/10.1038/sdata.2016.18
2. Stodden, V.: Reproducible research: Addressing the need for data and code sharing in computational science. Computing in Science & Engineering **12**(5), 8–12 (2010). https://doi.org/10.1109/MCSE.2010.113
3. Bouter, L.M.: Why research integrity matters and how it can be improved. Accountability in Research, 1–10 (2023). https://doi.org/10.1080/08989621.2023.2189332
4. LIBER Europe: Data Management Plans (DMPs) in a nutshell: Practical guide for research libraries (2025)
5. LIBER Europe: Improving access and reuse of research results: A LIBER action plan addressing copyright and data legislative issues (2025)
6. Borchert, M., Nizami, U., Stern, N., Lippincott, S.: Open science infrastructures speak about sustainability challenges. In: LIBER 2023 Annual Conference (2023)
7. Vines, T.H., Albert, A.Y.K., Andrew, R.L., Débarre, F., Bock, D.G., Franklin, M.T., Gilbert, K.J., Moore, J.-S., Renaut, S., Rennison, D.J.: The availability of research data declines rapidly with article age. Current Biology **24**(1), 94–97 (2014). https://doi.org/10.1016/j.cub.2013.11.014
8. Haven, T.L., Gopalakrishna, G., Tijdink, J.K., van der Schot, D., Bouter, L.M.: Promoting trust in research and researchers: How open science and research integrity are intertwined. BMC Research Notes **15**, 302 (2022). https://doi.org/10.1186/s13104-022-06169-y

9. National Academies of Sciences, Engineering, and Medicine: Reproducibility and replicability in science. The National Academies Press, Washington, DC (2019). https://doi.org/10.17226/25303
10. Hutter, F., Kotthoff, L., Vanschoren, J. (eds.): Automated Machine Learning: Methods, Systems, Challenges. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-05318-5
11. Theodorakopoulos, L., Theodoropoulou, A., Stamatiou, Y.: A state-of-the-art review in big data management engineering: real-life case studies, challenges, and future research directions. Eng **5**(3), 1266–1297 (2024). https://doi.org/10.3390/eng5030068
12. Munappy, A.R., Bosch, J., Holmström Olsson, H., Arpteg, A., Brinne, B.: Data management for production quality deep learning models: challenges and solutions. Journal of Systems and Software **191**, 111359 (2022). https://doi.org/10.1016/j.jss.2022.111359
13. Kotronoulas, G., Miguel, S., Dowling, M., Fernández-Ortega, P., Colomer-Lahiguera, S., Bağçivan, G., Pape, E., Drury, A., Semple, C., Dieperink, K.B., Papadopoulou, C.: An overview of the fundamentals of data management, analysis, and interpretation in quantitative research. Semin. Oncol. Nurs. **39**(2), 151398 (2023). https://doi.org/10.1016/j.soncn.2023.151398
14. M. Harita, A. Wong, R. Suppi, D. Rexachs, and E. Luque. A metaheuristic search algorithm based on sampling and clustering. *IEEE Access*, 12:15493–15508, 2024. https://doi.org/10.1109/ACCESS.2024.3354714.
15. M. Harita, A. Wong, D. Rexachs, E. Luque, E. Bruballa, and F. Epelde. Application of a sampling and clustering-based heuristic search algorithm to find an efficient staff configuration in an emergency department. *Expert Systems with Applications*, 295:128803, 2026. doi:https://doi.org/10.1016/j.eswa.2025.128803. https://www.sciencedirect.com/science/article/pii/S0957417425024212.
16. M. Harita, A. Wong, D. Rexachs, and E. Luque. KP01 solved by an n-dimensional sampling and clustering heuristic. In *Proceedings of the International Conference on Computational Science (ICCS)*, pages 229–236, Springer International Publishing, 2022. https://doi.org/10.1007/978-3-031-08754-7_31
17. M. Harita, A. Wong, D. Rexachs del Rosario, and E. Luque Fadón. Evaluation of a heuristic search algorithm based on sampling and clustering. In *IX Jornadas de Cloud Computing, Big Data & Emerging Topics*, 2021.
18. M. Harita, A. Wong, D. Rexachs, E. Luque, E. Bruballa, and F. Epelde. Replication data for: Application of a sampling and clustering-based heuristic search algorithm to find an efficient staff configuration in an emergency department. CORA.Repositori de Dades de Recerca, 2026. Version V1. https://doi.org/10.34810/data2951
19. M. Harita, A. Wong, R. Suppi, D. Rexachs, and E. Luque. Replication data for: A metaheuristic search algorithm based on sampling and clustering. CORA.Repositori de Dades de Recerca, 2026. Version V1. https://doi.org/10.34810/data3157
20. Z. Liu, E. Cabrera, M. Taboada, F. Epelde, D. Rexachs, and E. Luque. Quantitative evaluation of decision effects in the management of emergency department problems. *Procedia Computer Science*, 51:433–442, 2015.
21. Z. Liu, D. Rexachs, F. Epelde, and E. Luque. A simulation and optimization based method for calibrating agent-based emergency department models under data scarcity. *Computers & Industrial Engineering*, 2016. https://doi.org/10.1016/j.cie.2016.11.036.

Short papers

A Blockchain-Based MFA Model for Decentralization, Privacy, and Asymmetric-Key Security

Guillermo Daniel Romero Arregin[1], Miguel Méndez-Garabetti[1,2], Eduardo Piray[1]

[1] Universidad CAECE, Departamento de Sistemas, Mar del Plata, Argentina
gromeroarregin@gmail.com, {mmendezgarabetti,epiray}@ucaecemdp.edu.ar
[2] Free and Open Source Software/Hardware Research Laboratory (FOSSHLab), Mendoza, Argentina

Abstract. Authentication processes are fundamental for service providers and users alike. Currently, many users rely on centralized identity providers that verify identities to grant access to services. While these providers offer convenience, they require the disclosure of personal data, often leading to privacy concerns and third-party data access. Advancements in blockchain technology and smart contracts have enabled decentralized alternatives. Specifically, Sign-In with Ethereum (SIWE) leverages asymmetric cryptography, allowing users to authenticate via their public addresses and wallets. However, this approach introduces a significant risk: the loss or theft of private keys can result in the permanent loss of account control. This document proposes a multi-factor authentication (MFA) scheme that utilizes a smart contract to link a primary blockchain address with a secondary factor. This system provides an additional security layer against attackers without requiring personal information, effectively balancing account recovery resilience with user privacy.

Keywords: Blockchain · Multi-Factor Authentication · Sign-In with Ethereum · Smart Contracts · Privacy

1 Introduction

Authentication is essential for many applications and websites that provide Internet services, since it is the main mechanism used to protect both those services and the information they transmit or receive. Single-factor authentication (SFA) alone is often insufficient, and systems require an additional security layer. Multi-factor authentication (MFA) addresses this problem by combining multiple factors within the same authentication process, thereby providing a higher level of security than SFA alone [9].

Through the mechanism provided by Sign-In with Ethereum (SIWE), users can authenticate to web services using the same asymmetric keys with which they identify themselves on the blockchain. In this way, wallets can be used to

sign login requests and interact with services, replacing the traditional username-and-password combination [11, 2].

In this project an authentication model is proposed. This model makes use of multiple pairs of asymmetric keys to be used as MFA through the SIWE mechanism to provide users with a safe and robust method to associate their blockchain addresses with services and applications, while at the same time an additional layer of security is added by storing the necessary information in smart contracts to verify blockchain credentials. This research focuses not only on the scheme's security but also on ease of use for its users, monetary cost, storage cost, and processing cost in the blockchain, as all these factors are considered to be of great importance for this scheme's success and adoption.

To achieve this, the project considers the following specific objectives: to design an authentication mechanism that fulfills the proposed idea; to design the smallest possible smart contract that allows the implementation of the idea using only addresses as user attributes; to design a scheme based on SIWE and smart contracts and implement a prototype to test the proposal; and to evaluate the prototype by considering different qualitative and quantitative metrics.

2 Methodology

In order to prove the feasibility of the model an Ethereum-compatible prototype was implemented. Two authentication mechanisms were tested on this prototype, browsing the blockchain from the genesis block to locate the smart contract that contains the required main address or accessing directly through the smart contract's address. The factors used to evaluate this project are the existing vulnerabilities in the communications between components, storage space used by the smart contracts within the blockchain, the efficiency of the process to obtain the necessary data from the blockchain, the security of the smart contract's code and gas cost when deploying said smart contracts. The prototype mainly requires a SIWE-compatible service, a back-end that handles the logic of the model and verification of MFA data, a smart contract that contains said MFA data and a blockchain to store it. The implemented smart contract, written in Solidity [4], stores the contract type, the two addresses used to verify the user, and their respective getter methods declared as `external view` functions, which means that they do not modify blockchain state and therefore do not require a state-changing transaction when called. The IAM service Keycloak [6] was used to simplify interactions among the prototype components and to secure their communications. In addition, an OpenID Connect server was used as an identity provider in order to support SIWE and simplify prototype development. Consequently, the current implementation should be understood as partially centralized at the IAM layer, and a production deployment would require a different SIWE integration in order to reduce centralization and avoid additional security concerns.

For this prototype the authentication processes are handled through an Express.js [10] server. The server is responsible for completing a log-in attempt

by communicating with the blockchain to obtain the necessary smart contract and ensuring that the addresses entered by the user match those stored in the smart contract. Communication with the blockchain is carried out through the Ethers.js library [7], which is used to browse the blockchain, read transactions, and retrieve the information required by the authentication flow. While traversing the blockchain, the server queries transactions within blocks in order to identify smart contracts matching the model's contract type, and discards any other transaction as irrelevant to the model.

3 Results

This section presents the main results obtained from the prototype evaluation. For clarity, the results are organized into qualitative and quantitative findings.

3.1 Qualitative Results

User-to-contract binding Observation of the prototype reveals an important limitation. Because no user attributes other than blockchain addresses are stored, deployed smart contracts do not contain information that reliably links them to the legitimate users who rely on them, and such users may or may not coincide with the deployer address. As a result, the model cannot verify whether a contract was deployed by a legitimate user unless additional data are introduced or the process becomes partially centralized. This limitation reflects a central trade-off in the design: reducing metadata exposure improves privacy, but weakens the binding between user and contract. Within the current design, the only practical criterion available to distinguish a candidate contract from later conflicting deployments is deployment order. Under this assumption, the first smart contract associated with a given address is treated as valid and any later conflicting deployment is discarded, taking advantage of blockchain immutability. However, this means that the model does not cryptographically prevent an attacker from deploying a conflicting contract first; rather, it only makes such behavior economically less attractive because of deployment costs and the vast space of possible blockchain addresses. For this reason, deployment order should be understood as a practical heuristic rather than as a strong proof of legitimacy. Because the contract cannot be modified after deployment, there can only be one valid contract per address under the proposed interpretation. At the same time, the contract cannot be reliably linked to another contract for later modification without compromising some of the privacy, usability, or decentralization goals of the model.

Mechanism comparison Because the model ultimately relies on smart-contract deployment order for verification, the two mechanisms evaluated in the prototype become operationally similar. In practice, both require blockchain traversal or historical verification, either to locate the first contract associated with a given address or to verify that no earlier contract already exists. This requirement directly affects scalability and increases verification latency.

Architecture vulnerabilities Although the authentication flow is protected through Keycloak, communications to and from the blockchain must also be secured. In the current prototype, Ethers.js communicates with the blockchain through JSON-RPC over HTTP. Therefore, the architecture should enforce encrypted and authenticated transport, including HTTPS/TLS and, ideally, restricted exposure of JSON-RPC endpoints.

3.2 Quantitative Results

Storage costs within the blockchain Storage costs were estimated using the Hardhat Contract Sizer plugin [8], which analyzes Solidity code to report creation and runtime code size. According to the report, the proposed contract has a runtime size of 0.455 KiB and an init-code size of 0.806 KiB. Considering the Ethereum runtime bytecode limit of 24 KiB established in EIP-170 [1], the contract represents approximately 1.9% of the allowed maximum. The EVM optimizer was not enabled in this analysis because it was not expected to provide substantial benefits for this specific contract and could potentially increase gas consumption. Using these settings, the contract was compared against three widely known contracts obtained from Etherscan [5], as shown in Table 1. This comparison is intended only to contextualize code size and should not be interpreted as a comparison among functionally equivalent authentication mechanisms.

Table 1. Init and runtime code size of different contracts. Compiled using Solidity 0.8.28 and only considering the main contract of each code.

Contract	Init code size	Runtime code size
Used MFA contract	0.806 KiB	0.455 KiB
Tether USD	11.669 KiB	11.103 KiB
BNB	5.822 KiB	5.332 KiB
USDC	2.135 KiB	2.790 KiB

Security analysis of the smart contract The smart contract's security was tested with the tool Mythril [3], which analyzes EVM code to find vulnerabilities within it. Mythril did not report any errors or vulnerabilities.

Deployment cost Gas cost was estimated using the Hardhat Gas Reporter plugin [8]. Since the smart contract used in this model does not include state-changing methods intended for regular use, the most relevant factor for this analysis is deployment cost. Hardhat Gas Reporter indicates that deploying the contract with Solidity 0.8.28 and without EVM optimizations requires 159 404 gas units. According to Etherscan, the average gas price in January 2026 was approximately 1.44 Gwei, while according to Yahoo Finance [12] the average ETH

price in the same period was 2722.74 USD. Using these values, the estimated cost of deploying the contract is approximately 0.00023 ETH, which corresponds to roughly 0.62 USD per deployment.

Blockchain processing cost As discussed in the Mechanism Comparison section, both mechanisms require blockchain traversal, either from the genesis block until the first smart contract associated with an address is found, or from a known contract back toward the genesis block in order to verify that no previous contract exists for the same address. In the worst case, where the relevant contract is associated with the most recent transaction, the time complexity of the verification process is linear and can be expressed as $O(n)$, where n is the total number of verified transactions in the blockchain. This limitation makes the current verification strategy unsuitable for large-scale deployment unless it is complemented by indexing, caching, or an auxiliary lookup layer.

4 Conclusions and Future Work

The results obtained from the prototype indicate that the model should still be regarded as an exploratory solution. Its main limitations are the weak association between legitimate users and deployed smart contracts, and the poor scalability of a verification process that depends on blockchain traversal. Although the model provides a promising basis for privacy-oriented, blockchain-based multi-factor authentication, further improvements are required before production-level adoption.

A formal usability evaluation has not yet been carried out, so the adoption potential of the model should be interpreted cautiously. Another relevant obstacle is the monetary cost of deploying the smart contract. Although this cost is incurred only once and remains low, competing authentication models require no direct expense. At the same time, this cost is inherent to public blockchains and acts as a deterrent against spam, since attackers must likewise pay for every transaction submitted to the network.

Ongoing work addresses these limitations along several lines: improving verification scalability through indexing, caching, and alternative lookup strategies that reduce direct blockchain traversal while remaining consistent with on-chain data—taking care that such auxiliary components do not introduce new attack surfaces or undermine blockchain immutability; strengthening user-to-contract association without exposing personal information; replacing the partially centralized IAM configuration with a more decentralized, self-contained authentication flow; and supporting more flexible contract structures, including multiple secondary addresses, to reduce the impact of losing a single secondary key. Future evaluation stages will incorporate metrics on usability, latency, robustness, and deployment-oriented operation, so that the model can be assessed under conditions closer to real-world adoption.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Buterin, V.: EIP-170: Contract code size limit. Ethereum Improvement Proposals, https://eips.ethereum.org/EIPS/eip-170 (2016)
2. Chang, W., Rocco, G., Millegan, B., Johnson, N., Terbu, O.: ERC-4361: Sign-in with Ethereum. Ethereum Improvement Proposals, https://eips.ethereum.org/EIPS/eip-4361 (2021)
3. ConsenSys: Mythril: Security analysis tool for EVM bytecode. https://github.com/ConsenSysDiligence/mythril, accessed June 2026
4. Ethereum Foundation: Solidity documentation, version 0.8.28. https://docs.soliditylang.org (2024), accessed June 2026
5. Etherscan: Etherscan: The Ethereum blockchain explorer. https://etherscan.io, accessed June 2026
6. Keycloak Project: Keycloak: Open source identity and access management. https://www.keycloak.org, accessed June 2026
7. Moore, R.: Ethers.js: A complete Ethereum library and wallet implementation. https://docs.ethers.org, accessed June 2026
8. Nomic Foundation: Hardhat: Ethereum development environment. https://hardhat.org, plugins: hardhat-contract-sizer and hardhat-gas-reporter. Accessed June 2026
9. Ometov, A., Bezzateev, S., Mäkitalo, N., Andreev, S., Mikkonen, T., Koucheryavy, Y.: Multi-factor authentication: A survey. Cryptography **2**(1), 1 (2018). https://doi.org/10.3390/cryptography2010001
10. OpenJS Foundation: Express.js: Fast, unopinionated, minimalist web framework for Node.js. https://expressjs.com, accessed June 2026
11. Romero Arregin, G.D., Méndez-Garabetti, M.: Completar: titulo del trabajo previo de los autores (2025), entrada pendiente: copiar desde el biblio.bib original
12. Yahoo Finance: Ethereum USD (ETH-USD) price history. https://finance.yahoo.com/quote/ETH-USD/history, accessed June 2026

A Predictive Model on Percolation Theory and Machine Learning

Eduardo Miguel Zarate[1], Roman Bond[2], and Diego Encinas[2,3]

[1] Universidad Nacional de Catamarca, Facultad de Ciencias Exactas y Naturales.
[2] Proyecto de Investigación SimHPC - Programa TICAPPS - Instituto de Ingeniería y Agronomía - Universidad Nacional Arturo Jauretche.
[3] Instituto de Investigación en Informática LIDI (III-LIDI) - Facultad de Informática - UNLP - Centro Asociado CIC.
emzarate@exactas.unca.edu.ar, {rbond,dencinas}@unaj.edu.ar

Abstract. Early detection of vehicular faults remains a major challenge in the development of predictive maintenance strategies. Traditional diagnostic systems based on OBD2 (On-Board Diagnostics II) data are commonly used in a reactive manner, identifying failures only after they have already affected vehicle performance. To address this limitation, a predictive model based on percolation theory, OBD2 data, and Machine Learning is proposed. Real-world data were collected from an internal combustion vehicle using an ELM327 OBD2 scanner under different operating conditions. Next, a functional network was constructed in which the sensors were represented as nodes and their statistical relationships as edges. Percolation theory was applied to evaluate network robustness and estimate the critical connectivity threshold associated with system stability. In parallel, unsupervised learning techniques, particularly DBSCAN and Isolation Forest, were used to identify patterns and detect anomalous observations in the sensor data. The results indicate that the integration of network-based structural analysis and Machine Learning provides a valuable framework for the early detection of incipient vehicular faults and supports the design of data-driven predictive maintenance strategies.

Keywords: Percolation theory · OBD2 data · Predictive maintenance.

1 Introduction

In recent years, data-driven approaches have gained relevance in intelligent maintenance, particularly through Machine Learning techniques capable of identifying hidden patterns and anomalous behavior in complex systems. However, many of these approaches focus exclusively on the statistical behavior of sensor values without explicitly considering the structural relationships among monitored variables. In this regard, network-based modeling offers an alternative perspective by representing the system as an interconnected structure in which sensors interact dynamically.

Within this framework, percolation theory provides a useful mathematical basis for analyzing connectivity, robustness, and critical transitions in complex networks. Its application to vehicular systems makes it possible to evaluate how changes in sensor interactions may reflect states of instability associated with incipient faults. Based on this perspective, the present study proposes a predictive model for vehicular failure detection that integrates OBD2 data, percolation theory, and unsupervised machine learning. The aim is to identify anomalous patterns and assess structural vulnerability in order to support predictive maintenance strategies based on real operational data.

2 Theorical Background

Percolation theory provides a mathematical framework for studying connectivity, critical transitions, and robustness in complex systems. Originally developed to analyze propagation processes in disordered media, it was later extended to the study of complex networks, where the emergence or fragmentation of connected components can be examined through critical thresholds and structural resilience [3,7]. Previous studies have shown that network topology plays a central role in determining whether a system is robust under random failures yet vulnerable to disruptions affecting highly connected nodes [10,1,4,5].

From this perspective, a vehicular sensor system can be represented as a functional network in which sensors correspond to nodes and their statistical dependencies define the edges. This representation makes it possible to analyze the structural organization of the monitored system and to identify critical variables whose alteration may affect global connectivity. Complementarily, Machine Learning provides data-driven tools for detecting hidden patterns and anomalous behavior in high-dimensional sensor data [9,2]. In particular, unsupervised approaches are especially suitable in vehicular diagnostics when labeled fault data are limited or unavailable [8].

3 Methodology

The proposed methodology integrates real-time vehicular data acquisition, functional network modeling, percolation-based structural analysis, and unsupervised Machine Learning for predictive fault detection. Data were collected from an internal combustion vehicle using an ELM327 OBD2 scanner, which provided real-time measurements from sensors associated with engine and vehicle operation [6]. The recorded variables included parameters related to intake air temperature, manifold absolute pressure, ignition timing advance, engine coolant temperature, throttle position, fuel trim, and vehicle speed.

After acquisition, the data were preprocessed to ensure consistency and comparability across variables. This stage involved data cleaning, normalization, temporal organization, and the selection of representative operational observations. Based on these processed data, a functional network was constructed in which each sensor was represented as a node and the relationships between

sensors were established according to their statistical dependence. This representation enabled the analysis of the structural behavior of the vehicular system as a complex network.

Percolation theory was then applied to evaluate network robustness and to estimate the critical threshold associated with connectivity loss under node removal. In parallel, unsupervised Machine Learning techniques, specifically DBSCAN and Isolation Forest, were employed to identify latent groupings and detect anomalous observations in the sensor data.The combination of structural and data-driven analyses served as the basis for identifying patterns that could be linked to emerging faults in vehicles.

4 Results and Discussion

The most representative results of the proposed model were obtained under real driving conditions, where the vehicular system exhibited greater dynamic complexity and a richer pattern of interactions among sensors. From the normalized and consolidated dataset, a Pearson correlation matrix was computed (Table 1) in order to identify the statistical relationships among the monitored variables. The results revealed several relevant associations involving ignition timing advance, engine coolant temperature, intake air temperature, throttle position, fuel trim, and vehicle speed, providing an initial characterization of the functional dependencies within the vehicular system.

Table 1. Correlation matrix under real driving conditions (10-second time interval).

	ITA	**ECT**	**IAT**	**LTFT**	**MAP**	**STFT**	**TPS**	**VSS**
Ignition Timing Advance (ITA)	1	0.215	0.393	0.250	0.077	-0.286	0.592	0.714
Engine Coolant Temperature (ECT)	0.215	1	0.600	0.630	0.144	-0.891	0.261	0.330
Intake Air Temperature (IAT)	0.393	0.600	1	0.579	0.100	-0.664	0.385	0.472
Long-Term Fuel Trim (LTFT)	0.250	0.630	0.579	1	0.110	-0.817	0.239	0.338
Manifold Absolute Pressure (MAP)	0.077	0.144	0,100	0.110	1	-0.125	0.651	0.034
Short-Term Fuel Trim (STFT)	-0.286	-0.891	-0.664	-0.817	-0.125	1	-0.283	-0.377
Throttle Position Sensor (TPS)	0.592	0.261	0.385	0.239	0.651	-0.283	1	0.699
Vehicle Speed Sensor (VSS)	0.714	0.330	0.472	0.338	0.034	-0.377	0.699	1

Based on the correlation analysis, the strongest and most meaningful pairwise relationships between sensors were selected to define the edges of the functional graph. These connections summarize the most relevant dependencies among the monitored variables and provide the basis for representing the vehicular system as a functional network, as shown in Table 2.

The graph derived from these relationships revealed a more articulated and informative topology than those observed under cold-start and idle conditions. This result suggests that, during real driving, the vehicular system operates as a more interconnected structure in which changes in one variable may influence the global behavior of the network. The functional graph therefore provides a compact representation of the structural organization of the monitored system

Table 2. Edges of the functional graph under real driving conditions.

Sensor 1	Sensor 2	Correlation
Ignition Timing Advance	Throttle Position Sensor	0.592899639
Ignition Timing Advance	Vehicle Speed Sensor	0.714437897
Engine Coolant Temperature	Intake Air Temperature	0.600406744
Engine Coolant Temperature	Long-Term Fuel Trim	0.630406568
Engine Coolant Temperature	Short-Term Fuel Trim	-0.891561686
Intake Air Temperature	Long-Term Fuel Trim	0.579708494
Intake Air Temperature	Short-Term Fuel Trim	-0.664090326
Long-Term Fuel Trim	Short-Term Fuel Trim	-0.817774263
Manifold Absolute Pressure	Throttle Position Sensor	0.651033614
Throttle Position Sensor	Vehicle Speed Sensor	0.699886753

and makes it possible to identify the sensor interactions that are most relevant for predictive diagnosis.

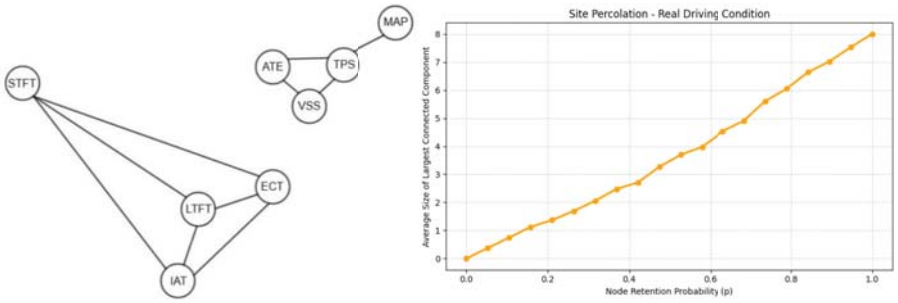

Fig. 1. On the left graph of vehicular sensors under real driving conditions. On the rigth site percolation curve of the vehicular sensor network under real driving conditions.

Once the functional network was constructed, percolation theory was applied to evaluate its structural robustness and to estimate the critical connectivity threshold (Figure 1). The percolation curve showed the transition from disconnected local failures to the emergence of large-scale connectivity, revealing the point at which local perturbations may evolve into system-wide structural instability. This finding is especially relevant from a diagnostic perspective, since it provides a mathematical criterion for distinguishing between isolated fluctuations and critical transitions that may compromise the global stability of the vehicular system.

Overall, the combined analysis of correlations, graph structure, and percolation behavior shows that the proposed model is capable of capturing both the internal organization and the structural vulnerability of the vehicular sensor system. In this sense, the real-driving condition offers the most informative scenario for evaluating predictive behavior, since it reflects the highest level of interaction among variables and allows the identification of connectivity patterns associated with possible incipient faults.

5 Conclusion

This study presented a predictive model for vehicular fault detection based on percolation theory, OBD2 data, and Machine Learning. The proposed approach combined correlation analysis, functional network modeling, and percolation-based structural evaluation in order to identify critical connectivity patterns in vehicular sensor behavior under real driving conditions. The results showed that the monitored system can be represented as an articulated network whose structural organization provides relevant information about potential instability and fault propagation.

The analysis also demonstrated that real driving conditions offer the most informative scenario for evaluating the predictive potential of the model, since they reveal stronger and more diverse interactions among the monitored variables. In this context, the functional graph and the percolation curve provided complementary evidence of structural robustness and critical transitions in system connectivity.

Overall, the findings support the usefulness of integrating network-based analysis and data-driven methods to move beyond conventional reactive diagnostics based on OBD2 systems. This perspective contributes to the development of predictive maintenance strategies and opens the possibility of extending the model to larger datasets, additional sensors, and more diverse vehicular operating conditions.

References

1. Albert, R., Jeong, H., Barabási, A.L.: Error and attack tolerance of complex networks. nature **406**(6794), 378–382 (2000)
2. Bishop, C.M., Nasrabadi, N.M.: Pattern recognition and machine learning. Springer (2006)
3. Broadbent, S.R., Hammersley, J.M.: Percolation processes: I. crystals and mazes. In: Mathematical proceedings of the Cambridge philosophical society. pp. 629–641. Cambridge University Press (1957)
4. Callaway, D.S., Newman, M.E., Strogatz, S.H., Watts, D.J.: Network robustness and fragility: Percolation on random graphs. Physical review letters **85**(25), 5468 (2000)
5. Cohen, R., Erez, K., Ben-Avraham, D., Havlin, S.: Resilience of the internet to random breakdowns. Physical review letters **85**(21), 4626 (2000)
6. Elm Electronics: ELM327: OBD to RS232 Interpreter. Elm Electronics Inc., Kingston, Canada (May 2020), available at https://www.elmelectronics.com/wp-content/uploads/2020/05/ELM327DSL.pdf
7. Grimmett, G.: What is percolation? In: Percolation, pp. 1–31. Springer (2012)
8. Rasheed, R., Qazi, F., Ahmed, A., Asif, A., Shams, H., et al.: Machine learning approaches for in-vehicle failure prognosis in automobiles: a review. VFAST Transactions on Software Engineering **12**(1), 169–182 (2024)
9. Samuel, A.L.: Some studies in machine learning using the game of checkers. IBM Journal of research and development **3**(3), 210–229 (1959)
10. Schwartz, N., Cohen, R., Ben-Avraham, D., Barabási, A.L., Havlin, S.: Percolation in directed scale-free networks. Physical Review E **66**(1), 015104 (2002)

Evaluating Large Language Models for Emergency Triage Prediction from Clinical Narratives: A Latin American Perspective

José M. Manríquez-Troncoso[1,2][0000−0001−7571−0048], Xaviera A. López-Cortés[1,2][0000−0002−7514−8777], and Felipe Tirado Maraboli[1,2][0000−0003−2001−658X]

[1] Departamento de Computación e Industrias, Universidad Católica del Maule, Talca, Chile
[2] Multidisciplinary Intelligent Data Science Lab, Talca, Chile
[3] Millennium Institute for Intelligent Healthcare Engineering (iHEALTH), Santiago, Chile

jose.manriquez@alumnos.ucm.cl

Abstract. Despite the growing interest in LLMs for clinical decision support, few studies have evaluated their ability to perform ESI-based triage using real clinical data from Latin American healthcare settings. This study systematically evaluates the capacity of LLMs to predict emergency triage levels from clinical narratives. Two models — Qwen3-30B and MedGemma-1.5-4B — were assessed across three progressive strategies: zero-shot prompting, Chain-of-Thought prompt engineering with ESI v4 logic, and LoRA fine-tuning on a synthetically constructed instruction-response dataset of 2,750 cases from Talca Regional Hospital, Chile. Fine-tuning yielded the most substantial performance gains, achieving accuracy above 0.55 and F1 scores up to 0.585, with the lightweight MedGemma model reaching performance comparable to the larger Qwen3-30B. These findings suggest that domain-specific models can be effectively adapted for triage classification tasks, supporting their potential as clinical decision-support tools and motivating future work on knowledge-graph-augmented approaches tailored to regional emergency protocols.

Keywords: Large language models· Emergency Severity Index

1 Introduction

Emergency departments are a cornerstone of modern healthcare systems, serving as the first line of care for a wide range of conditions. However, this critical component faces a state of perpetual crisis worldwide due to the constant increase in demand coupled with limited staff, space, and equipment, creating systemic overcrowding [2]. This directly compromises patient care quality and has a devastating impact on healthcare professionals, leading to burnout and decreased job satisfaction [7]. At the heart of this challenge is the triage process, which acts as the primary access control and resource allocation system in emergency departments [3]. Standardized protocols such as the Emergency Severity Index

(ESI) have been crucial in establishing a structured decision-making framework. However, their effectiveness largely depends on the subjective judgment of nursing staff, introducing significant variability and errors including under-triage and over-triage due to cognitive overload [1, 6]. Artificial intelligence, and more recently large language models (LLMs), emerge as promising tools to provide a more objective and consistent assessment, moving beyond simple prediction toward complex clinical reasoning tasks such as acuity assessment from narrative notes [5, 8]. However, these models are primarily trained on English corpora, tailored to realities very different from Latin America. This study addresses this gap by evaluating the ability of LLMs to predict ESI urgency levels using real clinical data from Talca Regional Hospital in Chile, comparing three progressive strategies: zero-shot prompting, Chain-of-Thought prompting with ESI v4 logic, and LoRA fine-tuning on a synthetically constructed instruction-response dataset.

2 Methodology

2.1 Dataset

This study utilized patient admission records from the Emergency Department of the Regional Hospital of Talca, Chile, covering the period from 2016 to 2018. These records include relevant clinical and administrative information, such as vital signs (e.g., blood pressure, temperature, oxygen saturation), demographic data (age, sex), primary reason for presentation, and triage level assigned according to the ESI system. Only patients over the age of 15 were included, and emergency dental care was excluded. Additionally, records without a reason for consultation, without an ESI level, or with missing data for most vital signs were discarded. After preprocessing, 650 cases were selected per ESI category, resulting in a balanced dataset of 3,250 samples for analysis.

2.2 Models evaluated

Two large-scale language models with complementary profiles were selected for the experiments. Qwen3-30B-A3B-Instruct-2507, developed by Alibaba Cloud, is a general-purpose model with a Mixture of Experts (MoE) architecture. It has a total of 30 billion parameters, with 3 billion active parameters per token. Its multilingual training and size position it as a benchmark model for complex reasoning tasks. On the other hand, MedGemma-1.5-4B-IT, developed by Google, is a lightweight language model with 4 billion parameters. It was pre-trained on biomedical corpora and fine-tuned using instructions for clinical tasks. Its specialization in medical terminology makes it a natural candidate for clinical decision support applications. Both models were run with 4-bit quantization using the NF4 scheme with double quantization (BitsAndBytes) and bfloat16 computational precision to reduce GPU memory consumption while maintaining minimal performance degradation.

2.3 Triage case generation

A triage case is defined as a structured textual description that simulates a patient's clinical presentation upon arrival at the emergency department. It typically includes symptoms, vital signs, relevant medical history, and a clinical first impression [4]. To generate these cases, the Qwen3-30B-A3B-Instruct-2507 and medgemma-1.5-4b-it models available on the Hugging Face platform were used.

Based on the hospital records described in the previous subsection, specific instructions (prompt engineering) were designed to obtain clinical narratives consistent with the format recommended in the ESI v4 implementation manual [4]. This procedure allowed us to enrich and standardize the information contained in the reason for visit, correcting spelling errors and expanding abbreviations commonly used in clinical settings that could hinder subsequent automatic analysis. From the 3,250 generated cases, 500 (100 per ESI category) were isolated via stratified sampling as a held-out test set, used exclusively for evaluation across all three experimental scenarios. The remaining 2,750 cases were used to construct the instruction-response dataset for LoRA fine-tuning, of which only those with both a valid narrative and successfully generated reasoning were retained in the final training corpus.

2.4 Instruction-response dataset construction

Based on the clinical narratives generated in the previous stage, an instruction-response dataset was constructed for fine-tuning models for triage tasks. In a second phase, each narrative was used as input to generate structured clinical reasoning, guided by the ESI v4 algorithm and conditioned by the true triage level assigned in the original record.

To this end, a fixed instruction was designed to determine if the patient requires immediate life-saving interventions, analyze the components of airway, breathing, and circulation, and finally assign the corresponding ESI level. In this scheme, the generated clinical narrative was used as input, while the output corresponded to a step-by-step clinical explanation followed by the final label in the format "Final Answer: ESI X".

In order to preserve an independent set for validation, the samples previously reserved for expert evaluation were excluded before the construction of the final dataset. Likewise, only those cases that simultaneously included clinical narrative and generated reasoning were retained, thus obtaining a structured corpus suitable for supervised training in classification and clinical reasoning tasks in triage.

2.5 Evaluation strategies

Three progressive experimental scenarios were designed to evaluate the models' capabilities. In the first scenario, a baseline was established using direct prompting (zero-shot), in which each model was asked to assign an ESI level from 1 to 5 based on the clinical narrative, without additional reasoning instructions. In the

second scenario, an advanced prompt engineering strategy was implemented that incorporates the conditional logic of the ESI v4 algorithm, structured in three sequential steps: assessment of immediate life-threatening intervention, identification of high-risk situations, and counting of expected diagnostic resources. This scheme forces the model to generate explicit clinical reasoning (Chain-of-Thought) before issuing its final prediction in the standardized ESI LEVEL: X format, which is extracted using regular expressions. In both scenarios, inference was performed in deterministic mode (greedy decoding, do_sample=False) with 4-bit quantization to optimize memory usage. In the third scenario, we explored model fine-tuning using low-rank adaptation (LoRA), utilizing the instruction–response dataset described in Section 2.4. This procedure allows the model weights to be specialized for the specific triage task without requiring full retraining, significantly reducing computational cost. The fine-tuning was performed using the PEFT library on the TRL framework. The LoRA configuration utilized a rank (r) of 16, an alpha of 32, and targeted the projection modules (q_proj, k_proj, v_proj, o_proj, gate_proj, up_proj, down_proj) with a dropout of 0.05. Training was conducted for 1 epoch with a learning rate of 3e-5 using the paged_adamw_8bit optimizer.

2.6 Evaluation metrics

Model performance was evaluated using four metrics calculated as a macro-level average, taking into account the five ESI categories. Accuracy measures the proportion of correctly classified cases out of the total. The macro F1-score corresponds to the harmonic mean of precision and recall calculated per class, making it particularly suitable for scenarios with a balanced distribution across categories. Recall quantifies the model's ability to correctly identify each ESI level, while specificity reflects its ability to rule out incorrect levels. This latter metric takes on particular clinical relevance in the context of triage, given that under-triage errors—classifying a critically ill patient at a lower level—can directly compromise patient safety

3 Results

Tables 1, 2, and 3 present the results obtained across the three experimental scenarios for each combination of triage note generator and prediction model. In the zero-shot scenario (Table 1), both models exhibit modest performance, with accuracy ranging from 0.293 to 0.353 and macro F1-scores between 0.193 and 0.225. Qwen3-30B consistently outperforms MedGemma-1.5-4B as a predictor regardless of the note source, while specificity remains high across all combinations ($\geq$0.898), suggesting that errors are distributed across levels rather than concentrated in critical misclassifications. The prompt engineering strategy (Table 2) produces marginal improvements over the baseline, with accuracy gains of approximately 2–3 percentage points and modest F1 increases. The most notable improvement is observed in specificity for Qwen3-30B predictions (up

to 0.946), indicating more conservative classification behavior under structured reasoning. Fine-tuning (Table 3) yields a substantial performance leap across all configurations, with accuracy increasing to approximately 0.54–0.56 and F1-scores reaching up to 0.585. Notably, the performance gap between models narrows considerably: MedGemma-1.5-4B fine-tuned on Qwen3-30B-generated notes achieves results comparable to the larger model (Acc=0.536, F1=0.560 vs. 0.559, F1=0.585). Additionally, the source of the clinical narratives shows minimal impact on prediction performance within the same scenario, suggesting that both generation models produce narratives of equivalent informativeness for the triage task. From a clinical safety perspective, the consistently high specificity (>0.910 for both fine-tuned models) acts as a critical safeguard against severe under-triage. Because specificity reflects the models' ability to successfully rule out incorrect levels , these results indicate that when misclassifications occur, they are less likely to systematically downgrade critically ill patients to non-urgent categories, thereby minimizing the most dangerous form of triage error.

Table 1. Direct prompting (zero-shot) results

Triage notes	Triage Prediction	Acc	F1	Spec.	Rec.
Qwen3 30B	Qwen3-30B	0.353	0.225	0.926	0.360
	Medgemma-1.5-4B	0.308	0.193	0.902	0.308
Medgemma-1.5 4B	Qwen3-30B	0.347	0.221	0.905	0.345
	Medgemma-1.5-4B	0.293	0.199	0.898	0.293

Table 2. Prompt engineering results

Triage notes	Triage Prediction	Acc	F1	Spec.	Rec.
Qwen3 30B	Qwen3-30B	0.375	0.248	0.946	0.378
	Medgemma-1.5-4B	0.318	0.206	0.908	0.328
Medgemma-1.5 4B	Qwen3-30B	0.372	0.243	0.939	0.367
	Medgemma-1.5-4B	0.302	0.228	0.900	0.302

Table 3. Model fine-tuning results

Triage notes	Triage Prediction	Acc	F1	Spec.	Rec.
Qwen3 30B	Qwen3-30B	0.559	0.585	0.935	0.564
	Medgemma-1.5-4B	0.536	0.560	0.910	0.543
Medgemma-1.5 4B	Qwen3-30B	0.557	0.580	0.933	0.559
	Medgemma-1.5-4B	0.540	0.577	0.919	0.545

4 Discussion

The results obtained across the three experimental scenarios reveal several relevant findings for the application of LLMs in clinical triage. First, the overall performance under zero-shot and prompt engineering conditions remained modest, with accuracy below 0.38 in the best cases, suggesting that general-purpose

LLMs are not yet ready for direct deployment in high-stakes clinical environments without further adaptation. This aligns with findings reported by Shen et al. [8], who observed similar limitations in prompt-engineered models prior to fine-tuning, and reinforces the notion that additional research is needed before LLMs can be reliably integrated into real-world emergency workflows. Second, a notable finding is the competitive performance between MedGemma-1.5-4B and Qwen3-30B across all scenarios, despite their substantial difference in parameter count. After fine-tuning, the performance gap between both models narrows considerably, with MedGemma achieving comparable accuracy and F1 scores. This suggests that domain-specific pretraining can partially compensate for model scale, making lightweight clinical models a viable option in resource-constrained settings. Third, the substantial performance improvement observed after LoRA fine-tuning with accuracy gains of approximately 20 percentage points over the prompt engineering baseline is consistent with the instruction tuning results reported by Shen et al. [8], where task-specific adaptation proved to be the dominant factor in performance improvement across model sizes. This evidence supports the hypothesis that LLMs can be effectively specialized for triage classification tasks through supervised fine-tuning on clinically grounded instruction-response datasets, even when constructed synthetically. Furthermore, addressing the directionality of errors is crucial for real-world deployment. While absolute accuracy remains moderate, the models' high specificity suggests a clinically protective behavior. In emergency medicine, over-triage (assigning a higher acuity than necessary) implies a temporary overallocation of resources, whereas under-triage directly jeopardizes patient survival. The fine-tuned models demonstrate a capacity to avoid catastrophic under-triage, reinforcing their potential utility as safe second-opinion tools that err on the side of caution when processing ambiguous clinical narratives. Finally, an important limitation of this study, shared with related work [8], is the dependence on data that may not fully reflect the sociocultural and operational particularities of Latin American emergency systems. As future work, we plan to explore the integration of knowledge graphs constructed from documentation specific to Chilean emergency services, including clinical protocols, resource availability, and regional epidemiological patterns. This approach aims to embed context-aware, country-specific knowledge into LLMs, potentially improving their alignment with local triage practice and reducing the gap between model predictions and real clinical decisions.

Acknowledgments. The authors gratefully acknowledge the Laboratory of Technological Research in Patterns Recognition (LITRP) at Universidad Católica del Maule (https://www.litrp.cl) for providing access to the computational infrastructure that made the experiments in this study possible.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ausserhofer, D., Zaboli, A., Pfeifer, N., Solazzo, P., Magnarelli, G., Marsoner, T., Siller, M., Turcato, G.: Errors in nurse-led triage: an observational study. International journal of nursing studies **113**, 103788 (2021)
2. Charan, G.S., Kalia, R., Dular, S.K., Kumar, R., Kaur, K.: Challenges faced by doctors and nurses in the emergency department: An integrated review. Journal of Education and Health Promotion **14**(1), 2 (2025)
3. FitzGerald, G., Jelinek, G.A., Scott, D., Gerdtz, M.F.: Emergency department triage revisited. Emergency medicine journal **27**(2), 86–92 (2010)
4. Gilboy, N., Tanabe, P., Travers, D., Rosenau, A.M., et al.: Emergency severity index (esi): a triage tool for emergency department care, version 4. Implementation handbook **2012**, 12–0014 (2012)
5. Jafari, M., Raee, A., Mohammad Karimi, N., Neshati, A., Modjallal, F., Zeinali, F.: Applications of machine learning models in optimizing triage accuracy and predicting patient outcomes in emergency care: A narrative review of current evidence. Journal of Surgery and Trauma pp. 0–0 (2025)
6. Nguyen, H., Meczner, A., Burslam-Dawe, K., Hayhoe, B.: Triage errors in primary and pre–primary care. Journal of medical Internet research **24**(6), e37209 (2022)
7. Phillips, K., Knowlton, M., Riseden, J.: Emergency department nursing burnout and resilience. Advanced emergency nursing journal **44**(1), 54–62 (2022)
8. Shen, Q., Zhang, X., Ren, H., Guo, Q., Yi, Z.: Knowledge-embedded large language models for emergency triage. Knowledge-Based Systems **318**, 113431 (2025)