INTELLI 2026

The Fifteenth International Conference on Intelligent Systems and Applications

ISBN: 978-1-68558-351-4

March 8th –12th, 2026

Valencia, Spain

INTELLI 2026 Editors

Gil Gonçalves, University of Porto, Portugal

Rui Pinto, Universidade do Porto, Portugal

INTELLI 2026

Forward

The Fifteenth International Conference on Intelligent Systems and Applications (INTELLI 2026), held between March 8-th, 2026 and March 12-th, 2026 in Valencia, Spain, continued a series of events on advances towards fundamental, as well as practical and experimental, aspects of intelligent systems and applications.

The information surrounding us is not only overwhelming but also subject to limitations of systems and applications, including specialized devices. The diversity of systems and the spectrum of situations make it almost impossible for an end-user to handle the complexity of the challenges. Embedding intelligence in systems and applications seems to be a reasonable way to move some complex tasks away from the user. However, this approach requires fundamental changes in designing the systems and applications, in designing their interfaces, and requires using specific cognitive and collaborative mechanisms. Intelligence has become a key paradigm, and its specific use takes various forms according to the technology or the domain a system or an application belongs to.

We take here the opportunity to warmly thank all the members of the INTELLI 2026 technical program committee, as well as all the reviewers. The creation of such a high-quality conference program would not have been possible without their involvement. We also kindly thank all the authors who dedicated much of their time and effort to contribute to INTELLI 2026. We truly believe that, thanks to all these efforts, the final conference program consisted of top-quality contributions. We also thank the members of the INTELLI 2026 organizing committee for their help in handling the logistics of this event.

We hope that INTELLI 2026 was a successful international forum for the exchange of ideas and results between academia and industry for the promotion of progress in the field of intelligent systems and applications.

INTELLI 2026 Chairs

INTELLI 2026 Steering Committee
Carsten Behn, Schmalkalden University of Applied Sciences, Germany
Stefano Berretti, University of Firenze, Italy
Marcin Paprzycki, Systems Research Institute, Polish Academy of Sciences – Warszawa, Poland
Gil Gonçalves, University of Porto, Portugal

INTELLI 2026 Publicity Chairs
Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain
José Miguel Jiménez, Universitat Politècnica de València, Spain
Sandra Viciano Tudela, Universitat Politècnica de València, Spain

INTELLI 2026
Committee

INTELLI 2026 Steering Committee

Carsten Behn, Schmalkalden University of Applied Sciences, Germany
Stefano Berretti, University of Firenze, Italy
Marcin Paprzycki, Systems Research Institute, Polish Academy of Sciences – Warszawa, Poland
Gil Gonçalves, University of Porto, Portugal

INTELLI 2026 Publicity Chairs

Francisco Javier Díaz Blasco, Universitat Politècnica de València, Spain
Ali Ahmad, Universitat Politècnica de València, Spain
José Miguel Jiménez, Universitat Politècnica de València, Spain
Sandra Viciano Tudela, Universitat Politècnica de València, Spain

INTELLI 2026 Technical Program Committee

Azizi Ab Aziz, Universiti Utara Malaysia, Malaysia
Lounis Adouane, Université de Technologie de Compiègne, France
Ashwani Kumar Aggarwal, Sant Longowal Institute of Engineering and Technology, India
Leo Aguilera, 33 Technologies LLC, USA
Ari Aharari, SOJO University, Japan
Bilal Ahmad, University of Warwick, UK
Mohd Ashraf Ahmad, Universiti Malaysia Pahang Al-Sultan Abdullah, Malaysia
Sathish Akula, Florida Polytechnic University, USA
Zaher Al Aghbari, University of Sharjah, UAE
Miltos Alamaniotis, University of Texas at San Antonio, USA
Antonios Alexos, University of California Irvine, USA
Sarra Alqahtani, Wake Forest University, USA
Rachid Anane, Coventry University, UK
Olugbenga Moses Anubi, Florida State University, USA
Arvind Bansal, Kent State University, USA
Suzanne Barber, The University of Texas at Austin, USA
Dariusz Barbucha, Gdynia Maritime University, Poland
Carmelo Bastos-Filho, University of Pernambuco, Brazil
Rafael Batres, Tecnologico de Monterrey, Mexico
Carsten Behn, Schmalkalden University of Applied Sciences, Germany
Fayçal Bensaali, Qatar University, Qatar
Lyes Benyoucef, Aix-Marseille University, France
Giuseppe Berio, IRISA | Université de Bretagne Sud, France
Stefano Berretti, University of Firenze, Italy
Mahdis Bisheban, National Research Council Canada (NRC), Canada
José Miguel Blanco, Masaryk University, Brno, Czech Republic
José Boaventura-Cunha, UTAD University / INESC-TEC- Technology and Science, Portugal

Copyright Information

For your reference, this is the text governing the copyright release for material published by IARIA.

The copyright release is a transfer of publication rights, which allows IARIA and its partners to drive the dissemination of the published material. This allows IARIA to give articles increased visibility via distribution, inclusion in libraries, and arrangements for submission to indexes.

I, the undersigned, declare that the article is original, and that I represent the authors of this article in the copyright release matters. If this work has been done as work-for-hire, I have obtained all necessary clearances to execute a copyright release. I hereby irrevocably transfer exclusive copyright for this material to IARIA. I give IARIA permission or reproduce the work in any media format such as, but not limited to, print, digital, or electronic. I give IARIA permission to distribute the materials without restriction to any institutions or individuals. I give IARIA permission to submit the work for inclusion in article repositories as IARIA sees fit.

I, the undersigned, declare that to the best of my knowledge, the article is does not contain libelous or otherwise unlawful contents or invading the right of privacy or infringing on a proprietary right.

Following the copyright release, any circulated version of the article must bear the copyright notice and any header and footer information that IARIA applies to the published article.

IARIA grants royalty-free permission to the authors to disseminate the work, under the above provisions, for any academic, commercial, or industrial use. IARIA grants royalty-free permission to any individuals or institutions to make the article available electronically, online, or in print.

IARIA acknowledges that rights to any algorithm, process, procedure, apparatus, or articles of manufacture remain with the authors and their employers.

I, the undersigned, understand that IARIA will not be liable, in contract, tort (including, without limitation, negligence), pre-contract or other representations (other than fraudulent misrepresentations) or otherwise in connection with the publication of my work.

Exception to the above is made for work-for-hire performed while employed by the government. In that case, copyright to the material remains with the said government. The rightful owners (authors and government entity) grant unlimited and unrestricted permission to IARIA, IARIA's contractors, and IARIA's partners to further distribute the work.

Table of Contents

Craft Beer Monitoring: A Cyber-Physical IoT System for Precision Brewing 1
Joao Teixeira, Joao Lourenco, Jose Felisberto, Leonardo Regadas, Muhammad Ibrahim, Rodrigo Figueiredo, Tomas Macedo, Rui Pinto, and Gil Goncalves

Edge-Based IoT and AI Framework for Real-Time Wastewater Potability Classification 7
Jose Isidro, Rafael Teixeira, Joao Costa, Carolina Goncalves, Diogo Ferreira, Pedro Azevedo, Pedro Simoes, Rui Pinto, and Gil Goncalves

An Edge-Centric IoT System for Smart Building Energy Management 13
Joao Silva, Joao Alves, Nuno Silva, Abdul Rauf, Victor Bongard, Victor Rodriguez, Nuno Moreira, Rui Pinto, and Gil Goncalves

Accelerating the Adoption of Asset Administration Shells through AI Agents 19
Camilo Velazquez-Rodriguez, Steven Kauffmann, Tom Pauwaert, Stijn Huysentruyt, and Axl Van Alboom

From Data Silos to Intelligent Operations: An AI-Based Approach to Ground Station Incident Investigation 25
Nieves Salor Moral, Dimitri Accad, and Andrea Di Luca

Residual Hybrid Motor Controller with Multi-Phase Learning 32
Johann Wiens and Christoph Reich

ENDURE: Ensemble Based Robust Reinforcement Learning with Reduced Sample Complexity 38
Sarra Alqahtani

Integrating Intuitive Interaction and Large Language Model Guidance for Efficient 3D Annotation 45
Haruya Ishigami, Kenji Iwata, and Yutaka Satoh

From Regulation to Relevance: Integrating User Values into Privacy Policy Scoring 51
Brian Kim and Suzanne Barber

Craft Beer Monitoring: A Cyber-Physical IoT System for Precision Brewing

João Teixeira[1], João Lourenço[1], José Felisberto[1], Leonardo Regadas[1], Muhammad Ibrahim[1], Rodrigo Figueiredo[1], Tomás Macedo[1], Rui Pinto[1,2], Gil Gonçalves[1,2]

Dept. de Engenharia Informática, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[1]

SYSTEC, ARISE, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[2]

Email: {up202108738, up202108864, up202108657, up202108144, up202301909, up202108870, up202008551}@up.pt

{rpinto, gil}@fe.up.pt

***Abstract*—In small-scale breweries, manual process monitoring is the main source of inconsistency, often leading to a significant waste of time, effort, and ingredients. This paper presents a Cyber-Physical Internet of Things (IoT) system designed to bring precision and predictability to the craft brewing process. Utilizing a 5-layer architecture, the system integrates an Arduino Nano ESP32 with sensors to monitor critical parameters, such as temperature, pH, and flow rate. While hardware limitations needed a strategic pivot to data simulation for pH and flow inputs, the core infrastructure was rigorously validated through the successful integration of a physical temperature sensor, Message Queuing Telemetry Transport (MQTT) communication protocols, and the Blynk cloud platform. Additionally, the system incorporates an external weather Application Programming Interface (API) to demonstrate extensibility. The resulting dashboard provides real-time data and proactive alerts, empowering brewers to prevent batch failure. By validating the soundness of this end-to-end blueprint, this work demonstrates how smart engineering can improve sustainability and quality control in artisanal production.**

Keywords-Cyber-Physical Systems; Internet of Things; Precision Brewing; Sustainability.

I. Introduction

The global brewing industry is currently navigating a structural transformation marked by the dual pressures of artisanal differentiation and industrial efficiency. While the craft beer segment has experienced explosive growth, the lack of process consistency due to the technological gap between industrial conglomerates and Small and Medium Enterprises (SMEs) is a critical challenge. While large-scale breweries utilize automated Supervisory Control and Data Acquisition (SCADA) systems to homogenize output, homebrewers and microbreweries frequently rely on manual monitoring methods.

The core problem addressed in this work is the "Brewer's Dilemma"—the fine line between art and waste [1]. In small-scale environments, manual process monitoring is the primary source of inconsistency. Because fermentation occurs within a closed vessel, small, untracked fluctuations in critical parameters, such as temperature and pH, often go undetected. Even minor deviations during key metabolic stages can irreversibly alter the flavor profile or stall the fermentation entirely. Consequently, this reliance on error-prone manual checks leads directly to waste: the loss of time, energy, and raw ingredients embedded in a failed batch. There is, therefore, an urgent need for a monitoring solution that eliminates this stochastic variability and ensures resource efficiency by design [2].

To address these challenges, this study aims to establish the theoretical and architectural foundations for the "Digital Brewmaster" [3]. The primary objective is to develop a low-cost, high-performance sensing ecosystem that brings industrial-grade precision to the artisanal sector. Specifically, this work focuses on validating the layers of an Internet of Things (IoT) architecture—communication, cloud processing, and user interfaces—to ensure that the system can reliably detect anomalies and alert the brewer before inconsistency turns into waste.

The proposed solution is a Cyber-Physical System (CPS) designed to democratize access to precision fermentation monitoring. The system utilizes a robust 5-layer IoT architecture anchored by the Arduino ESP32 microcontroller [4], the Message Queuing Telemetry Transport (MQTT) protocol [5], and the Blynk cloud platform [6]. To address the challenges of hardware validation, the project adopts a methodological pivot toward Data Simulation and Software-in-the-Loop (SIL) testing [7]. By generating synthetic data streams that mimic the non-linear kinetics of yeast fermentation, the system serves as a functional tool for brewers and a case study in resilient IoT systems engineering, capable of transitioning from reactive manual checks to proactive real-time monitoring.

The remainder of this paper is organized as follows: Section II establishes the fundamental knowledge and reviews related work, clustering recent research. Section III details the proposed Cyber-Physical System, outlining the 5-layer IoT architecture. Section IV presents the evaluation and results. Section V discusses the implications of these results, validating the hybrid architectural approach. Section VI concludes the paper with a summary of contributions and a set of future research directions.

II. Fundamental Knowledge and Related Work

Recent literature highlights a clear transition from offline measurements toward real-time, IoT-based monitoring in artisanal and small-scale production contexts. Kovačević *et al.* [8] present Winnie, a modular IoT system developed for small wineries that closely parallels the objectives of the Digital

Brewmaster concept. Their system employs distributed embedded sensing units installed on barrels to monitor temperature, humidity, and carbon dioxide (CO_2), using RS-485 communication to ensure robustness in harsh cellar environments. A key contribution of their work is the demonstration that low-cost sensors, when combined with integrity mechanisms, such as hash-based validation, can achieve reliability sufficient for process optimization. This finding supports the feasibility of commodity hardware for critical monitoring tasks and directly validates the design choices adopted in this work.

Hardware efficiency and energy sustainability are further addressed by Dzahir and Chia [9], who analyze the power consumption of ESP32 microcontrollers in MQTT-based monitoring systems. Their results show that, when deep-sleep modes are properly leveraged, ESP32-based nodes can operate continuously for extended periods on battery power. This is particularly relevant for craft brewing environments, where power outlets may not be readily available near fermenters or conditioning tanks. These findings reinforce the suitability of the ESP32 as a low-energy, scalable platform for continuous monitoring in artisanal production settings.

Beyond sensing and connectivity, several studies emphasize the growing role of the Digital Twin (DT) in fermentation and food-processing domains. Abdurrahman and Ferrari [10] provide a comprehensive review of DT applications in the food industry, classifying them into prognostic, reactive, and virtual commissioning models. While large-scale producers employ DTs primarily for energy optimization and process efficiency, their review identifies a notable lack of accessible DT solutions tailored to small producers. The authors argue that virtual models operating in parallel with physical processes enable safe exploration of “what-if” scenarios and represent a future cornerstone of quality control in food production.

This vision is further reinforced by Pierre *et al.* [11], who describe the engineering of a safety-critical Digital Twin for beer fermentation. Their system achieves continuous process sampling and reduces manual intervention by 91%, highlighting the potential of bidirectional digital–physical interaction. In such “true” Digital Twins, the cyber layer not only monitors but actively controls fermentation conditions, for example, by regulating cooling valves. While the Digital Brewmaster currently focuses on monitoring and alerting, its underlying architecture establishes the foundation required to evolve toward this level of closed-loop control.

A recurring challenge in the development of CPS and DT is the instability and unreliability of hardware during early prototyping. Balan *et al.* [12] note that this issue is often underreported in the literature and advocate for simulation-driven validation strategies. They propose a SIL methodology that decouples software validation from physical sensor availability. This approach directly supports the methodological pivot adopted in this project, where synthetic data streams were used to validate the cyber infrastructure when physical sensors proved unreliable.

Emerging research further extends these concepts toward data trust and virtual sensing. Blockchain technology [13] was integrated with multi-sensor IoT fermentation systems to ensure data integrity and traceability across supply chains [14]. By recording sensor data, such as temperature, pressure, and gas emissions, on an immutable ledger, their framework achieves full transaction reliability and enables verifiable provenance. While the Digital Brewmaster currently targets internal process control, this work suggests a future pathway in which small-scale brewers could leverage similar technologies to demonstrate product quality and compliance to consumers or regulators.

Complementarily, Ferrer *et al.* [15] and Zaidan *et al.* [16] explore the use of data-driven virtual sensors to estimate difficult-to-measure variables using correlations among inexpensive measurements. In the context of fermentation, this approach could enable the estimation of specific gravity—traditionally measured with a hydrometer—by combining temperature data with CO_2 release dynamics. Such techniques indicate that future iterations of the "Digital Brewmaster" could replace specialized instrumentation with AI-driven virtual sensing, further reducing system cost and lowering adoption barriers for craft brewers.

A synthesis of the current state-of-the-art reveals two specific voids that the "Digital Brewmaster" addresses.

1) **Accessibility for SMEs**. While large producers rely on advanced SCADA and DT systems, equivalent solutions remain largely inaccessible to craft brewers. As noted by Abdurrahman and Ferrari [10], there is a lack of "middle-ground" architectures combining industrial-grade monitoring with affordable commodity hardware [8]. Consequently, SMEs remain constrained to manual practices, leading to increased variability, waste, and inefficiency [17].
2) **Validation bottlenecks**. A defining challenge in developing custom IoT solutions is the reliability of physical hardware during prototyping, where harsh brewery environments and sensor availability can derail software validation. Most IoT literature assumes fully functional hardware. This work contributes to the field by practically applying the SIL methodology [12], demonstrating that a brewing CPS can be validly engineered even when physical sensor integration is obstructed.

III. Proposed Solution

The proposed solution is a CPS designed to enable continuous, automated, and real-time monitoring of critical parameters in the craft beer brewing process. The system follows a layered IoT architecture, which promotes modularity and scalability.

A. System Design and Architecture

The architecture is structured into five distinct layers, as represented in Figure 1: i) *Physical*; ii) *Hardware*; iii) *Communication*; iv) *Cloud*; and v) *Application*. Each layer fulfills a specific role in the data acquisition, transmission, processing, and visualization pipeline, enabling a clear abstraction of responsibilities and facilitating system extensibility.

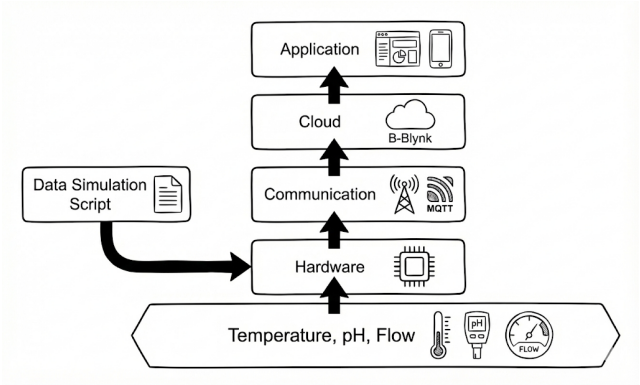

Figure 1. System Architecture.

The *Physical Layer* interfaces directly with the brewing environment by acquiring key process variables. Three sensors—temperature, pH, and flow rate—were selected due to their critical influence on fermentation dynamics, yeast metabolism, and final product quality. Temperature governs enzymatic reactions and yeast activity, while pH affects enzyme efficiency and microbial stability. Flow rate monitoring enables the detection of abnormal liquid behavior, such as leaks or blockages. Together, these measurements provide a representative snapshot of the brewing process state.

The *Hardware Layer*, implemented using an Arduino Nano ESP32 microcontroller [4], bridges both physical and digital processes. The ESP32 integrates a dual-core processor and embedded Wi-Fi, offering sufficient computational capacity for edge-level data acquisition and preprocessing while maintaining low power consumption [9]. This layer performs sensor sampling, basic data conditioning, and message formatting before forwarding data to higher layers.

The *Communication Layer*, based on the MQTT protocol [5], ensures efficient and reliable data exchange between the embedded node and cloud services. MQTT is particularly well-suited for IoT systems characterized by constrained resources and intermittent connectivity [9][12]. It's a lightweight publish–subscribe model, which decouples data producers from consumers and enhances scalability and fault tolerance. Sensor data are published to predefined topics, allowing cloud services and user applications to subscribe as needed.

The *Cloud Layer* acts as the central hub for data aggregation, processing, and system management. The Blynk IoT platform [6] was used to manage data ingestion, virtual devices, and rule-based logic. This layer aggregates sensor streams, stores historical data, and triggers alerts when predefined thresholds are exceeded. Cloud-based deployment ensures persistent storage and remote accessibility, supporting longitudinal analysis and operational traceability [8].

The *Application Layer* provides the human–machine interface for system interaction. A mobile dashboard was developed to visualize real-time measurements, historical trends, and system status indicators. By presenting deviations through visual cues and push notifications, this layer supports proactive decision-making. The interface was designed to be lightweight and intuitive, ensuring accessibility for non-technical users, such as homebrewers and small-scale producers.

A key strength of the proposed architecture is its extensibility. Its modular design enables the integration of additional sensors, predictive modules, or external data sources without requiring architectural redesign. This capability is illustrated through the integration of a weather Application Programming Interface (API), the OpenWeather Map API [18][19], which enriches brewing data with ambient environmental context. Such extensibility aligns with recent advances in digital twin and context-aware CPS architectures [10][11].

B. Development Process

The development of the proposed CPS followed an incremental and iterative engineering methodology. The initial goal was to implement a fully integrated proof of concept incorporating three physical sensors—temperature, pH, and flow rate—real-time data transmission, cloud-based processing, and a functional user dashboard.

During early prototyping, several hardware-related challenges emerged, particularly in sensor calibration, signal stability, and component reliability. These issues hindered the acquisition of consistent and reliable measurements from the pH and flow sensors. Such limitations are common in IoT prototyping environments, where hardware constraints and deployment conditions can delay or impede higher-layer validation [12].

To avoid constraining overall system validation, a methodological pivot was adopted toward a SIL strategy. Synthetic data streams were generated to emulate realistic sensor behavior, allowing continued system development despite partial hardware unavailability. A simulation module was implemented directly on the ESP32 to produce time-varying pH and flow values, enabling the validation of communication, cloud, and application layers independently of physical sensors.

In parallel, the temperature sensor was fully integrated and operated as a real physical input. This hybrid validation approach, combining physical sensing with simulated data, enabled end-to-end testing of the cyber infrastructure. It verified correct MQTT message formation, cloud ingestion, data parsing, historical storage, and dashboard visualization.

System development progressed incrementally, with individual subsystems validated in isolation before full integration. This approach reduced debugging complexity and improved traceability of system-level issues. The final system underwent functional testing under varying data conditions to assess responsiveness, stability, and real-time performance.

Overall, this methodology aligns with contemporary CPS engineering practices that emphasize simulation-driven validation and decoupled testing to mitigate hardware dependency risks and accelerate development cycles [10][12].

C. Technological Stack

The embedded component of the system is built around the Arduino Nano ESP32 microcontroller [4]. This platform was selected due to its integrated Wi-Fi connectivity, low power consumption, and mature development ecosystem. Prior work

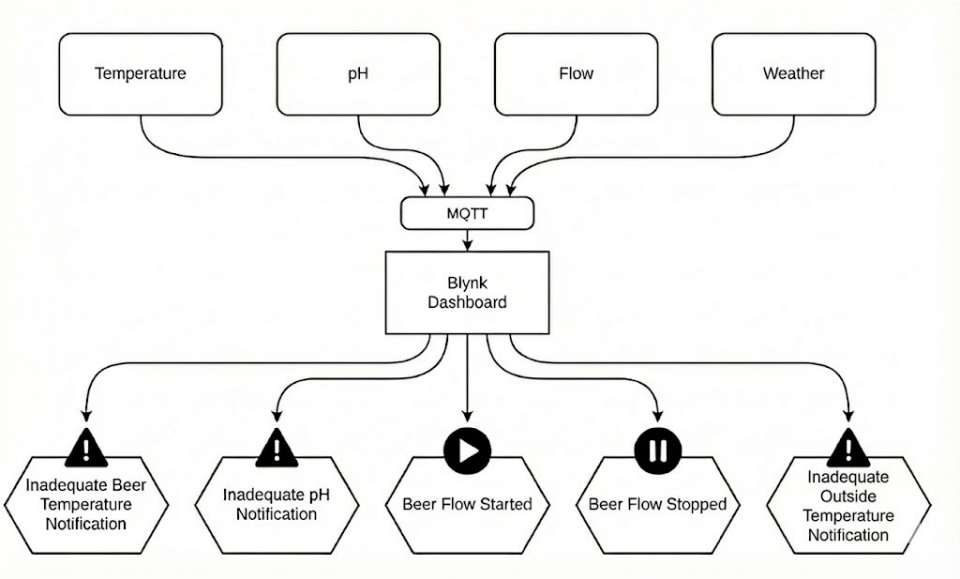

Figure 2. Blynk Notification Logic.

has shown the ESP32 to be well suited for continuous MQTT-based monitoring in resource-constrained IoT deployments [9]. In the proposed system, the microcontroller performs sensor acquisition, basic preprocessing, and data transmission to the cloud layer.

MQTT [5] was adopted as the communication protocol due to its lightweight design, asynchronous message handling, and robustness under unstable network conditions. Its publish–subscribe paradigm decouples data producers from consumers, simplifying integration of multiple subscribers and enabling future scalability without requiring changes to the embedded node [12].

The cloud layer was implemented using the Blynk IoT platform [6]. Blynk provides integrated support for device management, virtual pin mapping, data visualization, and alert generation. This abstraction of backend infrastructure enabled rapid prototyping and allowed development efforts to focus on CPS logic rather than cloud deployment. The user interface was implemented using Blynk's mobile user interface components and includes real-time gauges, historical plots, and threshold-based notifications, as shown in Figure 2. This interface allows brewers to monitor process conditions remotely and receive early warnings of abnormal trends, which is critical for preventing batch failures [8].

To demonstrate system extensibility, the OpenWeather Map API [18][19] was integrated into the architecture. It retrieves ambient temperature and humidity data, which may influence fermentation behavior, and triggers notifications when environmental conditions fall outside predefined ranges. Beyond immediate monitoring, this integration establishes a foundation for context-aware analytics and future predictive models, in line with recent DT-based CPS research trends [10][11].

In parallel, a simulation module was implemented as a Python-based SIL component executed within the development environment rather than on the microcontroller. This simulator interacts directly with the cloud layer through the Blynk HTTP REST API and generates bounded, time-dependent synthetic signals that approximate realistic sensor dynamics. This approach enables systematic stress testing of the cyber layers under controlled conditions, supporting validation of alert logic, dashboard responsiveness, and data handling mechanisms. Such decoupled validation strategies are consistent with emerging practices in virtual sensing and DT-based CPS development [12][15][16].

IV. Evaluation and Results

The system evaluation was conducted in two phases: 1) Physical Validation, which focused on the Negative Temperature Coefficient (NTC) thermistor and MQTT latency; and 2) Logic Validation, utilizing the Virtual Edge Device to test system responsiveness and connectivity.

A. Physical Sensor Calibration

The KY-028 NTC thermistor [20] was integrated into the Hardware Layer, as illustrated in Figure 3.

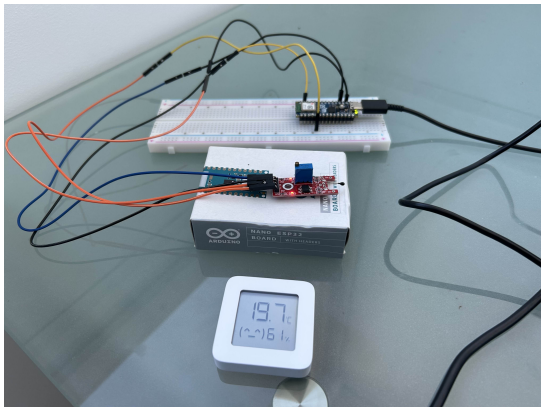

Figure 3. Arduino Nano ESP32, connected to the KY-028 thermistor [20].

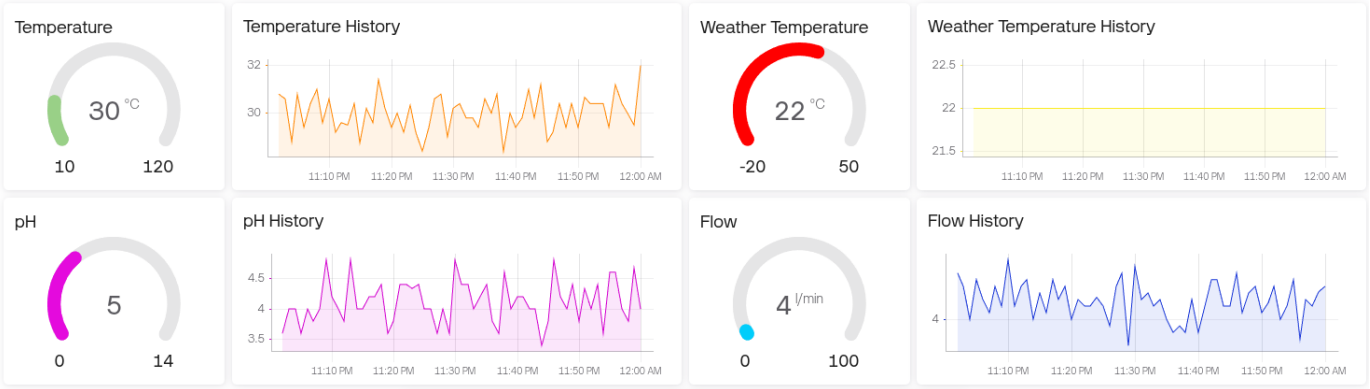

Figure 4. Blynk Dashboard.

To convert the raw analog signal (0-4095) from the ESP32 into degrees Celsius, the Steinhart-Hart equation was implemented within the firmware, as represented in (1), where: i) T is the temperature in Kelvin; ii) R is the resistance in ohms; and iii) A, B, and C are the Steinhart-Hart coefficients.

$$\frac{1}{T} = A + B\ln(R) + C\ln(R)^3 \quad (1)$$

Calibration was performed by waterproofing the sensor and submerging it in a water bath alongside a standard analog fluid thermometer. To address the specific characteristics of the NTC module used, the series resistance parameter (R_{series}) in the voltage divider calculation was adjusted to $49k\Omega$. Post-calibration comparisons confirmed that the sensor readings consistently remained within $\pm 1°$C of the reference thermometer, validating its suitability for the brewing process proof of concept.

B. Logic Validation via Data Simulation

To validate the alert mechanisms and dashboard visualization without the latency of physical brewing processes, the Python automation script described before was executed. Figure 4 represents the Blynk dashboard. The testing protocol involved two data injection strategies:

1) **Nominal State Simulation:** The script injected random values within standard fermentation ranges (e.g., Weather Temperature $16-22°$C, pH $4.0-6.0$) to verify the stability of the dashboard visualization.
2) **Critical State Simulation:** The script injected specific boundary values defined in the code (Temperature=$105°$C, pH=7, Flow=3 L/min) to force-trigger the system's alarm states.

The cloud platform successfully identified the out-of-bound values, making the mobile application trigger the respective high-priority push notifications for each. This confirmed that the conditional logic resides correctly within the cloud layer and is not dependent on the physical sensor's sampling rate.

Regarding connectivity, system latency was measured as the time difference between the transmission of an MQTT packet from the ESP32 (or the Virtual Edge Device) and the visual update on the Dashboard. Over 100 trials, the average end-to-end latency was observed to be $<$ 200ms. Given that fermentation is a slow-moving biochemical process (changing over hours or days), this responsiveness is orders of magnitude faster than required, confirming the suitability of the MQTT protocol for this application.

V. Discussion

The results demonstrate that the proposed five-layer IoT architecture provides a resilient and modular framework for monitoring artisanal brewing processes. Throughout the evaluation period, the system successfully performed data acquisition, MQTT-based transmission, cloud-side processing, and dashboard visualization without observable latency or packet loss. These results confirm the architectural soundness of the proposed CPS and its suitability for small-scale, resource-constrained brewing environments.

In contrast to existing work that prioritizes fully integrated hardware deployments in harsh cellar conditions [8][11], this study contributes a distinctive methodological advantage through the adoption of a SIL validation strategy. A recurring limitation in IoT and CPS research is the prototype-to-deployment bottleneck, where hardware instability or slow biological processes, such as fermentation, significantly delay the validation of cyber components. By decoupling software validation from physical sensor availability [12], this work demonstrated that communication, data handling, alerting logic, and user interfaces can be stress-tested and refined independently of final hardware integration. This approach positions the Digital Brewmaster as a more agile and practical framework for small and medium-sized enterprises, enabling iterative refinement before committing to full sensor deployment.

Despite these positive results, the study faced notable challenges related to hardware procurement and integration. Limited access to high-precision pH and flow sensors constrained the physical scope of the prototype, while the integration of available sensors with the ESP32 required extensive circuit conditioning and troubleshooting. Combined with component delivery delays, these constraints prevented long-duration val-

idation under real fermentation conditions and limited the assessment of sensor drift and environmental noise effects.

VI. Conclusion and Future Work

Overall, this work designed and validated a Digital Twin–oriented CPS for monitoring key brewing parameters and confirmed that a layered IoT architecture can reliably support end-to-end data flow, even under partial hardware availability. By combining real sensor inputs with simulated data streams, the full cyber pipeline was validated in a controlled yet realistic manner, establishing the system as a robust proof of concept rather than a production-ready deployment.

Future work will focus on resolving the identified hardware limitations and completing full sensor integration to enable long-term validation in operational brewing environments. This will allow the evaluation of sensor accuracy, drift, and robustness under real process conditions. In addition, future iterations will extend the current rule-based monitoring toward predictive analytics and machine learning capabilities, enabling early fault detection and autonomous decision support, as envisioned in recent CPS and digital twin research [15]. Through these extensions, the Digital Brewmaster aims to evolve from a monitoring platform into a comprehensive, intelligent assistant for craft brewing operations.

Acknowledgment

This work is financially supported by national funds through the FCT/MCTES (PIDDAC), under the Associate Laboratory Advanced Production and Intelligent Systems – ARISE LA/P/0112/2020 (DOI 10.54499/LA/P/0112/2020) and the Base Funding (UIDB/00147/2020) and Programmatic Funding (UIDP/00147/2020) of the R&D Unit Center for Systems and Technologies – SYSTEC.

References

[1] F. A. Rizos, "Contract law-brewing a solution: An argument for fairness in massachusetts beer franchise laws," *W. New Eng. L. Rev.*, vol. 43, pp. 47–77, 2021.

[2] M. C. Gruba, D. Denes, R. C. G. Lobo, and A. J. Isaak, "Circular economy initiatives: Strategic implications, resource management, and entrepreneurial innovation in a brazilian craft beer ecosystem during the covid era," *Sustainability*, vol. 14, no. 19, 2022, ISSN: 2071-1050. DOI: 10.3390/su141911826.

[3] S. Violino, S. Figorilli, C. Costa, and F. Pallottino, "Internet of beer: A review on smart technologies from mash to pint," *Foods*, vol. 9, no. 7, 2020, ISSN: 2304-8158. DOI: 10.3390/foods9070950.

[4] M. Banzi and M. Shiloh, *Getting Started with Arduino*, 3rd. Sebastopol, CA, USA: Maker Media, Inc., 2014, ISBN: 978-1449363338.

[5] A. Banks and R. Gupta, *MQTT Version 3.1.1*, OASIS Standard, retrieved: February, 2026, 2014. [Online]. Available: https://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html.

[6] D. Babić, I. Jovović, T. Popović, N. Kovač, and S. Čakić, "An internet of things system for environmental monitoring based on esp32 and blynk," in *2022 26th International Conference on Information Technology (IT)*, 2022, pp. 1–5. DOI: 10.1109/IT54280.2022.9743538.

[7] G. S. R. Umang, K. K. Shivanandaswamy, S. Pallavi, and S. Rajagopal, "Software-in-the-loop (sil) method," in *Proceedings of the 10th International Conference on Mechanical, Automotive and Materials Engineering: CMAME 2023, 20–22 December, Da Nang, Vietnam*, Springer Nature, 2024, pp. 243–257. DOI: 10.1007/978-981-97-4806-8_21.

[8] I. Kovačević, I. Aleksi, T. Keser, and T. Matić, "Winnie: A sensor-based system for real-time monitoring and quality tracking in wine fermentation," *Applied Sciences*, vol. 15, no. 21, 2025, ISSN: 2076-3417. DOI: 10.3390/app152111317.

[9] M. A. Syahmi Md Dzahir and K. Seng Chia, "Evaluating the energy consumption of esp32 microcontroller for real-time mqtt iot-based monitoring system," in *2023 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT)*, 2023, pp. 255–261. DOI: 10.1109/3ICT60104.2023.10391358.

[10] E. E. M. Abdurrahman and G. Ferrari, "Digital twin applications in the food industry: A review," *Frontiers in Sustainable Food Systems*, vol. Volume 9 - 2025, 2025, ISSN: 2571-581X. DOI: 10.3389/fsufs.2025.1538375.

[11] P.-E. Goffi, R. Tremblay, and B. Oakes, *Engineering a digital twin for the monitoring and control of beer fermentation sampling*, 2025. arXiv: 2508.18452 [cs.SE].

[12] G. Balan et al., "A perspective on software-in-the-loop and hardware-in-the-loop within digital twin frameworks for automotive lighting systems," *Applied Sciences*, vol. 15, no. 15, 2025, ISSN: 2076-3417. DOI: 10.3390/app15158445.

[13] M. A. Ferrag et al., "Blockchain technologies for the internet of things: Research issues and challenges," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2188–2204, 2019. DOI: 10.1109/JIOT.2018.2882794.

[14] J. Lee, J. Jeong, and S. Kim, "Pressure-guided lstm modeling for fermentation quantification prediction," *Sensors*, vol. 25, no. 17, 2025, ISSN: 1424-8220. DOI: 10.3390/s25175251.

[15] P. Ferrer-Cid, J. Paredes-Ahumada, J. M. Barcelo-Ordinas, and J. Garcia-Vidal, "Virtual sensor-based proxy for black carbon estimation in iot platforms," *Internet of Things*, vol. 27, p. 101 284, 2024, ISSN: 2542-6605. DOI: https://doi.org/10.1016/j.iot.2024.101284.

[16] M. A. Zaidan et al., "Virtual sensors: Toward high-resolution air pollution monitoring using ai and iot," *IEEE Internet of Things Magazine*, vol. 6, no. 1, pp. 76–81, 2023. DOI: 10.1109/IOTM.001.2200103.

[17] M. C. Ravanal, J. P. Doussoulin, and B. Mougenot, "Does sustainability matter in the global beer industry? bibliometrics trends in recycling and the circular economy," *Frontiers in Sustainable Food Systems*, vol. Volume 8 - 2024, 2024, ISSN: 2571-581X. DOI: 10.3389/fsufs.2024.1437910.

[18] S. Chadha, N. Gupta, and R. Chauhan, "Development of weather forecast application using api," in *2024 1st International Conference on Advances in Computing, Communication and Networking (ICAC2N)*, 2024, pp. 653–658. DOI: 10.1109/ICAC2N63387.2024.10895829.

[19] OpenWeather Ltd., *Weather API - OpenWeatherMap*, retrieved: February, 2026, 2026. [Online]. Available: https://openweathermap.org/api.

[20] Arduino Modules Info, *KY-028 Digital Temperature Sensor Module*, retrieved: February, 2026. [Online]. Available: https://arduinomodules.info/ky-028-digital-temperature-sensor-module/.

Edge-Based IoT and AI Framework for Real-Time Wastewater Potability Classification

José Isidro[1], Rafael Teixeira[1], João Costa[1], Carolina Gonçalves[1], Diogo Ferreira[1], Pedro Azevedo[1], Pedro Simões[1], Rui Pinto[1,2], Gil Gonçalves[1,2]
Dept. de Engenharia Informática, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[1]
SYSTEC, ARISE, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[2]
Email:{up202108831, up202006485, up202108714, up202108781, up202205295, up201905966, up202403063} @edu.fe.up.pt
{rpinto, gil}@fe.up.pt

***Abstract*—Access to safe drinking water is a critical public health requirement. However, conventional water quality monitoring approaches remain labor-intensive, intermittent, and heavily dependent on delayed laboratory analyses or cloud-centric infrastructures. Such architectures introduce latency, connectivity dependencies, and limited operational resilience in decentralized or bandwidth-constrained systems. This paper presents a proof-of-concept edge-based Internet of Things (IoT) framework for real-time wastewater potability classification. The proposed system integrates an Arduino-based sensing node for physicochemical data acquisition with a Raspberry Pi edge gateway executing local machine learning inference. By relocating data processing and decision-making from the cloud to the edge, the system enables low-latency, autonomous classification while preserving data locality and operational continuity under limited connectivity. Experimental results demonstrate that resource-constrained edge hardware is capable of supporting real-time water quality assessment and classification, validating the feasibility of edge-centric architectures as a scalable and cost-effective alternative for resilient water quality monitoring systems.**

Keywords-Internet of Things; Edge Computing; Water Quality Monitoring; Artificial Intelligence; Wastewater Treatment.

I. INTRODUCTION

Access to safe drinking water is critical for public health, yet water quality monitoring systems remain largely reactive, relying on labor-intensive sampling procedures and delayed laboratory analyses. In wastewater and industrial water infrastructures, contamination events can evolve rapidly, posing significant risks to environmental safety, economic stability, and community health if not promptly detected and mitigated [1][2]. Consequently, the core problem addressed in this work is the inability of existing Smart Water Quality Monitoring Systems (SWQMS) [3] to support timely, autonomous decision-making under strict latency, reliability, and data-governance constraints.

Recent advances in the Industrial Internet of Things (IIoT) have improved the granularity and frequency of water quality data acquisition. However, the majority of deployed SWQMS architectures remain fundamentally cloud-centric, outsourcing data processing and decision logic to remote data centers. This design introduces a critical latency bottleneck, defined by the round-trip time required for sensor data to reach the cloud and for control actions to be issued in response [4]. In safety-critical wastewater scenarios, such delays can prevent the system from reacting within the narrow temporal window required to contain contamination events, undermining its ability to function as a real-time "digital reflex arc" [5][6].

In addition to latency issues, cloud-based SWQMS solutions suffer from limited operational resilience. The reliance on continuous connectivity creates a single point of failure: when communication with the cloud is disrupted, intelligent monitoring and control capabilities are degraded or lost entirely. Furthermore, the continuous transmission of raw sensor data to external infrastructures raises concerns related to data sovereignty, privacy, and regulatory compliance within industrial environments. Collectively, these limitations reveal a misalignment between the architectural assumptions of cloud-centric SWQMS and the real-time, reliability-critical requirements of water quality management.

To address this gap, this paper proposes an Edge-based SWQMS grounded in edge analytics, where Machine Learning (ML) inference is performed directly at the network perimeter. By relocating the intelligence layer from Cloud to Edge, the proposed system enables low-latency, autonomous decision-making while preserving full functionality under intermittent or absent external connectivity. This architectural shift directly targets the identified problem by prioritizing speed, reliability, and data locality as first-class design objectives.

Accordingly, the main contributions of this work are:

1) The design and implementation of an Arduino-based sensor node for continuous, real-time acquisition of physicochemical water quality parameters.
2) The deployment of a localized ML inference engine, hosted on a Raspberry Pi and exposed through a FastAPI backend, enabling autonomous and low-latency decision-making at the Edge.
3) The development of an integrated local dashboard that provides real-time operational insights while ensuring that all sensitive telemetry remains confined to the on-premise network.

The remainder of this paper is organized as follows: Section II reviews Related Work; Section III details the Methods and System Architecture; Section IV describes the Experimental Setup; Section V presents Results and Discussion; and Section VI concludes the study and outlines future work.

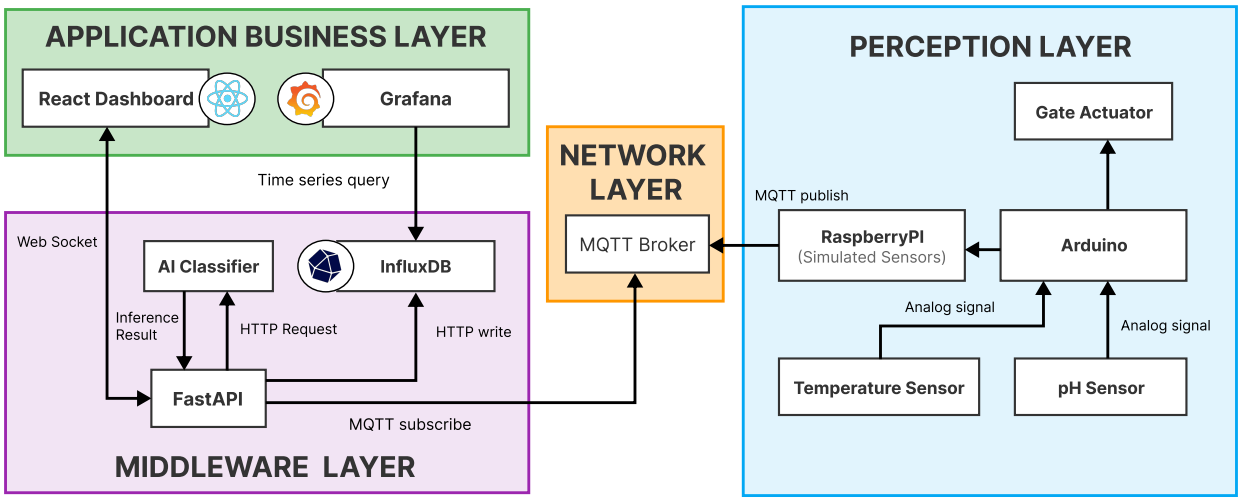

Figure 1. High-level architecture diagram illustrating the overall system design.

II. Related Work

The management of industrial and municipal wastewater is subject to strict environmental regulations, requiring frequent and reliable spatiotemporal monitoring to ensure public health protection and environmental sustainability [7]. Within this context, *Water Quality Monitoring* (WQM) refers to the systematic assessment of physicochemical and biological parameters—such as pH, temperature, turbidity, conductivity, and dissolved solids—to determine water safety and compliance with regulatory thresholds. Traditionally, WQM has played a central role in pollution detection, regulatory enforcement, and operational control in water treatment and distribution systems.

Historically, WQM has relied predominantly on manual sampling followed by laboratory-based analysis. While this approach provides high measurement accuracy, it presents several limitations [8]. Sampling campaigns are typically sporadic and spatially sparse, resulting in limited coverage and reduced sensitivity to short-lived contamination events. Moreover, the inherent delay between sample collection and laboratory results introduces significant latency, preventing timely responses in critical situations where rapid intervention is required [9]. In addition to these constraints, laboratory-based monitoring is labor-intensive and costly, requiring specialized equipment and trained personnel, which limits its feasibility for continuous or large-scale deployment [10].

To address these shortcomings, recent research and industrial practice have increasingly embraced the IIoT paradigm, giving rise to SWQMS [3]. SWQMS leverage distributed sensing, embedded computing, and networked communication to enable continuous, automated, and near real-time monitoring of water quality parameters. By replacing manual sampling with in situ sensors and automated data acquisition, these systems improve temporal resolution, enhance situational awareness, and reduce operational costs. Furthermore, IIoT-based approaches promote system decentralization, interoperability, and scalability, supporting heterogeneous devices and facilitating integration with higher-level data analytics and decision-support tools [11].

From a systems perspective, SWQMS typically build upon the generic layered architecture widely adopted in IoT applications. In its most common form, this architecture consists of five conceptual layers [12]–[15]. The *Sensing Layer* interfaces directly with the physical environment and includes sensors and embedded devices responsible for data acquisition. The *Network or Gateway Layer* ensures data transmission and aggregation, using wired or wireless communication technologies, such as Wi-Fi, cellular networks, or industrial field-buses. Collected data are then handled by a *Data or Cloud Layer*, where information is ingested via protocols, such as Message Queuing Telemetry Transport (MQTT) or Hypertext Transfer Protocol (HTTP), and stored, often in scalable Not Only SQL (NoSQL) databases. On top of this, the *Analytics Layer* processes raw telemetry into actionable knowledge, performing tasks such as anomaly detection, threshold-based alerts, and performance indicator computation [16]. Finally, the *Presentation Layer* provides user-facing interfaces for visualization, monitoring, and control, enabling operators to observe system status in real time and, in some cases, trigger remote actuation [17][18].

While this layered architecture is effective, most implementations remain cloud-centric. Recent efforts have begun exploring edge computing for environmental monitoring. For instance, Yadav *et al.* [19] proposed a framework for simultaneous air and water monitoring, emphasizing the reduction of network load. Similarly, Ren *et al.* [20] demonstrated how edge nodes can perform preliminary data filtering to enhance system longevity. Unlike these approaches, which often still treat the edge as a pre-processing buffer for the cloud, our solution proposes a fully autonomous edge gateway capable of localized Machine Learning (ML) inference and decision-making, ensuring operation even during total backhaul failure. This architectural shift prioritizes speed and data locality.

III. Methods and System Architecture

Unlike traditional cloud-centric approaches, this proof of concept adopts a local edge-based architecture, illustrated in Figure 1. The system operates on a Raspberry Pi coupled with an Arduino Nano acting as the sensor node, enabling low-latency data processing and AI inference without reliance on external cloud services [21][22]. This design prioritizes responsiveness, privacy, and operational continuity in environments with limited or unreliable internet connectivity.

The Arduino Nano serves as the primary interface with the physical environment, handling analog signal acquisition and actuator control. It performs analog-to-digital conversion for connected sensors and directly actuates elements, such as LED indicators, based on classification feedback received from higher layers. While the pH sensor is fully integrated into the hardware, the system accounts for missing sensing equipment—namely, Total Dissolved Solids (TDS), turbidity, and hardness sensors—through a custom simulation engine. This engine, hosted on the Raspberry Pi, generates bounded synthetic sensor data, enabling full system validation and end-to-end testing despite partial hardware availability.

Data exchange between the embedded hardware and processing components is achieved using the MQTT protocol via a Mosquitto broker [23]. This event-driven communication model was selected due to its lightweight overhead and robustness under unstable network conditions. The Arduino Nano publishes raw telemetry data to predefined topics, which are asynchronously consumed by backend services. This decoupled design supports scalability and flexibility, allowing additional sensor nodes or parallel dashboards to be integrated with minimal changes to the communication layer.

At the core of the system, a middleware component functions as the central intelligence layer, bridging data acquisition and user interaction. Implemented using FastAPI [24], this layer enables high-throughput data ingestion and low-latency processing. Incoming data streams are analyzed in real time using a local scikit-learn classification model, which evaluates water chemistry and produces potability predictions within seconds. Simultaneously, all measurements are persisted in InfluxDB [25], a time-series database optimized for high-frequency sensor data. By performing AI inference locally and avoiding cloud-based computation, the system ensures data privacy and remains fully operational in offline scenarios.

The user interface was designed to address the needs of both technical and non-technical stakeholders. Grafana [26] supports in-depth engineering analysis through detailed visualizations of raw sensor data and system behavior over time. In parallel, a lightweight React-based web application [27], shown in Figure 2, provides an intuitive interface for end-users. This application uses WebSockets to maintain a persistent connection with the middleware, ensuring that sensor updates, classification results, and alerts are reflected instantly on the dashboard without requiring manual page refreshes.

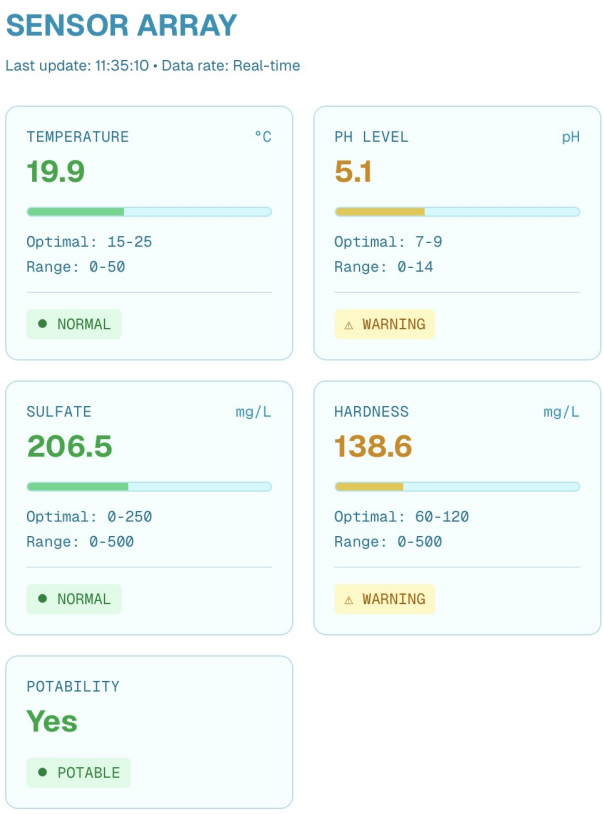

Figure 2. React Web App Dashboard sample.

IV. Experimental Setup

The physical layer of the system comprises two main computational units: an ESP32-based Arduino Nano and a Raspberry Pi 3, as illustrated in Figure 3. The Arduino Nano operates as the sensing node, interfacing directly with the physical environment through analog data acquisition. It is connected to a set of analog sensors, including a pH probe, enabling real-time measurement of key physicochemical parameters from the water source. The Raspberry Pi functions as the edge gateway, aggregating sensor data received via MQTT and executing the Python-based AI inference engine.

The prototype enclosure was implemented using a low-cost, easily reproducible design intended to manage and route water flow based on the system's classification outcome. A cut plastic bottle serves as the primary container for receiving incoming water samples and housing the sensing process.

At the bottle outlet, a flexible hose segment connects to a manual two-way valve that acts as a physical bifurcation point. This valve is linked to two additional hose segments, each leading to a separate container. One container is designated for water classified as potable, while the other collects water identified as non-potable. The routing decision is conceptually driven by the inference result: samples classified as safe are directed to the potable container, whereas unsafe samples are routed to the alternative container.

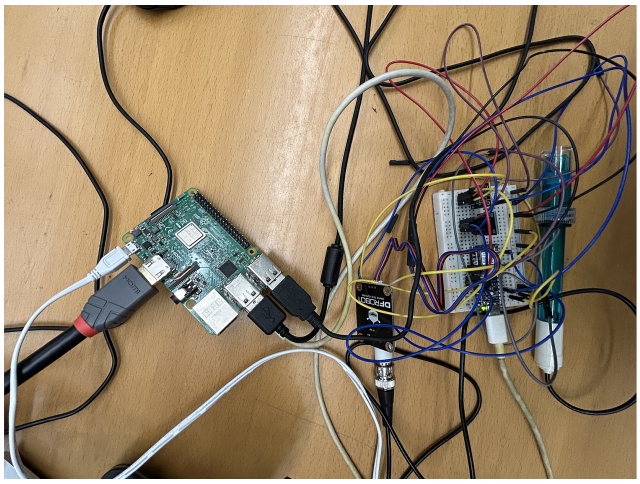

Figure 3. Overall implemented system integrating the Arduino and Raspberry Pi with all sensors.

This design provides a clear physical separation between potable and non-potable water, demonstrating the practical applicability of the proposed system in real-world scenarios. At the same time, it intentionally avoids permanent structural housing or refined mechanical fabrication, reinforcing the prototype's focus on functional validation, affordability, and reproducibility rather than industrial-grade deployment.

A. Dataset and Metrics

The proposed models were trained and evaluated using the Water Quality dataset [28], which comprises 3,276 samples described by the physicochemical indicators listed in Table I. This dataset is well-suited to the study for three main reasons: (i) its multivariate structure reflects the complexity of water potability assessment; (ii) the weak linear correlations among features capture chemical interdependencies, motivating learning-based approaches over fixed threshold laboratory checks; and (iii) its alignment with global regulatory standards supports future scalability and deployment.

The dataset contains nine numerical features and a binary target variable. Missing values were identified in several attributes, most notably pH (14.99%), sulfate (23.84%), and trihalomethanes (4.95%). Based on feature distribution analysis, mean imputation was applied to address these gaps.

A primary limitation of this study is the synthetic origin of the dataset. The high density of unique values across multiple features suggests limited underlying source diversity, which may reduce the granularity of learned decision boundaries and affect model generalization.

As the physical prototype currently supports only pH sensing, the remaining chemical parameters required by the model were synthetically generated to enable end-to-end testing. These values were constrained to World Health Organization (WHO) recommended ranges, ensuring complete feature vectors for inference. Model performance was evaluated using complementary classification metrics [29], including accuracy (ACC), precision ($PREC$), recall ($Recall$), $F1$-score, and the area under the Receiver Operating Characteristic (ROC) curve (AUC), providing a comprehensive assessment under class imbalance and non-linear decision boundaries. The binary potability label follows standard encoding, where 0 denotes unsafe water, and 1 indicates water suitable for human consumption.

TABLE I. CHARACTERISTICS OF WATER QUALITY MEASUREMENTS.

Material	Description	WHO Range
pH	Acidity or alkalinity	6.5 to 8.5
Hardness	Ca/Mg salts	-
Solids (TDS)	Dissolved minerals	500 to 1000 mg/L
Chloramines	Disinfectants	$\leq$ 4 mg/L
Sulfate	Mineral compound	3 to 30 mg/L
Conductivity	Electrical ability	$\leq$ 400 μS/cm
Organic Carbon	Pollutant indicator	$<$ 4 mg/L
Trihalomethanes	Chlorination by-product	$\leq$ 80 ppm
Turbidity	Suspended solids	$\leq$ 5 NTU
Potability	1=Safe; 0=Unsafe	-

B. ML models and code architecture

The proposed system requires an ML model capable of assessing water potability; however, limitations related to data availability, data quality, and the gap between synthetic and real-world data prevent strong guarantees on model performance. To address these constraints and ensure long-term adaptability, a flexible model management architecture was implemented to support the seamless introduction of new datasets and alternative models as they become available.

The architecture is structured around two core components: *Model_Class* and *Controller_Models*. The *Model_Class* encapsulates model-specific functionality, including training, evaluation, persistence, and metadata management. Each model, along with its configuration and performance metrics, can be serialized and restored from a JSON representation. The *Controller_Models* component orchestrates model usage by loading available models from a predefined repository, selecting a designated primary model, and executing inference requests. This separation of concerns enables transparent model replacement and simplifies future system evolution.

To explore water potability from diverse algorithmic perspectives, the model selection strategy spans linear baselines to non-linear ensemble methods. Logistic Regression (LR) and a Stochastic Gradient Descent (SGD) classifier with log loss were implemented as baseline models to evaluate whether a global linear decision boundary can effectively separate the nine-dimensional chemical feature space.

Complementing these global approaches, the k-Nearest Neighbors (kNN) algorithm was employed to capture local data structure, operating under the assumption that chemically similar samples exhibit similar potability outcomes. To further

TABLE II. COMPREHENSIVE PERFORMANCE METRICS AND TRAINING TIME (IN SECONDS) FOR TESTED MODELS.

Model	$F1$	ACC	AUC	$PREC$	$Recall$	$Time$
LR	0.48	0.63	0.52	0.39	0.63	0.424
SGD	0.48	0.63	0.50	0.39	0.63	0.013
kNN	0.49	0.63	0.51	0.50	0.62	0.0117
CNB	0.54	0.53	0.53	0.55	0.53	0.002
RF	**0.65**	**0.68**	**0.69**	**0.67**	**0.68**	**0.884**

address class imbalance—a common characteristic of environmental datasets—Complement Naive Bayes (CNB) was included. Although the current dataset exhibits limited imbalance, CNB is architecturally suited to mitigate the dominance of the majority (non-potable) class and is expected to be particularly relevant in real-world deployments.

Finally, a Random Forest (RF) model was used to capture complex, non-linear relationships among features. By aggregating multiple decision trees through ensemble learning, RF enables the identification of intricate decision boundaries and mitigates variance introduced by data imputation. This multi-model strategy provides a comprehensive assessment of whether water safety can be adequately described by linear thresholds or instead emerges from highly specific, non-linear chemical interactions.

V. RESULTS AND DISCUSSION

This section evaluates the system from two perspectives: the predictive accuracy of the localized AI models and the operational performance of the integrated edge architecture.

A. Model Performance Evaluation

The comparative analysis of the tested algorithms is summarized in Table II.

The empirical results demonstrate that water potability classification is a significantly non-linear task, as evidenced by the failure of linear architectures. Both LR and SGD converged to an AUC of approximately 0.5, indicating that linear hyperplanes possess near-zero discriminatory power for this feature space and effectively default to majority-class guessing. While CNB achieved a marginally higher $F1$ score of 0.54, its poor accuracy of 0.53 reveals a high false-positive rate, failing to reliably distinguish chemical safety margins.

The RF emerged as the best architecture, achieving an AUC of 0.69 and a balanced $F1$ score of 0.65. Its success is tied to its ensemble nature, which mitigates the impact of noise and the variance introduced by missing value imputation, factors that crippled the distance-based kNN. Using recursive partitioning, the RF successfully captured high-order interactions between variables. This performance gap confirms that water potability is not governed by independent feature thresholds but by a multifaceted chemical synergy that only the ensemble-based decision trees could effectively resolve.

B. System Integration and Latency Analysis

Beyond algorithmic accuracy, the system's ability to function as an integrated control unit was validated through real-time testing. The Edge AI model demonstrated the capability to process incoming MQTT packets and return the classification within milliseconds. By hosting the inference engine locally via FastAPI, the system eliminated the "Cloud Round-Trip Time", which in industrial IoT environments can fluctuate between 100ms and several seconds depending on network congestion [4].

The React-based dashboard successfully maintained a persistent WebSocket connection, reflecting potability changes instantly. This responsiveness is critical for the "Control Rule" aspect of the project, as any delay in detecting non-potable water could lead to the contamination of downstream reservoirs before a manual or cloud-based intervention could occur.

VI. CONCLUSION AND FUTURE WORK

This work presents a proof of concept for a wastewater detection system combining local IoT sensing with edge-level AI inference. The results demonstrate that relocating intelligence from the cloud to the edge improves responsiveness and operational reliability, providing deterministic latency suitable for control logic. The study further validates that low-cost platforms, such as the Raspberry Pi 3, can execute complex ensemble models, confirming that non-linear water quality interactions can be effectively resolved at the edge in a scalable and accessible manner.

However, several limitations define the current boundary conditions. The system relies on a partially synthetic dataset, as several chemical parameters were generated to compensate for incomplete hardware integration. In addition, missing values in the source data required mean imputation, potentially simplifying learned decision boundaries. While sufficient for validating the edge architecture and middleware logic, the resulting model may not fully reflect the variability, noise, and stochastic behavior of real wastewater streams.

Future work will prioritize full sensor integration to eliminate reliance on simulated inputs and improve real-world robustness. The architecture is also designed to evolve from a monitoring solution into a closed-loop control system through the replacement of manual routing with electronically actuated solenoid valves. Validating automated actuation within the existing framework will enable real-time, autonomous water diversion, enhancing system resilience and applicability in decentralized environments.

From a broader socio-economic perspective, such edge-enabled systems may support municipal authorities in reducing operational costs and preventing contamination through rapid, localized decision-making. The modular design and use of standardized protocols, including MQTT, facilitate integration with existing industrial infrastructures without requiring large-scale system overhauls.

ACKNOWLEDGMENTS

This work is financially supported by national funds through the FCT/MCTES (PIDDAC), under the Associate Laboratory Advanced Production and Intelligent Systems – ARISE LA/P/0112/2020 (DOI 10.54499/LA/P/0112/2020) and the Base Funding (UIDB/00147/2020) and Programmatic Funding

(UIDP/00147/2020) of the R&D Unit Center for Systems and Technologies – SYSTEC.

REFERENCES

[1] A. Oros, "Bioaccumulation and trophic transfer of heavy metals in marine fish: Ecological and ecosystem-level impacts," *Journal of Xenobiotics*, vol. 15, no. 2, pp. 59–72, 2025, [retrieved: February, 2026], ISSN: 2039-4713. DOI: 10.3390/jox15020059.

[2] Y. Gelaye, "Public health and economic burden of heavy metals in ethiopia: Review," *Heliyon*, vol. 10, no. 19, Oct. 2024, ISSN: 2405-8440. DOI: 10.1016/j.heliyon.2024.e39022.

[3] R. Martínez et al., "On the use of an iot integrated system for water quality monitoring and management in wastewater treatment plants," *Water*, vol. 12, no. 4, 2020, ISSN: 2073-4441. DOI: 10.3390/w12041096. [Online]. Available: https://www.mdpi.com/2073-4441/12/4/1096.

[4] P. Hu, C. He, Y. Zhu, and T. Li, "The product quality inspection scheme based on software-defined edge intelligent controller in industrial internet of things," *Journal of Cloud Computing*, vol. 12, no. 1, pp. 113–128, 2023, [retrieved: February, 2026], ISSN: 2192-113X. DOI: 10.1186/s13677-023-00487-7.

[5] R. Wiryasaputra, C.-Y. Huang, Y.-J. Lin, and C.-T. Yang, "An iot real-time potable water quality monitoring and prediction model based on cloud computing architecture," *Sensors*, vol. 24, no. 4, 2024, ISSN: 1424-8220. DOI: 10.3390/s24041180.

[6] F. Tosi et al., "Enabling image-based streamflow monitoring at the edge," *Remote Sensing*, vol. 12, no. 12, 2020, ISSN: 2072-4292. DOI: 10.3390/rs12122047.

[7] U. Zhalmagambetova, D. Assanov, A. Neftissov, A. Biloshchytskyi, and I. Radelyuk, "Implications of water quality index and multivariate statistics for improved environmental regulation in the irtysh river basin (kazakhstan)," *Water*, vol. 16, no. 15, pp. 2203–2220, 2024, [retrieved: February, 2026], ISSN: 2073-4441. DOI: 10.3390/w16152203.

[8] H. H. Lou et al., "A new area of utilizing industrial internet of things in environmental monitoring," *Frontiers in Chemical Engineering*, vol. 4, p. 842514, 2022.

[9] H. Cao et al., "Advancing clinical biochemistry: Addressing gaps and driving future innovations," *Front Med (Lausanne)*, vol. 12, p. 1521126, Apr. 2025.

[10] A. T. Chafa, G. Chirinda, and S. Matope, "Design of a real–time water quality monitoring and control system using internet of things (iot)," *Cogent Engineering*, vol. 9, 2022.

[11] A. Benis, O. Tamburis, C. Chronaki, and A. Moen, "One digital health: A unified framework for future health ecosystems," *J Med Internet Res*, vol. 23, no. 2, e22189, Feb. 2021, ISSN: 1438-8871. DOI: 10.2196/22189. [Online]. Available: http://www.ncbi.nlm.nih.gov/pubmed/33492240.

[12] Y. Dai, Z. Huang, N. Khan, and M. S. Labbo, "Smart water management: Governance innovation, technological integration, and policy pathways toward economic and ecological sustainability," *Water*, vol. 17, no. 13, 2025, ISSN: 2073-4441. DOI: 10.3390/w17131932.

[13] J. Salgado, C. Pizarro, L. Wong, and J. Castillo, "Iot watercare: Water quality control system in unofficial settlements of peru based in an iot architecture," in *2022 31st Conference of Open Innovations Association (FRUCT)*, 2022, pp. 277–288. DOI: 10.23919/FRUCT54823.2022.9770881.

[14] J. Li, X. Yang, and R. Sitzenfrei, "Rethinking the framework of smart water system: A review," *Water*, 2020.

[15] R. M. M. Salem, M. S. Saraya, and A. M. T. Ali-Eldin, "An industrial cloud-based iot system for real-time monitoring and controlling of wastewater," *IEEE Access*, vol. 10, pp. 7096–7121, 2022.

[16] W. Zhang, F. Ma, M. Ren, and F. Yang, "Application with internet of things technology in the municipal industrial wastewater treatment based on membrane bioreactor process," *Applied Water Science*, vol. 11, no. 52, p. 52, 2021.

[17] A. G. Orozco-Lugo et al., "Monitoring of water quality in a shrimp farm using a fanet," *Internet of Things*, vol. 18, p. 100170, 2022, ISSN: 2542-6605. DOI: https://doi.org/10.1016/j.iot.2020.100170.

[18] A. Mamani-Saico and P. R. Yanyachi, "Implementation and performance study of the micro-ros/ros2 framework to algorithm design for attitude determination and control system," *IEEE Access*, vol. 11, pp. 128451–128460, 2023. DOI: 10.1109/ACCESS.2023.3330441.

[19] S. Yadav and A. Kumar, "Integrating edge computing and IoT for real-time air and water quality monitoring systems," *International Journal of Engineering Science*, vol. 11, pp. 1507–1511, 2025, [retrieved: February, 2026]. [Online]. Available: https://theaspd.com/index.php/ijes/article/view/4588.

[20] J. Ren, Q. Zhu, and C. Wang, "Edge computing for water quality monitoring systems," *Mobile Information Systems*, vol. 2022, no. 1, p. 5056606, 2022, [retrieved: February, 2026]. DOI: https://doi.org/10.1155/2022/5056606. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2022/5056606. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1155/2022/5056606.

[21] E. Upton and G. Halfacree, *Raspberry Pi User Guide*, 3rd. Indianapolis, IN, USA: Wiley, 2014, ISBN: 978-1118921661.

[22] M. Banzi and M. Shiloh, *Getting Started with Arduino*, 3rd. Sebastopol, CA, USA: Maker Media, Inc., 2014, ISBN: 978-1449363338.

[23] R. A. Light, "Mosquitto: Server and client implementation of the mqtt protocol," *Journal of Open Source Software*, vol. 2, no. 13, p. 265, 2017.

[24] M. Lathkar, "Getting started with fastapi," in *High-Performance Web Apps with FastAPI: The Asynchronous Web Framework Based on Modern Python*. Berkeley, CA: Apress, 2023, pp. 29–64, ISBN: 978-1-4842-9178-8. DOI: 10.1007/978-1-4842-9178-8_2.

[25] InfluxData, *InfluxDB Documentation*, retrieved: February, 2026, 2025. [Online]. Available: https://docs.influxdata.com/influxdb/.

[26] S. Kirešová et al., "Grafana as a visualization tool for measurements," in *2023 IEEE 5th International Conference on Modern Electrical and Energy System (MEES)*, IEEE, 2023, pp. 1–5.

[27] Meta Platforms, Inc., *React: A JavaScript Library for Building User Interfaces*, retrieved: February, 2026, 2025. [Online]. Available: https://reactjs.org/.

[28] Aditya Kadiwal and Kaggle Contributors, *Water Potability Dataset*, retrieved: February, 2026, 2020. [Online]. Available: https://www.kaggle.com/datasets/adityakadiwal/water-potability.

[29] G. Naidu, T. Zuva, and E. M. Sibanda, "A review of evaluation metrics in machine learning algorithms," in *Artificial Intelligence Application in Networks and Systems*, R. Silhavy and P. Silhavy, Eds., Cham: Springer International Publishing, 2023, pp. 15–25, ISBN: 978-3-031-35314-7.

An Edge-Centric IoT System for Smart Building Energy Management

João Silva[1], João Alves[1], Nuno Silva[1], Abdul Rauf[1], Victor Bongard[1],
Victor Rodríguez[2], Nuno Moreira[2],
Rui Pinto[1,3], Gil Gonçalves[1,3]
Dept. de Engenharia Informática, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[1]
Dept. de Ciência de Computadores, Faculdade de Ciências, Universidade do Porto, Porto, Portugal[2]
SYSTEC, ARISE, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal[3]
Email: {up202108713, up202108670, up202005501, up202502193, up202502847}@edu.fe.up.pt
{up202105114, up202104873}@edu.fc.up.pt
{rpinto, gil}@fe.up.pt

***Abstract*—Rising energy costs and suboptimal environmental control in industrial and service facilities are frequently associated with limited real-time visibility, static control strategies, and low resilience in conventional building management systems. Many existing solutions rely on centralized architectures, which increases latency and reduces robustness under connectivity disruptions, while offering limited support for proactive optimization. This paper presents a modular, edge-centric Internet of Things (IoT) platform for multi-zone environmental monitoring and control, designed to preserve local autonomy under partial failure conditions. Environmental and occupancy data are acquired by microcontroller-based sensor nodes and disseminated via Message Queuing Telemetry Transport (MQTT) to a local edge gateway, where data validation, aggregation, and coordination are performed before selective synchronization with backend services for persistence and visualization. The platform integrates lightweight forecasting mechanisms that combine historical occupancy patterns with external weather information to support anticipatory control decisions. The proposed architecture enforces a clear separation between data and control planes, supports scalable multi-zone deployments, and enables energy-efficient operation with low deployment overhead, making it suitable for small- and medium-scale facilities.**

Keywords-IoT; CPS; MQTT; InfluxDB; SQLite; Grafana; Edge Computing; AI Forecast.

I. Introduction

The increasing pressure to reduce energy consumption and meet sustainability targets has intensified the need for more efficient management of lighting and Heating, Ventilation, and Air Conditioning (HVAC) systems in industrial and service facilities [1][2]. Despite advances in automation technologies, many buildings still rely on static schedules and coarse-grained control strategies, providing limited real-time visibility into environmental conditions and system behavior. Consequently, energy usage is often misaligned with actual occupancy and operational patterns, leading to waste and reduced efficiency [3]. Although Internet of Things (IoT) and Cyber-Physical Systems (CPS) approaches enable fine-grained sensing and adaptive control [4], many existing solutions remain centralized and cloud-dependent, reducing robustness under network disruptions and limiting suitability for resource-constrained environments. Moreover, fragmented deployments frequently result in transient or poorly persisted data, hindering long-term analysis and informed decision-making.

To address these limitations, this work proposes a resilient, edge-centric IoT/CPS platform for multi-zone environmental monitoring and control. The system connects sensors, microcontrollers, edge services, and backend components through a scalable architecture that separates telemetry from control execution. It supports real-time data acquisition, persistence, and visualization, exposes well-defined APIs for integration, and enables closed-loop control of lighting and HVAC systems. Lightweight forecasting mechanisms based on occupancy patterns and weather data allow proactive, energy-aware operation [5], while preserving local autonomy under connectivity failures.

The main contributions of this paper are threefold: (1) the design of a modular, edge-first IoT/CPS architecture that maintains operation during gateway or network disruptions; (2) the implementation of a clear separation between telemetry and control flows to enhance robustness and scalability; and (3) the integration of lightweight, edge-based decision support for occupancy- and weather-aware environmental control within a practical proof-of-concept deployment.

The remainder of this paper is organized as follows. Section II reviews related work and identifies research gaps. Section III details the proposed architecture and system design. Section IV presents implementation and evaluation results. Finally, Section V discusses limitations and outlines future research directions.

II. Related Work

Smart building energy management has evolved from traditional, proprietary Building Management Systems (BMS) toward IoT- and CPS-based architectures integrating distributed sensing, communication, and actuation. Conventional BMS solutions are typically centralized and offer limited interoperability and adaptability, motivating the adoption of modular IoT-based approaches as energy efficiency and sustainability requirements increase.

Recent surveys show that contemporary Building Energy Management Systems (BEMSs) increasingly rely on distributed IoT sensors, wireless protocols, and data-driven analytics to enable real-time monitoring and adaptive control. Reviews by Akbulut *et al.* and Poyyamozhi *et al.* highlight

TABLE I. ARCHITECTURAL COMPARISON WITH REPRESENTATIVE ENERGY MANAGEMENT PARADIGMS

Feature	Centralized IoT/BEMS	AI/DL Cloud HVAC	Microgrid EMCS	This Work
Primary processing location	Cloud	Cloud	Hierarchical/Central	Edge-first
Local autonomy under backend failure	Limited	Limited	Partial	Yes
Explicit data/control separation	Rare	Not emphasized	Layered control	Yes
Edge telemetry buffering	Uncommon	Uncommon	Application-specific	Yes
Acknowledgment-based actuation	Not typical	Not typical	Application-specific	Yes
Lightweight edge deployment	Variable	High compute	Grid-scale infra	Yes (RPi/MCU)
Edge-level predictive logic	Rare	Cloud-based	Optimization-based	Yes (lightweight)
Target deployment scale	Building	Building	Grid/Microgrid	Multi-zone building

the convergence of IoT, Artificial Intelligence (AI), and energy management, reporting consistent reductions in HVAC and lighting energy consumption while preserving occupant comfort [3][4]. These studies confirm the effectiveness of fine-grained sensing and analytics but also reveal a strong dependence on centralized processing infrastructures.

Occupancy-aware control has become a key research direction, replacing static schedules with dynamic adaptation based on inferred or predicted presence. Probabilistic occupancy prediction models have demonstrated improved HVAC optimization compared to binary detection approaches [5]. In parallel, predictive strategies such as Model Predictive Control (MPC) have shown measurable energy savings when combined with real-time sensor data and forecasts [2]. AI- and Deep Learning-based approaches further enhance adaptability under dynamic environmental and occupancy conditions [1]. However, many of these methods assume centralized computation, large datasets, and substantial processing resources, limiting applicability in lightweight or resource-constrained environments.

Energy-management optimization is also widely investigated beyond buildings. Microgrid Energy Management and Control Systems (EMCS) integrate layered control and optimization objectives for grid stability [6], while model-free wind-farm control and learning-augmented predictive control in hybrid and plug-in hybrid electric vehicles demonstrate the broader applicability of data-driven and predictive strategies across energy domains [7]–[9]. These cross-domain applications reinforce the trend toward predictive and model-based optimization but often target large-scale infrastructures rather than building-level deployments.

Despite significant progress in sensing, analytics, and predictive control, architectural limitations remain. Many IoT/BEMS solutions retain strong cloud dependency for persistence, analytics, and decision-making, increasing latency and reducing robustness under network disruptions [10]. Edge components are frequently treated as data relays rather than autonomous coordination layers, and explicit separation between telemetry and control flows is seldom emphasized [11]. Moreover, resilience mechanisms such as telemetry buffering, acknowledgment-based actuation tracking, and fault isolation across architectural layers are not systematically addressed in building-scale systems. While Digital Twin and AI-driven approaches represent the state of the art, they often introduce architectural complexity and computational overhead that hinder practical deployment in small- and medium-scale environments [12]. Additionally, reproducible and privacy-preserving evaluation strategies—such as the controlled use of synthetic data—remain underexplored in applied IoT/CPS prototypes [13].

Table I summarizes the architectural positioning of the proposed system relative to representative paradigms in the literature. In contrast to centralized BEMS and AI-driven HVAC solutions, which prioritize optimization but depend on continuous backend availability, the proposed architecture emphasizes edge autonomy, explicit separation of telemetry and control planes, acknowledgment and buffering mechanisms, and low-overhead deployment suitable for small- and medium-scale facilities.

III. PROPOSED SOLUTION

The proposed system is an IoT-based BEMS for intelligent environmental monitoring and control in indoor spaces. Although motivated by industrial and office scenarios, the architecture is generic and adaptable to residential or mixed-use environments. It combines distributed sensing, edge coordination, and backend services to automate lighting and HVAC control, targeting energy efficiency while maintaining comfort constraints.

The design follows a layered IoT/CPS architecture that separates sensing, communication, processing, and user interaction. Embedded devices perform local data acquisition

and actuation; an edge gateway coordinates communication and autonomy; backend services provide persistence and visualization without being required for core operation. The architecture prioritizes resilience, modularity, and scalability, enabling continued local functionality under connectivity disruptions and incremental system expansion through standard protocols and well-defined interfaces. In addition to reactive control, the system integrates lightweight, edge-based decision support mechanisms—such as occupancy- and weather-aware logic—to demonstrate anticipatory control within a practical proof-of-concept deployment.

A. Layered IoT Architecture

The system adopts a layered IoT/CPS architecture (Figure 1) that separates physical, communication, processing, and user-facing concerns, aligning with CPS principles of sensing, computation, actuation, and feedback.

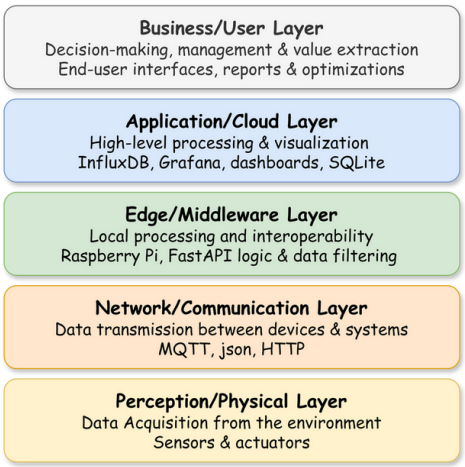

Figure 1. High-level overview of the adopted architecture.

- Perception/Physical Layer: ESP32-based devices connected to environmental sensors and actuators perform local data acquisition and immediate actuation.
- Network/Communication Layer: MQTT [14] supports publish–subscribe device-to-edge communication, while HTTP enables interaction among edge, backend, and frontend services.
- Edge/Middleware Layer: Implemented on a Raspberry Pi [15], this layer hosts the MQTT broker and services for device coordination, aggregation, buffering, and local autonomy.
- Application/Cloud Layer: Provides persistence and system state management. In the proof of concept, services are containerized and deployed locally but remain architecturally decoupled from the edge.
- Business/User Layer: Exposes visualization, configuration, and supervisory functions through a web-based interface.

Although the current implementation supports a single gateway, the architecture is designed to scale across multiple zones and multi-gateway deployments.

B. Data and Control Flow

System operation is structured around two explicit runtime paths (Figure 2):

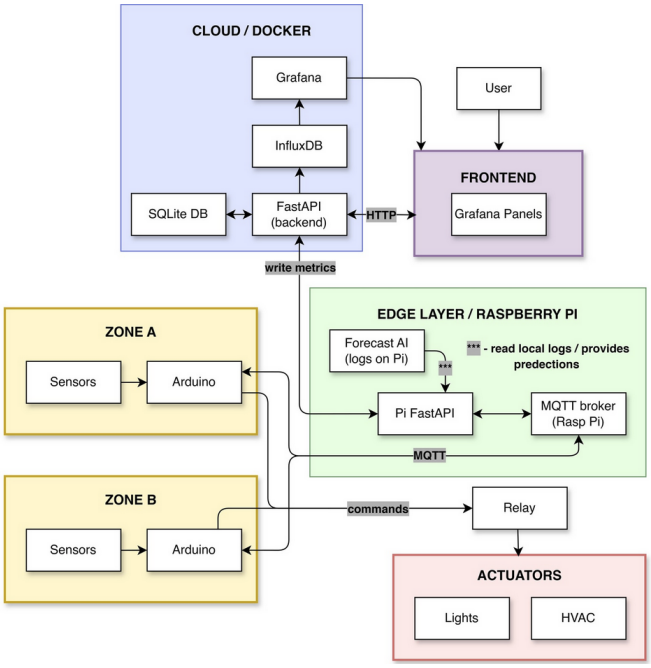

Figure 2. Overview of the data and control flows across system layers.

Data Flow. Embedded devices periodically publish telemetry via MQTT to the edge gateway. The gateway aggregates and enriches measurements before forwarding structured data to backend services for persistence and visualization. Time-series data and metadata are stored separately to support both historical analysis and state tracking. During backend unavailability, telemetry is buffered at the edge and synchronized upon recovery.

Control Flow. Control commands originate either from user interactions or edge-level decision logic. Commands propagate from frontend to backend and then to the edge gateway, which dispatches them to devices via MQTT. An acknowledgment-based mechanism tracks commands with a pending status and confirms completion upon execution, enabling reliable actuation and safe retry under transient failures.

This explicit separation improves robustness, fault isolation, and scalability.

C. Edge Gateway and Reliability Mechanisms

The edge gateway plays a central role in coordinating communication, enforcing local autonomy, and supporting resilience under partial failures. By hosting critical coordination and control logic at the edge, the system reduces dependence on continuous backend availability and avoids single points of failure. At runtime, the gateway hosts the MQTT broker and executes services for device management, data aggregation, buffering, and command dispatch. Sensing and actuation continue locally during backend disruptions, with telemetry and logs stored at the edge until synchronization is restored.

Reliability is further enhanced through acknowledgment-based command handling. Control commands are applied immediately at the edge and confirmed back to the backend, ensuring that delayed or missed transmissions can be retried without duplication. Complementary mechanisms include telemetry buffering during disconnections, periodic resynchronization after recovery, and basic reconnection strategies at the device level. Collectively, these mechanisms enable continued operation under common failure scenarios while preserving consistency and recoverability.

D. Intelligent Decision Support at the Edge

The system integrates lightweight, edge-based decision support mechanisms to enable anticipatory control without imposing significant computational overhead.

Weather forecast data retrieved from the OpenWeatherMap API [16] are used to estimate short-term temperature trends, including direction and magnitude. These trends inform pre-conditioning decisions, mitigating thermal inertia effects and improving comfort while reducing peak energy demand.

Occupancy forecasting is formulated as a probabilistic classification problem, producing continuous occupancy likelihood estimates rather than binary labels. Due to the absence of real occupancy traces, models are trained using synthetically generated behavioral data that capture temporal patterns and stochastic variability. This approach supports reproducibility and privacy-preserving evaluation. At runtime, the trained model is used for inference at the edge, generating near-future occupancy probabilities that are integrated into control decisions.

Control logic follows a hybrid strategy combining predictive inputs with deterministic, rule-based constraints. Thermal comfort is managed through context-dependent temperature bands (e.g., occupied, unoccupied, peak tariff periods), while pre-conditioning decisions are based on predicted occupancy probability, deviation from comfort targets, and external temperature trends. Stability mechanisms such as minimum action intervals are employed to prevent oscillatory behavior and reduce actuator wear.

IV. Evaluation and Results

A. Implementation Scope and Experimental Setup

The system was implemented as a proof of concept using lightweight, widely adopted technologies suitable for IoT/CPS prototyping. ESP32 microcontrollers [17] were used for sensing and actuation, while a Raspberry Pi served as the edge gateway. Communication relied on MQTT, with edge and backend services implemented using FastAPI [18]. System configuration and state metadata were persisted in SQLite [19], while time-series telemetry was stored in InfluxDB [20] and visualized through Grafana dashboards [21]. A web-based frontend implemented in React [22] provided supervision and interaction. Lightweight decision support was implemented using a logistic regression model trained on synthetically generated occupancy data.

The evaluation focused on validating architectural decisions and runtime behavior rather than production readiness or quantitative energy savings. The system was deployed locally, without reliance on external cloud infrastructure, and supports a single edge gateway managing multiple logical zones. While the architecture is designed to scale to multi-zone and multi-gateway deployments, these aspects were not fully exercised in the current implementation. Figure 3 illustrates the implemented supervisory interface.

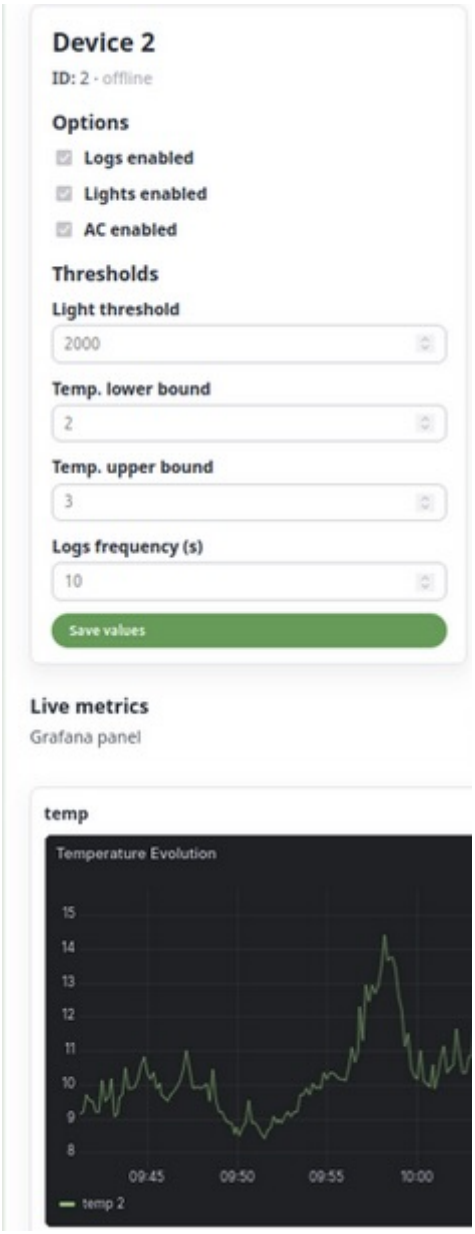

Figure 3. Frontend implementation (device page).

B. Validation Methodology and Scenarios

Evaluation was structured around end-to-end validation of the system's data and control paths under representative operating conditions. The data pipeline—from MQTT ingestion

at the edge to persistence in InfluxDB and visualization in Grafana—was validated using both real sensor data and mocked telemetry streams. Dashboard correctness was verified across different time ranges and filtering conditions, and displayed values were cross-validated against stored records.

Control paths (frontend → backend → edge → MQTT → devices) were exercised to confirm correct propagation of commands and state updates. Command handling relied on acknowledgment-based mechanisms between backend and edge components, enabling verification of reliable actuation and safe retry behavior under transient failures. All services were deployed locally using containerized environments [23], enabling functional validation while abstracting network latency and remote availability effects.

Two complementary validation scenarios were considered. First, a physical edge deployment was evaluated using one Raspberry Pi and two ESP32-based nodes, each representing an independent zone. Real-time telemetry (temperature, luminosity, and binary state variables) was published via MQTT, and control actions were issued from the frontend. In the absence of physical HVAC or lighting equipment, actuation was emulated using onboard indicators, allowing verification of message routing, state transitions, and bidirectional communication without hardware-specific dependencies.

Second, a mixed-data scenario was used to assess scalability and observability under higher logical load. Multiple logical devices were emulated using mocked telemetry to validate MQTT topic handling, backend routing, and dashboard filtering. This scenario enabled evaluation of persistence, visualization, and system behavior independently of physical device availability.

Edge-level decision logic integrating weather forecast data was evaluated by inspecting inference outputs and control decisions generated on the gateway. Due to abstracted actuation and the coexistence of manual and automated control paths, this component was assessed primarily in terms of computational correctness and system integration rather than observable physical effects.

C. Resilience Observations and Results

Table II summarizes the main resilience-related observations obtained during evaluation. The results confirm that the proposed architecture supports reliable telemetry ingestion, persistence, and visualization, maintaining consistent supervisory views once time-series data is available.

From a resilience perspective, the system demonstrated effective fault isolation across device, edge, and backend layers. Backend restarts did not affect previously persisted data, and buffered telemetry was successfully forwarded after recovery. Similarly, edge restarts resulted in brief service interruptions, after which devices automatically reconnected and normal data and control flows were restored. These behaviors validate the edge-centric design objective of preserving local autonomy under partial failures.

Observed data loss was confined to the device–edge boundary during edge downtime, reflecting an intentional design trade-off favoring lightweight embedded devices without persistent buffering. Importantly, such losses did not propagate upstream, preserving system stability and observability. Embedded dashboards consistently reflected the recovered system state following disruptions, reinforcing the suitability of the platform for supervisory monitoring.

TABLE II. OBSERVED RESILIENCE-RELATED BEHAVIORS DURING EVALUATION.

Metric / Scenario	**Observed Behavior**
Physical devices tested	2 ESP32 + 1 Raspberry Pi
Mocked devices tested	Up to 10 logical devices
Backend restart recovery time	∼2 seconds
Edge restart recovery time	Few seconds (boot + services)
Data loss (ESP32 → edge)	During edge downtime
Data loss (edge → cloud)	Not observed; buffered data forwarded
Device reconnection	Automatic, without manual intervention
Command acknowledgment	Implemented cloud-edge only

Overall, the evaluation demonstrates that the proposed architecture achieves its primary objectives of modularity, observability, and resilience in a multi-zone, edge-centric setting, while highlighting device-level buffering and acknowledgment mechanisms as natural directions for future enhancement.

V. CONCLUSION AND FUTURE WORK

This work presented a modular, edge-centric IoT/CPS platform for multi-zone environmental monitoring and control, designed to preserve local autonomy while enabling scalable supervision and centralized observability. By combining distributed sensing with edge-level coordination and backend persistence, the proposed architecture supports reliable real-time and historical monitoring, explicit separation between data and control flows, and resilient operation under partial connectivity failures. The evaluation confirmed that the system maintains consistent behavior across device, edge, and backend layers, validating the architectural choices with respect to modularity, observability, and fault isolation.

While the proof-of-concept demonstrates the feasibility of the proposed approach, several limitations remain. Visualization relies on embedded dashboard components, which introduce authentication and browser policy constraints that may complicate operational deployment. Data quality is dependent on sensor calibration and timestamp consistency, and the current implementation does not include automated validation or anomaly detection mechanisms. In addition, although predictive components based on occupancy and weather context are integrated at the edge, their influence on control policies remains limited and has not yet been validated using real occupancy data.

Future work will focus on extending resilience through device-level buffering and backfilling mechanisms, strengthening authentication and access control across backend services and visualization components, and introducing automated monitoring and alerting. The decision-support layer will be further integrated into control logic and evaluated with real-world occupancy data. Finally, enhanced device management, calibration workflows, and deployment automation (e.g., continuous integration and continuous deployment (CI/CD) and container orchestration) will be explored to improve maintainability and support larger-scale deployments.

ACKNOWLEDGMENT

This work is financially supported by national funds through the FCT/MCTES (PIDDAC), under the Associate Laboratory Advanced Production and Intelligent Systems – ARISE LA/P/0112/2020 (DOI 10.54499/LA/P/0112/2020) and the Base Funding (UIDB/00147/2020) and Programmatic Funding (UIDP/00147/2020) of the R&D Unit Center for Systems and Technologies – SYSTEC.

REFERENCES

[1] S. A. Aghili et al., "Artificial Intelligence Approaches to Energy Management in HVAC Systems: A Systematic Review," *Buildings*, vol. 15, no. 7, Mar. 2025, ISSN: 2075-5309. DOI: 10.3390/buildings15071008. Accessed: Jan. 13, 2026.

[2] M. K. Pasupuleti, "Model Predictive Control for Smart HVAC Systems in Green Buildings," *International Journal of Academic and Industrial Research Innovations (IJAIRI)*, vol. 5, pp. 1–11, Jun. 2025. DOI: 10.62311/nesx/rphcrefcs1.

[3] L. Akbulut et al., "A Systematic Review of Building Energy Management Systems (BEMSs): Sensors, IoT, and AI Integration.," *Energies (19961073)*, vol. 18, no. 24, p. 6522, Dec. 2025, ISSN: 1996-1073. DOI: 10.3390/en18246522. Accessed: Jan. 13, 2026.

[4] M. Poyyamozhi et al., "IoT—A Promising Solution to Energy Management in Smart Buildings: A Systematic Review, Applications, Barriers, and Future Scope," *Buildings*, vol. 14, no. 11, Oct. 2024, ISSN: 2075-5309. DOI: 10.3390/buildings14113446. Accessed: Jan. 13, 2026.

[5] I. Khan, O. Zedadra, A. Guerrieri, and G. Spezzano, "Occupancy Prediction in IoT-Enabled Smart Buildings: Technologies, Methods, and Future Directions," *Sensors (Basel, Switzerland)*, vol. 24, no. 11, p. 3276, May 2024, ISSN: 1424-8220. DOI: 10.3390/s24113276. Accessed: Jan. 13, 2026.

[6] S. Ahmad, M. Shafiullah, C. B. Ahmed, and M. Alowaifeer, "A Review of Microgrid Energy Management and Control Strategies," *IEEE Access*, vol. 11, pp. 21 728–21 757, Feb. 2023. DOI: 10.1109/ACCESS.2023.3248511.

[7] M. A. Ahmad, M. R. Hao, R. M. T. R. Ismail, and A. N. K. Nasir, "Model-free Wind Farm Control Based on Random Search," in *2016 IEEE International Conference on Automatic Control and Intelligent Systems (I2CACIS)*, Oct. 2016, pp. 131–134. DOI: 10.1109/I2CACIS.2016.7885302.

[8] J. J. Jui, M. A. Ahmad, M. I. Milla, and M. I. M. Rashid, "Optimal Energy Management Strategies for Hybrid Electric Vehicles: A Recent Survey of Machine Learning Approaches," *Journal of Engineering Research*, vol. 12, pp. 454–467, 2024. DOI: 10.1016/j.jer.2024.01.016.

[9] X. Sun, J. Fu, H. Yang, M. Xie, and J. Liu, "An Energy Management Strategy for Plug-in Hybrid Electric Vehicles Based on Deep Learning and Improved Model Predictive Control," *Energy*, vol. 269, 2023. DOI: 10.1016/j.energy.2023.126772.

[10] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016. DOI: 10.1109/JIOT.2016.2579198.

[11] A. Sivanathan, H. Habibi Gharakheili, and V. Sivaraman, "Managing iot cyber-security using programmable telemetry and machine learning," *IEEE Transactions on Network and Service Management*, vol. 17, no. 1, pp. 60–74, 2020. DOI: 10.1109/TNSM.2020.2971213.

[12] A. Ghaemi, Y. Rezgui, I. Petri, T. Beach, and A. Ghoroghi, "Ai and digital twin applications in building energy management: A state-of-the-art review," in *2025 IEEE International Conference on Engineering, Technology, and Innovation (ICE/ITMC)*, 2025, pp. 1–11. DOI: 10.1109/ICE/ITMC65658.2025.11106601.

[13] N. Harth, C. Anagnostopoulos, and D. Pezaros, "Predictive intelligence to the edge: Impact on edge analytics," *Evolving Systems*, vol. 9, no. 2, pp. 95–118, Jun. 2018, ISSN: 1868-6486. DOI: 10.1007/s12530-017-9190-z. [Online]. Available: https://doi.org/10.1007/s12530-017-9190-z.

[14] A. Banks and R. Gupta, "MQTT Version 3.1.1," 2014, Accessed: Jan. 19, 2026. [Online]. Available: https://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html.

[15] S. E. Mathe, H. K. Kondaveeti, S. Vappangi, S. D. Vanambathina, and N. K. Kumaravelu, "A comprehensive review on applications of raspberry pi," *Computer Science Review*, vol. 52, p. 100 636, 2024.

[16] OpenWeather Ltd., "Weather API - OpenWeatherMap," 2025, Accessed: Jan. 19, 2026. [Online]. Available: https://openweathermap.org/api.

[17] N. Cameron, "Esp32 microcontroller," in *ESP32 Formats and Communication: Application of Communication Protocols with ESP32 Microcontroller*. Berkeley, CA: Apress, 2023, pp. 1–54, ISBN: 978-1-4842-9376-8. DOI: 10.1007/978-1-4842-9376-8_1. [Online]. Available: https://doi.org/10.1007/978-1-4842-9376-8_1.

[18] M. Lathkar, "Getting started with fastapi," in *High-Performance Web Apps with FastAPI: The Asynchronous Web Framework Based on Modern Python*. Berkeley, CA: Apress, 2023, pp. 29–64, ISBN: 978-1-4842-9178-8. DOI: 10.1007/978-1-4842-9178-8_2. [Online]. Available: https://doi.org/10.1007/978-1-4842-9178-8_2.

[19] J. Kreibich, *Using SQLite*. " O'Reilly Media, Inc.", 2010.

[20] InfluxData, "InfluxDB Documentation," 2025, Accessed: Jan. 19, 2026. [Online]. Available: https://docs.influxdata.com/influxdb/.

[21] S. Kirešová et al., "Grafana as a visualization tool for measurements," in *2023 IEEE 5th International Conference on Modern Electrical and Energy System (MEES)*, IEEE, 2023, pp. 1–5.

[22] Meta Platforms, Inc., "React: A JavaScript Library for Building User Interfaces," 2025, Accessed: Jan. 19, 2026. [Online]. Available: https://reactjs.org/.

[23] I. Docker, "Docker Documentation," 2023, Accessed: Jan. 19, 2026. [Online]. Available: https://docs.docker.com/.

Accelerating the Adoption of Asset Administration Shells through AI Agents

Camilo Velázquez-Rodríguez*, Steven Kauffmann*, Tom Pauwaert*, Stijn Huysentruyt†, Axl Van Alboom†
*Flanders Make, Leuven, Belgium; email: {camilo.velazquez, steven.kauffmann, tom.pauwaert}@flandersmake.be
†Universiteit Gent, Ghent, Belgium; email: {stijn.huysentruyt, axl.vanalboom}@ugent.be

***Abstract*—The adoption of Industry 4.0 practices has grown worldwide, driven by the need for standardised, transparent, and interoperable data exchange. The Asset Administration Shell (AAS) provides a standardised template for asset information, enabling companies to share and integrate data across systems. However, AAS creation from existing datasheets remains a manual and time-consuming process, hindering large-scale adoption. In this paper, we propose an Artificial Intelligence (AI) agent-based approach that automates the transformation of Portable Document Format (PDF) datasheets into AAS instances, which are then serialised into AAS files. The agents extract, structure, and validate asset information against the Industrial Digital Twin Association (IDTA) guidelines to ensure compliance with industry standards. We demonstrate the approach in a use case scenario, illustrating its potential to streamline the creation of AAS files and facilitate their adoption in a manufacturing environment.**

Keywords-Intelligent agents; Datasheet; AAS; Industry 4.0.

I. Introduction

The Industry 4.0 paradigm has revolutionised manufacturing processes by integrating advanced technologies, such as Internet of Things (IoT), Cyber-Physical Systems, and Artificial Intelligence [1]. This transformation has led to the emergence of smart factories. Nowadays, interconnected systems can communicate and make autonomous decisions to optimise production efficiency and flexibility [2]. A key enabler of this paradigm is the Asset Administration Shell (AAS), which serves as a digital representation of physical assets in the manufacturing environment [3].

The AAS provides multiple benefits. These include a standardised framework for managing and exchanging information about assets, facilitating interoperability among diverse systems and devices. By encapsulating all relevant data, functionalities, and services related to an asset, the AAS enables seamless integration into Industry 4.0 ecosystems. Additionally, the AAS supports advanced functionalities, such as predictive maintenance, real-time monitoring, and data analytics, which are crucial for optimising manufacturing processes [4]. Despite its potential, the adoption of AAS in industrial settings faces several challenges.

One of the main challenges is the complexity of the AAS implementation. New submodel templates released by the IDTA [5] and other standardisation bodies require significant effort to understand and apply correctly. In this regard, companies often struggle with the initial setup, configuration, and customisation of AAS for their specific use cases. Even though tools and frameworks exist to facilitate AAS creation, they often demand a steep learning curve and technical expertise. Therefore, manual processes, often chosen for AAS deployment, can be time-consuming and error-prone, hindering widespread adoption.

The lack of awareness among stakeholders is another barrier to AAS adoption. Many organisations are still unfamiliar with the benefits and functionalities of AAS, leading to resistance in embracing this relatively new technology. This is accompanied by the need for skilled personnel to manage and maintain AAS-based systems [6]. Many of the companies, although interested in Industry 4.0, lack the necessary expertise to effectively implement submodel templates and further utilise AAS in their operations. For example, Small and Medium-sized Enterprises (SMEs) may find it particularly challenging to allocate resources for training and development in this area [7]. In addition, many industries still have plenty of machinery and equipment not Industry 4.0-ready, making the integration of AAS even more complex [8].

To address some of the mentioned challenges, there is a growing interest in leveraging automated solutions to streamline the adoption process. In this context, Artificial Intelligence (AI) agents have emerged as promising tools to facilitate the deployment and management of AAS in manufacturing environments. AI agents can autonomously perform tasks, such as data collection, analysis, and decision-making, thereby reducing the burden on human operators [9]. By integrating AI agents with AAS, it is possible to enhance the efficiency and effectiveness of asset management processes, ultimately accelerating the adoption of AAS in Industry 4.0 settings.

This paper explores the potential of AI agents to support the adoption of AAS in Industry 4.0 environments. More specifically, we investigate how AI agents can assist in the automatic creation of AAS instances from PDF datasheets. Generated AAS instances can then be later integrated into a component library to speed up the commissioning of physical assets. PDF datasheets are commonly used to document technical specifications of assets. Oftentimes, they contain valuable information that can be utilised to populate AAS submodel templates. However, manually extracting and structuring this information can be labour-intensive and prone to errors. Our paper makes the following contributions:

- Automatic generation of AAS from PDF datasheets using AI agents and knowledge graphs.
- A case study demonstrating the application of our approach in generating AAS instances for an industrial robot arm, highlighting the practical implications and benefits.
- The discussion of challenges faced during implementation and future directions in this area.

While our approach shows promise, certain limitations are acknowledged: generated AAS quality depends on input data

completeness, manual intervention remains necessary for eClass classification, and organisations with heavily customised legacy systems may face adoption barriers.

The remainder of the paper is organised as follows: In Section II, we review the related work on AAS adoption and AI agents in Industry 4.0. Section III provides background information on AAS, AI agents, knowledge graphs, etc. Section IV presents our approach for automatic AAS generation based on the automated information extraction from PDF datasheets. Section V describes the case study, experimental setup, and faced challenges during implementation. Section VI discusses the results, their implications, and potential limitations of our approach. Finally, Section VII concludes the paper and outlines future research directions.

II. Related Work

A. Traditional Agents for AAS Adoption

The interaction of agents with assets in industry has been broadly studied over the past few years. Precisely, the definition of Multi-Agent Systems (MAS) by Wooldridge [10] laid the foundation for agents in industry. As per [10], a MAS is a "*society of intelligent, cooperative, proactive, and autonomous entities*" (i.e., agents). These agents rely on symbolic reasoning, explicit knowledge representations (e.g., logic-based systems), and planning algorithms to make decisions. They are designed for specific tasks, such as negotiation, coordination, or resource allocation in MAS. Their intelligence is rule-based or logic-based, with limited adaptability outside predefined domains. Communication is structured, often relying on predefined protocols (e.g., Foundation for Intelligent Physical Agents - Agent Communication Language (FIPA-ACL)) and ontologies for inter-agent interactions. Therefore, we will refer to these as *traditional agents* for the remainder of the paper.

Based on the previous MAS definition, a group of works applied it to facilitate the adoption of AAS in industry [6, 11–17]. They leverage MAS for the implementation, digitisation, and adoption of AAS. These works describe the use of different traditional agents in conjunction with AAS to enhance standardisation and communication between physical assets and their digital twin counterparts. Many assume that AAS files are already filled with the required information in a specific submodel and leverage their approach on this assumption. In reality, the instantiation of an AAS submodel might take several hours (depending on the asset and submodel complexities), making it a labour intensive process.

Our approach makes use of Large Language Model (LLM)-based models, in contrast to traditional agents, such as in the previous works. Modern agents interact with a transformer-based neural network trained on massive datasets, also known as an LLM. This can be general-purpose or fine-tuned for a specific task. LLM-based agents can operate autonomously but often rely on human prompts or predefined workflows to initiate actions. These agents excel in natural language communication, interacting through conversational interfaces without requiring formal protocols. Unlike Wooldridge's agents [10], LLM-based agents demonstrate emergent general-purpose reasoning, enabling them to handle a wide range of tasks without domain-specific programming.

B. LLMs (Agents) for AAS Adoption

The use of LLMs to facilitate the adoption of AAS in industry is a novel approach. Prompts, as well as agents powered by LLMs, have been used to generate AAS files from unstructured data [18]. Xia et al., [18] propose a framework that leverages LLMs to interpret text-based data and generate AAS-compliant files. They based their approach on the capabilities of LLMs to match semantically similar concepts of syntactically different terminologies. This can be useful when dealing with unstructured data from various sources, as LLMs can understand context and meaning beyond rigid schemas.

Although we also use LLMs to syntactically map different terminologies, our approach goes beyond that of Xia et al., [18] in several ways. First, our approach automatically extracts text information from PDF datasheets, instead of relying on text-based data as input. Second, our process involves a knowledge database creation and a querying step to prevent hallucinations from the LLM. Third, we also map eClass IDs to the extracted data to ensure that the generated AAS files are compliant with industry standards. We are not aware of any other work that uses LLMs to facilitate the adoption of AAS in industry apart from the work of Xia et al., [18].

III. Background

A. AI Approaches and Techniques

1) Vector databases: Vector databases are specialised systems for storing and indexing high-dimensional embeddings, enabling efficient approximate similarity search (e.g., nearest neighbour) over data, such as text, images or audio [19].

2) Large Language Models: LLMs are transformer-based neural networks pre-trained on massive corpora, which scale in parameter count, data, and compute, exhibiting emergent capabilities without task-specific tuning. While their applications span across many domains, main limitations include data bias, hallucination, high inference/training costs, and challenges in evaluation and governance [20, 21].

3) AI Agents: AI agents are systems that not only generate responses but perceive, plan, decide and act (often using tools, memory, and environment interactions) to achieve goals with some degree of autonomy [22].

4) Knowledge graphs: Knowledge graphs are directed, labelled, multi-relational graphs in which nodes represent entities and edges represent semantic relationships. These graphs are often enriched with an ontology and are used to integrate, query, and reason over structured and semi-structured data. They provide explainability and support complex relational queries and inference. Some of their challenges include scalability, schema alignment, entity/edge extraction, and maintaining/updating large, heterogeneous graphs [23–25].

B. Asset Administration Shell

The Asset Administration Shell (AAS) is a standardised digital representation of industrial assets throughout their

lifecycle. An AAS file can contain one or more AAS instances, where each instance represents a digital twin of a specific physical or logical asset. AAS files are organised via modular submodels, which capture technical, operational, semantic, and communication aspects enabling interoperability across heterogeneous systems and stakeholders [17, 18, 26–29]. The Industrial Digital Twin Association (IDTA) [30] publishes, maintains and distributes AAS submodel template files publicly. Currently, there are various submodels available for specific target domains (e.g., digital nameplate for industrial equipment, functional safety, technical data, etc.) [5], while others are intended to be released in the near future (e.g., carbon footprint, maintenance instructions, etc.).

Each AAS submodel defines specific fields to be filled in, which correspond to the characteristics of the asset being digitalised. Fields also have types (e.g., date, integer, string) depending on the data to be stored. An asset might have more characteristics than those specified by the submodel template. To handle these differences, some submodel templates (e.g., digital nameplate) can be extended with custom fields. These fields are of string type to be completed with generic information, and not specific to the submodel template.

IV. Approach

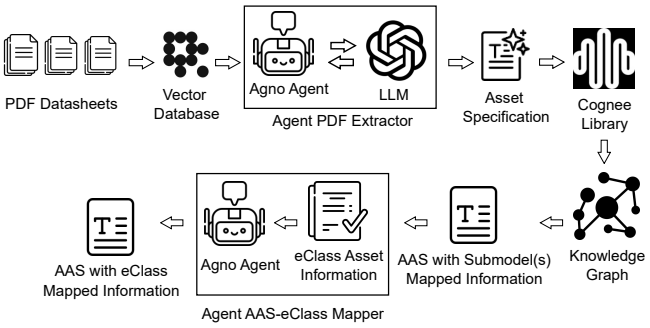

Figure 1. Approach to automatically move from PDF datasheets to AAS files.

Figure 1 shows the implemented approach. Datasheets in the form of PDF files represent the input to the approach. These are distributed as part of the documentation for different assets (e.g., robot arm, industrial pump, etc.). Additionally, they describe the main asset characteristics to be extracted later on.

A vector database is used to process the documents by extracting text and assigning a vector to each extracted token. Extracted words and their corresponding vectors will later be used to map characteristics that might be syntactically different but semantically similar. Currently, multiple vector database tools exist providing help to automatic information extraction and vector matching. We selected LanceDB [31] due to its integration with the agent framework we use. Our approach is not limited to this specific vector database, and other tools can be used as well.

Once the documents are stored, processed and vectorised, an AI agent will interact with this information and a pre-established Large Language Model (LLM). The selection of LLMs also depends on many factors related to the task at hand. For example, some models are better suited for code generation, while others excel at natural language understanding. We selected the OpenAI models [32] to interact with the agent, as they have demonstrated state-of-the-art performances for many different tasks. Specifically, we used the *Generative Pre-trained Transformer (GPT)-4.1* model to extract assets' characteristics from the documents.

The agent framework Agno [33] provides the necessary tools for our approach. Agno's integration with the LanceDB vector database and other services and models (e.g., OpenAI) allowed us to focus on the agents' configuration, instead of implementation details. Previously processed datasheet documents serve as a knowledge base, whilst *GPT-4.1* is used as a language model. We also make use of the internal reasoning tool of the Agno framework in the agent configuration. This tool instructs the agent to solve the task at hand as a series of logical steps and to trace back the reasoning process. Reasoning includes analysis and revision of the subtasks' performance. The configured prompt instructs the agent to extract technical information about one specific asset at a time.

The result of running the agent, given the previous configuration specifications, is a file containing extracted technical information about the specified asset. As the information of different assets may be heterogeneous and diverse, the language model might tend to hallucinate. This phenomenon has been identified as a drawback for these models [34, 35]. Therefore, we opted to address this issue by relying solely on the extracted information from the datasheets. We assume that the datasheets contain the necessary information to populate the AAS templates.

The Cognee library [36] provides an approach to solve the hallucination related to LLMs. The library enforces agents to rely on the information from a knowledge graph which is created from a set of pre-defined sources (e.g., PDF datasheets). This knowledge graph is generated in the form of a Neo4J database [37] where nodes represent automatically extracted entities and edges their corresponding relationships. Once the knowledge graph is created (cf., Figure 1), responses to queries are limited to only the scope of the information within the graph. Thus, a hallucination is not possible by construction, as the agent cannot invent information not present in the graph.

Figure 1 also shows that after the Neo4J knowledge graph creation, queries are requested to the database. Cognee enables natural language queries, which are transformed into the Cypher graph query language. Our approach further leverages this feature by mapping extracted asset characteristics to the fields in an AAS template file. Asset characteristics and AAS fields may differ syntactically while conveying similar semantics.

An instantiation of the specified AAS templates is the outcome of the previous mapping step. The AAS template files are based on the IDTA guidelines. More specifically, in this work, we used the AAS template for Digital Nameplate, Technical Data, and Functional Safety. The three instantiations are then combined into one AAS file per asset. These can be found in the IDTA AAS Template Repository on GitHub [38].

The inclusion of additional templates is part of future work.

A final mapping step is performed to enrich the AAS instantiation with additional information. eClass [39] IDs are standard identifiers for asset characteristics that might not be part of the default information in templates (i.e., custom-added fields). Therefore, we mapped such custom-added characteristics to their corresponding eClass IDs. For this purpose, we used another AI agent configured similarly to the previous one.

The outcome of our approach is an AAS file per asset, which is automatically generated from the corresponding datasheets. These AAS files can then be used in Industry 4.0 applications to represent assets in a standardised manner. They contain the main technical information as well as additional eClass IDs for custom-added fields.

V. Use Case

We present a use case that illustrates the practical application of our proposed method. The use case focuses on a real-world scenario where PDF datasheets are processed by our approach to generate AAS files and be used in an Industry 4.0 context. We describe the specific steps taken in this use case, the encountered challenges, and the achieved results.

A. Scenario Description and Resolution

1) Description: The use case involves a manufacturing company that aims to digitise its asset information by converting existing PDF datasheets into AAS files. The company has a large repository of datasheets for various components and equipment used in its production processes. We will focus on an asset type to demonstrate the effectiveness of our approach.

We focus our approach on robotic arms used in assembly lines. We consider a 6-axis articulated robotic arm commonly used for pick-and-place operations. More specifically, we use the datasheets provided by the manufacturer Stäubli for their TX2-90XL model [40]. This robotic arm is equipped with various sensors and actuators, and its datasheets contain detailed technical specifications, performance data, and operational guidelines. The datasheets provide information for three variants of the TX2-90 model (e.g., the TX2-90, TX2-90L, and TX2-90XL), each with different payload capacities and reach. Precisely, we selected these datasheets as they have a complex structure, including multiple tables, images, and technical diagrams, which pose challenges for automated data extraction.

Characteristics

	TX2-90	TX2-90L	TX2-90XL
Load capacity	14 kg	12 kg	7 kg
Reach (between axis 1 and 6)	1000 mm	1200 mm	1450 mm
Number of degrees of freedom	6	6	6
Repeatability - ISO 9283	± 0.02 mm	± 0.02 mm	± 0.02 mm
Weight	114 kg	117 kg	119 kg
UL certification	✓	✓	✓

Figure 2. Screenshot of the Stäubli TX2-90XL robotic arm datasheet.

Figure 2 shows a table extracted from the datasheet, which lists the technical specifications of the TX2-90 robotic arm variants. Since this use case focuses on the robot arm TX2-90XL, our approach needs to correctly identify and extract the relevant information from the table in Figure 2. Additionally, other relevant information that might be in multiple documents also needs to be extracted. This can include the robot's degrees of freedom, maximum speed, repeatability, and operating conditions.

2) Resolution: As previously described in Section IV, our approach involves several steps to generate AAS files from PDF datasheets. First, the PDF datasheets are processed using a vector database to extract and vectorize the textual information. Next, a smart agent configured with the Agno framework interacts with the vector database and a pre-established LLM (GPT-4.1) to extract the relevant technical specifications of the TX2-90XL robotic arm. The agent is prompted to focus on the specific asset and extract information, such as payload capacity, reach, degrees of freedom, maximum speed, repeatability, and operating conditions. A manual review of the extracted information is performed to ensure accuracy and completeness. The extracted information is then mapped to the corresponding fields in an AAS template file using the Cognee library. A second mapping step is performed to link the extracted information to the appropriate eClass IDs. Such a mapping is performed using the eClass database to ensure that each field in the AAS file is correctly identified and classified. The mapped information is written into the AAS templates, resulting in a complete AAS file for the TX2-90XL robotic arm. Finally, the generated AAS file is validated against the AAS schema to ensure compliance with Industry 4.0 standards.

B. Challenges

During the execution of the use case, several challenges were encountered. One of the main challenges was dealing with the heterogeneity of the datasheet formats. The datasheets contained various layouts, tables, and images, which made it difficult for the agent to extract the relevant information. To address this challenge, we leveraged the capabilities of the vector database to effectively index and retrieve the textual information. Additionally, the heterogeneity of the language used in the datasheets posed a challenge in case this process would be performed manually. Therefore, we relied on the LLM's ability to understand and process technical language and similar terminologies.

Each of the mappings (to the AAS template and to the eClass IDs) also presented challenges. The AAS template needed to be analysed to accommodate the specific attributes of the TX2-90XL robotic arm. The mapping to eClass IDs required access to an up-to-date eClass database and understanding of the classification system. To overcome these challenges, we ensured that the AAS template was flexible enough to allow for assets that may not appear in the standard. Also, we manually focus on the robotic arm domain to ensure that the correct eClass IDs were assigned. In this regard, future work could explore automating the eClass mapping process further.

Figure 3. Screenshot of the Neo4J knowledge graph created from the available datasheets.

Lastly, another challenge was ensuring the accuracy of the extracted information. As we relied on official datasheets, and also based on a knowledge base with the extracted information, we mitigated the risk of inaccuracies. Despite this, we still performed a manual review of the extracted information to ensure its correctness.

C. Results

The results of the use case proved the effectiveness of our approach in generating AAS files from PDF datasheets. One of the first results we obtained was the knowledge graph in the Neo4J database. This graph served as a valuable resource for querying and retrieving information about the robotic arm.

Figure 3 shows a screenshot of the knowledge graph in Neo4J, illustrating the relationships between different technical specifications of the TX2-90XL robotic arm. Queries can be in natural language which are then translated into Cypher queries to retrieve specific information. For example, a query such as *"What is the maximum payload of the TX2-90XL robotic arm?"* would return this specific information from the graph.

The generated AAS instance for the TX2-90XL robotic arm contained the technical specifications extracted from the datasheets. Figure 4 shows a screenshot of the generated AAS instance in the visualisation tool AASX Package Explorer [41].

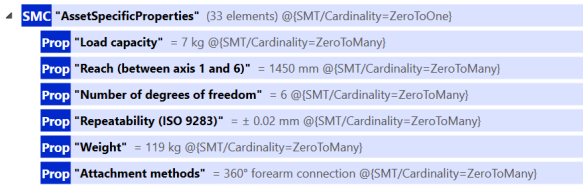

Figure 4. Screenshot of the generated AAS instance for the Stäubli TX2-90XL robotic arm in the AASX Package Explorer.

Our approach also includes the datasheet PDFs as part of the AAS file, ensuring that the original documentation is available for reference. A visualisation tool supporting PDF rendering can display the datasheets directly from the generated file.

VI. Discussion

A. Implications

The proposed approach has several implications for the adoption of AAS in industrial settings. By leveraging AI agents to automate the creation of AAS instances and files, we can reduce the time and effort required for their implementation. This can lead to faster integration of AAS into existing systems, promoting interoperability and standardisation across different platforms. Previously considered tedious and error-prone tasks can now be streamlined, allowing organisations to focus on leveraging the benefits of AAS.

The integration of industries into the Industry 4.0 paradigm could be further accelerated by this approach. With the ability to quickly generate AAS instances and files, industries can more readily adopt digital twins and other Industry 4.0 technologies, leading to improved efficiency in manufacturing processes. The use of AI agents also opens up new possibilities for customisation and scalability, as AAS can be tailored to specific use cases and easily adapted to changing requirements.

B. Limitations

Despite the promising results, there are several limitations to our approach that need to be addressed in future work. First, the quality of the generated AAS instances is heavily dependent on the input data and the capabilities of the AI agents. Inaccurate or incomplete data can lead to suboptimal AAS instances, which may not fully meet the requirements of the intended application. Second, the current implementation may not fully capture the complexity of certain industrial scenarios. While AI agents are capable of generating AAS files for a wide range of use cases, there may exist specific cases that require more sophisticated modelling and representation. Finally, there is still a need for manual intervention in the eClass classification process, which may introduce delays.

Beyond mentioned technical limitations, practical challenges exist regarding adoption and use by existing organisations. Many industrial environments have heavily customised legacy systems and established workflows deeply integrated into their operations. Introducing AAS-based approaches may require significant organisational change management. Furthermore, organisations with investments in proprietary data management systems may face resistance to adopting AAS formats, particularly if migration costs are perceived as prohibitive.

VII. Conclusion and Future Work

In this paper, we presented an approach leveraging AI agents to facilitate AAS creation from PDF datasheets in industrial settings. We utilize the Cognee library to create a knowledge graph from datasheet information, which AI agents use to generate and populate the AAS structure. A use case involving an industrial robot arm demonstrates the potential of this approach to automate the AAS creation process. Future work will focus on extensive evaluations in real-world industrial environments and integration of generated AAS files into component libraries for flexible manufacturing systems. This enables dynamic asset assignment, reducing manual programming effort and increasing production flexibility.

References

[1] H. Lasi, P. Fettke, H.-G. Kemper, T. Feld and M. Hoffmann, 'Industry 4.0', *Business & information systems engineering*, vol. 6, no. 4, pp. 239–242, 2014.

[2] Y. Lu, 'Industry 4.0: A survey on technologies, applications and open research issues', *Journal of industrial information integration*, vol. 6, pp. 1–10, 2017.

[3] E. Tantik and R. Anderl, 'Integrated data model and structure for the Asset Administration Shell in Industrie 4.0', *Procedia Cirp*, vol. 60, pp. 86–91, 2017.

[4] C. Wagner et al., 'The role of the Industry 4.0 Asset Administration Shell and the Digital Twin during the life cycle of a plant', in *2017 22nd IEEE international conference on emerging technologies and factory automation (ETFA)*, IEEE, 2017, pp. 1–8.

[5] IDTA, 'IDTA Submodel Templates', Accessed: 2025-01-06, 2024. [Online]. Available: https://industrialdigitaltwin.org/en/content-hub/submodels.

[6] A. Sidorenko et al., 'The MAS4AI framework for human-centered agile and smart manufacturing', *Frontiers in Artificial Intelligence*, vol. 6, p. 1 241 522, 2023.

[7] F. Faiz, V. Le and E. K. Masli, 'Determinants of digital technology adoption in innovative SMEs', *Journal of Innovation & Knowledge*, vol. 9, no. 4, p. 100 610, 2024.

[8] A. Alqoud, D. Schaefer and J. Milisavljevic-Syed, 'Industry 4.0: a systematic review of legacy manufacturing system digital retrofitting', 2022.

[9] B. Wang, X. Li, T. Freiheit and B. I. Epureanu, 'Learning and intelligence in human-cyber-physical systems: Framework and perspective', in *2020 Second International Conference on Transdisciplinary AI (TransAI)*, IEEE, 2020, pp. 142–145.

[10] M. Wooldridge, *An Introduction to Multiagent Systems*. John wiley & sons, 2009.

[11] L. Sakurada and P. Leitão, 'Multi-Agent Systems to Implement Industry 4.0 Components', in *2020 IEEE conference on industrial cyberphysical systems (ICPS)*, IEEE, vol. 1, 2020, pp. 21–26.

[12] A. López, O. Casquero and M. Marcos, 'Design patterns for the implementation of Industrial Agent-based AASs', in *2021 4th IEEE International Conference on Industrial Cyber-Physical Systems (ICPS)*, IEEE, 2021, pp. 213–218.

[13] B. Vogel-Heuser, F. Ocker and T. Scheuer, 'An approach for leveraging Digital Twins in agent-based production systems', *at-Automatisierungstechnik*, vol. 69, no. 12, pp. 1026–1039, 2021.

[14] S. Jungbluth et al., 'Dynamic Replanning using Multi-Agent Systems and Asset Administration Shells', in *2022 IEEE 27th International Conference on Emerging Technologies and Factory Automation (ETFA)*, IEEE, 2022, pp. 1–8.

[15] L. Sakurada, P. Leitao and F. De La Prieta, 'Engineering a Multi-agent Systems Approach for Realizing Collaborative Asset Administration Shells', in *2022 IEEE International Conference on Industrial Technology (ICIT)*, IEEE, 2022, pp. 1–6.

[16] L. Sakurada, P. Leitao and F. De la Prieta, 'Agent-Based Asset Administration Shell Approach for Digitizing Industrial Assets', *Ifac-Papersonline*, vol. 55, no. 2, pp. 193–198, 2022.

[17] L. Sakurada, F. De La Prieta and P. Leitao, 'Ten Years of Asset Administration Shell: Developments, Research Opportunities, and Adoption Challenges', *IEEE Access*, 2025.

[18] Y. Xia, Z. Xiao, N. Jazdi and M. Weyrich, 'Generation of Asset Administration Shell with Large Language Model Agents: Toward Semantic Interoperability in Digital Twins in the Context of Industry 4.0', *IEEE Access*, vol. 12, pp. 84 863–84 877, 2024.

[19] Y. Han, C. Liu and P. Wang, 'A Comprehensive Survey on Vector Database: Storage and Retrieval Technique, Challenge', *arXiv preprint arXiv:2310.11703*, 2023.

[20] W. X. Zhao et al., 'A Survey of Large Language Models', *arXiv preprint arXiv:2303.18223*, vol. 1, no. 2, 2023.

[21] K. S. Kalyan, 'A survey of GPT-3 family large language models including ChatGPT and GPT-4', *Natural Language Processing Journal*, vol. 6, p. 100 048, 2024.

[22] P. Zhao, Z. Jin and N. Cheng, 'An In-depth Survey of Large Language Model-based Artificial Intelligence Agents', *arXiv preprint arXiv:2309.14365*, 2023.

[23] A. Hogan et al., 'Knowledge Graphs', *ACM Computing Surveys (Csur)*, vol. 54, no. 4, pp. 1–37, 2021.

[24] S. Ji, S. Pan, E. Cambria, P. Marttinen and P. S. Yu, 'A Survey on Knowledge Graphs: Representation, Acquisition, and Applications', *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 2, pp. 494–514, 2021.

[25] C. Peng, F. Xia, M. Naseriparsa and F. Osborne, 'Knowledge Graphs: Opportunities and Challenges', *Artificial intelligence review*, vol. 56, no. 11, pp. 13 071–13 102, 2023.

[26] S. Rongen, N. Nikolova and M. van der Pas, 'Modelling with AAS and RDF in Industry 4.0', *Computers in Industry*, vol. 148, p. 103 910, 2023.

[27] J. Arm et al., 'Automated Design and Integration of Asset Administration Shells in Components of Industry 4.0', *Sensors*, vol. 21, no. 6, p. 2004, 2021.

[28] L. Sakurada, F. De la Prieta and P. Leitao, 'The Role of Multi-Agent Systems in Realizing Asset Administration Shell Type 3', *Future Internet*, vol. 17, no. 7, p. 270, 2025.

[29] A. Alexopoulos, G. Kalogeras, K. Koutras and A. Kalogeras, 'Why Asset Administration Shells: A Survey on Uses and Challenges', *IEEE Access*, 2025.

[30] IDTA, 'Industrial Digital Twin Association', Accessed: 2025-01-06, 2024. [Online]. Available: https://industrialdigitaltwin.org/en/.

[31] LanceDB, 'LanceDB: Developer-friendly, serverless vector database', Accessed: 2025-01-06, 2024. [Online]. Available: https://lancedb.com/.

[32] OpenAI, 'OpenAI Platform Models', Accessed: 2025-01-06, 2024. [Online]. Available: https://platform.openai.com/docs/models.

[33] Agno, 'Agno: AI Agent Framework', Accessed: 2025-01-06, 2024. [Online]. Available: https://docs.agno.com/introduction.

[34] Z. Ji et al., 'Towards Mitigating Hallucination in Large Language Models via Self-Reflection', in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 1827–1843.

[35] A. Martino, M. Iannelli and C. Truong, 'Knowledge Injection to Counter Large Language Model (LLM) Hallucination', in *European Semantic Web Conference*, Springer, 2023, pp. 182–185.

[36] Cognee, 'Cognee: Knowledge Graph Generation', Accessed: 2025-01-06, 2024. [Online]. Available: https://www.cognee.ai/.

[37] Neo4j, Inc., 'Neo4j Graph Database', Accessed: 2025-01-06, 2024. [Online]. Available: https://neo4j.com/.

[38] IDTA, 'IDTA AAS Submodel Templates Repository', Accessed: 2025-01-06, 2024. [Online]. Available: https://github.com/admin-shell-io/submodel-templates.

[39] eCl@ss e.V., 'eCl@ss: International Standard for the Classification and Description of Products and Services', Accessed: 2025-01-06, 2024. [Online]. Available: https://eclass.eu/en/.

[40] Stäubli Robotics, 'Stäubli TX2-90XL Industrial Robot', Accessed: 2025-01-06, 2024. [Online]. Available: https://www.staubli.com/us/en/robotics/products/industrial-robots/tx2-90.html.

[41] Eclipse AASPE, 'AASX Package Explorer', Accessed: 2025-01-06, 2024. [Online]. Available: https://github.com/eclipse-aaspe/package-explorer.

From Data Silos to Intelligent Operations: An AI-Based Approach to Ground Station Incident Investigation

Dimitri Accad and Nieves Salor Moral
Ground Segment & Operations Solutions
Starion Deutschland GmbH
Darmstadt, Germany
e-mail: {d.accad | n.salor}@stariongroup.eu

Andrea Di Luca
Fondazione Bruno Kessler
Trento, Italy
e-mail: adiluca@fbk.eu

***Abstract*—The rapid increase in data generated by modern space systems has introduced significant challenges related to the scale, diversity, and fragmentation of operational information. Operators are increasingly overwhelmed by the daily influx of information, which is often loosely structured or spread across isolated systems. Although the integration of Artificial Intelligence (AI) into space ground segment operations offers promising opportunities, it demands addressing the inherent operational siloed contexts, security concerns, the diverse and evolving information sources, and the high-risk constraints of space missions. This work presents an Incident Investigation AI system that assists operators in analysing, understanding, and resolving ground station incidents threatening space missions. Traditional incident investigation relies on subjective expert judgment, manual log searches, and fragmented document access. This leads to data duplication, inconsistent formats, and knowledge loss. This paper addresses these limitations through dedicated data-processing pipelines that continuously extract, ingest, and harmonise information. The resulting knowledge powers agentic systems that assist operators throughout the incident lifecycle, including similarity identification, correlation analysis, root cause analysis, and automated conclusion generation. These components create an extensible framework that, beyond immediate operational gains, establishes a foundation for future collaborative agentic ecosystems capable of reasoning over complex operational data and supporting more autonomous, data-driven decision-making in space operations.**

Keywords-artificial intelligence; generative AI; agentic systems; space operations; incident investigation; data pipelines; vector databases; root cause analysis; ground stations; ground segment.

I. INTRODUCTION

Space systems generate a large amount of daily data. These systems can sometimes run into issues classified as incidents. These need to be investigated by operational personnel to take the relevant action in a timely manner. To perform this task, operators rely on experience, technical documentation, and reports on previous incident investigations. Accessing these sources is not a trivial task because of the data being scattered across siloed systems, being inconvenient to access, or becoming unavailable due to staff turnover and organizational memory loss. The investigation itself is a daunting manual process due to the overload of information to corroborate.

We designed a solution to support this investigation work by providing a suite of AI tools to empower operators with insights and sources they could hardly access before. All these tools are accessible through a User Interface built in React [1] where they can access past and current investigations, previous incident similarity graphical insights, chat with Engi (i.e., an AI agent with access to all the knowledge gathered over years of previous investigations), or review initial conclusions drafted by an autonomous AI assistant upon investigation creation.

The implementation of these agents presented technical challenges due to context length limitations of Large Language Models (LLMs) [2],[3] and available hardware constraints. Loading all available data into the model was therefore not viable, while external proprietary models were precluded by confidentiality concerns [4].

Recent advances in the AI landscape, such as Retrieval Augmented Generation (RAG) systems [5] and vector databases [6] enabled us to encode our data sources as vectors. Using similarity search methods (keyword, vector, or hybrid), relevant documents can be retrieved at run-time [7]. The adoption of this technology circumvented part of the software and hardware limitations described previously. It also provided opportunities to lay out the similarities generated in an interactive graph users can navigate to discover new insights based on investigations closely related to the one at hand. To improve user adoption, this approach was combined with the existing metadata of the documents, comprising manual correlations created by operators in previous investigations. The result is a knowledge graph visualisation [8] where the information is structured as nodes and relationships as edges between documents.

As a whole, the system is based on well-defined and well-scoped agents providing different features. These agents are powered by state-of-the-art methods to provide quality responses despite the small model size able to fit on the available hardware (i.e., 8 Billion parameter dense models). The designed network of expert agents was tailored to different scenarios with access to distinct but sometimes overlapping data sources. Space operations being a sensitive field where accuracy is paramount, multiple quality control techniques were deployed, such as re-ranking, hallucination detection, and relevance checks, to ensure integrity and transparency in the results generated by our system. These controls are distributed throughout the agentic workflow, from ensuring the relevance of the retrieved documents, to providing safeguards against potential model hallucinations [9] and checking the quality of

 ISBN: 978-1-68558-351-4

the final generated answer.

While RAG-based systems have been explored for industrial troubleshooting in manufacturing [10] or aerospace-specific RAG evaluation (e.g., NASA's VALOR [11]), no deployed system currently addresses the specific challenges of ground station incident investigation, namely the combination of heterogeneous operational data sources, security-constrained on-premise deployment, and the need for multi-expert reasoning over years of accumulated institutional knowledge.

The main contributions of this work are:

- A unified data pipeline architecture that continuously harmonizes heterogeneous operational data sources (logs, reports, communications) into queryable knowledge bases.
- A multi-expert agentic system combining retrieval-augmented generation with quality control mechanisms (re-ranking, hallucination detection) tailored for mission-critical environments.
- An "API-in-the-loop" protocol enabling autonomous agents to operate within network-restricted secure environments.
- A deployed, operational system validated through real-world use at ESOC ground stations.

This paper is structured as follows: Section II discusses the design of the data pipelines to continuously extract and structure operational data from diverse sources while touching on the construction of the similarity graphs. Section III presents the agentic AI system architecture, detailing the agent-assisted incident investigation workflow with its quality control mechanisms, as well as the automated conclusion generation process. Finally, conclusions and future work are drawn in Sections IV and V, respectively.

II. Knowledge Base Construction

Building an effective knowledge base relies on two essential components: properly processing incoming data and structuring it for practical use. The following subsections outline the technical architecture and methodologies employed to achieve both objectives.

A. Pipeline Architecture and Processing

Ground stations generate large amounts of data daily, stored in different but isolated environments. At the European Space Operations Centre (ESOC), these include Station Terminal Computer logs (STCs), Anomaly Report Traces (ARTs), technical documents, and communication records. To retrieve all this information and make it available to process through our pipelines, scheduled cron jobs execute daily retrieval operations, extracting data produced within the last 24 hours, or immediately, depending on the source. The purpose of the pipeline described in Figure 1 is to unify the data by extracting structured information and generating embeddings from their content and storing it in either a relational or vector database.

The pipeline first categorizes files by type, as each category requires a distinct extraction methodology, and subsequently stores the processed data in separate vector store collections. Source documents undergo deduplication at two levels: file-level checks using hash-based algorithms [12], and record-level checks for CSV files where individual rows may duplicate or update previous entries.

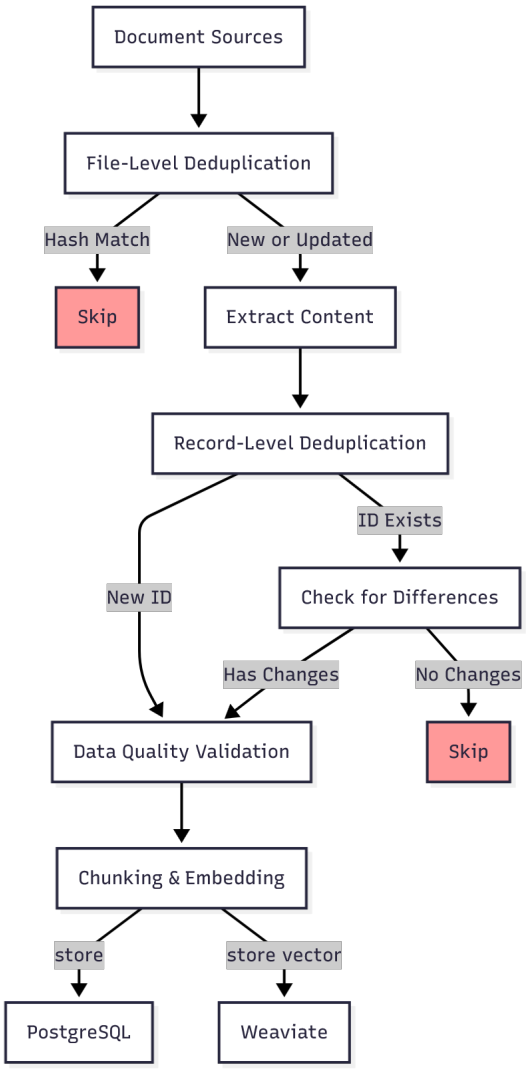

Figure 1. Pipeline architecture and workflow.

Each entry is analysed for missing information, as the data have mostly been created by humans over multiple years. Inconsistencies in notation, schema changes or mistakes are common and need to be handled. Next, a rule-based approach discards entries missing critical fields (e.g., main description) while retaining those with missing non-critical fields, such as cross-references or metadata.

To store the vector representation of each entry, the first step is chunking the information to ensure better retrieval performance [13]. Ensuring temporal and context coherence through these fragments is key to enabling transparent sourcing of the information used for each query. Chunks are obtained by overlapping them and enriching their metadata with timestamps, source identifiers, and their position in the overall text. After segmentation and embedding, the vector representation

are stored in the chosen vector database, Weaviate [14], along with the positional metadata described above, while the structured information and remaining metadata are stored in a relational database, PostgreSQL [15]. The information from one database can be traced back to the other thanks to a unique identifier. The simplified pipeline workflow is described in Figure 1.

Additionally, users can manually upload new data (e.g., files of similar structure to the one described above, or images and PDFs that may contain illustrations) to the system. As the essence of this information, where content and structure are unknown, is not efficiently captured by traditional methods, a multimodal approach using ColPali [16] was adopted, which avoids lossy text-only segmentation and allows images, diagrams, and tables to be retrieved alongside text passages. The open-source Gemma3 [17] vision Large Language Model (vLLM) [18] was selected to understand the information and create embeddings from them, which are then stored in Weaviate.

B. *Structuring Data for Practical Use*

The vector database is segregated into different collections for each document type, enabling similarity searches on specific entries. The results are enriched using the metadata link to the relational database described previously. The resulting information is leveraged to produce a similarity graph using t-Distributed Stochastic Neighbors Embedding (t-SNE) [19], a dimensionality reduction technique that visualizes high-dimensional data by projecting each data point onto a 2D map. This processing also clusters related ideas, boosting the insights of the similarity graph by grouping similar concepts together.

In order to offer more context to the users, the solution also generates a second type of plot similar to a knowledge graph. It is built using metadata on related incidents as links to other files to enrich the graphic, as demonstrated in Figure 2.

This view is fully browsable and interactive; clicking on each node provides more information about the incident and empowers users with new ways to investigate by traversing the knowledge graph, supporting rapid and evidence-driven investigations across missions. Additional metadata related to dates or ground station origin are also present, which the AI assistant will leverage to further filter the results to investigate. Figure 3 summarises the post-processing data storage architecture and the flow of information from these sources to the visualisation services and AI agents.

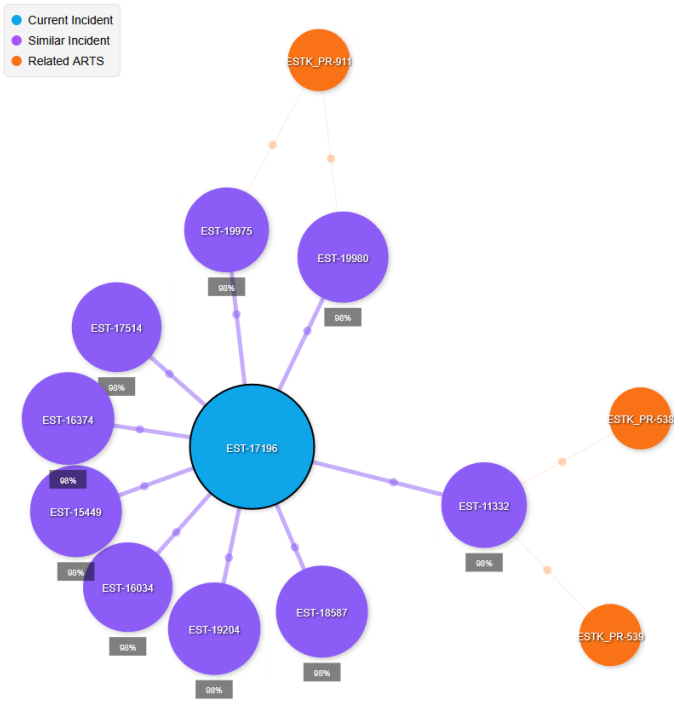

Figure 2. Knowledge Graph Visualisation.

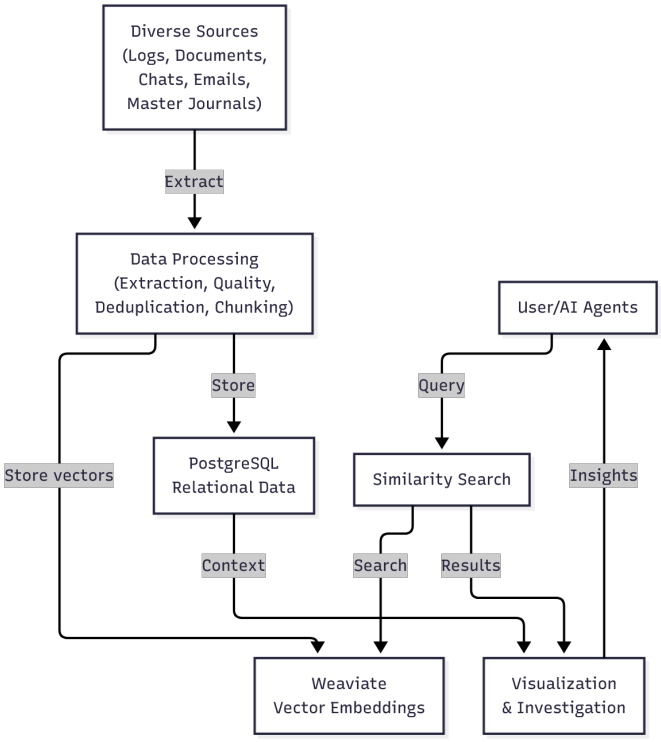

Figure 3. Data Lifecycle.

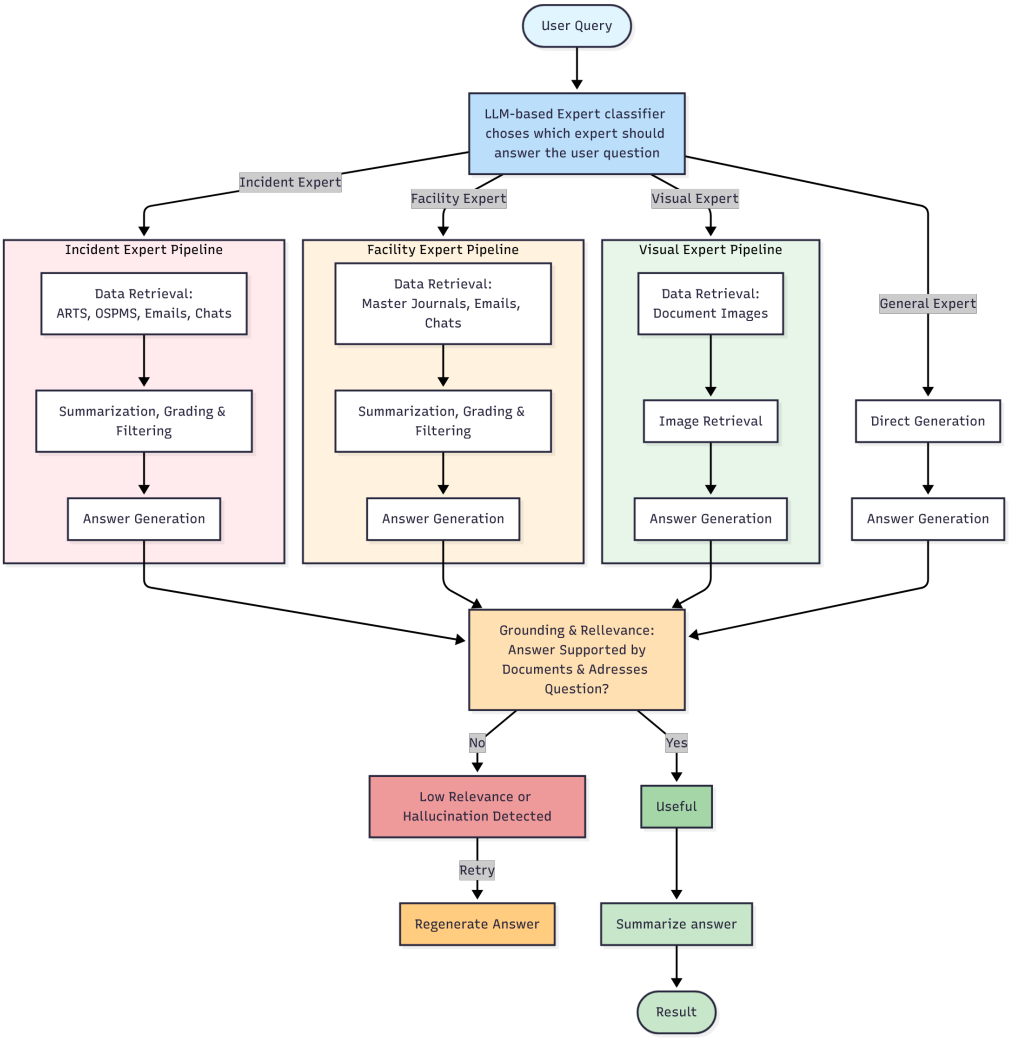

Figure 4. AI Architecture diagram.

III. Agentic AI for Incident Investigation

The assistants are designed using open-source software and models, with LangGraph [20] orchestrating the agent steps and qwen3:8b [21] serving as an expert model. Model selection was driven by practical constraints: most open-source LLMs support fewer than 250K tokens, compared to 1M tokens in some closed-source alternatives [22],[23]. Furthermore, open-source models with larger context windows require multiple GPUs to store model weights and the Key-Value (KV) cache [24], which were unavailable during development. These hardware, architectural, and licensing constraints made qwen3:8b the optimal choice, balancing performance with size.

A. Agent-Assisted Incident Investigation

Our approach employs the multi-expert workflow architecture shown in Figure 4. The first step identifies and expands acronyms in the user query using a predefined dictionary. This step is essential because domain-specific terminology can be misinterpreted by general-purpose LLMs unfamiliar with the specialized lexicon. The updated query is routed by a model which selects the most appropriate experts from a pool of four specialists with access to different data sources and able to run in parallel.

- **Incident Expert**: Focused on incident reports and communications for anomaly-centric analyses. It can answer questions about similar past incidents and anomalies.
- **Ground Station Expert**: Specialized in ground station logs and communications for timeline reconstruction and subsystem-level investigations.
- **Visual Document Expert**: Interprets visual content including diagrams, charts, scanned pages, and embedded images in documents.
- **General Expert**: Handles general questions outside the scope of other specialists.

The answers from the specialists are then gathered and graded for quality before being summarised and returned to the user.

The retrieval part of the workflow is run against the vector database described in previous sections, and uses hybrid search [25] and automatic filtering based on the user query.

Hybrid search is used for textual sources. It is a combination of semantic similarity (vector search) with keyword-based scoring, specifically Best Match 25 Model with Extension to Multiple Weighted Fields (BM25F) [26]. This method strikes a balance between natural language queries and technical identifiers, such as error codes or timestamps. Filters are created autonomously by the agent using dspy [27] to provide consistent formatting of the filtering arguments. Filters enable agents to execute focused searches based on the relevant metadata (e.g., dates, stations, etc.) to improve the relevance of the retrieved sources.

The visual document expert is a special case as it needs to search through multi-modal documents, such as PDF manuals, schematics, or diagrams. Retrieval is performed using the ColPali framework. This allows the assistant to ground its reasoning on figures and tables, as well as textual evidence.

Once the relevant chunks are identified, their original document is retrieved to put each chunk in context before a summarisation step condenses each entry into concise representations preserving key information, such as timestamps, subsystem identifiers, and error codes. These representations are then re-ranked using another LLM to grade them against the user query, this method has been shown to improve answers by refining the quality of the subset [28]. If not enough documents are deemed useful, the filters are adapted and the user query is rewritten for performance inside a retry loop.

After retrieval, each expert generates a response to the user query. The response then undergoes two validation steps using LLM-as-a-judge methods [29]: a hallucination check to ensure the answer is grounded in the provided documents, and a quality check to verify that it addresses the user question. If either validation fails, query parameters are adjusted, and the retrieval is retried.

Once the system's answers pass all quality checks, it is returned to the user with direct links to the documents used for the generation in the User Interface (UI) thanks to the metadata. All these steps ensure that the quality meets the standards of this mission-critical domain while providing clear and transparent answers within seconds.

B. Automated Conclusion Generation

The last developed agent is focused on generating the investigation closure. Its objectives are to propose conclusions, a root cause, the impact on systems, and mitigation actions for each incident. However, the implementation contrasts with previous agents. This one is a fully autonomous agentic loop with tool calling [30], thus, it mimics how issues would be investigated by real users. The reason for the different implementation was simply because users started to trust the system and felt more comfortable with a pure agentic loop. Its architecture is described in Figure 5.

The main agent node represented in purple in Figure 5 uses qwen3:8b in thinking mode. By enabling this feature, the model is allowed more time to reflect on its actions and current state of its internal investigation; great improvements have been observed in the way the model handles the tools at its disposal to steer the investigation in a meaningful direction at the cost of increased inference time.

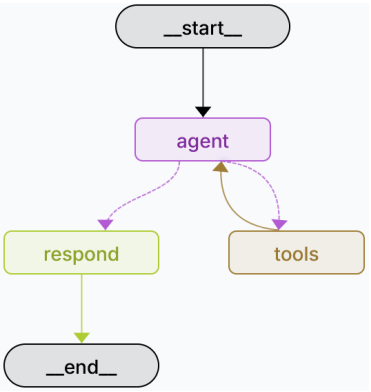

Figure 5. Conclusion Agent Architecture.

A single search tool is made available for the model to use in autonomy. The agent decides how to use it, with vector, keyword, or hybrid searches being performed independently based on what is thought to be the relevant step to perform. The output of the tool is post-processed and stored in the agent state [31] in a way that filters out duplicate results in subsequent queries. This state is ingested into the model prompt at each loop to prevent context expansion and information dilution after a few loops, maintaining the agent performance in long-running sessions; this is called context engineering [32].

With the agent being deployed on a secure server limiting outgoing requests, a networking issue was encountered, preventing the model from enriching the data retrieved from the vector database (running in the same environment) with the metadata of the relational database (running on another server). The nominal flow would have been for a GET request to be made against the relational database for the documents retrieved so far to enrich them with added context. The Human-In-The-Loop (HITL) protocol [33],[34] was used to address this issue by redirecting the call to the API endpoint of our database instead of asking for user feedback. The attempt was successful, as the HITL protocol does not rely on emitting a request from the server, but instead our request is embedded in the response body of the initial call to the model. By processing the response body on the other end of the system, which has unrestricted access to PostgreSQL, the requested data can be sent back to the agent by resuming the conversation where it was left off with the protocol. Hence, "API-in-the-loop" was implemented to ensure the agent could access the relevant data sources for the investigation.

At the start of each loop, the agent's partial findings are reflected upon, and a decision is made whether to investigate further in a set direction or if enough information has been gathered to draw conclusions. In the latter case, the proposed conclusions, root cause, impact, and mitigation strategies derived from the investigation are written by the main agent,

making sure all sources used to generate the response are cited. Additionally, the model is prompted to score its confidence in the provided answers based on how well sources correlate. This feature builds trust and focuses the user's attention on low scoring answers to encourage further manual or assisted investigation. However, this free-form text cannot be used as-is by the system as structured JSON output is expected to populate the UI fields with this information. A response agent is subsequently called after conclusions are drawn and is set up to answer in structured JSON format following the expected schema of the calling service. It is instructed to simply reformulate the main agent findings without changing its content, separating the content of the answer from the confidence score it self-evaluated. This approach provides flawless accuracy in the generated JSON schema, ensuring the system does not fail due to a formatting mistake.

The resulting assistant is able to adapt the amount of time spent investigating issues based on their complexity and to reach documents that would not have been considered in the investigation agent described in Section III-A, thanks to the ability to investigate step by step, delving deeper into the analysis of relevant documents uncovered at each loop. This design implements a human-in-the-loop paradigm, as the generated conclusions are presented to users for review before being added to the incident database, creating a powerful hybrid approach where the system conducts comprehensive, autonomous investigations while human review ensures that only verified findings enter the organizational knowledge base.

IV. Conclusion

This paper presents an already-in-use AI-driven incident investigation framework designed to mitigate the fragmentation of operational data and the loss of institutional knowledge in ground station environments. The framework integrates automated data processing pipelines that continuously aggregate and standardize information across heterogeneous sources, interactive visualization interfaces revealing relationships and similarities among data sources, a multi-expert agentic system for supporting user investigations, and an autonomous reasoning agent formulating preliminary conclusions for each incident. The system incorporates rigorous validation mechanisms and transparency protocols to maintain accuracy and explainability while operating within hardware and security constraints inherent to these critical environments.

The practical deployment revealed innovative solutions including the "API-in-the-loop" protocol for network-restricted environments and context engineering techniques for managing open-source LLM constraints. Initial results have been promising, with users expressing satisfaction with the system's capabilities. Ongoing work consists in further refinement of the quality of the answers and mitigation of hallucinations inherent in LLM-based systems that require continuous monitoring and improvement.

Beyond immediate operational gains, this work establishes a foundation for future collaborative agentic ecosystems capable of reasoning over complex data and supporting more autonomous, data-driven decision-making in space operations.

V. Future Work

Future work will expand the existing solution, where two main directions are already identified and ongoing:

Unified Agentic Architecture: The multi-expert agent (Section III-A) and the conclusion agent (Section III-B) currently use different architectures. Migrating the multi-expert system to a fully autonomous tool-calling architecture, as implemented in the conclusion agent, would enable deeper investigation capabilities and more flexible reasoning across all user interactions.

Real-Time Collaborative Investigation: The current system provides answers after investigation completes. Future iterations will surface the agent's reasoning process in real-time, allowing operators to observe which documents are being consulted and why. Implementing human-in-the-loop interaction would enable operators to guide, correct, or redirect the investigation mid-analysis. This transparency and control would transform the system from a tool that delivers answers into a "co-pilot" experience that builds more trust with the user and could define the future vision for autonomous space operations.

References

[1] Meta, *React: A JavaScript library for building user interfaces*, 2025. Accessed: 2025-01-29. [Online]. Available: https://react.dev/.

[2] A. Vaswani *et al.*, *Attention is all you need*, 2023. arXiv: 1706.03762 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/1706.03762.

[3] W. X. Zhao *et al.*, *A survey of large language models*, 2025. arXiv: 2303.18223 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2303.18223.

[4] S. Chen, S. Wong, L. Chen, and Y. Tian, *Extending context window of large language models via positional interpolation*, 2023. arXiv: 2306.15595 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2306.15595.

[5] P. Lewis *et al.*, *Retrieval-augmented generation for knowledge-intensive nlp tasks*, 2021. arXiv: 2005.11401 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2005.11401.

[6] L. Ma *et al.*, *A comprehensive survey on vector database: Storage and retrieval technique, challenge*, 2025. arXiv: 2310.11703 [cs.DB]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2310.11703.

[7] W. Li *et al.*, *Approximate nearest neighbor search on high dimensional data — experiments, analyses, and improvement (v1.0)*, 2016. arXiv: 1610.02455 [cs.DB]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/1610.02455.

[8] A. Hogan *et al.*, "Knowledge graphs," *ACM Computing Surveys*, vol. 54, no. 4, pp. 1–37, Jul. 2021, ISSN: 1557-7341. DOI: 10.1145/3447772. [Online]. Available: http://dx.doi.org/10.1145/3447772.

[9] Z. Ji *et al.*, "Survey of hallucination in natural language generation," *ACM Computing Surveys*, vol. 55, no. 12, pp. 1–38, Mar. 2023, ISSN: 1557-7341. DOI: 10.1145/3571730. [Online]. Available: http://dx.doi.org/10.1145/3571730.

[10] A. Narimani and S. Klarmann, "Integration of Large Language Models for Real-Time Troubleshooting in Industrial Environments based on Retrieval-Augmented Generation (RAG)," in *Proceedings of the 7th European Industrial Engineering and Operations Management Conference*, Augsburg, Germany: IEOM Society International, 2024, ISBN: 979-8-3507-1737-2. DOI: 10.46254/EU07.20240085. [Online]. Available: https://ieomsociety.org/proceedings/2024germany/85.pdf.

[11] K. S. Prakash *et al.*, "VALOR: Validation for Aerospace LLM Output and Reasoning," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, California, NASA Technical Memorandum NASA/TM–20260000076, 2026. [Online]. Available: https://ntrs.nasa.gov/citations/20260000076.

[12] J. Alakuijala, B. Cox, and J. Wassenberg, *Fast keyed hash/pseudo-random function using simd multiply and permute*, 2017. arXiv: 1612.06257 [cs.CR]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/1612.06257.

[13] R. Schwaber-Cohen and A. Patel, *Chunking Strategies for LLM Applications*, 2025. Accessed: 2025-01-29. [Online]. Available: https://www.pinecone.io/learn/chunking-strategies/.

[14] Weaviate, *Weaviate: Vector Database*, 2025. Accessed: 2025-01-29. [Online]. Available: https://weaviate.io/.

[15] PostgreSQL Global Development Group, *PostgreSQL: The World's Most Advanced Open Source Relational Database*, 2025. Accessed: 2025-01-29. [Online]. Available: https://www.postgresql.org/.

[16] M. Faysse *et al.*, *Colpali: Efficient document retrieval with vision language models*, 2025. arXiv: 2407.01449 [cs.IR]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2407.01449.

[17] G. Team *et al.*, *Gemma 3 technical report*, 2025. arXiv: 2503.19786 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2503.19786.

[18] F. Lin, *Vision language models: A survey of 26k papers*, 2025. arXiv: 2510.09586 [cs.CV]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2510.09586.

[19] T. T. Cai and R. Ma, *Theoretical foundations of t-sne for visualizing high-dimensional clustered data*, 2022. arXiv: 2105.07536 [stat.ML]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2105.07536.

[20] LangChain, *LangGraph: Build resilient language agents as graphs*, 2025. Accessed: 2025-01-29. [Online]. Available: https://langchain-ai.github.io/langgraph/.

[21] A. Yang *et al.*, *Qwen3 technical report*, 2025. arXiv: 2505.09388 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2505.09388.

[22] OpenAI, *OpenAI: Introducing GPT 4.1*, 2025. Accessed: 2025-01-29. [Online]. Available: https://platform.openai.com/docs/models/gpt-4.1.

[23] Anthropic, *Claude Sonnet 4 now supports 1M tokens of context*, 2025. Accessed: 2025-01-29. [Online]. Available: https://claude.com/blog/1m-context.

[24] R. Pope *et al.*, *Efficiently scaling transformer inference*, 2022. arXiv: 2211.05102 [cs.LG]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2211.05102.

[25] Weaviate, *Weaviate: Hybrid search*, 2025. Accessed: 2025-01-29. [Online]. Available: https://docs.weaviate.io/weaviate/search/hybrid.

[26] S. Robertson and H. Zaragoza, "The probabilistic relevance framework: Bm25 and beyond," *Foundations and Trends in Information Retrieval*, vol. 3, pp. 333–389, Jan. 2009. DOI: 10.1561/1500000019.

[27] O. Khattab *et al.*, *Dspy: Compiling declarative language model calls into self-improving pipelines*, 2023. arXiv: 2310.03714 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2310.03714.

[28] X. Ma, Y. Gong, P. He, H. Zhao, and N. Duan, *Query rewriting for retrieval-augmented large language models*, 2023. arXiv: 2305.14283 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2305.14283.

[29] J. Gu *et al.*, *A survey on llm-as-a-judge*, 2025. arXiv: 2411.15594 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2411.15594.

[30] LangChain, *LangChain: Tool calling documentation*, 2025. Accessed: 2025-01-29. [Online]. Available: https://docs.langchain.com/oss/python/langchain/tools.

[31] LangChain, *LangChain: State documentation*, 2025. Accessed: 2025-01-29. [Online]. Available: https://docs.langchain.com/oss/python/langgraph/use-graph-api.

[32] L. Mei *et al.*, *A survey of context engineering for large language models*, 2025. arXiv: 2507.13334 [cs.CL]. Accessed: 2025-01-29. [Online]. Available: https://arxiv.org/abs/2507.13334.

[33] X. Wu *et al.*, "A survey of human-in-the-loop for machine learning," *Future Generation Computer Systems*, vol. 135, pp. 364–381, Oct. 2022, ISSN: 0167-739X. DOI: 10.1016/j.future.2022.05.014. [Online]. Available: http://dx.doi.org/10.1016/j.future.2022.05.014.

[34] LangChain, *LangChain: Human in the loop protocol*, 2025. Accessed: 2025-01-29. [Online]. Available: https://docs.langchain.com/oss/python/langchain/human-in-the-loop.

Residual Hybrid Motor Controller with Multi-Phase Learning

Johann Wiens, Christoph Reich
Institute of Data Science, Cloud Computing and IT-Security
Furtwangen University of Applied Sciences
Furtwangen, Germany
Email: johann-wiens@outlook.com, christoph.reich@hs-furtwangen.de

***Abstract*—Direct-Current (DC) motors serve as the fundamental actuators in emerging intelligent systems, including advanced robotics, autonomous vehicles, and smart home infrastructure. Precise and energy-efficient control of these motors is paramount, as control quality directly dictates system safety, movement fluidity, and operational longevity. However, conventional Proportional-Integral-Derivative (PID) controllers in compact motor drives often struggle with repeatable nonlinearities and unmodeled dynamics. Examples are friction, cogging torque, and saturation effects. These effects can reduce stability margins and require frequent manual retuning. This paper proposes a hybrid motor controller. A tuned PID controller provides the baseline loop. A Long–Short-Term-Memory (LSTM) module adds a residual path. The baseline loop guarantees stability, transparency, and hard actuator limits. The compact three-layer LSTM adds a bounded corrective voltage to compensate repeatable nonlinearities. Learning runs in two phases. (1) Behavior Cloning: the LSTM first learns to imitate the PID controller. This fixes the control direction and creates a safe starting point for Phase 2. (2) Advantage-Weighted Regression (AWR): a task-specific advantage function rewards residual actions only when they reduce trajectory error. A conservative gate discards degrading updates. For realistic and reproducible evaluation, the geared DC drive is modeled as an averaged Pulse-Width-Modulation (PWM) Brushless-Direct- Current (BLDC) motor model. The model includes Stribeck friction, iron/core losses, cogging torque, current and voltage saturation, and a lumped thermal winding model. Tracking tasks follow quintic S-curve trajectories with output-side load disturbances. One trajectory is held out to test generalization. Across all scenarios, the residual path lowers position Mean-Squared-Error (MSE) compared to pure PID and reduces overshoot at similar rise times. Comparing the MSEs, PID+LSTM reduces the error by about 98.5% on average relative to PID only.**

Keywords-Hybrid controller; Smart System AI-Controlled; Proportional–integral–derivative controller; LSTM controller; Behavior Cloning; Advantage-Weighted Regression; DC motor; motor modeling.

I. Introduction

PID control is still the de facto standard in industrial motion systems. It is simple, interpretable, and when properly tuned robust to moderate uncertainty [1]. Modern compact drives, however, increasingly operate in regimes with strong nonidealities. These include static and Stribeck friction, iron and core losses, cogging torques, saturations, and thermal drift. In such regimes, purely linear control often faces a trade-off between steady-state accuracy, overshoot, and actuator stress. Laborious retuning for each operating point is then common [2].

A promising alternative is a residual learned policy on top of a proven baseline controller. The baseline loop guarantees reasonable behavior. A bounded data-driven term compensates repeatable, state-dependent errors [3]. This pattern is attractive for safety-relevant actuation. The learned component can be range-limited, while the stabilizing structure stays in place.

This study considers a geared DC drive and adds a compact LSTM residual to a strongly tuned PID. Learning proceeds in two stages. First, Behavior-Cloning (BC) aligns the LSTM with the PID to give a conservative starting point. Then, Advantage-Weighted Regression (AWR) uses a simple task-centered advantage. A conservative policy-improvement gate prevents regressions [4]. The BLDC model includes key multiphysics effects (Stribeck friction, iron-loss torque, cogging, thermal dynamics) and hard limits. Smooth quintic S-curves serve as references. They are a canonical setup for point-to-point motions [5].

The contributions of this work are:

- Realistic, reproducible motor model of a geared DC drive with multiphysics effects.
- Compact hybrid PID+LSTM controller with two-phase training (BC $\rightarrow$ AWR) and conservative update gating.
- Systematic evaluation on S-curve trajectories including load disturbances and a hold-out trajectory for generalization.

The remainder of the paper is as follows. Section II reviews related work. Section III explains the controller and learning scheme. Section IV describes the model and scenarios. Section V reports results and limitations. Section VI concludes.

II. Related Work

Recent work couples classical PI/PID loops with small learned compensators in a parallel or residual path. The goal is to handle repeatable nonlinearities while keeping the baseline structure. On a DC servo test stand, the authors of [6] use a PID baseline and add an online Artificial-Neural-Network (ANN) / Recurrent-Neural-Network (RNN) precompensator gated by a fuzzy system. They report improved overshoot and steady-state accuracy in hardware. In contrast, the present work uses a direct LSTM residual with two-phase off-policy learning. This simplifies the architecture and training compared to the fuzzy-gated RNN in [6], For Permanent-Magnet-Synchronous-Motor (PMSM) Field-Oriented-Control (FOC), the authors of [7] add a small feedforward network that corrects PI transients. After pruning and quantization, the network runs on Microcontroller Unit (MCU) hardware and reduces overshoot. Both approaches share the same safety idea used here: a bounded learned term adds to the conventional loop, instead of replacing it.

Recurrent networks, especially LSTMs, can capture temporal dependencies in torque ripple, flux dynamics, or friction memory. In Direct-Torque-Control (DTC) of induction machines, the authors of [8] replace the switching table with a ConvLSTM selector and improve low-speed behavior in simulation. For synchronous machines, the authors of [9] use an LSTM-driven predictive current controller. Two-degree-of-freedom designs (feedforward + feedback) also profit from recurrent models. Yin et al. [10] train an LSTM inversion model as a feedforward compensator and combine it with linear feedback for nanopositioning. Broader evaluations of ML-based PMSM drive controllers also support the use of recurrent architectures [11].

In this paper, the learned component is deliberately residual and strictly bounded. The two-stage training (BC $\rightarrow$ AWR) is used to preserve stability and interpretability of the baseline.

Key differences to prior work: Compared to hybrid PI/PID+network approaches such as [7], this study differs in three main aspects:

- *Model and task:* The focus is a geared DC drive modeled as an averaged PWM BLDC. The model includes friction (with Stribeck), iron/core losses, cogging torque, saturation, and a lumped thermal model. The evaluation uses S-curve point-to-point references and output-side load disturbances. Prior work often uses simpler friction models or different machines with less detailed multiphysics.
- *Controller architecture:* The controller uses a small LSTM as a residual voltage path in parallel to a fixed, tuned PID. The PID is never turned off. The residual voltage is hard-bounded. Thus, authority and stability margins remain with the classical loop, unlike approaches that replace switching tables or MPCC [8].
- *Learning procedure:* Many earlier works rely on direct supervised training or problem-specific optimization. Here, learning uses two phases. First, Behavior Cloning imitates the PID and fixes the control direction. Second, Advantage-Weighted Regression updates the residual, with improvements measured against the PID baseline and protected by a conservative gate.

Scope of this paper: This paper:

- designs a compact three-layer LSTM residual that augments a tuned PID position loop for a realistic geared DC drive;
- trains this residual in two phases: PID-mimicking Behavior Cloning, followed by Advantage-Weighted Regression using a simple, trajectory-error-based advantage and conservative policy-improvement gating; and
- compares the hybrid controller against the pure PID on multiple S-curve trajectories, including a hold-out trajectory and load disturbances. Metrics include position MSE, overshoot, rise time, and actuator limits.

The results show that a strictly bounded residual LSTM can consistently improve tracking over a strong industrial baseline without changing the underlying control structure.

III. Hybrid Control Approach

The proposed hybrid approach combines a robust PID baseline with a small residual LSTM. The PID provides a well-understood, stabilizing foundation, while the LSTM is trained in two phases (Behavior Cloning, then AWR) to compensate repeatable, state-dependent nonlinearities that a fixed-parameter PID cannot fully address.

Phase 1: Behavior Cloning (BC)

In Phase 1 (cf. Figure 1), the LSTM is trained by supervised learning to imitate the PID output $u_{\mathrm{PID}}(t)$ [12]. Let $u_{\mathrm{LSTM}}(t;\theta)$ denote the LSTM output. The horizon $\mathcal{T}$ is used to minimize

$$\min_{\theta} \; \mathcal{L}_{\mathrm{BC}}(\theta) \; = \; \sum_{t\in\mathcal{T}} \left\| u_{\mathrm{LSTM}}(t;\theta) - u_{\mathrm{PID}}(t) \right\|_2^2. \qquad (1)$$

Only improving updates are accepted; degrading ones are rejected. This yields a conservative starting point.

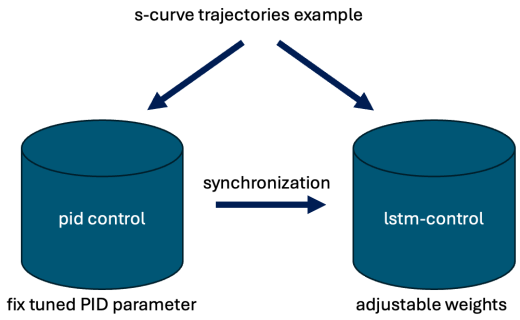

Figure 1. Phase 1: Behavior Cloning (BC).

Phase 2: Reinforcement Learning with AWR

In Phase 2 (cf. Figure 2), the full closed-loop drive is used. Advantage-Weighted Regression (AWR) [13] acts as the policy-improvement method. PID and LSTM both receive the control error $e(t)$; PID parameters remain fixed. The LSTM is updated by weighted regression. Only updates that improve performance are accepted (conservative policy improvement [4]).

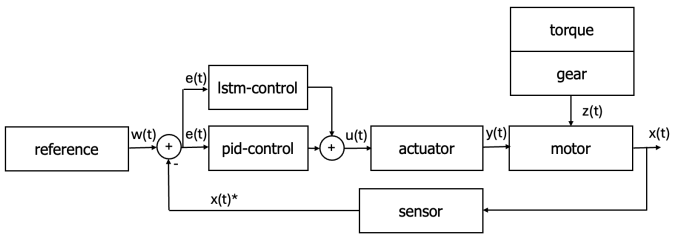

Figure 2. Phase 2: Reinforcement Learning (AWR).

Implementation Notes

- **Safety bounds:** The PID path is always active. The LSTM output is bounded via saturation or a blend gain.
- **Curriculum:** BC starts with simple references (step, ramp, sine) and then uses richer trajectories before AWR.
- **Validation:** Trajectory error, control effort, and robustness (parameter variations) are evaluated systematically.

IV. Simulation Environment and Modeling

The goal is a realistic and reproducible assessment of the hybrid PID+LSTM controller on S-curves.

Simulation instead of hardware: Evaluation starts in simulation to make disturbances and drifts reproducible and to avoid risk and wear during early policy iterations. The setup follows the idea of a digital twin: a high-fidelity dynamic model used as a virtual test bench before deployment. The model captures dominant nonidealities: Stribeck friction, iron losses, cogging, saturations, and thermal effects. Hardware validation is planned as a next step.

References: smooth S-curves: Quintic S-curves with zero boundary conditions (position, velocity, acceleration) are used as references. They have low jerk and are common in industry. They also isolate controller performance from high-frequency artifacts [5].

A. Overview and Time Discretization

All dynamics are integrated with explicit Euler at step $dt{=}10^{-4}\,\mathrm{s}$. PWM is modeled as an *averaged* actuation voltage with small ripple. The core structure is residual:

$$V_{\mathrm{tot}}(k) \;=\; V_{\mathrm{PID}}(k) \;+\; V_{\mathrm{LSTM}}(k). \tag{2}$$

B. Reference Trajectories and Scenarios

The numbers are not chosen arbitrarily, but are the mathematically compelling result when searching for a curve that starts at 0 and ends at 1 and has no velocity or acceleration at either end. For $0{\le}t{\le}T$,

$$s(t) = 10\big(\tfrac{t}{T}\big)^3 - 15\big(\tfrac{t}{T}\big)^4 + 6\big(\tfrac{t}{T}\big)^5 . \tag{3}$$

Ten scenarios with different reach times T and end positions are evaluated; see Table I. An external load $T_{\mathrm{Load}}{=}0.8\,\mathrm{N\,m}$ acts by default from $40\,\%$ to $80\,\%$ of the scenario duration.

C. Motor Model: BLDC with Gearbox, Averaged PWM

The BG 42x15 with planetary gearbox PLG 42S is represented as a DC drive with states current i, angular velocity ω, angle θ, and winding temperature T_{cu}. The fundamental equations are:

$$v = R(T)\,i + L\,\dot{i} + K_e^{*}\,\omega, \tag{4}$$

$$J\,\dot{\omega} = K_t^{*}\,i - T_{\mathrm{fric}}(\omega,T) - T_{\mathrm{iron}}(\omega) - T_{\mathrm{cog}}(\theta) - T_{\mathrm{load}}. \tag{5}$$

Cf. [14]. Parameters K_e^{*} and K_t^{*} vary slightly with current and temperature.

Nonidealities:

- **Friction** (viscous, Coulomb, Stribeck) cf. [15]:

$$T_{\mathrm{fric}}(\omega) \;=\; B\,\omega + T_c\,\mathrm{sgn}(\omega) + (T_s{-}T_c)\,\exp\!\Big(-\Big(\frac{\omega}{\omega_s}\Big)^2\Big)\,\mathrm{sgn}(\omega). \tag{6}$$

- **Iron losses** cf. [16]:

$$P_{\mathrm{iron}}(\omega) = k_h\,|\omega| + k_e\,\omega^2, \tag{7}$$

$$T_{\mathrm{iron}}(\omega) = \frac{P_{\mathrm{iron}}(\omega)}{\max(|\omega|,\varepsilon)}. \tag{8}$$

- **Cogging torque:**

$$T_{\mathrm{cog}}(\theta) \;=\; T_{\mathrm{cog,\,amp}}\sin(n_{\mathrm{el}}\,\theta). \tag{9}$$

The harmonic is included as in [17].

- **PWM artifacts:** averaged actuation voltage with small $v_{\mathrm{ripple}}(t)$ at $f_{\mathrm{PWM}}{=}20\,\mathrm{kHz}$.
- **Saturations:** current and voltage limits

$$|i|{\le}i_{\mathrm{peak}}, \qquad |v|{\le}V_{\mathrm{bus}}. \tag{10}$$

- **Thermal model:** A lumped RC model is used cf. [18]:

$$C_{\mathrm{th}}\,\dot{T}_{\mathrm{cu}} \;=\; i^2R(T_{\mathrm{cu}}) + |T_{\mathrm{iron}}\omega| - h\,(T_{\mathrm{cu}} - T_{\mathrm{amb}}). \tag{11}$$

The output torque is reflected through gear ratio r and efficiency η:

$$T_{\mathrm{load,m}} \;=\; \frac{T_{\mathrm{load,out}}}{r\,\eta}. \tag{12}$$

Linear position follows from $\theta_{\mathrm{out}}{=}\theta/r$ and shaft radius r_s:

$$p \;=\; \frac{\theta}{r}\,r_s, \tag{13}$$

cf. the arc-length relation [19].

D. Baseline Controller (PID)

With position error $e(k) = p_{\mathrm{ref}}(k) - p(k)$:

$$V_{\mathrm{PID}}(k) \;=\; K_p\,e(k) + K_i\sum_{j\le k} e(j)\,dt + K_d\,\frac{e(k)-e(k-1)}{dt}. \tag{14}$$

A moderately aggressive tuning ($K_p{=}180$, $K_i{=}200$, $K_d{=}10$) is employed; V_{PID} is limited to $|V| \le V_{\mathrm{bus}}$ [1].

E. Residual Controller (LSTM)

The LSTM (3 layers, hidden size 12) receives sequences of the last $T_{\mathrm{seq}}{=}10$ samples with features

$$\left[\frac{t}{T_{\mathrm{end}}},\ \frac{i}{i_{\mathrm{peak}}},\ \frac{e}{e_{\mathrm{scale}}}\right] \in [-1,1]^3, \qquad e_{\mathrm{scale}}{=}0.30\,\mathrm{m}. \tag{15}$$

The output is mapped to a physical voltage:

$$V_{\mathrm{LSTM}}(k) \;=\; \tfrac{1}{2}V_{\mathrm{bus}}\cdot\tanh\big(\mathrm{Head}(\mathrm{LSTM}(\mathrm{Seq}))\big), \tag{16}$$

ensuring $|V_{\mathrm{LSTM}}| \le \frac{1}{2}V_{\mathrm{bus}}$ (authority remains primarily with the PID) [20]; the combination follows Equation 2.

F. Two-Stage Learning Scheme

Phase 1: Behavior Cloning (BC): Datasets from PID rollouts (S-curves, load window) provide supervision. The target is the normalized PID voltage; optimization uses MSE with soft clipping through $\tanh(\cdot)$ [21].

Phase 2: Advantage-Weighted Regression (AWR): Hybrid rollouts (PID+LSTM) are then generated. Advantage relative to PID is defined as

$$A_k \;=\; \frac{|e_{\mathrm{PID},k}| - |e_{\mathrm{hyb},k}|}{e_{\mathrm{scale}}}, \qquad w_k \;=\; \exp(\beta A_k),\ \beta{=}4. \tag{17}$$

Regression on the residual actions is performed, weighted by w_k; updates are conservatively accepted using BC validation [4].

V. Results and Discussion

Among the ten S-curve scenarios in Table I, scenarios 1–9 are used both for training and evaluation, whereas scenario 10 is held out and used only for testing. We, therefore, treat scenario 10 as the primary generalization case: the controller must track a reference it has never seen during training.

TABLE I. Reference trajectory scenarios.

Scenario number	Position reach time [s]	Position move [m]	Simulation time [s]	**MSE PID only**	**MSE PID+LSTM**	Mode
1	0.600	0–0.30	0–0.600	**2.23e-05**	**3.79e-07**	Train+Test
2	0.725	0–0.10	0–0.725	**2.14e-05**	**4.18e-07**	Train+Test
3	0.550	0–0.15	0–0.550	**2.27e-05**	**2.54e-07**	Train+Test
4	0.300	0–0.15	0–0.300	**2.13e-05**	**2.72e-07**	Train+Test
5	0.350	0–0.05	0–0.350	**2.14e-05**	**3.60e-07**	Train+Test
6	0.650	0–0.30	0–0.600	**2.28e-05**	**4.13e-07**	Train+Test
7	0.450	0–0.10	0–0.725	**1.96e-05**	**3.69e-07**	Train+Test
8	0.350	0–0.15	0–0.550	**2.05e-05**	**2.41e-07**	Train+Test
9	0.250	0–0.15	0–0.300	**1.88e-05**	**2.56e-07**	Train+Test
10	**0.475**	**0–0.15**	**0–0.475**	2.26e-05	2.35e-07	**Test only**

The hybrid controller (PID+LSTM) improves trajectory tracking on all tasks compared to the standalone tuned PID. It shows lower error energy and faster settling at similar rise times, while respecting the same actuator limits.

A. Training Dynamics and Phase Contributions

Validation MSE: Figure 3 shows a "big step + fine-tuning" pattern: Behavior Cloning (BC) reduces validation MSE from 1.049×10^{-2} to 4.22×10^{-4} (epoch 7). Advantage-Weighted Regression (AWR) further refines to 4.34×10^{-4} (epoch 4). A conservative gate rejects later degrading epochs (epochs 5–6). *Interpretation:* BC brings the LSTM close to the PID policy

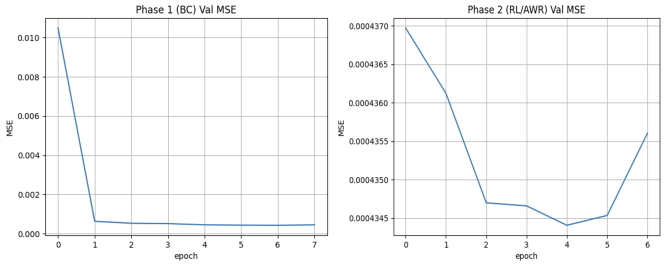

Figure 3. Validation MSE: Phase 1 (BC, left) and Phase 2 (AWR, right).

and avoids risky exploration. AWR then uses a task-centered advantage relative to the PID baseline. It shifts the residual action where it reduces trajectory error. Conservative policy improvement prevents large regressions. This explains why most of the gain appears in BC, while AWR adds robustness and small refinements, for example near friction transitions.

B. Tracking on S-Curves and Error Statistics

Position and velocity profiles: In Figure 4 and Figure 5, the hybrid controller tracks the reference more tightly and with less overshoot. Velocity is smoother, with less ripple. Peak velocities stay on the PID level ($\approx 80\,\mathrm{rad/s}$). Thus, the hybrid keeps timing but reduces error energy. *Error energy*

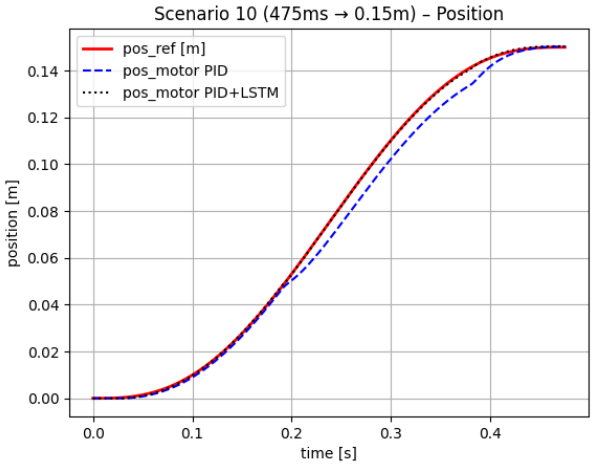

Figure 4. Test scenario 10 position: hybrid vs. PID control.

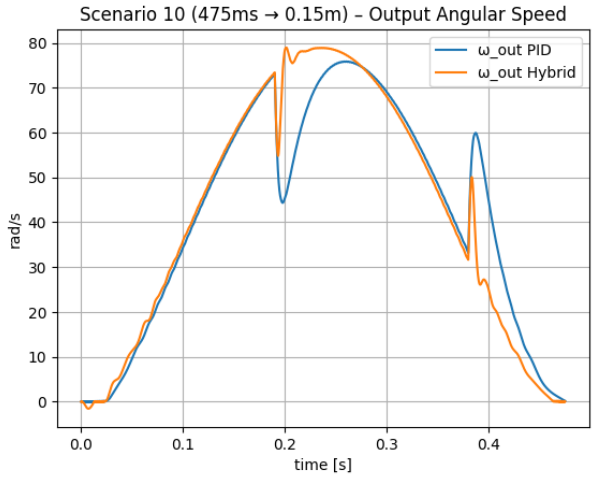

Figure 5. Test scenario 10 speed: hybrid vs. PID control.

and overshoot: Error energy and overshoot.: Across scenarios, the residual path consistently reduces position MSE relative to pure PID. In the test scenario, the reduction is roughly one order of magnitude, see Figure 6.

C. Actuation Effort and Saturations

Voltages and currents: The smoother trajectories appear in the torque and voltage profiles as well. The LSTM output is strictly limited to $|V_{\mathrm{LSTM}}| \leq \frac{1}{2} V_{\mathrm{bus}}$ Thus, the PID retains the main authority and the residual cannot cause extreme actions. Transient stresses are reduced overall. *Interaction with saturations:* The total voltage $V_{\mathrm{tot}} = V_{\mathrm{PID}} + V_{\mathrm{LSTM}}$ is still clipped by the bus limit. No new saturation regimes appear. The

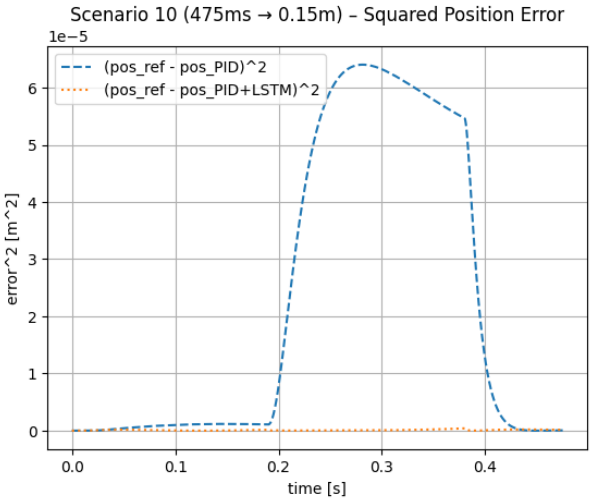

Figure 6. Test scenario 10 mean-squared error: hybrid vs. PID control.

residual mainly compensates repeatable nonlinearities (Stribeck friction, iron-loss torque, cogging) instead of creating extra peaks.

D. Robustness to Disturbances and Drift

Load window: The reference tasks include an output-side load window ($T_{\mathrm{Load}} = 0.8\,\mathrm{Nm}$, 40–80 % of thee scenario duration). Figure 7 shows the torque in test scenario 10. The hybrid shows smaller error peaks in this interval. It leverages recurring, state-dependent patterns in the nonidealities, while the PID ensures stability. *Parametric uncertainties:* The motor

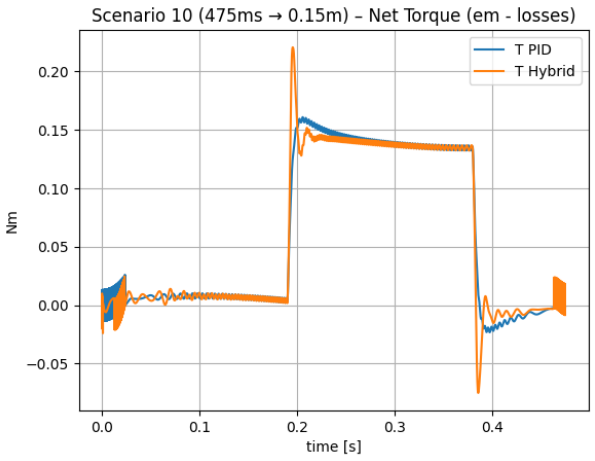

Figure 7. Test scenario 10 torque: hybrid vs. PID control.

model includes temperature-dependent winding parameters, iron losses, cogging, and saturations. Parameters drift during the trajectory, for example due to heating. Because the residual is bounded and the PID remains active, stability margins are preserved. Robustness comes from this separation: the PID stabilizes, the LSTM compensates.

E. Generalization and Hold-out

One trajectory is used as a *hold-out* and appears only in testing; cf. Table I). On this unseen reference, the hybrid controller again achieves tighter tracking, smoother velocity, and lower error energy. This suggests that the residual mainly exploits repeatable, state-dependent patterns. It does not simply memorize specific trajectories.

F. Residual Contribution and Interpretability

Small recurrent architecture: The LSTM is intentionally compact (3 layers, hidden size 12) and receives a short history ($T_{\mathrm{seq}} = 10$) of normalized features. A final *tanh* enforces a hard bound on V_{LSTM}. This keeps the PID "in charge" and improves interpretability and safety. The division of labor is clear: PID handles linear and broadband effects; the LSTM corrects structured nonlinearities.

Conservative updates: During AWR, an update is kept only if it improves validation over the BC baseline. This avoids sudden policy jumps. Table II) shows that epochs 5 and 6 are rejected, even though their MSE change is small.

TABLE II. TRAINING RESULTS ACROSS PHASES AND EPOCHS.

Phase	Epoch	Val-MSE	Status
PHASE 1: BEHAVIOR CLONING			
Phase 1/BC	01	0.010490	improved
Phase 1/BC	02	0.000625	improved
Phase 1/BC	03	0.000525	improved
Phase 1/BC	04	0.000507	improved
Phase 1/BC	05	0.000445	improved
Phase 1/BC	06	0.000430	improved
Phase 1/BC	07	0.000422	improved
Phase 1/BC	08	0.000444	kept
PHASE 2: RL FINE-TUNING (AWR)			
Phase 2/RL	00	0.000437	baseline before updates
Phase 2/RL	01	0.000436	improved
Phase 2/RL	02	0.000435	improved
Phase 2/RL	03	0.000435	improved
Phase 2/RL	04	0.000434	improved
Phase 2/RL	05	0.000435	rejected
Phase 2/RL	06	0.000436	rejected

G. Computational Load and Deployment Readiness

The small LSTM and bounded residual authority are chosen with embedded hardware in mind. The PID loop remains unchanged. The LSTM adds one saturated summation path. Real-time feasibility will depend on the target platform but is aided by the compact model size.

H. Limitations and Implications

Simulation domain: All results are obtained in simulation. Simplifications, such as averaged PWM and a single cogging harmonic, may smooth some hardware effects. Metrics focus on position MSE; energy and thermal objectives are not optimized. Architecture and hyperparameters were tuned manually. Systematic ablations (e.g., over residual gain, hidden size, or sequence length) are left for future work.

Practical significance: Despite these limits, the study suggests that residual learning is a practical way to gain

performance without discarding established control structures. This is important in safety-critical and certified systems. The consistent error reduction under strict limits and a gentle learning pipeline supports future hardware or Hardware-in-the-Loop (HiL) tests.

VI. Conclusion and Future Work

This paper presented a hybrid controller that augments a tuned PID with a compact, strictly bounded LSTM in a residual path. In a realistic simulation of an averaged PWM BLDC, including friction, iron-loss, cogging, saturation, and thermal effects, the hybrid consistently outperforms pure PID. It reduces trajectory errors (about an order-of-magnitude reduction of error energy in test scenario 10), smooths velocity and torque, and maintains similar rise times without extra peaks. The two-stage learning scheme is simple to use: BC gives most of the gain, and AWR refines it conservatively.

Limitations: All experiments are in simulation. Simplified models of iron losses, cogging harmonics, and PWM may hide some hardware details. Metrics focus on position MSE; energy and temperature are treated as constraints, not optimization targets. Architecture and hyperparameters are not yet systematically explored. Generalization is tested only on S-curves with one hold-out trajectory.

Outlook: Future work will address: (i) hardware or HiL validation with the same residual bounds and safety limits, (ii) other motion profiles (trapezoidal velocity, jerk-limited), speed or torque control, and contact/backlash effects, (iii) multi-objective cost functions that include error, energy, and temperature, and (iv) more formal guarantees for safe policy updates.

Acknowledgment

This work was funded by the state of Baden-Württemberg within the MoDe ProBio project (FEIH_PAN_2685068).

References

[1] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA, 2006.

[2] B. Armstrong-Hélouvry, P. Dupont, and C. C. de Wit, "A survey of models, analysis tools and compensation methods for the control of machines with friction", *Automatica*, vol. 30, no. 7, pp. 1083–1138, 1994.

[3] T. Johannink et al., "Residual reinforcement learning for robot control", *arXiv:1812.03201*, 2019.

[4] S. Kakade and J. Langford, "Approximately optimal approximate reinforcement learning", in *Proceedings of ICML*, 2002.

[5] T. Flash and N. Hogan, "The coordination of arm movements: An experimentally confirmed mathematical model", *Journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.

[6] Z. Huang et al., "Fuzzy inference system enabled neural network feedforward compensation for position leap control of dc servo motor", *Scientific Reports*, vol. 14, no. 20814, 2024. DOI: 10.1038/s41598-024-71647-1.

[7] M. J. M. Elele, S. Pau Danilo und Zhuang, and T. Facchinetti, "Compensating pi controller's transients with tiny neural network for vector control of permanent magnet synchronous motors", *World Electric Vehicle Journal*, vol. 16, no. 4, p. 236, 2025. DOI: 10.3390/wevj16040236.

[8] S. Potturi et al., "Direct torque control of induction motor using convlstm based on gaussian pillbox surface", *Mathematical Problems in Engineering*, 2022. DOI: 10.1155/2022/4408271.

[9] I. Hammoud, S. Hentzelt, T. Oehlschlaegel, and R. Kennel, "Learning-based model predictive current control for synchronous machines: An lstm approach", *European Journal of Control*, 2022, Online first.

[10] R. Yin, Y. Chen, Z. Gong, and J. Ren, "Lstm-inversion-based feedforward–feedback nanopositioning control", *Machines*, vol. 12, no. 11, p. 747, 2024. DOI: 10.3390/machines12110747.

[11] A. M. Tom and J. L. Febin Daya, "Design of machine learning-based controllers for speed control of pmsm drive", *Scientific Reports*, vol. 15, no. 17826, 2025. DOI: 10.1038/s41598-025-02396-y.

[12] *Imitation: Documentation (behavior cloning overview)*, https://imitation.readthedocs.io/, Accessed for background on supervised behavior cloning, 2025.

[13] X. B. Peng, A. Kumar, G. Zhang, and S. Levine, "Advantage-weighted regression: Simple and scalable off-policy reinforcement learning", *arXiv preprint arXiv:1910.00177*, 2019.

[14] D. Tilbury, W. Messner, and R. Hill, "Dc motor speed: System modeling", Accessed: Oct. 19, 2025, Control Tutorials for MATLAB and Simulink (CTMS), 2025, Accessed: Oct. 19, 2025. [Online]. Available: https://ctms.engin.umich.edu/CTMS/index.php?example=MotorSpeed&section=SystemModeling.

[15] C. Canudas de Wit, H. Olsson, K. J. Åström, and P. Lischinsky, "A new model for friction", *IEEE Transactions on Automatic Control*, vol. 40, no. 3, pp. 419–425, 1995.

[16] C. Oliver, "A new core loss model", Ridley Engineering App Note, 2011, Accessed: Oct. 19, 2025. [Online]. Available: https://ridleyengineering.com/images/phocadownload/new%20core%20loss%20model.pdf.

[17] N. S. Samala, "Cogging torque and torque ripple analysis on permanent magnet synchronous machines", 2019, Accessed: Oct. 19, 2025. [Online]. Available: https://www.politesi.polimi.it/handle/10589/148988.

[18] P. H. Mellor, D. Roberts, and D. R. Turner, "Lumped parameter thermal model for electrical machines with tefc design", *IEEE Proceedings B (Electric Power Applications)*, vol. 138, no. 5, pp. 205–218, 1991.

[19] "Arc length and area of a sector", Relation $s = r\theta$, Accessed: Oct. 19, 2025, LibreTexts, 2022, Accessed: Oct. 19, 2025. [Online]. Available: https://courses.lumenlearning.com/ccbcmd-math-1/chapter/arc-length-and-area-of-sector/.

[20] T. P. Lillicrap et al., "Continuous control with deep reinforcement learning", *arXiv:1509.02971*, 2015.

[21] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network", in *Advances in Neural Information Processing Systems (NIPS 1)*, 1988.

ENDURE: Ensemble Based Robust Reinforcement Learning with Reduced Sample Complexity

Sarra Alqahtani
Department of Computer Science
Wake Forest University
Winston-Salem, United States of America
e-mail: alqahtas@wfu.edu

***Abstract*—Intelligent systems deployed in the real world must be both *robust* to distributional shifts and *efficient* in how they learn from interaction. Reinforcement Learning (RL) has delivered strong results in controlled settings, but practical adoption is often limited by high sample costs and brittle behavior on rarely seen states. We present ENDURE, an ensemble-based method for model-free RL that targets intelligent-systems requirements: reliability, data efficiency, and simple integration into existing stacks. ENDURE reuses policies saved from a *single* training run and selects a compact, diverse subset via an *on-policy state diversity* metric that requires no extra environment interaction. At execution time, ENDURE applies a *risk-aware voting* rule that chooses the action with the lowest estimated short-horizon failure probability, improving safety without retraining. Across benchmark control tasks, ENDURE reaches optimal performance with up to $10\times$ fewer samples (CartPole) and about $2\times$ fewer samples (InvertedPendulum-v2) than single-policy baselines, while maintaining strong behavior on underrepresented states. We discuss engineering considerations—guardrails, metrics, and integration points making ENDURE a practical component for real-world intelligent systems.**

Keywords*: *Reinforcement Learning*; *Ensemble Methods*; *Robust Control*; *Risk Estimation*; *Sample Efficiency.

I. INTRODUCTION

Modern intelligent systems increasingly rely on autonomous decision-making under uncertainty. While RL offers a principled framework for sequential control, two constraints frequently hinder deployment: (i) *sample efficiency*: collecting interactions is slow, costly, or risky—and (ii) *robustness*: policies can degrade on states that were rare during training. For intelligent systems operating in dynamic environments, these constraints translate directly into engineering risks, higher operating costs, and difficult validation/assurance cycles.

Deep RL has demonstrated impressive results in games and control [1]–[3], yet many approaches achieve peak performance only after large-scale interaction budgets and may exhibit brittle behavior on distributional edges [4]. Ensemble methods are a natural fit for the intelligent-systems goal of reliability: different policies often specialize in different regions of the state space, and combining them can improve coverage [5]–[8]. However, training many agents in parallel or sequence substantially *increases* sample and compute budgets at odds with the efficiency imperative.

This paper introduces ENDURE, an ensemble approach designed around intelligent-systems constraints. Instead of training multiple agents, ENDURE treats *policies saved across a single training run* as a pool of candidates. We select a small, diverse subset using an *on-policy state diversity* criterion that leverages where each snapshot tends to operate well, avoiding additional rollouts. At execution time, we replace majority/averaging with a *risk-estimation voting* rule: for each candidate policy, estimate the short-horizon probability of failure from the current state and choose the least risky action. The result is a drop-in mechanism that increases robustness and reduces interaction cost, while fitting standard RL stacks and metrics.

From an engineering standpoint, ENDURE emphasizes: (a) *resource awareness*—no extra environment interaction to build the ensemble; (b) *risk-aware control*—an explicit failure-probability proxy used at decision time; and (c) *measurable outcomes*—sample complexity reduction, coverage of hard states, and simple integration points for guardrails and human-in-the-loop review.

Contributions. This work makes the following contributions to intelligent systems design and operation:

1) **Single-run ensembles for data efficiency.** We form diverse policy ensembles by reusing checkpoints from a *single* training run, avoiding the interaction and compute overhead of multi-agent training.
2) **On-policy state diversity for selection.** We introduce a practical diversity metric that identifies complementary policies based on their on-policy state distributions, guiding ensemble selection without additional rollouts.
3) **Risk-aware voting at inference.** We propose a failure-probability–based action selector that prefers low-risk actions at run time, improving robustness on underrepresented states with minimal overhead.
4) **Evidence on control benchmarks.** On CartPole and InvertedPendulum-v2, ENDURE attains optimal performance with up to $10\times$ and about $2\times$ fewer samples, respectively, than single-policy baselines, illustrating readiness for resource-constrained intelligent systems.

Paper organization. Section II provides background and notation. Section III describes ENDURE, including on-policy state diversity (Section III-A1) and risk-estimation voting (Section III-B) with precise labeling and interaction accounting. Section IV presents experiments and discusses the validation budget, checkpoint-frequency sensitivity, and robustness evaluation protocols. Section V concludes.

 ISBN: 978-1-68558-351-4

II. Background

RL studies how an intelligent *agent* learns to make a sequence of decisions by interacting with an *environment*. At each step, the agent observes the current *state* (what the world looks like), takes an *action* (what to do next), and receives a *reward* (a number that indicates how desirable the outcome was). Over time, the agent adjusts its *policy*—its rule for choosing actions from states—to maximize the total (discounted) reward it expects to collect in the future.

A. Core Concepts (Plain Language)

State. The information the agent uses to decide (e.g., the pole angle and cart position in CartPole).
Action. The control the agent applies (e.g., push cart left or right; set a continuous torque).
Reward. Immediate feedback that encodes the task goal (e.g., $+1$ for each balanced time step).
Policy. A mapping from states to actions. It can be a table, a set of rules, or a neural network.
Episode/Trajectory. One run from a start state until termination (success, failure, or timeout).
Return. The total future reward from now onward; commonly written $G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k+1}$ with $0 \leq \gamma < 1$ a discount factor.
Value / Q-Value. Predict how good a state (or state–action) is, in terms of expected return, if you keep following the current policy.

Formally, many RL problems are modeled as a Markov Decision Process (MDP) with states $\mathcal{S}$, actions $\mathcal{A}$, transition dynamics $P(\cdot \mid s, a)$, reward function $R(s, a)$, and discount γ. The objective is to find a policy $\pi(a \mid s)$ that maximizes expected return.

B. Why RL Is Powerful—but Also Costly and Brittle

RL excels when it is hard to write explicit rules but easy to evaluate outcomes (games, control, operations). However, two practical issues often limit real-world use:
Sample complexity. Learning a good policy can require millions of interactions. In physical systems (robots, vehicles) each sample is slow or expensive.
Generalization and robustness. Policies learn from the states they *happen* to visit during training. When they later face unusual or infrequent situations, performance can drop sharply. In RL, we say the policy is strong on its *on-policy* states (the distribution it visits while acting) but may be weak on underrepresented parts of the state space.

C. How Ensembles Help in RL

Ensembles combine multiple policies and choose an action by "voting." In supervised learning this improves accuracy and robustness. In RL, ensembles can also broaden coverage: different policies tend to specialize in different regions of the state space or phases of the task. The challenge is that training many separate RL agents makes sample complexity even worse.

Our work avoids that cost by reusing *policies from a single training run* (e.g., checkpoints saved at different times). These snapshots already differ in what they have learned and where they are strong, giving us diversity without extra environment interactions.

D. On-Policy State Diversity (Intuition)

To select a small but useful subset of snapshot policies, we measure how different their *on-policy* state visitations are. Intuitively, if Policy A routinely visits states that Policy B rarely sees—and vice versa—then combining them should cover more situations. We call this property *on-policy state diversity* and use it to guide ensemble selection before we spend any budget evaluating candidates online.

E. Risk Estimation for Safer Action Selection

Even within an ensemble, we still need to pick one action at each step. Majority or averaging can work, but they ignore *risk*. We instead estimate, for each candidate policy, the probability that taking its suggested action will lead to failure within a short horizon (e.g., the next H steps). The ensemble then chooses the *least risky* action. Practically, we train a lightweight predictor that maps the current state to an estimated failure probability for each policy. This *risk-estimation voting* prefers actions that are both competent and less likely to unravel, improving robustness without retraining.

F. How This Connects to Our Contributions

This paper operationalizes the above ideas as ENDURE. We (i) harvest diverse policies from one training run, (ii) score candidate ensembles using on-policy state diversity to minimize evaluation budget, and (iii) deploy a risk-aware voting rule that picks the action with the lowest predicted short-term failure probability. The result is an ensemble that achieves competitive (often optimal) performance with far fewer samples and better behavior on hard or underrepresented states.

III. ENDURE: Robust Reinforcement Learning Through Ensemble

This work proposes ENDURE, a robust RL framework for learning and selecting an ensemble of diverse policies from *a single training run*. ENDURE has two components: (1) Ensemble policy selection via on-policy state diversity, and (2) Ensemble action selection via risk estimation voting. The ensemble policy selection component finds the best subset of K policies from the entire set of N policies collected during one training run. The policy selection is based on a metric called *on-policy state diversity* to be described in Section III-A1. The voting technique is developed based on risk estimation of each policy failing the task in the next H timesteps as described in Section III-B.

A. Ensemble Policy Selection via On-Policy State Diversity

The RL agent's policy changes throughout training and its optimality in a state is heavily correlated with the frequency the agent visits that state [4]. Motivating by this fact, we split the training process into N different training periods creating N different policies. Given the N policies for the agent throughout training, we here consider the problem of policy selection, or

selecting K out of N policies to be in the final ensemble. In supervised learning, the problem of selecting K models to be in an ensemble from a set of N possible models is known as *model selection*. It has been shown in supervised learning settings that a necessary and sufficient criteria for an ensemble to perform better than its counterpart individual models is for the individual models to be both accurate and diverse [9]. In RL settings, we empirically found that accuracy, as defined by high reward, is insufficient for producing a good ensemble. Figure 1a shows the evaluation results of ensemble RL agents made up of the best policies, according to their rewards. However, those policies are not optimal in comparison to the randomly selected policies in Figure 1b. Further experiments in Section IV show that these policies are not optimal.

It is more challenging to quantify and optimize *diversity* in RL settings than in the supervised learning settings [6]. Prior works follow an ad-hoc approach. Particularly, they set $N{=}K$, and train these K policies end-to-end to include in the ensemble, using methods to ensure the resulting K policies are diverse [5]–[8]. However, in practice there is no guarantee of diversity once these policies are trained. Since these approaches only have K policies to choose from after training, it is infeasible to choose a new set of policies, should this be the case. Supervised learning researchers have long attempted to solve this problem. Model selection methods have been developed to pick the most accurate and diverse classifiers through examining model performance on a validation set, and the difference in model outputs within this validation set [10]–[12]. In RL settings, an analogous approach would require rolling out the potential ensemble for some number of episodes to gather information about the reward the ensemble achieves. These rollouts would contribute heavily to the sample complexity. We show in Figure 1b that, without considering the additional rollout sample complexity, we can find high-performing ensembles by simply testing ensembles with random assortments of policies. However, we want to minimize the amount of possible ensembles we have to test for reward. To this end, we develop a metric, *on-policy state diversity*, to measure the diversity of individual policies in a more efficient way.

1) On-Policy State Diversity: In [4], an agent's on-policy states are defined as the states which the agent encounters during a standard training episode, and off-policy states as those states which are not encountered as frequently. They found that the percentage of the maximum return the policy achieves, starting in a state s, is correlated with how often the agent visits s [4]. We aim to choose a subset of $K{\leq}N$ policies which have diverse expertise within the set of possible environment states.

Given a set $S_i = \{s_{i,1}, \ldots, s_{i,T}\}$ for each policy i, containing a small ($T \leq 5000$) set of vector-valued on-policy states encountered by that policy, we summarize each S_i by its empirical mean μ_i and elementwise standard deviation σ_i:

$$\mu_i = \frac{1}{T}\sum_{t=1}^{T} s_{i,t}, \qquad \sigma_i = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(s_{i,t}-\mu_i)^2}. \tag{1}$$

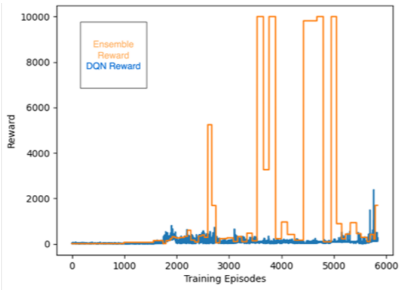

(a) Ensemble vs DQN reward with the 4 best policies according to their rewards.

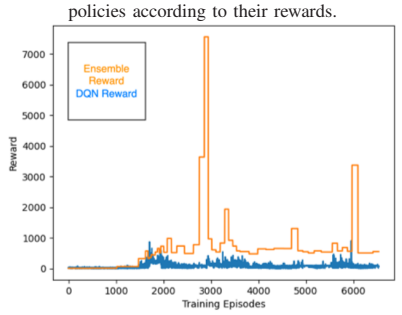

(b) Ensemble vs DQN reward with four random policies in the ensemble.

Figure 1. Performance of DQN [13] and ensemble policies across a training run in the CartPole environment. Plots differ in ensemble selection mechanisms.

We then define the pairwise on-policy state diversity metric between policy i and policy j as a normalized distance:

$$\lambda_{i,j} = \frac{1}{d}\left(\|\mu_i - \mu_j\|_2^2 + \|\sigma_i - \sigma_j\|_2^2\right), \tag{2}$$

where d is the state dimension. This yields a fast-to-compute proxy for differences in typical visited states and spread.

We then define on-policy state diversity of a sampled ensemble of policies as the average pairwise state diversity among the individual policies:

$$\lambda(\mathcal{E}) = \frac{2}{K(K-1)}\sum_{i\in\mathcal{E}}\sum_{\substack{j\in\mathcal{E}\\ j>i}}\lambda_{i,j}. \tag{3}$$

This metric can be computed across large S_i within seconds after precomputing (μ_i, σ_i) for each checkpoint. Computing all pairwise $\lambda_{i,j}$ costs $O(N^2 d)$ time, and evaluating $\lambda(\mathcal{E})$ for a candidate ensemble is $O(K^2)$ given the pairwise table.

Concrete values and feasibility. In our continuous-control experiments (Section IV), we save checkpoints uniformly across training: for CartPole we train for 1,000,000 timesteps and use $N = 20$ checkpoints, and for InvertedPendulum-v2 we train for 250,000 timesteps and use $N = 10$ checkpoints. We use ensemble size $K = 4$. For CartPole, full enumeration yields $\binom{20}{4} = 4845$ candidate ensembles, which is feasible to score by Equation 3. We then evaluate only a subset under a validation budget of M ensembles (below).

We propose to perform policy selection in a way that minimizes the number of ensemble rollouts we need to perform. Assume we have a validation budget such that we can test the

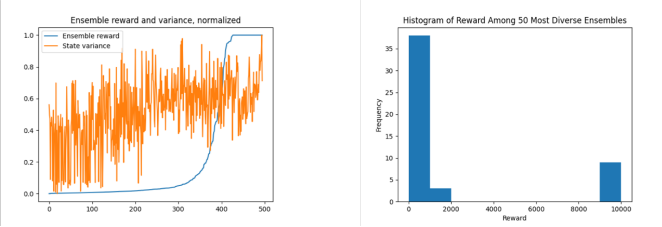

(a) Average pairwise on-policy state diversity versus reward for 500 candidate ensembles.

(b) Histogram of the reward of the 50 most diverse ensembles according to on-policy state diversity.

Figure 2. Experiments conducted with the same training run in the CartPole environment.

performance of M ensembles. We sort all $\binom{N}{K}$ ensembles by their mutual on-policy state diversity values and choose the top M ensembles with respect to this metric. We then test these candidate ensembles in the environment, and choose the best ensemble according to the observed reward. Figure 2b shows how many 'good' ensembles we can expect from this method in the CartPole environment. Figure 2b shows that in the top 50 ensembles with respect to $\lambda(\mathcal{E})$ (Equation 3), 10 ensembles achieve what we consider a high reward, whereas the remaining 40 achieve a reward similar to the individual policies within the ensemble. In these experiments, we use $M = 50$ as the validation budget.

B. Ensemble Action Selection via Risk Estimation Voting

Given an ensemble with K policies, we aim to develop a method for selecting one action among the selected actions by the ensemble policies which will be fed into the environment. Specifically, given a set of actions $A_t = \{\pi_i(s_t) \mid 0 \leq i < K\}$, we wish to develop a function f such that $\hat{a}_t = f(A_t)$. A standard voting scheme is the majority vote approach, which takes an action according to:

$$\hat{a}_t = \arg\max_{a \in \mathcal{A}} \left(\sum_{i=0}^{K-1} \mathbb{1}[A_t[i] = a] \right) \tag{4}$$

where $\mathbb{1}$ is the indicator function. This voting scheme only works well when a_t is discrete. In the case of a continuous, vector-valued action space, we can use average voting. In this scheme, the action is selected as

$$\hat{a}_t = \frac{1}{K} \sum_{i=0}^{K-1} A_t[i]. \tag{5}$$

As an alternative to these approaches, we propose risk estimation voting. The intuition behind this approach is to pick the action from the policy that is the least likely to fail the task, from the current state.

Risk definition and labels. We assume failure and success are well-defined via the environment termination condition (e.g., pole falls, pendulum diverges), distinct from time-limit truncation. For checkpoint policy i, we define the short-horizon failure probability conditioned on *(state, action)*:

$$g_i(s,a) = \Pr\Big(\text{failure occurs within the next } H \text{ steps} \,\Big|\, s_t = s,\ a_t = a,\ \pi_i\Big). \tag{6}$$

We obtain binary risk labels from trajectories collected during training *while checkpoint policy i is active* (or within its training window), introducing no additional environment interaction beyond the single training run. For each logged transition $(s_{i,t}, a_{i,t})$, we label

$$y_{i,t}^{(H)} = \mathbb{1}\Big[\text{failure occurs in } \{t+1, \ldots, t+H\} \text{ within the same episode}\Big]. \tag{7}$$

This labeling uses only episode termination signals already produced during the single training run and does not use data from later checkpoints to label earlier ones (avoiding "future-data leakage" that could inflate sample-efficiency claims).

Suppose we have an oracle function, g_i, which computes the probability that policy i will fail the task in the next H timesteps from state s_t. Then, risk estimation voting can be derived as

$$\hat{a}_t = A_t \left[\arg\min_{0 \leq i < K} \{g_i(s_t)\} \right]. \tag{8}$$

Action-conditioned risk voting (implementation). To address action-dependence, we use the action proposed by policy i, $a_t^{(i)} = \pi_i(s_t)$, and select:

$$\hat{a}_t = a_t^{(i^\star)}, \quad i^\star = \arg\min_{0 \leq i < K} \hat{g}_i\left(s_t, a_t^{(i)}\right), \tag{9}$$

where $\hat{g}_i$ approximates g_i using supervised learning on (s,a) and labels $y_{i,t}^{(H)}$.

Given the risk estimation voting scheme, all we need to do is approximate the function g_i, such that $\hat{g}_i(s_t) \approx g_i(s_t)$ $\forall s_t$. Since s_t can potentially take infinite values, [14] chooses to approximate g_i through a neural network, trained with supervised learning. We follow their methodology, which involves first collecting a dataset $\{(s_t, g_i(s_t))\}$ that contains around $100,000$ samples. This dataset can be collected from observing the agent as it trains. We use a neural network with two fully-connected layers of 64 neurons each, and train it to predict $g_i(s_t)$ from s_t.

Overhead note. For larger systems, training a separate risk estimator per checkpoint can be expensive; in practice, we can train risk estimators only for the K selected ensemble policies rather than all N checkpoints, and we can reuse the same logged data collected during training (no extra interaction).

IV. Experiments

We perform experiments in discrete and continuous environments. First, we analyze risk estimation ensemble voting in a gridworld environment. Then, we examine the effect of risk estimation voting combined with on-policy state diversity model selection in continuous control environments CartPole and InvertedPendulum-v2.

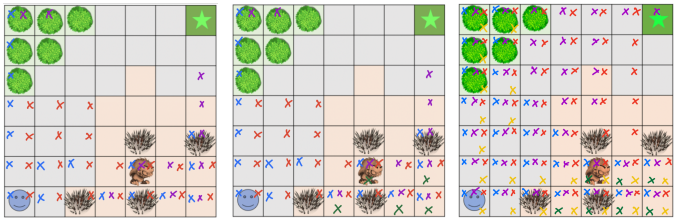

(a) Three phases of the RL agent. The blue, purple, and red x's are for the beginning, middle, and end of training.

(b) Three phases of the RL agent along with an ensemble (green) of the three; the ensemble uses risk-estimation voting.

(c) Three phases of a suboptimal RL agent, its risk-estimation ensemble (green), and its majority-vote ensemble (yellow).

Figure 3. Performance of individual and ensemble RL policies in the gridworld environment. X's mark states from which the corresponding policy fails to complete the task.

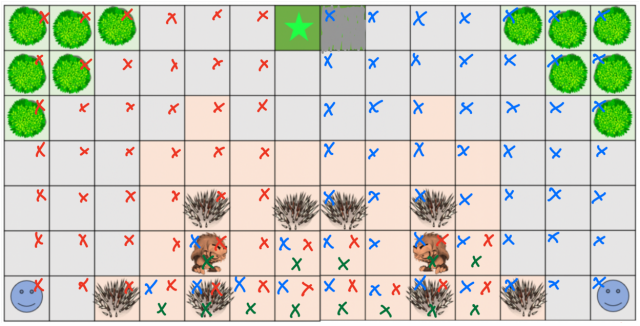

Figure 4. Two RL agents, the blue one starts in the bottom-left, the red one starts in the bottom-right, and their ensemble (green), using risk estimation voting where x's represent states from which the corresponding agent fails to complete the task.

A. Discrete Environments

We adapt the gridworld from [15], shown in Figure 3. In this environment, the agent starts in the bottom-left corner, and the goal is to navigate to the top-right corner using the actions up, left, down, and right. If the agent gets too close to the monster, or if it takes too many timesteps to reach the goal state, the agent fails and receives a large negative reward.

In Figure 3(a), we show the adeptness of three phases of an RL agent. We train an RL agent for a fixed number of timesteps in this environment and save its weights at the beginning, middle, and end of training. In Figure 3, we denote the states from which an agent fails to solve the task with an appropriately colored 'X'. We see that the agent at the beginning, middle, and end of training has different sets of states from which it can solve the task.

In Figure 3(b), we include the performance of an ensemble agent using risk estimation voting. Note that instead of using a neural network to approximate g_i in this environment, we used a tabular representation, since the gridworld environment has a finite set of possible states. We see that the ensemble agent only fails the task from the states in which all individual agents also fail the task, indicating that the ensemble can correctly choose which agent to act based on individual specialization. Figure 3(c) shows the same experiment, but with a worse base RL agent and a majority vote baseline. We observe that there are many states from which the individual policies fail the task, but the risk estimation ensemble method is able to successfully complete the task from many of these states. This shows that the risk estimation ensemble policies have the potential to perform much better than any individual policy. We also observe that the majority vote policy does better than the individual policies, but still has many states from which it fails the task. This shows that the improvement in the policy is due in part to risk estimation voting, not just the nature of ensembles.

To further visualize the benefit of ensembles when policies specialize in different states, we created an alternate gridworld, which is the original gridworld concatenated with a flipped version of itself (Figure 4). We trained one RL agent (blue) which starts in the bottom-left corner and one agent which starts in the bottom-right (red). When we combine these two agents in an ensemble with risk estimation, the ensemble agent can solve the task from all states except a handful from which the task is impossible.

B. Continuous Environments

Next, we examine the performance of the risk estimation ensemble in two environments with continuous state spaces, CartPole and InvertedPendulum-v2 [3]. Both these environments could continue forever if the control policy is good enough, so we cut the episodes off at 10,000 timesteps. The agent receives a reward of 1 for every timestep it is left standing, so the maximum reward is also 10,000.

We trained the agents for both environments using Deep Q-Learning (DQN) [13] and Deep Deterministic Policy Gradient (DDPG) [16], but we present only the optimal performance of each of them. The agent for the CartPole environment was trained with DQN for 1,000,000 timesteps; the DDPG agent was trained on InvertedPendulum-v2 for 250,000 timesteps. For testing, we perform ensemble policy selection at set intervals and run the ensemble with risk estimation voting, to give an idea of what kind of ensemble performance is possible at various numbers of training timesteps.

Checkpointing parameters. In these experiments, we save $N = 20$ checkpoints for CartPole (every 50,000 timesteps) and $N = 10$ checkpoints for InvertedPendulum-v2 (every 25,000 timesteps). We use ensemble size $K = 4$. We use a validation budget of $M = 50$ candidate ensembles ranked by on-policy state diversity (Section III-A1) and selected by online evaluation.

Figure 5(a) shows the results for the CartPole environment. We show the average DQN agent reward, the maximum reward that the DQN agent ever achieves, and the reward of the best ensemble, which our algorithm finds. The risk estimation ensemble can achieve the maximum reward of 10,000 after only 100,000 timesteps of training, which is about a $10\times$ decrease in sample complexity (measured against the single-run training budget).

Figure 5(b) shows the results for the same experiment in the InvertedPendulum-v2 environment, except with DDPG RL

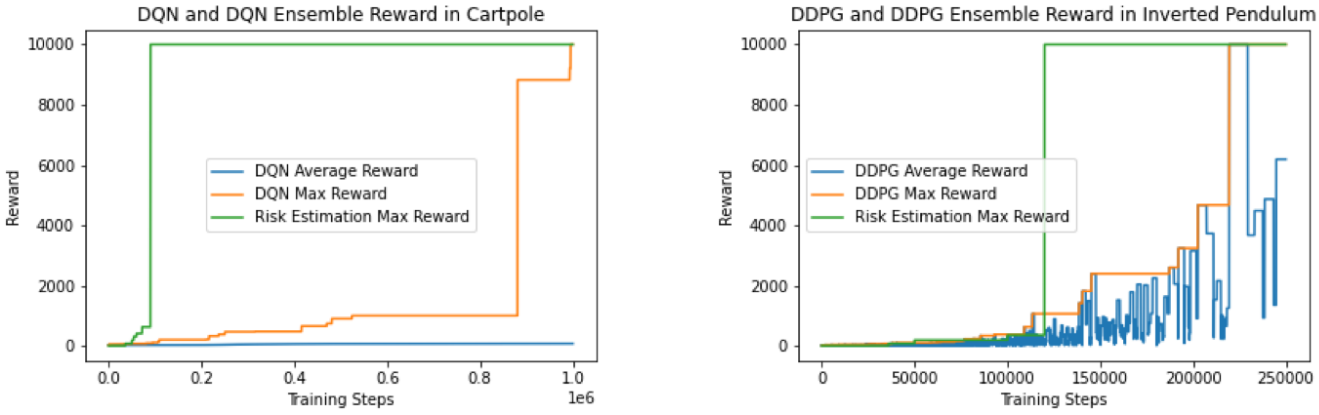

(a) CartPole (DQN) **(b)** InvertedPendulum (DDPG)

Figure 5. Average single-policy performance, best checkpoint, and best ensemble. (a) CartPole; (b) InvertedPendulum.

training algorithm. We see that the risk estimation ensemble reaches the maximum reward at around 125,000 timesteps, about a $2\times$ decrease in sample complexity (against the single-run training budget).

Checkpoint frequency sensitivity (engineering consideration). Ensemble quality depends on N and checkpoint spacing. If checkpoints are too frequent, adjacent policies may be redundant and reduce effective diversity; if too sparse, the pool may miss complementary behaviors. Practical guidance is to save checkpoints across distinct learning phases and optionally use non-uniform spacing (denser early training, sparser late training) to increase candidate diversity at fixed N.

V. Conclusion

We presented ENDURE, an ensemble method for RL tailored to intelligent-systems constraints: robustness, efficiency, and ease of integration. By harvesting policy snapshots from a *single* training run, selecting a compact subset via on-policy state diversity, and applying a risk-aware voting rule at execution time, ENDURE improves behavior on underrepresented states while substantially reducing the interactions needed to reach optimal performance. On standard control tasks, ENDURE achieves up to $10\times$ (CartPole) and about $2\times$ (InvertedPendulum-v2) reductions in sample complexity relative to single-policy baselines.

For practitioners, ENDURE is a drop-in augmentation to existing RL stacks that aligns with common assurance practices: it exposes explicit selection criteria, admits guardrails (e.g., thresholds, escalation), and supports measurement via standard control metrics. Future work will study deployment aspects including uncertainty calibration for the risk estimator, integration with formal checks for safety constraints, extension to multi-agent settings, and evaluation on hardware-in-the-loop or edge scenarios where interaction budgets are tight.

References

[1] D. Silver et al., "Mastering the game of Go with deep neural networks and tree search", *Nature*, vol. 529, no. 7587, pp. 484–489, 2016. DOI: 10.1038/nature16961.

[2] O. Vinyals et al., "Starcraft II: A new challenge for reinforcement learning", arXiv preprint arXiv:1708.04782, 2017, arXiv: 1708.04782. [Online]. Available: https://arxiv.org/abs/1708.04782.

[3] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.

[4] J. Uesato et al., "Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures", arXiv preprint arXiv:1812.01647, 2018, arXiv: 1812.01647. [Online]. Available: https://arxiv.org/abs/1812.01647.

[5] K. Lee, M. Laskin, A. Srinivas, and P. Abbeel, "Sunrise: A simple unified framework for ensemble learning in deep reinforcement learning", in *Proceedings of the 38th International Conference on Machine Learning*, M. Meila and T. Zhang, Eds., ser. Proceedings of Machine Learning Research, vol. 139, PMLR, 2021, pp. 6131–6141.

[6] R. Saphal, B. Ravindran, D. Mudigere, S. Avancha, and B. Kaul, "Seerl: Sample efficient ensemble reinforcement learning", arXiv preprint arXiv:2001.05209, 2020, arXiv: 2001.05209. [Online]. Available: https://arxiv.org/abs/2001.05209.

[7] Q. He, C. Gong, Y. Qu, X. Chen, X. Hou, and Y. Liu, "Mepg: A minimalist ensemble policy gradient framework for deep reinforcement learning", arXiv preprint arXiv:2109.10552, 2021, arXiv: 2109.10552. [Online]. Available: https://arxiv.org/abs/2109.10552.

[8] X. Chen, L. Cao, C. Li, Z. Xu, and J. Lai, "Ensemble network architecture for deep reinforcement learning", *Mathematical Problems in Engineering*, vol. 2018, pp. 1–6, 2018. DOI: 10.1155/2018/2129393.

[9] T. G. Dietterich, "Ensemble methods in machine learning", in *Multiple Classifier Systems*, ser. Lecture Notes in Computer Science, vol. 1857, Berlin, Heidelberg: Springer, 2000, pp. 1–15.

[10] R. Caruana, A. Niculescu-Mizil, G. Crew, and A. Ksikes, "Ensemble selection from libraries of models", in *Proceedings of the Twenty-First International Conference on Machine Learning*, ser. ICML '04, New York, NY, USA: Association for Computing Machinery, 2004, p. 18. DOI: 10.1145/1015330.1015432.

[11] T. Pang, K. Xu, C. Du, N. Chen, and J. Zhu, "Improving adversarial robustness via promoting ensemble diversity", arXiv preprint arXiv:1901.08846, 2019, arXiv: 1901.08846. [Online]. Available: https://arxiv.org/abs/1901.08846.

[12] X. Dong, Z. Yu, W. Cao, Y. Shi, and Q. Ma, "A survey on ensemble learning", *Frontiers of Computer Science*, vol. 14, pp. 241–258, 2020. DOI: 10.1007/s11704-019-8208-z.

[13] V. Mnih et al., "Playing atari with deep reinforcement learning", arXiv preprint arXiv:1312.5602, 2013, arXiv: 1312.5602. [Online]. Available: https://arxiv.org/abs/1312.5602.

[14] E. A. Abolfathi, J. Luo, P. Yadmellat, and K. Rezaee, "Coachnet: An adversarial sampling approach for reinforcement learning", arXiv preprint arXiv:2101.02649, 2021, arXiv: 2101.02649. [Online]. Available: https://arxiv.org/abs/2101.02649.

[15] J. van der Waa, J. van Diggelen, K. van den Bosch, and M. Neerincx, "Contrastive explanations for reinforcement learning in terms of expected consequences", arXiv preprint arXiv:1807.08706, 2018, arXiv: 1807.08706. [Online]. Available: https://arxiv.org/abs/1807.08706.

[16] T. P. Lillicrap et al., "Continuous control with deep reinforcement learning", in *International Conference on Learning Representations (Poster)*, OpenReview.net, 2016.

 ISBN: 978-1-68558-351-4

Integrating Intuitive Interaction and Large Language Model Guidance for Efficient 3D Annotation

Haruya Ishigami†‡, Kenji Iwata† and Yutaka Satoh†‡
†: National Institute of Advanced Industrial Science and Technology, Tsukuba, Japan
‡: University of Tsukuba, Tsukuba, Japan
e-mail: {ishigami.h2002, kenji.iwata, yu.satou} @aist.go.jp

***Abstract*— In this paper, we propose a method to intuitively annotate and analyze 3D models of disaster scenes by integrating 3D Gaussian Splatting (3DGS) and a Large Language Models (LLM). However, most existing systems are limited to visual reproduction, and it remains challenging to sufficiently utilize the generated data for practical problem-solving in the real world. Using aerial data captured by drones, we construct a two-layer structure that combines high-quality visual representation through 3DGS with surface mesh geometry. This framework enables the intuitive and flexible placement of annotations, such as panels and lines, at intended positions in 3D space in response to user operations. Furthermore, we developed a method where a LLM analyzes disaster risks based on user instructions and visual information, and subsequently displays the analysis results at arbitrary locations within the 3D environment. By fusing intuitive information sharing with AI-driven analysis support, the proposed system enhances the efficiency and quality of decision-making in disaster management.**

Keywords-3D Gaussian Splatting; Annotation System; Large Language Model.

I. INTRODUCTION

Recently, the emergence of innovative 3D scene reconstruction technologies, exemplified by 3DGS [1], has enabled the rapid generation of high-quality and high-density 3D models for large-scale and complex environments. Building on this development, numerous efforts are being made across various fields to high-fidelity reconstruct and visualize real-world spaces using images acquired by drones and other platforms [2]. However, most existing systems primarily focus on visual reproduction and viewing, and it remains challenging to sufficiently utilize the generated 3D data for practical decision-making and problem-solving in the real world [3].

Data utilization is particularly indispensable in disaster response scenarios to rapidly assess damage and consider rescue plans or evacuation routes. While the use of 3D models via drones has proven effective for gaining a bird's-eye view of an entire site [4], simple visualization is insufficient to support complex decision-making in environments characterized by high uncertainty. To achieve rapid and accurate information sharing, a mechanism is required to directly annotate crucial information—such as hazardous locations, navigable paths, and damage status—onto the 3D model. Furthermore, to supplement limited human resources on-site and ensure reliable rescue operations, it is considered highly effective to combine objective situational analysis and decision-making support provided by AI [5].

3DGS is well-suited for the purpose of this study due to its ability to generate models with high speed and excellent visual quality. However, 3DGS is a volumetric rendering method that represents each point as 3D Gaussian distributions, which are then projected into screen space to synthesize images. Consequently, it lacks explicit geometric boundary information, such as distinct surfaces and edges, presenting a technical challenge in that depth information cannot be directly obtained. To associate and display manual user input or AI analysis results at accurate positions in 3D space, obtaining reliable depth information is essential.

In this study, we constructed a system that integrates 3D model generation technology with a LLM to improve operational efficiency and realize smooth information sharing on-site. Specifically, the process begins with capturing site images and constructing a data foundation. Next, mesh information is generated using Surface-Aligned Gaussian Splatting (SuGaR) [6] after initial processing with COLMAP [7]. By leveraging accurate depth information derived from this mesh, we enable intuitive 3D spatial annotation by the user. Furthermore, we implemented advanced analysis support functions by integrating an LLM. This allows the LLM to verbalize damage situations and propose countermeasures based on user-provided information, visual features of the site, and instruction prompts.

The objective of this system is to enhance the quality of decision-making in disaster scenarios by fusing intuitive information sharing with AI-driven analysis support. This approach, which allows for rapid information annotation while seamlessly navigating between bird's-eye and immersive perspectives, is expected to drastically improve information transmission efficiency among rescue teams and related organizations, thereby contributing to more reliable rescue operations.

II. RELATED WORK

3D Gaussian Splatting (3DGS) is a method that takes multi-view images as input, represents and optimizes a 3D point cloud as Gaussian distributions, and generates novel view images with high speed and high precision. It acquires a 3D point cloud and camera parameters from RGB images using COLMAP's Structure-from-Motion [7]. Using this point

 ISBN: 978-1-68558-351-4

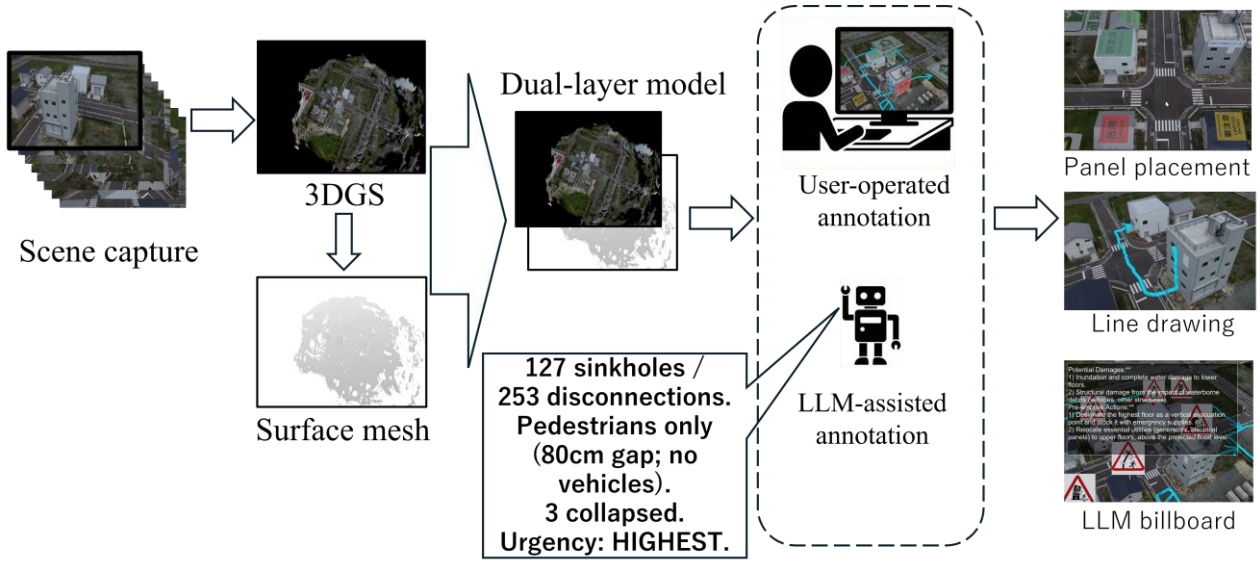

Figure 1. Overview of our proposed LLM-assisted 3D annotation system.
Drone-captured imagery is used to generate a Dual-layer model (integrating 3DGS and SuGaR surface mesh). The system combines user-driven operations (panels/lines) with LLM-assisted analysis (disaster risk/mitigation) to embed spatial information. The resulting annotations are viewable from arbitrary perspectives.

cloud as an initial Gaussian distribution, it optimizes various parameters (3D position, opacity, covariance, and SH coefficients representing color) through a machine learning approach. During optimization, Gaussians are also added or removed. Numerous presentations have been made regarding this method, including applications to diverse fields, such as robotics and VR, as well as speed and accuracy improvements [2].

Generally, 3DGS requires images captured from several dozen different viewpoints. If there are few input images, the shooting and calculation costs are reduced, but there is a problem that rendering quality drops significantly. In response to this, Few-shot View Gaussian Splatting [9] estimates monocular depth from rendered images and uses that information to complement the shape of unobserved regions, achieving high-quality view synthesis even from about three images.

Furthermore, Instant Splat [10] is a method that generates an initial point cloud from a large-scale pre-trained geometric model without relying on SfM. It reconstructs both camera parameters and the scene with high speed and high precision through self-supervised optimization. By using Gaussian-based bundle adjustment, it achieves a speedup of more than 30 times compared to conventional 3DGS while realizing high rendering quality (SSIM) even in sparse-view environments.

On the other hand, while 3DGS enables high-quality image generation, there is a challenge that the optimized Gaussian distribution is not structurally organized, making it difficult to extract a clear mesh structure. Therefore, it is difficult to use the generation results directly for editing or annotation, and additional processing is required to apply it to operations where users specify arbitrary positions in 3D space to add information. To address this problem, SuGaR, proposed by Guédon et al., introduces a regularization term to align Gaussian distributions with the scene surface. Using this alignment, it achieves fast and high-precision mesh extraction via Poisson reconstruction. Furthermore, through a mechanism that simultaneously optimizes the mesh and Gaussians, it enables high-quality rendering in a shorter time compared to conventional Neural SDF-based methods, as well as flexible operations, such as editing, animation, and lighting adjustment.

In recent years, the utilization of drone technology in disaster response has attracted attention, and comprehensive organization of these trends is proceeding. A survey [4] investigated 52 research papers published between 2009 and 2020, classifying drone applications in disasters into four categories: Mapping/Disaster Management, Search and Rescue, Transportation, and Training [8]. The contribution to the mapping field is particularly notable, and its effectiveness for situational awareness and decision-making support at disaster sites has been confirmed. On the other hand, discussions regarding use in post-disaster areas, such as victim identification and medical support, are not yet sufficient and are pointed out as future challenges.

Applications that realize an interactive 3D manipulation environment using pen input and touch operations on mobile devices include Feather [11] and Cozy Blanket [12]. While these applications possess high operability and convenience, they are primarily aimed at 3D content creation and do not support annotation uses for adding and sharing information on 3D models of real spaces.

Research on the systematization of interaction design in Virtual Reality (VR) and Augmented Reality (AR) environments is also progressing [13]. Research is being conducted on immersive systems that allow interactive manipulation [14], such as deforming 3D content generated by 3D Gaussian Splatting, and frameworks that can perform shape changes, color adjustments, and style transfers on 4D scenes [15]. However, many of these focus on model

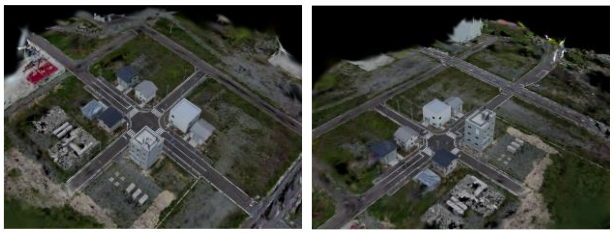

Figure 2. 3DGS

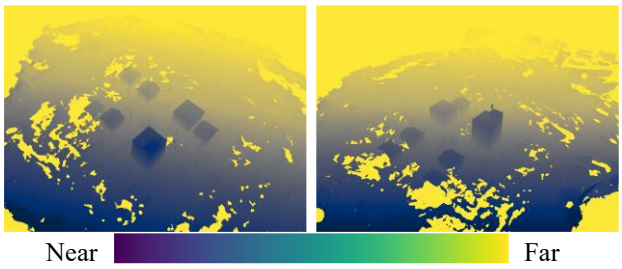

Figure 3. Example of Depth Information Computation (Color map)

generation or interaction design, and there is room for consideration regarding mechanisms to support information addition and sharing on generated 3D models. Therefore, in this study, we propose a method to realize intuitive and accurate annotation operations using 3D model generation technology to address such issues.

III. PROPOSED METHOD

A. System Overview

This system is an interactive annotation tool, and its overview is illustrated in Figure 1. From the input images, a 3DGS model is generated for visual display (Figure 2), and a surface mesh is constructed using SuGaR. Figure 3 shows the result of computing depth values from the surface mesh and assigning a color map to them.Users can freely navigate the 3D scene from arbitrary viewpoints, place pictograms representing disaster risks analyzed by the LLM (Google Gemini), and draw line objects through mouse-drag operations.

B. Determination of Target Coordinates

Since 3DGS models do not explicitly represent surface geometry, a surface mesh generated by SuGaR is introduced to identify the 3D coordinates corresponding to a click position on the screen.

When a user performs an annotation operation, the depth value d corresponding to the input screen coordinates (x, y) is determined based on the surface mesh. The screen coordinates (x, y) and the depth value d are then converted into normalized device coordinates (NDC). The conversion is defined as follows:

$$x_{ndc} = \frac{2x}{ScreenWidth} - 1\,,$$

$$y_{ndc} = \frac{2y}{ScreenHeight} - 1\,, \tag{1}$$

$$z_{ndc} = 2d - 1\,,$$

Here, d is a normalized depth value where 0 is the nearest and 1 is the farthest within the range from the camera's Near Clip to Far Clip.Next, using the camera's projection matrix P, the coordinates (x_c, y_c, z_c) in the camera coordinate system are obtained from the NDC coordinates $(x_{ndc}, y_{ndc}, z_{ndc})$ as follows:

$$\begin{bmatrix} x_c \\ y_c \\ z_c \\ w \end{bmatrix} = P^{-1} \begin{bmatrix} x_{ndc} \\ y_{ndc} \\ z_{ndc} \\ 1 \end{bmatrix} \tag{2}$$

As a correction for perspective transformation, the obtained camera coordinates are normalized using w:

$$x'_c = x_c / w, \qquad y'_c = y_c / w, \qquad z'_c = z_c / w \tag{3}$$

Finally, by using the view matrix V, the world coordinates $(X_w, Y_w, Z_w)$are obtained:

$$\begin{bmatrix} X_w \\ Y_w \\ Z_w \\ w \end{bmatrix} = V^{-1} \begin{bmatrix} x'_c \\ y'_c \\ z'_c \\ 1 \end{bmatrix} \tag{4}$$

Based on the computed world coordinates, annotations are placed at positions in the 3D space corresponding to the user's operations. During this process, to prevent misdetections caused by noise inherent to 3DGS or occlusions such as power lines, offset processing is applied. In addition, for line drawing operations, instead of directly using the depth value corresponding to the simple click position, the median depth value within a neighboring region is used. This approach enables stable target coordinate specification even in the presence of outliers.

C. Panel Placement Method

In the panel placement process (Figure 3), the user first selects the type of panel to be placed, and the panel is then positioned at the user-specified location via a click operation. Since the orientation of the panel must be adjusted for each specified location, the proposed system determines the orientation of the panel object based on the local structure of the 3D scene around the target coordinates and the user's viewpoint at the time of annotation input.

To uniquely define an orientation in 3D space, two axes are required: the Z-axis (forward vector) and the Y-axis (upward vector). In this system, the Z-axis direction is estimated by applying principal component analysis (PCA) to infer the orientation of the surface in the vicinity of the target coordinates, while the Y-axis direction is determined using the upward vector of the camera at the time of annotation input.

By defining the Z- and Y-axis directions in this manner, the panel object is naturally aligned with the surface geometry, adapting to local surface irregularities and inclinations, thereby improving visual consistency. In addition, using the camera's upward direction for the Y-axis preserves the expected notion of the "top" of the panel relative to the user's

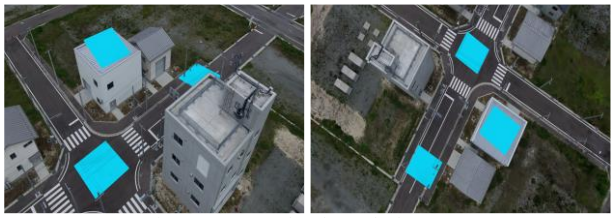

Figure 4. Example of Panel Placement

Figure 5. Example of Line Drawing

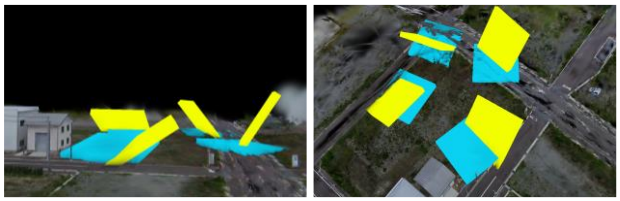

Figure 6. Comparison of Panel Placement Methods

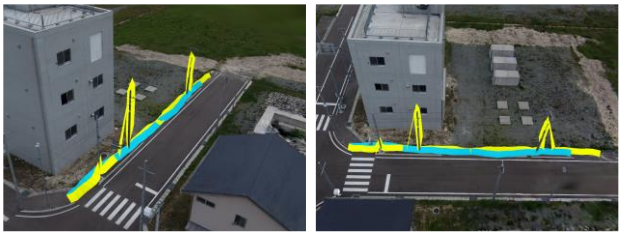

Figure 7. Comparison of Line Drawing Methods

viewpoint at the time of annotation, enabling panel placement that is consistent with the user's operational intent.

To obtain the surface normal of the 3D scene around the target coordinates, the user specifies a single point in screen space. Screen-space coordinates within an n-pixel neighborhood around the selected point are then collected. The corresponding depth values are retrieved and converted into 3D coordinates. Finally, principal component analysis (PCA) is applied to these local 3D points to estimate a normal vector representing the local surface orientation.

D. Line Drawing Method

In the line drawing process (Figure 5), coordinates sampled during the user's drag operation are connected by cylindrical segments to form a polyline. If input coordinates are sampled every frame, excessive sampling density causes the system to overly react to minor surface irregularities (e.g., small bumps on the ground), resulting in unintended zigzag trajectories. Conversely, excessive smoothing may lead to the loss of important terrain features such as steps or steep slopes.

To address this trade-off, sampling is triggered based on a fixed movement distance threshold. This approach preserves essential terrain characteristics while reducing noise caused by small surface irregularities, enabling smooth line drawing. In addition, to prevent drawn points from being embedded within the model surface, a constant offset is applied in the direction toward the camera.

When drawing lines, objects located in the foreground may interfere with the intended placement on the target surface. To mitigate this issue, in addition to the depth value corresponding to the user-specified screen coordinate (x, y), depth values of neighboring screen coordinates are also computed. By using the median of these depth values, the system enables stable line placement on the intended target surface even when small foreground objects, such as utility poles, are present.

E. Integration of LLM Analysis and Human Operation

In the LLM-assisted functionality, risk factors and mitigation measures related to structures displayed on the LLM interface are generated at arbitrary timings specified by the user. Although LLMs are capable of advanced situational reasoning, they are not well-suited for precise 3D localization. Therefore, the proposed system adopts a process in which the LLM generates semantic information and the user explicitly specifies its placement location.

The generated information is placed as billboard-style pictograms in the 3D space. When a pictogram is clicked, detailed analysis results are displayed in a pop-up window, allowing the user to interactively review the generated content.

IV. EXPERIMENT

We conducted an evaluation using disaster simulation data obtained at the Fukushima Robot Test Field [8]. In the comparison of panel placement methods (Figure 6), principal component analysis (PCA) was employed to determine the surface normal direction, and panel placement was performed while varying the number of points used for normal estimation. Specifically, 9 points (yellow panels) and 225 points (light blue panels) were used. Compared to the yellow panels, the light blue panels were less affected by local surface irregularities in the surface mesh at the placement location and were more consistently aligned with the ground plane, indicating more stable placement.

In the comparison of line drawing methods (Figure 7), using the click position coordinates directly (yellow line in Figure 7) resulted in interference with obstacles such as utility poles. In contrast, when using the median of depth values (light blue line in Figure 7), stable line drawing was achieved along the ground context without interference. These results confirm that lines can be accurately drawn along object surfaces, such as the ground and building walls.

Regarding the evaluation of the LLM-assisted functionality, the objective of this experiment was to integrate and visualize analysis results obtained from Gemini within the 3D space, and to verify its operation. For each target structure,

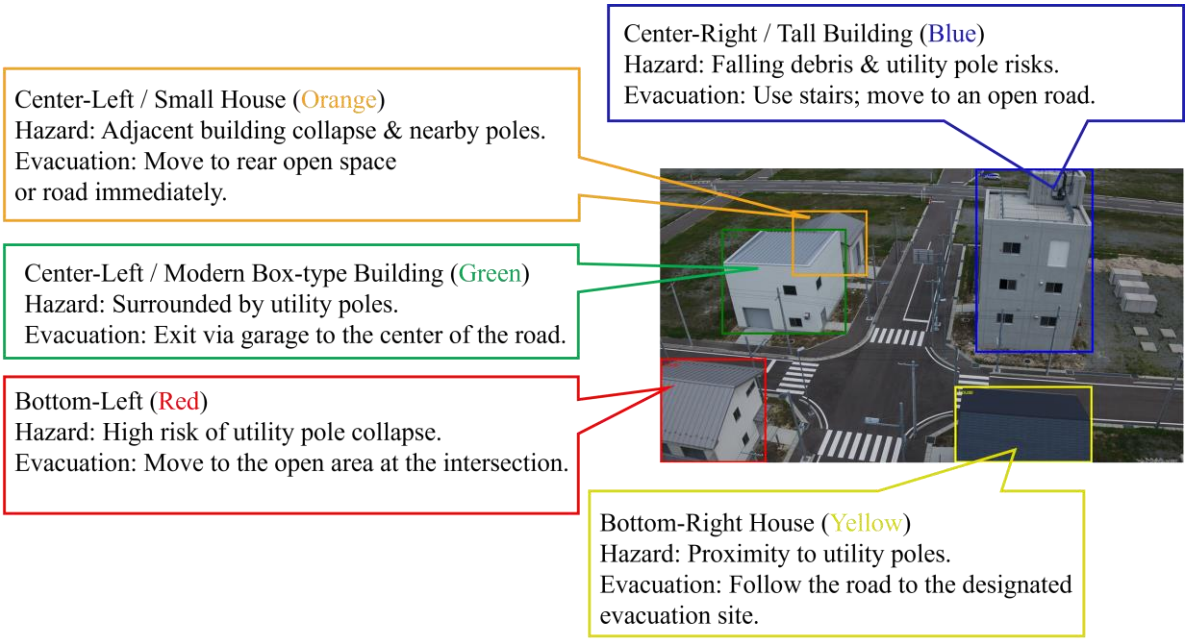

Figure 8. LLM output result

Figure 9. Use Case of the System

the LLM analyzes potential risks that may arise during a disaster and provides guidance on evacuation actions.

Figure 8 shows an example of output generated by the LLM. For each detected building, the results are categorized into risk information and evacuation guidance. Even when buildings are spatially close to each other, the results are distinguished by color-coded representations, and positional information and simple structural features are provided, making it easy to confirm the correspondence between the scene and the generated textual descriptions. For example, the result shown in the center-left (Orange) recognizes that the building is adjacent to a large house and recommends evacuating to a vacant lot behind the building rather than to an intersection. The result shown in the center-right (Blue) generates statements such as “tall structure” and “caution regarding falling rooftop equipment,” confirming that relevant structural characteristics are identified. In addition, relatively small objects such as utility poles are also recognized, and appropriate warnings indicating danger or caution during evacuation are generated. These results demonstrate that the system can assess risks and propose evacuation actions based on the identified risks.

The generated results can be placed as pictograms at appropriate locations in the 3D space through user interaction. Detailed analysis results output by the LLM can also be reviewed via pop-up displays (Figure 9, left). These findings demonstrate that users can rapidly and comprehensively add information by simply reviewing the LLM’s proposals and specifying placement locations, without the need to compose textual descriptions from scratch. Figure 9 also shows the results of annotation operations performed in the 3D space. As demonstrated, user-generated annotations can be reviewed from various viewpoints, enabling flexible and comprehensive spatial understanding.

V. CONCLUSION

In this study, we proposed a system that enables accurate annotation placement by combining the high-quality visual rendering of 3D Gaussian Splatting (3DGS) with depth information derived from a mesh model. By incorporating a surface mesh, the proposed approach overcomes the inherent

limitation of 3DGS in depth acquisition. Furthermore, the use of median depth values and offset processing allows for noise-robust and highly visible annotations.

The proposed mechanism, in which AI and humans collaboratively enrich information in a 3D space, provides a foundational framework for future disaster response systems. As for future work, we plan to conduct usability evaluations through questionnaires and user studies assuming real-world field operations, to assess the operability of the system and the effectiveness of information presentation. In addition, we will explore the integration of LLMs fine-tuned with disaster-related data to improve the reliability of generated proposals and to further automate expert-level situational assessment.

By seamlessly connecting the information organized within this system to on-site operational support and enabling end-to-end assistance from analysis to field activities, we aim to contribute to faster and more effective disaster response.

REFERENCES

[1] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian splatting for real-time radiance field rendering," ACM Trans. Graph., vol. 42, no. 4, pp. 139:1-139:14, 2023.

[2] B. Fei, J. Xu, R. Zhang, Q. Zhou, W. Yang, and Y. He, "3D Gaussian as a New Vision Era: A Survey," arXiv preprint arXiv:2402.07181, 2024.

[3] F. Nex, D. Roca, J. Fritsch, M. Gerke, and N. Kerle, "Seismic Damage Semantics on Post-Earthquake LOD3 Building Models Generated by UAS," ISPRS Int. J. Geo-Inf., vol. 10, no. 5, p. 345, 2021.

[4] S. M. S. M. Daud, M. Y. P. M. Yusof, C. C. Heo, L. S. Khoo, M. K. C. Singh, and M. S. Mahmood, "Applications of drone in disaster management: A scoping review," Science & Justice, vol. 62, no. 1, pp. 30-42, 2022.

[5] T. Seidl, C. Heckman, and S. Nikolaidis, "3D Gaussian Splatting for Human-Robot Interaction," in Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24), pp. 986-990, 2024.

[6] A. Guédon and V. Lepetit, "Sugar: Surface-aligned gaussian splatting for efficient 3D mesh reconstruction and high-quality mesh rendering," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5354-5363, 2024.

[7] J. L. Schönberger and J. M. Frahm, "Structure-from-Motion Revisited," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 4104-4113, 2016.

[8] Fukushima Robot Test Field, https://www.fipo.or.jp/robot/en/ 2025.12.01.

[9] Z. Zhu, Z. Fan, Y. Jiang, and Z. Wang, "FSGS: Real-Time Few-shot View Synthesis using Gaussian Splatting," in European conference on computer vision, Cham: Springer Nature Switzerland, pp. 145-160, 2024.

[10] Z. Fan et al., "InstantSplat: Sparse-view SfM-free Gaussian Splatting in Seconds," arXiv preprint arXiv:2403.20309, 2024.

[11] Feather, https://www.feather.art/ 2025.12.01.

[12] CozyBlanket, https://sparseal.com/cozyblanket/ 2025.12.01.

[13] M.-X. Chen, H. Hu, R. Yao, L. Qiu, and D. Li, "A Survey on the Design of Virtual Reality Interaction Interfaces," Sensors, vol. 24, no. 19, pp. 6204, 2024.

[14] Y. Jiang et al., "VR-GS: A Physical Dynamics-Aware Interactive Gaussian Splatting System in Virtual Reality," arXiv preprint arXiv:2401.16663, 2024.

[15] L. Liu, C. Wang, Z. Chen, and D. Xu, "4DGS-Craft: Consistent and Interactive 4D Gaussian Splatting Editing," arXiv preprint arXiv:2510.01991, 2025.

From Regulation to Relevance: Integrating User Values into Privacy Policy Scoring

Brian Kim and Suzanne Barber
The Center for Identity
The University of Texas at Austin
Austin, Texas, United States of America (USA)
e-mail: briankim31415@gmail.com, sbarber@identity.utexas.edu

***Abstract*—Despite increasing regulatory efforts to protect user data, such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), privacy policies—the main way users encounter these protections—remain difficult to understand and largely ineffective in guiding informed consent. Existing automated tools for evaluating these policies prioritize legal compliance but often overlook what users personally value in data privacy. This paper proposes a user-centered approach that integrates individual privacy values into policy scoring systems. A survey capturing user comfort levels with various types of Personally Identifiable Information (PII) reveals wide variability in privacy concerns. These insights are embedded into the PrivacyCheck™ tool to enable personalized policy evaluations. The results show that incorporating user values can significantly shift policy scores, especially in areas related to user control, highlighting a gap between regulatory standards and real user priorities. This work moves toward privacy tools that better reflect what users actually care about, supporting more meaningful and transparent data governance.**

Keywords–Privacy policy evaluation; User privacy values; Personalized scoring; Personally Identifiable Information.

I. Introduction

Digital platforms increasingly collect, store, and process personal data, making the protection of PII a critical public concern. Governments and regulatory bodies have responded with major privacy regulations, such as the European Union's GDPR, the CCPA, and its amendment, the California Privacy Rights Act (CPRA).

While such regulations shape data protection norms, users experience them primarily through privacy policies—the official documents describing how websites and applications collect, use, and share personal data. Ideally, policies should empower informed decisions, but studies show they are long, written at a college reading level, and filled with legal or technical jargon, making them inaccessible to the general public [1]. As a result, most users skip them or misunderstand their content, leading to uninformed consent and reduced transparency in data practices [2][3].

To address this shortcoming, automated tools interpret and evaluate privacy policies using machine learning [4]–[6]. PrivacyCheck™ [7] rates policies against regulatory frameworks, such as the GDPR and Fair Information Practice Principles (FIPPs), but these tools often apply uniform criteria and overlook variation in what users value. As a result, two users with different sensitivities can receive the same policy score, even when the policy aligns well with one user's priorities and poorly with another's.

We argue that privacy policy evaluation should go beyond regulatory compliance to reflect personal values. We propose a user-centered approach that integrates individual privacy preferences into automated policy scoring systems. We conducted a user study to assess comfort with sharing different types of PII and privacy-related practices. Results show general caution—especially for possession-based information—but substantial variation across individuals, highlighting the need for personalized evaluation frameworks. This variation suggests that a single "good" score can mask trade-offs across factors, such as data retention and third-party sharing, that users weigh differently.

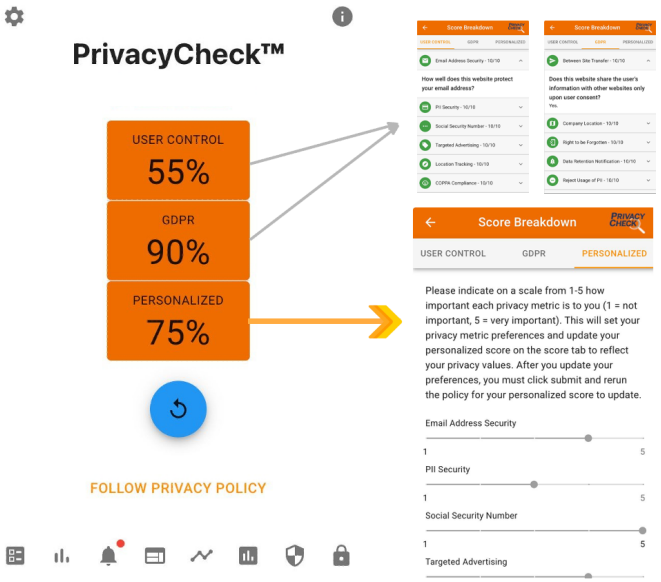

Figure 1. An example of PrivacyCheck™ scoring with the personalized scoring mechanism.

We map the survey insights onto the PrivacyCheck™ scoring model, enabling personalization based on aggregated or individual values (see Figure 1). Incorporating user values shifts policy scores—particularly for user control—revealing gaps between regulatory benchmarks and user priorities. By embedding user values into policy assessments, this work bridges legal compliance and meaningful user understanding.

The remainder of this paper is organized as follows. Section 2 reviews related work, Section 3 describes the survey, Section 4 details the value-aligned scoring and evaluation, and Section 5 concludes.

II. Related Work

A. Tools for Interpreting Privacy Policies

Automated tools analyze privacy policies using Natural Language Processing (NLP) and machine learning to improve transparency, regulatory compliance, and user understanding. Early systems, such as Privee [5], extract key practices from policy text. Polisis [4] uses deep learning with hierarchical attention to classify policy content and present it through an interactive user interface.

More recent tools include PrivacyGuide [6] and PolicyChecker [8]. They assess GDPR compliance using NLP and supervised learning

but often overlook individual user values. The Value-Centered Privacy Assistant (VcPA) [9][10] addresses this gap by helping users make app decisions aligned with personal value profiles derived from survey data.

B. User Values Regarding Privacy Policies

Surveys and experiments have examined user attitudes toward privacy policies and privacy-enhancing technologies. Ibdah et al. [11] show that despite concern about data privacy, participants often avoid policies due to complexity, length, and helplessness, prioritizing convenience over informed consent. Choi et al. [12] and Ebbers et al. [13] report similar themes in smart speakers and digital assistants, emphasizing transparency and control with features, such as customizable data collection and meaningful consent. These studies collectively suggest that usability and perceived control, not just formal compliance, shape whether privacy tools influence real behavior.

Wang & Li [14] propose a machine learning framework to infer privacy preferences from demographic and behavioral data, enabling personalized recommendations with limited user input. Building on this foundation, we directly integrate user-reported comfort levels with different types of PII into an automated policy evaluation system, aligning assessments with user values.

C. Privacy Policy Rating Systems PrivacyCheck™

We focus on PrivacyCheck™ [7], a machine learning–based privacy policy rating system. Implemented as a Chrome browser extension, PrivacyCheck™ identifies and scores key privacy practices from policy webpages, providing structured assessments aligned with frameworks, such as the GDPR and FIPPs.

Originally launched in 2015, the initial version of PrivacyCheck™ detected policy content across ten core dimensions, including PII collection (e.g., email, Social Security Number (SSN)), data aggregation, and law enforcement sharing. It used keyword detection and parsing to assign categorical risk scores reflecting GDPR and Federal Trade Commission (FTC) standards. Later iterations added scalability, version tracking, and user engagement.

In its current implementation, PrivacyCheck™ rates privacy policies across 20 factors derived from regulatory guidelines [15], as shown in Table I. Factors span disclosures, such as data usage, retention, third-party transfers, user access rights, and consent mechanisms. Each factor is scored and aggregated into two overall scores:

- GDPR Score: Reflects the presence or absence of disclosures relevant to specific GDPR articles
- User Control Score: Indicates the extent to which a policy outlines mechanisms for user agency and data governance based on FIPPs

In this work, PrivacyCheck™ provides the baseline for comparing regulatory compliance to user privacy value coverage, supporting an empirical investigation of transparency and data governance in online privacy policies.

III. User Attitudes Toward the Collection and Processing of PII

A. Survey Design

To evaluate users' comfort levels with digital privacy practices, we conduct a survey to capture individual attitudes toward the collection and processing of various types of PII by websites and mobile applications. This study was approved by the Institutional Review Board (IRB). The survey is structured around six thematic categories as listed below.

The first two categories address data policies and processing practices, assessing participants' perceptions of broader organizational behaviors, such as data retention, third-party data sharing, and transparency in privacy policies. The other four categories focus on participants' attitudes toward specific types of PII. These categories are informed by foundational principles in information security and user authentication, specifically the types of data individuals have, are, do, and know. Together, these six categories were designed to comprehensively reflect real-world contexts of digital data collection and usage.

1) Data Policies (17 questions)
2) Data Processing (8 questions)
3) What You HAVE (16 questions): possession-based identifiers (e.g., device IDs)
4) What You ARE (10 questions): biometric information (e.g., fingerprints)
5) What You DO (10 questions): behavioral data (e.g., browsing history)
6) What You KNOW (20 questions): knowledge-based credentials (e.g., passwords)

The final survey included a total of 81 questions. Each question in the survey is phrased in a standardized format to ensure clarity and consistency across categories. The wording generally follows the structure: "How comfortable are you with websites or mobile apps... [question topic]." The question topic typically begins with an action verb, such as "collecting" or "sharing," depending on whether the item relates to data acquisition or dissemination, and is followed by a specific type of PII; for example, "How comfortable are you with websites or mobile apps collecting your online shopping patterns?" or "How comfortable are you with websites or mobile apps sharing your data for academic or research purposes?" This consistent phrasing helps participants easily interpret the intent of each question while maintaining focus on the specific privacy-related scenario being assessed.

Participants respond to each item using a five-point Likert scale: (1) very comfortable, (2) slightly comfortable, (3) no difference, (4) slightly uncomfortable, and (5) very uncomfortable. The selected scale enables nuanced expression and fine-grained analysis of user comfort levels, revealing both the types of PII users are most sensitive about and the intensity of that sensitivity.

To ensure the quality of the responses and detect inattentive participation, one question from each of the six main categories is deliberately duplicated in another category where the context still made the question relevant. These duplicated questions serve as internal consistency checks. A response is flagged as inconsistent if a participant answered a duplicated question pair in a way that suggested a reversal in sentiment—for example, if an answer shifted from either very uncomfortable or slightly uncomfortable to either very comfortable or slightly comfortable, or vice versa. Additionally, if a participant changed their response from an extreme stance—defined as very uncomfortable or very comfortable—to no difference, or from no difference to an extreme stance, this is also treated as an inconsistency. Submissions containing more than one of these response reversals are considered unreliable and excluded from the dataset. Using this criterion, nine participant submissions are removed.

Demographics. In the survey, we also collect participants' demographic information, including gender, age, ethnicity, occupation, and education level. Additionally, we ask five questions about participants' general attitudes toward privacy practices and their daily use of websites and applications, as listed below. This information is used to explore potential correlations between user profiles and their

TABLE I. PRIVACYCHECK™ SCORING FACTORS.

User Control Factors
How well does this website protect your email address?
How well does this website protect your credit card information and address?
How well does this website handle your Social Security Number?
Does this website use or share your PII for marketing purposes?
Does this website track or share your location?
Does this website collect PII from children under 13?
Does this website share your information with law enforcement?
Does this website notify or allow you to opt-out after changing their privacy policy?
Does this website allow you to edit or delete your information from its records?
Does this website collect or share aggregated data related to your identity or behavior?
GDPR Compliance Factors
Does this website share the user's information with other websites only upon user consent?
Does this website disclose where the company is based or where the user's PII will be processed and transferred?
Does this website support the right to be forgotten?
If they retain PII for legal purposes after a user request to be forgotten, will they inform the user?
Does this website allow the user to reject the use of their PII?
Does this website restrict the use of PII of children under the age of 16?
Does this website advise the user that their data is encrypted even while at rest?
Does this website ask for the user's informed consent before processing data?
Does this website implement all principles of data protection by design and by default?
Does this website notify the user of security breaches without undue delay?

comfort levels with digital privacy, enabling analysis of trends across different demographic groups.

1) On average, how many hours per day do you spend online or on mobile apps?
2) How often do you read each website's or mobile application's privacy policy?
3) How comfortable are you with technology?
4) Do you agree with the statement that "I feel that I get more accomplished because of technology?"
5) List 3 to 5 of your most used websites and mobile apps.

Survey participants are recruited from the undergraduate and graduate student populations in the Electrical and Computer Engineering (ECE) and Computer Science (CS) departments at the University of Texas at Austin. Participation is entirely voluntary. In total, we collect 99 valid responses. Of the participants, 70.7% identify as male, 23.2% as female, and 6.1% as other gender identity. In terms of ethnicity, 59.6% identify as Asian, 17.2% as White, and 16.2% as Hispanic or Latino. The average age of participants is 21 years.

A third of participants (33.3%) report that they never read privacy policies, while 42.4% say they rarely do. Only 24.2% indicate that they read them sometimes or often. Regarding technological proficiency, 43.4% of participants consider themselves experts, and 46.5% report an advanced level of comfort with technology. Participants also report their average daily internet usage, with responses spanning a broad range, as illustrated in Figure 2. Finally, participants list their top 3 to 5 most frequently used websites and applications. Aggregated results reveal the 20 most-used platforms, shown in Figure 2. Social media platforms dominate usage patterns, with Instagram emerging as the most frequently mentioned.

B. Survey Results

The overall mean score across all survey questions is 3.85 on a five-point Likert scale. As shown in Figure 3, questions in the "Data Policies" and "What You HAVE" categories receive higher scores, suggesting that users are generally cautious about websites and apps collecting their PII, particularly possession-based information. In contrast, scores are lower in the "What You KNOW" and "Data Processing" categories, indicating more users are comfortable with sharing knowledge-based PII and are less concerned about how their data is processed. The questions with the highest and lowest scores are shown in Table II.

We also calculate the standard deviation for each survey question to assess variability in participants' responses, as shown in Figure 4. The average standard deviation across all questions is 1.06. We observe that questions with higher mean scores—indicating lower comfort levels—tend to have lower variance, suggesting stronger consensus when participants feel uncomfortable about specific data practices or sharing certain types of information. In contrast, questions with lower mean scores—reflecting higher comfort—exhibit greater variance, indicating a wider range of opinions. This pattern reveals a general agreement on what makes users uncomfortable, while comfort tends to be more subjective and varies significantly across individuals.

In Table III, we list questions with the highest and lowest variance. Notably, all five questions with the highest variance belong to the "What You KNOW" category, whereas the questions with the lowest standard deviations are from the "Data Policies" and "What You HAVE" categories.

Additionally, to assess the relationships between users' responses across questions, we compute the Pearson correlation coefficient (r). The results, shown in Figure 5, summarize the number of strongly correlated question pairs ($r > 0.6$) across category combinations. Few questions have negative correlations, and none are strongly negative. Most strong correlations occur within the same category, likely reflecting thematic consistency. In particular, questions in the "What You HAVE" and "What You KNOW" categories exhibit high internal correlations. Notably, the "What You KNOW" category is the only group to show strong correlations with questions from other categories, suggesting it may capture broader user attitudes that extend beyond its specific theme.

To ensure that correlation patterns are not merely the result of universally agreeable questions, a subset of questions is selected based on higher distributional variability (standard deviation greater than 1). Pearson correlation results for the refined question set are shown in Figure 6. Even with this stricter filtering, the "What You KNOW" category remains prominent, exhibiting a high number of strong correlations, especially within its own group.

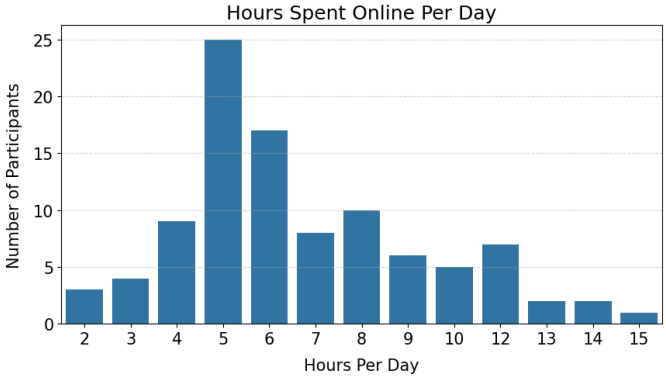

(a) Distribution of hours spent online per day.

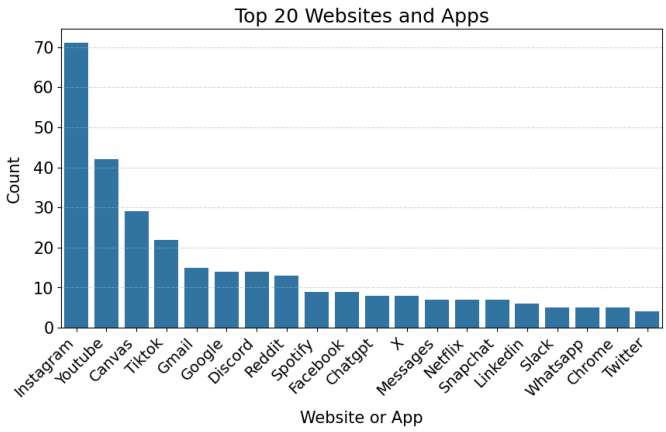

(b) Top 20 most frequently used websites and apps.

Figure 2. Demographic Information.

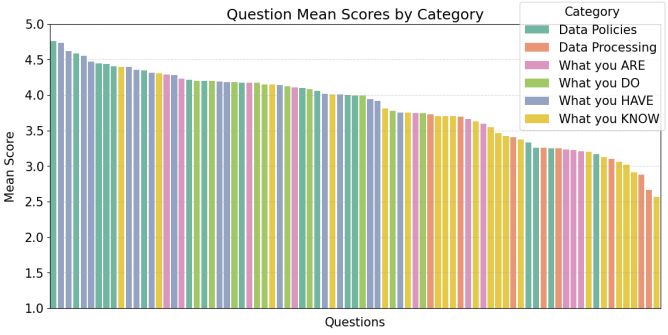

Figure 3. Distribution of question mean scores color-coded by category.

Finally, we calculate correlations between participants' average scores on high-distribution questions and their demographic information. For continuous demographic variables, we used Pearson correlation; for categorical variables, we used analysis of variance (ANOVA). None of the correlations are strong, suggesting no meaningful relationship between demographic factors and their responses to more polarizing or nuanced questions.

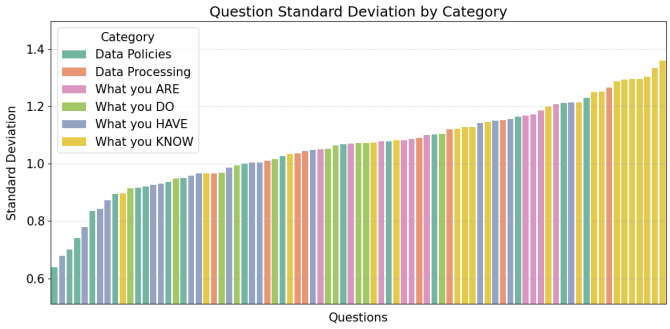

Figure 4. Distribution of question score standard deviations color-coded by category.

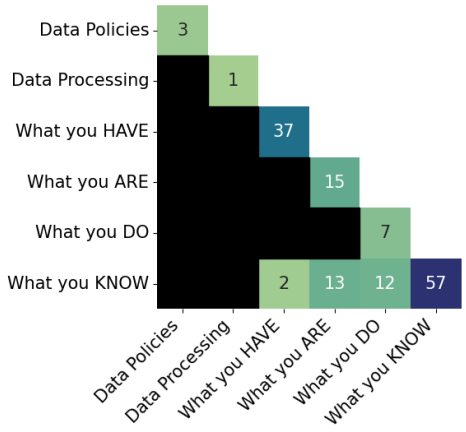

Figure 5. Count of Strongly Correlated Question Pairs across Categories.

Figure 6. Counts of Strongly Correlated High-Variance Question Pairs across Categories.

TABLE II. QUESTIONS WITH THE HIGHEST AND LOWEST MEAN SCORES.

Question	Category	Mean
... not notifying you if a data breach has occurred?	Data Policies	4.76
... having access to and/or collecting copies of your SSN?	What You HAVE	4.74
... having access to and/or collecting copies of your birth certificate?	What You HAVE	4.62
... collecting information from children under the age of 13?	Data Policies	4.59
... having access to and/or collecting copies of your passport?	What You HAVE	4.56
... collecting your email?	What You KNOW	3.02
... collecting your gender?	What You KNOW	2.91
... collecting your data for curated recommendations?	Data Processing	2.88
... collecting your data for developer bug detection and analytics?	Data Processing	2.67
... collecting your preferred language?	What You KNOW	2.57

TABLE III. QUESTIONS WITH THE HIGHEST AND LOWEST STANDARD DEVIATIONS IN SCORE DISTRIBUTIONS.

Question	Category	St.D	Mean
... collecting your ethnicity?	What You KNOW	1.36	3.13
... collecting your name?	What You KNOW	1.33	3.20
... collecting your phone number?	What You KNOW	1.30	3.46
... collecting your preferred language?	What You KNOW	1.30	2.57
... collecting your gender?	What You KNOW	1.30	2.91
... having access to and/or collecting copies of your birth certificate?	What You HAVE	0.78	4.62
... sharing your data with business partners that are not explicitly listed?	Data Policies	0.74	4.40
... not notifying you if a data breach has occurred and your information may be compromised?	Data Policies	0.70	4.76
... having access to and/or collecting copies of your SSN?	What You HAVE	0.68	4.74
... collecting information from children under the age of 13?	Data Policies	0.64	4.59

IV. PERSONALIZED PRIVACYCHECK™ WITH VALUE-ALIGNED SCORING

Using the collected data on users' attitudes toward privacy practices, we incorporate values into the privacy policy rating system—PrivacyCheck™. We map each survey question to corresponding factors and re-weight these factors based on aggregated responses to better align PrivacyCheck™ scores with users' values. We also implement a personalization feature that allows individual users to adjust factor weights and compare the resulting scores with the original regulatory scores.

A. Aligning PrivacyCheck™ Scoring with User Value

As described in subsection 2.3, PrivacyCheck™ evaluates privacy policies using two scoring dimensions: GDPR Compliance and User Control. Each dimension comprises a set of factors. For each factor, the model assigns 0/1/2 for not addressed, unclear, or clear. These raw scores are normalized to a 0–10 scale per factor and summed to produce two overall scores out of 100—one reflecting GDPR Compliance and the other User Control. The list of corresponding questions is provided in Table I.

At a high level, if $x_f \in \{0, 1, 2\}$ is the raw score for factor f, then the base score for dimension d (with factors F_d) is computed as S_d. To incorporate user values, we map survey responses to scoring factors; each survey question links to one or more criteria used in PrivacyCheck™. We compute average responses per question and aggregate them by factor to assign a user value weight, then normalize these weights within each dimension to sum to 1. For population-level analysis, v_f is the average response across participants; for deployment, v_f can be computed per user. We normalize v_f to survey weights w_f, which initialize slider values. Users adjust slider values s_f, normalized via a softmax function to user weights $\mathbf{u}$ (with u_f as the fth element). The base score sums normalized factor scores across F_d, scaling each x_f to a 0–10 range. The weighted score is

$$w_f = \frac{v_f}{\sum_{k \in F_d} v_k}, \quad S_d^{(u)} = 10 \cdot \sum_{f \in F_d} u_f \cdot \frac{x_f}{2}$$

We integrate this feature into the PrivacyCheck™ browser extension via a new "Personalization" tab, as shown in Figure 1. The interface presents adjustable sliders for each factor, initialized with survey-derived weights. Users can modify these sliders to reflect their preferences, and the values are normalized to sum to 1, allowing users to increase the influence of specific factors and receive customized scoring.

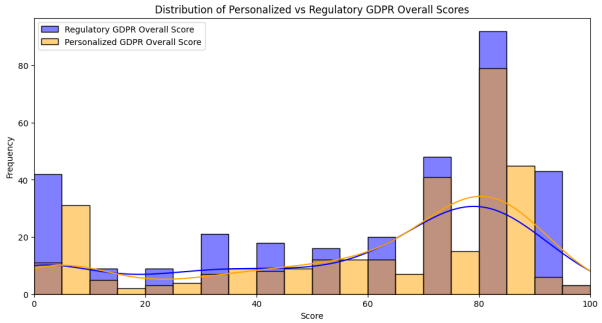

Figure 7. Distribution of personalized vs regulatory GDPR scores.

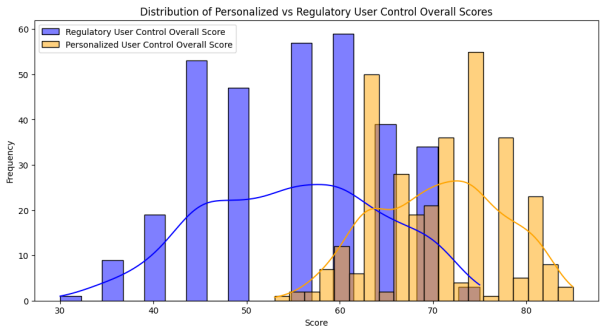

Figure 8. Distribution of personalized vs regulatory User Control scores.

B. Evaluation of the Alignment

To evaluate the value-aligned scoring system, we run the aligned PrivacyCheck™ on 321 privacy policies from the original 392-policy corpus [7], excluding invalid URLs or scoring errors. We compare value-aligned scores—generated using survey-based weights (without user-adjusted sliders)—against the original regulatory scores. Figure 7 and Figure 8 present the results for GDPR Compliance and User Control, respectively.

GDPR Compliance scores remain relatively similar under value-aligned weighting, whereas User Control scores increase. This suggests users emphasize control over how their data is collected and used, especially for sensitive "What You HAVE" information, such as Social Security Numbers and location data. This shift indicates regulatory assessments may underrepresent user priorities, and value-aligned scoring can surface concerns de-emphasized in traditional evaluations.

V. Conclusion and Discussion

This work addresses whether privacy policies reflect the values and comfort levels of the users they aim to protect. Through a user-centered survey and a personalized policy evaluation mechanism in PrivacyCheck™, we provide an empirical and technical foundation for answering that question.

Our findings show that strong, consistent discomfort with opaque data practices, particularly those involving sensitive identifiers, such as government-issued documents and breach notifications. Responses to less invasive practices, such as email or language collection, are more tolerant and varied. This pattern suggests contextual and value-driven interpretations of privacy. While regulations, such as the GDPR, provide a foundation for protecting user rights, they do not fully capture individual privacy concerns.

By integrating user preferences into PrivacyCheck™, we show that personalized scoring shifts evaluations, particularly in user control. Conventional regulatory assessments may underestimate areas that users deem critical. At the individual level, differences in priorities can shift factor emphasis and change policy rankings, supporting user-specific decision-making, such as comparing apps against stated privacy priorities. This opens the door to interfaces that let users tune scores to their values without requiring them to parse long policy text.

A. Limitations

The survey sample is limited to 99 respondents recruited from one university's ECE and CS populations, which may not represent broader user populations. Survey responses are self-reported and may not fully capture how users behave when making real privacy decisions. The mapping from 81 survey items to the 20 PrivacyCheck™ factors also introduces modeling assumptions, and aggregated factor weights may mask nuanced preferences that are not yet expressed in the survey.

The policy evaluation uses a subset of the original 392-policy corpus after excluding invalid URLs and scoring errors, so results may differ on other corpora or newly updated policy text. Finally, the current evaluation focuses on score shifts rather than downstream outcomes such as whether personalized scoring changes users' consent behavior or long-term trust. Broader demographic sampling, additional policy datasets, and behavioral validation would strengthen external validity.

To translate these findings into practice, organizations should enhance transparency by explaining, in clear language, how personal data is collected, used, and shared—especially for sensitive identifiers. It may also be helpful to provide more intuitive and granular options for users to manage privacy preferences around consent and data sharing. While legal compliance remains essential, aligning practices with user expectations can foster trust and engagement. Tools like the personalized version of PrivacyCheck™ show how user feedback can bridge regulatory intent and lived experience.

In closing, we aim to narrow the gap between what privacy policies claim and what users value. Centering user preferences supports tools and frameworks that are compliant yet intuitive and respectful. As digital ecosystems expand, integrating user values into privacy design can help build more transparent and trustworthy data governance.

References

[1] J. Tang, H. Shoemaker, A. Lerner, and E. Birrell, "Defining privacy: How users interpret technical terms in privacy policies," *Proceedings on Privacy Enhancing Technologies*, vol. 2021, no. 3, pp. 70–94, 2021. DOI: 10.2478/popets-2021-0038.

[2] C. McClain, M. Faverio, M. Anderson, and E. Park, *How Americans view data privacy*, Pew Research Center, Report 18 [retrieved: January, 2026], 2023. [Online]. Available: https://www.pewresearch.org/internet/2023/10/18/how-americans-view-data-privacy/.

[3] B. Auxier et al., *Americans' attitudes and experiences with privacy policies and laws*, Pew Research Center: Internet, Science & Tech, [retrieved: January, 2026], 2019. [Online]. Available: https://www.pewresearch.org/internet/2019/11/15/americans-attitudes-and-experiences-with-privacy-policies-and-laws/.

[4] H. Harkous et al., "Polisis: Automated analysis and presentation of privacy policies using deep learning," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 531–548.

[5] S. Zimmeck and S. M. Bellovin, "Privee: An architecture for automatically analyzing web privacy policies," in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 1–16.

[6] W. B. Tesfay, P. Hofmann, T. Nakamura, S. Kiyomoto, and J. Serna, "Privacyguide: Towards an implementation of the EU GDPR on internet privacy policy evaluation," in *Proceedings of the fourth ACM international workshop on security and privacy analytics*, 2018, pp. 15–21.

[7] R. N. Zaeem, R. L. German, and K. S. Barber, "Privacycheck: Automatic summarization of privacy policies using data mining," *ACM Transactions on Internet Technology (TOIT)*, vol. 18, no. 4, pp. 1–18, 2018.

[8] A. Xiang, W. Pei, and C. Yue, "Policychecker: Analyzing the GDPR completeness of mobile apps' privacy policies," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 3373–3387.

[9] S. E. Carter, "A value-centered exploration of data privacy and personalized privacy assistants," *Digital Society*, vol. 1, no. 3, p. 27, 2022.

[10] S. E. Carter et al., *In pursuit of privacy: The value-centered privacy assistant*, [retrieved: January, 2026], 2023. [Online]. Available: https : / / api . semanticscholar . org / CorpusId : 260775437.

[11] D. Ibdah, N. Lachtar, S. M. Raparthi, and A. Bacha, ""Why should I read the privacy policy, I just need the service": A study on attitudes and perceptions toward privacy policies," *IEEE Access*, vol. 9, pp. 166 465–166 487, 2021.

[12] H. Choi, J. Park, Y. R. Choi, and Y. Jung, "User preferences of privacy-enhancing attributes of a smart speaker," *International Journal of Human–Computer Interaction*, vol. 39, no. 18, pp. 3649–3662, 2023.

[13] F. Ebbers, J. Zibuschka, C. Zimmermann, and O. Hinz, "User preferences for privacy features in digital assistants," *Electronic Markets*, vol. 31, pp. 411–426, 2021.

[14] W. Wang and B. Li, "Learning personalized privacy preference from public data," *Information Systems Research*, pp. 1–20, 2024, Articles in Advance. DOI: 10.1287/isre.2023.0318.

[15] R. N. Zaeem et al., "Privacycheck v3: Empowering users with higher-level understanding of privacy policies," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 2022, pp. 1593–1596.

 ISBN: 978-1-68558-351-4