

THE REALISE PROJECT – RESPONSIBLE EDUCATION AND LEARNING WITH ARTIFICIAL INTELLIGENCE IN SYSTEMS AND ELECTRONICS

Fernando Fernández-Martínez, Iván Martín-Fernández, Sergio Esteban-Romero, Pedro Malagón-Marzo, Manuel Gil-Martín, & Juan Manuel Montero-Martínez

Department of Electronics Engineering, Information Processing and Telecommunications Center, ETSI Telecomunicación, Universidad Politécnica de Madrid (Spain)

Abstract

The widespread availability of generative artificial intelligence (AI) tools is transforming university teaching practices, particularly in technical disciplines centered on programming, system design, and experimental validation. While these tools offer new opportunities for learning support and creativity, they also raise concerns regarding students' reasoning, autonomy, and professional ethics. The REALISE project—Responsible Education and Learning with Artificial Intelligence in Systems and Electronics—addresses this challenge through a pedagogical framework that promotes ethical, reflective, and responsible use of AI in engineering education. Developed at the School of Telecommunications Engineering of the Universidad Politécnica de Madrid (UPM), REALISE is implemented in three experimentally intensive courses: Laboratory of Electronic Systems (LSEL) in the Master's in Electronic Systems Engineering (MUISE), and Sectorial Applications (APSE) and Projects in Data and Systems Engineering (PIDS) in the Bachelor's in Data and Systems Engineering (GISD). The methodology combines project-based learning with explicit documentation of students' interactions with AI tools, requiring justification of prompts, generated outputs, corrections, and critical reflections. In LSEL, authentic learning is reinforced through real sensor measurements, diversified hardware configurations, practical constraints, and debugging of intentionally faulty code. In APSE and PIDS, students actively build and curate datasets for AI model development, fostering awareness of data quality, representativeness, and bias. REALISE aims to integrate "Responsible AI" as a transversal competence, strengthen analytical and documentation skills, and generate transferable educational materials for broader adoption. The project contributes to the modernization of engineering education by promoting a human-centered and ethically grounded digital transformation.

Keywords: *Artificial intelligence, responsible education, educational innovation, critical thinking, ethics, assessments and measurements in education.*

1. Introduction

Generative artificial intelligence (AI) tools are rapidly reshaping higher education, particularly in engineering disciplines where programming, system design, and experimental validation are central. Although AI can support learning, prototyping, and documentation, it also challenges the development of reasoning skills, autonomy, and professional ethics. In technical contexts, syntactically correct code generation does not guarantee conceptual understanding or engineering competence.

The REALISE project (Responsible Education and Learning with Artificial Intelligence in Systems and Electronics) proposes a formative and ethical integration of AI into engineering education. Instead of restricting AI use, REALISE redesigns learning activities and assessment strategies so that AI functions as cognitive support while responsibility for reasoning, validation, and decision-making remains with the student. Developed at the School of Telecommunications Engineering of the Universidad Politécnica de Madrid (UPM), REALISE builds on prior challenge-based learning experiences that integrated AI into electronics education and demonstrated improvements in engagement and practical competence (Bravo-Pacheco et al., 2025). The framework is grounded in the Semantic Bridge paradigm, which provides a theoretical basis for understanding the limitations of AI in experimental engineering contexts and for designing learning scenarios that foster deep understanding and critical thinking.

2. The semantic bridge paradigm

The Semantic Bridge paradigm distinguishes between syntactic information and semantic meaning in engineered systems. Digital systems operate on symbols, data structures, and signals according to formal rules, but functional meaning emerges only when these elements are interpreted within a physical or

operational context. Similar notions appear in ontology-based integration, where semantic bridges connect heterogeneous conceptual spaces (Ambalavanan et al., 2025), and in enterprise AI through semantic data contracts that formalize shared meaning across data pipelines (Kuthuru, 2025).

Generative AI models primarily operate at the syntactic level, identifying statistical regularities in symbolic representations. Although their outputs may appear coherent, these models lack direct access to physical reality and cannot autonomously establish causal relationships between digital representations and real-world behavior. This limitation is particularly evident in experimental domains such as embedded systems and electronics, where validation against physical constraints is essential.

Recent research has explored algorithmic realizations of semantic bridges through structured representations such as Abstract Meaning Representation (AMR) graphs (Semantic Bridge, 2025). These approaches provide valuable insights into how semantic structure can be formally represented and manipulated, particularly for symbolic reasoning and interpretability. In the context of engineering education, however, our work extends this perspective by emphasizing semantic grounding through experimental interaction with physical systems.

From an educational perspective, meaningful learning occurs when students actively construct or reconstruct the semantic bridge linking abstract representations (code, models, data) to their physical or functional consequences. By intentionally designing tasks where this bridge is incomplete or context-dependent, students are required to reason, experiment, and interpret results. In such settings, AI may assist with syntax or documentation, but it cannot replace the student's role in constructing meaning.

3. Educational context and objectives

Previous implementations of embedded AI technologies in laboratory subjects have shown that combining hardware experimentation with AI-based components effectively supports the development of practical AI skills in electronics students (Gil-Martín et al., 2023). REALISE is implemented across three courses with strong experimental components:

- Laboratory of Electronic Systems (LSEL), Master's in Electronic Systems Engineering (MUISE).
- Sectorial Applications (APSE) and Projects in Data and Systems Engineering (PIDS), Bachelor's in Data and Systems Engineering (GISD).

These courses involve open-ended problems using hardware platforms, sensor data, and system-level constraints. Integrating AI in this context requires careful pedagogical design to preserve learning objectives related to reasoning, validation, and ethical responsibility.

The project objectives are:

1. To integrate responsible and reflective AI use as a transversal competence.
2. To promote critical thinking, autonomy, and self-regulation.
3. To design activities resistant to fully automated resolution.
4. To strengthen technical documentation and objective evaluation skills.
5. To generate transferable methodologies and educational materials.

4. Methodology and instructional design

The instructional approach combines project-based learning, experimental validation, and personalized problem settings. Each student or group interacts with unique hardware configurations, datasets, or constraints, limiting the effectiveness of generic AI-generated solutions. Activities include real sensor data acquisition, hardware diversification, and realistic constraints related to energy, timing, or environmental conditions.

A core methodological element is the explicit documentation of AI usage. Students must record prompts, generated outputs, applied corrections, and a critical assessment of AI suggestions. This shifts assessment from final results to learning processes, fostering metacognition and ethical awareness while positioning AI as an assistive tool rather than a substitute for reasoning.

5. Implementation scenarios

5.1. Laboratory of electronic systems (LSEL)

In LSEL, learning activities are explicitly designed to reinforce semantic grounding through interaction with real hardware, heterogeneous configurations, and experimental validation. Students work with physical sensors, acquire real measurements from their environment, and deploy solutions under realistic constraints related to timing, energy consumption, and environmental variability. Hardware

configurations are intentionally diversified across students or groups, ensuring that system behavior cannot be fully anticipated from abstract specifications alone.


A central strategy in LSEL is the deliberate introduction of semantic friction between abstract representations and observable system behavior. Students are frequently required to debug intentionally faulty code fragments, where errors are not purely syntactic but contextual—for example, incorrect assumptions about sensor behavior, communication timing, or peripheral configuration. These tasks force students to reason about causality, validate hypotheses experimentally, and iteratively refine their designs, thereby constructing the Semantic Bridge between code and physical behavior.

5.1.1. Exam design and semantic grounding. As a concrete implementation of this paradigm, LSEL adopts an individual, hands-on examination model centered on MQTT-based communication, a widely used protocol in embedded and IoT systems (HiveMQ, 2023). The exam requires students to design, implement, and execute a publish-subscribe system that interacts with a live MQTT broker, emphasizing operational correctness, interpretation of runtime behavior, and system-level reasoning.

Rather than assessing the reproduction of known code patterns, the exam evaluates the student's ability to translate abstract protocol concepts—such as topics, publishers, subscribers, and message semantics—into a functioning cyber-physical system.

Figure 1 illustrates the hierarchical MQTT data model provided to students as part of the exam scenario. While the syntactic structure of sensing, alerting, and control domains is explicitly defined, the operational semantics of each branch must be inferred, implemented, and validated through system execution. This partial specification deliberately requires students to construct the Semantic Bridge between abstract communication structures and observed system behavior.

Figure 1. Hierarchical MQTT data model used in the exam scenario.



The assessment context is intentionally underspecified at the syntactic level, requiring students to interpret requirements, justify design decisions, and validate system behavior through execution and observation. Table 1 summarizes the exam context, communication infrastructure, and semantic constraints imposed on data exchange, while Table 2 outlines the structure of the exam exercises and their associated reasoning focus.

Table 1. MQTT Exam Context and Data Model.

Element	Description
Scenario	Smart greenhouse monitoring and control system
Broker	IP: 10.8.42.100, Port: 1883
Topic hierarchy	Sensors / Alerts / Control
Personalization	Sector ID embedded in topics
Payload constraints	Numerical (sensors); {ON, OFF} (control); categorical (alerts)

Table 2. Exam Exercises and Assessment Focus.

Exercise	Core Task	Reasoning Focus
1) Data aggregation	Subscribe and compute rolling averages	Interpretation of live data
2) Alert generation	Detect anomalies and publish alerts	Semantic threshold reasoning
3) Hierarchical control	React to macro-sector alerts	System-level contextual reasoning

The effectiveness of this exam model lies in the fact that successful completion depends on the alignment of three elements: syntactic correctness (valid code and protocol usage), semantic interpretation (meaning of data and messages), and contextual validation (observed system behavior). While students may consult AI tools for syntax reminders or library usage, AI-generated solutions alone are insufficient, as they

cannot infer runtime behavior, diagnose context-dependent faults, or validate decisions against experimental outcomes.

From a responsible AI perspective, this exam design naturally constrains misuse while explicitly allowing AI as a support tool. Students remain accountable for interpretation, debugging, and validation, ensuring that AI augments—rather than replaces—engineering judgment.

Table 3. Comparison between traditional assessment models and the REALISE exam model.

Dimension	Traditional Exam Model	REALISE Exam Model
Nature of the task	Predominantly theoretical or template-based coding exercises	Individual, hands-on implementation in a real execution environment
Focus of assessment	Correctness of final answer or code output	Reasoning process, system behavior, and experimental validation
Contextual grounding	Abstract or simulated scenarios	Real communication scenarios with live brokers and runtime interaction
Role of code	End product to be evaluated	Medium for demonstrating understanding and decision-making
Treatment of AI tools	Often implicit or uncontrolled	Explicitly allowed as support, but insufficient for task completion
Resistance to automation	Low: solutions can often be generated directly by AI	High: requires semantic grounding, execution, and debugging
Evaluation perspective	Result-oriented	Process-oriented and reflective
Skills emphasized	Memorization, syntax, procedural reproduction	System-level reasoning, interpretation, autonomy, and critical judgment

Table 3 summarizes the key differences between conventional assessment approaches and the exam model adopted in the REALISE framework. While traditional exams tend to prioritize syntactic correctness and reproducibility of expected solutions, the REALISE model deliberately relocates assessment to the construction of meaning at the system level. By requiring students to execute, observe, and interpret the behavior of a live MQTT-based system, the exam enforces the active reconstruction of the Semantic Bridge between abstract protocol descriptions and their functional manifestation in a real environment. This shift not only increases resistance to automated solution generation, but also aligns assessment with the core learning objectives of experimental engineering: reasoning under uncertainty, validation through observation, and responsible use of computational tools.

5.2. Sectorial applications (APSE) and projects in data and systems engineering (PIDS)

In the Sectorial Applications (APSE) and Projects in Data and Systems Engineering (PIDS) courses, the REALISE framework is applied at a higher level of abstraction, focusing on data-centric and system-level AI design. Students actively participate in the creation, curation, and analysis of datasets used to develop AI-based systems, fostering concrete awareness of data quality, representativeness, and bias. This hands-on involvement makes explicit the ethical and technical implications of data-driven decisions and mirrors professional practices such as the use of semantic data contracts to ensure accountability and shared understanding across AI pipelines (Kuthuru, 2025).

The central project in APSE and PIDS involves the design of gesture-based Human–Computer Interaction systems, aligned with current research in embedded hand-gesture recognition and Edge AI (Esteban-Romero et al., 2025; Gil-Martín et al., 2025). Students implement complete processing pipelines—including data acquisition using MediaPipe landmarks, model training, and embedded deployment—on resource-constrained platforms such as the Raspberry Pi. Inference is performed locally, reinforcing privacy-by-design principles and exposing students to real deployment constraints.

From a pedagogical standpoint, this scenario operationalizes the Semantic Bridge paradigm at the data level. Students must define what their data means—how gestures are labeled, how variability is handled, and how preprocessing choices affect downstream behavior. While generative AI can assist with scripting, automation, or exploration of alternative architectures, it cannot replace the physical act of data collection nor the contextual interpretation of model behavior under real-world conditions.

A key learning outcome is the explicit analysis of trade-offs between model accuracy, computational cost, and system responsiveness. Students are required to justify architectural choices and assess whether performance gains justify increased complexity under embedded constraints. This reflective process reinforces responsible AI use by making limitations, assumptions, and consequences visible and subject to critical evaluation.

6. Evaluation framework

REALISE currently adopts a formative evaluation strategy based on student questionnaires, reflective reports, and instructor observations, focusing on autonomy, reasoning effort, and AI usage patterns. Qualitative analysis of technical documentation emphasizes justification, validation, and critical assessment of AI outputs. Future extensions include AI-assisted analysis of reflective artefacts and interaction logs to identify patterns in reasoning and AI usage. These tools are intended to augment, not replace, human judgment and will follow the same principles of transparency and responsibility promoted in learning activities.

7. Discussion

The REALISE framework contributes to ongoing discussions on how engineering education can remain pedagogically meaningful in the presence of increasingly capable generative AI tools. Rather than treating AI as a threat to assessment integrity, this work demonstrates that learning activities and evaluations can be redesigned so that understanding, experimentation, and interpretation cannot be trivially automated.

The Semantic Bridge paradigm provides a unifying lens across courses, revealing a structural limitation of AI in experimental and system-oriented domains: the inability to autonomously ground abstract representations in physical or contextual meaning. By positioning assessment tasks at this boundary, REALISE re-centers evaluation on human judgment, validation, and responsibility. At the same time, the framework highlights practical challenges, particularly the increased demand for reflective documentation and process-oriented assessment. Addressing scalability through carefully designed AI-assisted evaluation tools—used transparently and responsibly—emerges as a necessary complement to the pedagogical approach rather than a contradiction of it.

8. Conclusions and future work

This paper has presented the conceptual foundations and instructional design of the REALISE project, an educational framework for responsible and reflective AI integration in engineering education. Grounded in the Semantic Bridge paradigm, REALISE demonstrates how learning activities and assessments can remain meaningful and resistant to trivial automation.

Future work will focus on systematic empirical evaluation, refinement of assessment instruments—including controlled AI-based analysis—and transfer of the methodology to other engineering disciplines and institutional contexts.

References

- Ambalavanan, R., Snead, R. S., Marezika, J., Towett, G., Malioukis, A., & Mbogori-Kairichi, M. (2025). Ontologies as the semantic bridge between artificial intelligence and healthcare. *Frontiers in Digital Health*, 7(1668385). <https://doi.org/10.3389/fdgth.2025.1668385>
- Bravo-Pacheco, J. M., Esteban-Romero, S., Martín-Fernández, I., Gil-Martín, M., & Fernández-Martínez, F. (2025). Transforming electronics education with AI: A challenge-based learning approach. In *Proc. of the Education and New Developments 2025* (pp. 161-165). Budapest, Hungary.
- Esteban-Romero, S., Marini, M. R., Lanzino, R., & Gil-Martín, M. (2025). Towards multi-view hand pose recognition using a fusion of image embeddings and leap 2 landmarks. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (Volume 3: ICAART)* (pp. 918-925). Setúbal, Portugal: SciTePress.
- Gil-Martín, M., Luna-Jiménez, C., San-Segundo, R., Pardo, J. M., & Fernández-Martínez, F. (2023). Human activity recognition systems based on artificial intelligence for enhancing student motivation. In *16th Annual International Conference of Education, Research and Innovation (ICERI2023)* (pp. 7731-7740). Seville, Spain: IATED.
- Gil-Martín, M., Marini, M. R., Martín-Fernández, I., Esteban-Romero, S., & Cinque, L. (2025). Hand gesture recognition using MediaPipe landmarks and deep learning networks. In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (Volume 3: ICAART)* (pp. 24-30). Setúbal, Portugal: SciTePress.
- HiveMQ. (2023). *MQTT essentials – All core concepts explained*. Retrieved January 21, 2026, from <https://www.hivemq.com/info/mqtt-essentials/>
- Kuthuru, A. (2025). Semantic data contracts: A new integration paradigm for enterprise AI and database systems. *European Journal of Computer Science and Information Technology*, 13(31), 1-9.
- Semantic bridge: Universal multi-hop question generation via AMR-driven graph synthesis. (2025). *arXiv preprint arXiv:2508.10013*.