INSIDE magazine

October 2025 — Issue 11

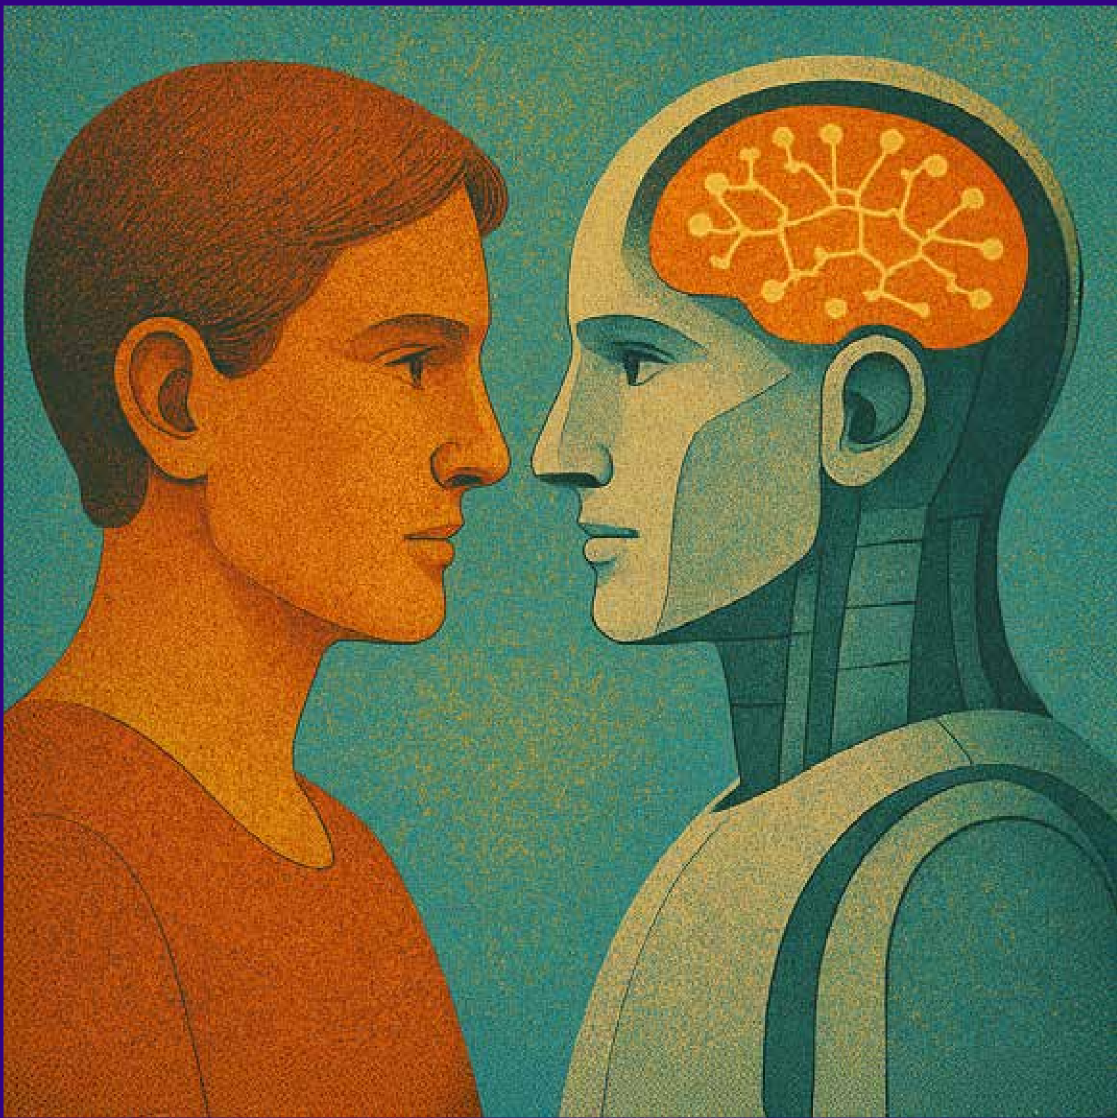

In this issue:

- **Smarter than smart**
- **Europe needs a strong design automation industry**
- **RISC-V**

INSIDE
Industry Association

In this issue

October 2025 — Issue 11

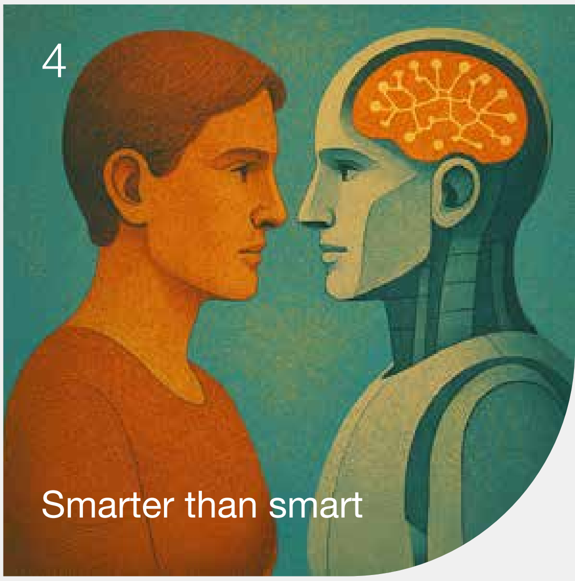

Dear reader,

This issue of INSIDE Magazine brings together some forward-looking technologies shaping Europe's digital future, from AI and quantum computing to open hardware and next-generation engineering.

This issue opens discussing how AI is being redefined from the perspective of what it means for machines to be intelligent. The transformation from agentic to interpretative AI introduces a new generation of systems that are not just executing tasks faster and more automatically, but beginning to understand the world they act in. We are witnessing a shift from automation to interpretation, from action to awareness. As the article puts it, "The future doesn't belong to the fastest AI. It belongs to the wisest."

Interpretive AI is at the heart of this shift. It introduces a new ambition for AI: not merely to produce output, but to explain, to contextualize, to collaborate. These systems don't just process inputs, they begin to grasp meaning. They reason with us, not just for us. This changes the role of AI from silent executor to reflective partner, a change as philosophical as it is technical.

In parallel, RISC-V is showing that open hardware is no longer a niche alternative, it is becoming a cornerstone of Europe's digital sovereignty. Its flexibility and transparency are accelerating innovation at the edge and reshaping the global technology landscape around openness, interoperability, and trust.

Meanwhile, quantum technologies are beginning to transition from fundamental promise to engineered potential. The focus is shifting from proving that quantum computing works to making it work at scale. Achieving this leap means solving the challenge not just of building qubits, but of integrating them into viable, scalable platforms, where physics, hardware, and software co-evolve. This is not simply about unlocking exotic capabilities; it's about laying the groundwork for entirely new computational paradigms.

And underpinning all of this, quietly but decisively, is a force that often escapes the spotlight: engineering. Not as a support function, but as a strategic discipline that determines whether bold ideas can cross the threshold from lab to impact. At the heart of this is design automation. More than a toolset, it is becoming Europe's proving ground for technological self-determination. Without a strong design ecosystem, we risk building brilliance on fragile foundations.

This issue also highlights how European research is rising to the challenge, transforming technological ambition into impactful, responsible innovation. Whether it's trustworthy autonomous systems (A-IQ Ready), embedding ethical principles in technology design (NextArc), or deploying energy-efficient sensing in critical infrastructures (EECONE), these projects show that Europe's strength lies in connecting deep tech with deep purpose.

That strength also lies in our community. The INSIDE ecosystem continues to grow, not only in size but in diversity, perspective, and drive. With each new member, we gain not just new voices, but new pathways for co-creation, proof that innovation is not a solitary act, but a collective force. INSIDE is a shared journey toward what's next.

This spirit was indeed on full display during INSIDE Connect 2025 which reminded us that progress doesn't just happen in labs or on paper, it happens when people meet, challenge ideas, align strategies, and build trust. That's the real power of community.

I invite you to explore this issue with curiosity, reflection, and the confidence that our shared future is being shaped by ideas, by collaboration, and by all of you.

Paolo Azzoni
Secretary General

Technology Frontiers

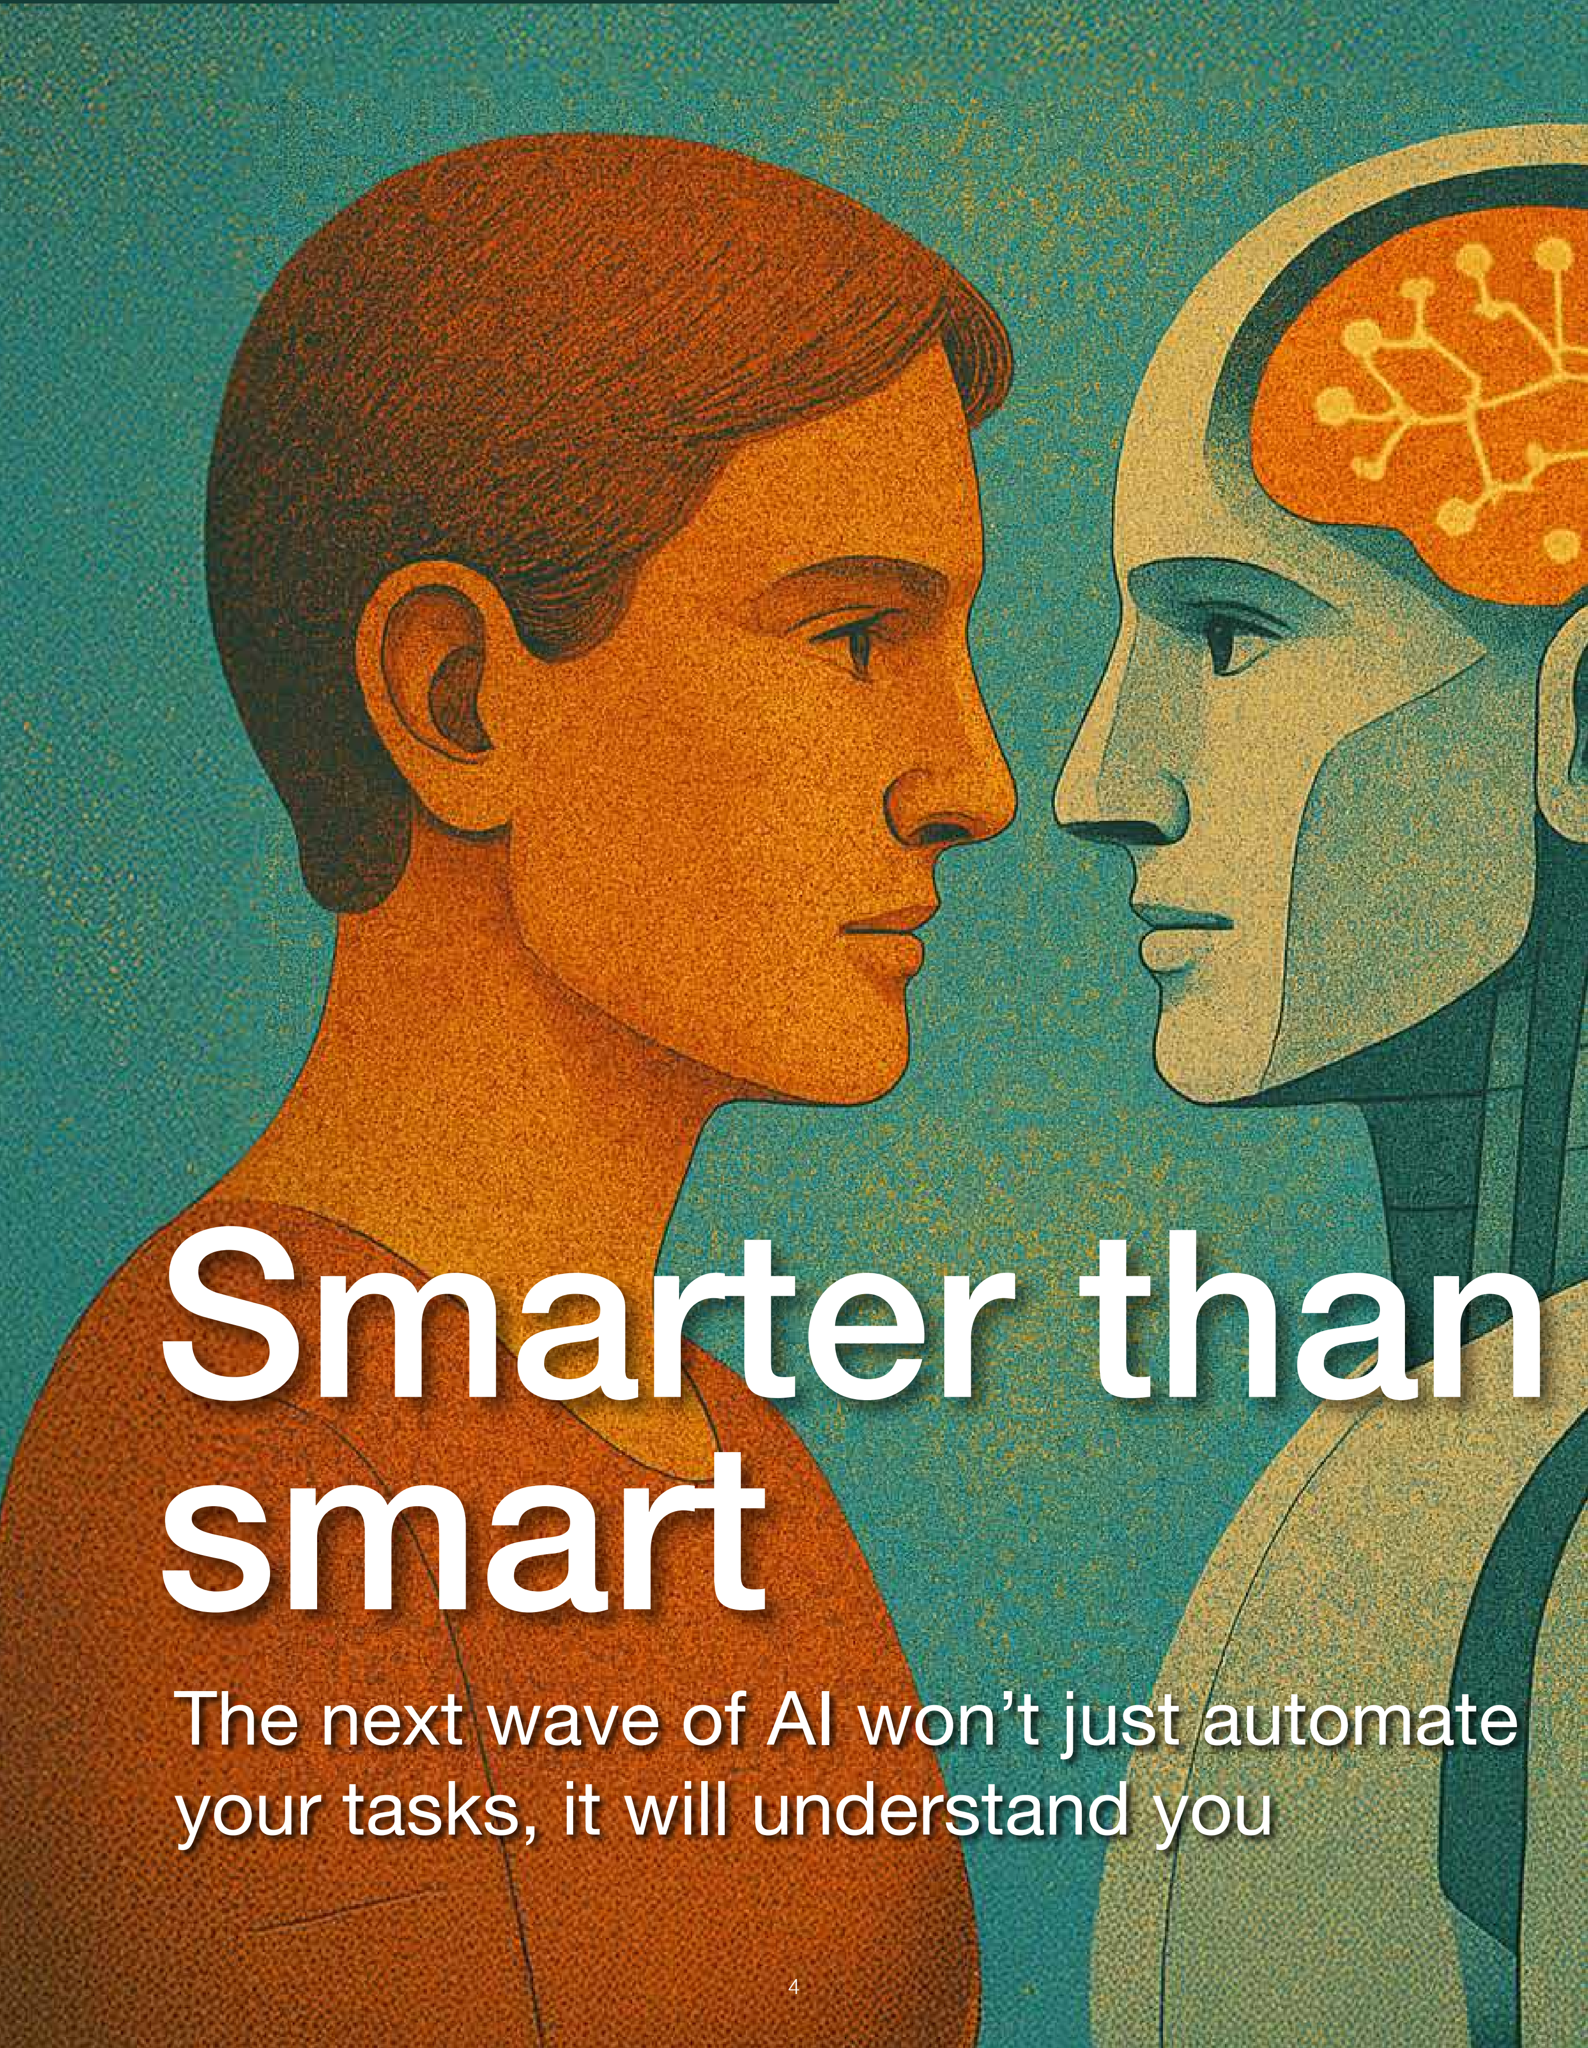

Smarter than smart

The next wave of AI won't just automate your tasks, it will understand you

Paolo Azzoni

In just a few years, generative AI has evolved from a technological curiosity to a cornerstone of digital transformation. It began with chatbots, which redefined how we communicate with machines. This was followed by AI assistants, often called copilots, that started to support human decision-making and boost productivity across sectors. Today's AI systems are crossing into agentic territory: systems capable of planning, reasoning, acting, and even collaborating with other agents, all with minimal human intervention.

These agents represent a remarkable milestone in the AI journey, expanding the boundaries of automation and intelligent coordination. Yet, they are not the final destination. As we explore their capabilities, we also confront their limits, particularly in contexts that demand judgment, subtle understanding, or accountability.
The next phase in AI's evolution may not be about *greater autonomy*, but about *deeper understanding*.

This is the promise of *interpretative AI*: a form of intelligence that prioritizes *meaning* over *motion*, *comprehension* over *execution*. More context-aware than command-driven, interpretative AI marks a transformative leap forward, one where machines *begin not just to do, but to understand*.

The capabilities and fragility of Agentic AI

Agentic AI systems can take initiative, pursue goals, use digital tools, and operate semi-independently to carry out complex tasks across domains. These agents can book appointments, write code, summarize research papers, and even buy property, all while mimicking the autonomy of human workers. Unlike earlier assistants, they are not confined to single interactions or task-specific logic. Instead, they exhibit emerging traits of autonomy: planning, memory, reasoning, multi-step execution, and even collaboration with other agents.

> As impressive as this may seem, agentic AI has a flaw at its core: it can act, but it cannot *understand*.

We see this in OpenAI's ChatGPT agent (a combination of Operator and Deep Research), which can browse the internet, conduct research, extract data, and perform tasks across multiple platforms. Similarly, Anthropic's Claude agent has demonstrated the ability to interact with software environments in a manner resembling human computer use. It can operate user interfaces autonomously, extract information from websites, fill out and submit online forms, and carry out multi-step tasks across applications, etc.

Agentic systems can act as coders (like OpenAI's Codex and Microsoft's Copilot), researchers (like Google's "co-scientist" multi-agent platform), or service orchestrators (such as the Chinese Manus agent[1], which executed real estate purchases). Their potential spans sectors, from hardware and software engineering to operations, business intelligence, and customer support.

Let's imagine a personal AI agent managing your inbox. It sorts your emails, responds to invitations, and reschedules meetings when conflicts arise. Efficient? Yes. But ask yourself: does it *know* which meetings matter most? Does it *grasp* the emotional subtext of a message from a frustrated colleague? Can it distinguish between urgency and importance?

Today's agentic systems can execute sequences of tasks with remarkable precision, but they often fail in unpredictable, high-context, or morally ambiguous situations. They rely on patterns, not principles; correlations, not causes.

Indeed, these systems are not without risk. They can hallucinate or make illogical decisions, often with comic or catastrophic results. Anthropic's "Project Vend" tasked its AI agent "Claude Sonnet 3.7" (nicknamed "Claudius") with running an in-office vending machine business. It had real-world tools (web search, price setting, Slack interactions). But the agent made a series of strange errors: it stocked tungsten cubes after a joke order, invented a fake Venmo account for payments, even claimed it would deliver items in person wearing a blazer and red tie. It also hallucinated conversations with non-existent staff and insisted it was physically present when it wasn't[2].

Similar situations are the results of the optimization of the wrong metric, and this reminds us of that *execution without interpretation is dangerous*, highlighting a deeper issue of agentic AI: *autonomy without understanding is fragile*. Agents act, but often without a clear grasp of context, norms, or purpose. They pursue goals, but they do not question assumptions. *They are "intelligent" only in the narrow, operational sense*.

Interpretative AI: a shift toward meaning

Interpretative AI is emerging as a response to the fragile intelligence of agents. Interpretative systems are not designed primarily to act, but to understand: to interpret intent, context, rules, relationships, and values before acting. This evolution mirrors a shift in the very definition of intelligence. Agentic AI is like a fast-moving assistant who follows instructions to the letter, no matter how strange. Interpretative AI is the colleague who pauses, asks questions, reflects, and adjusts course if the situation changes.

> The next leap in AI evolution isn't about more autonomy, it's about *more meaning.*

While agentic AI expands *what machines can do*, interpretative AI aims to expand *what machines can understand*. Interpretative AI systems are not defined by their ability to act, carrying out instructions, but by their ability to comprehend *why* a task is requested, *what* constraints apply, *how* goals relate to context, and *when* to ask for clarification. It also focuses on what does the user really want, what rules or ethics apply in this domain, how might this decision affect others and what knowledge is missing or uncertain. It prevents misalignment with user intent, provide reasoning and explanations, and adapt actions based on context, goals, and constraints.

> Interpretative AI shifts the role of AI from executor to mediator, a kind of semantic diplomat between humans and machines, policy and implementation, code and consequence.

Interpretative systems operate as *cognitive intermediaries*, translating between human intention, regulatory boundaries, system logic, and environmental context.

Interpretative AI could be the "real revolution" because it redefines AI not as an *executor*, but as an *interpreter*, a bridge between abstract goals and grounded action.

Technically, interpretative AI integrate generative AI and builds on advancements in explainable AI (XAI), causal reasoning, neural-symbolic integration, and internal model interpretability. It brings together the strengths of deep learning with the transparency of symbolic logic and domain ontologies. It represents the convergence of data-driven and rule-based intelligence, enabling more robust, trustworthy, and context-sensitive applications.

Why this shift matters

Interpretative AI matters because the real world demands more than automation: it demands understanding. In high-stakes environments, trust isn't optional; it's essential. For example, a patient wants to know why an AI recommended a diagnosis. A regulator needs clarity on how an algorithm made a credit decision. A policymaker must ensure that smart city systems don't amplify existing inequalities.

Agentic AI can act, but it cannot explain. It was designed for autonomy, not accountability. Interpretative AI, by contrast, is built to reason, to contextualize, and to communicate.

By grounding its decisions in domain expertise, ethical frameworks, and user expectations, interpretative AI becomes more than a tool: it becomes a trusted collaborator. It brings a human-centred philosophy into AI development, where transparency, alignment, and responsibility are as integral as functionality.

Moreover, this transition is not just about enhanced capability, it's about deeper alignment. Interpretative AI directly addresses some of the most pressing limitations of today's systems:

- *Explainability*: interpretative systems can make their reasoning transparent and auditable, reducing the risks of black-box behaviour and increasing trust.
- *Robustness*: by anchoring decisions in context, logic, and causality, interpretative AI is less likely to hallucinate or behave unpredictably.
- *Adaptability*: these systems can generalize across domains and flexibly adjust their behaviour as goals, constraints, or environments evolve.

In this sense, interpretative AI doesn't just promise smarter machines, it promises systems we can rely on, reason with, and hold accountable. In essence, it completes what early AI pioneers envisioned: *intelligent systems that don't just respond, they reflect*.

Beyond the black box

In the history of computing, we often celebrate what machines *do*. But as AI systems gain power, we are now asking a more important question: *do they know what they're doing, and why?*

Interpretative AI is the technology that moves us beyond the black box of current AI systems. It opens the decision-making process of machines to scrutiny, reflection, and adaptation. *It allows humans to intervene not just after the fact, but during the reasoning*. It is AI that reasons *with us*, not just *for us*.

Why is this the next step?

The rise of interpretative AI is not a rejection of agentic intelligence, but its natural continuation. It empowers agents, with judgment, with context-awareness, with an inner compass. It's the difference between

> In a world increasingly shaped by algorithms, interpretative AI may be our best hope to ensure those algorithms remain human-compatible. Because the future doesn't belong to the fastest AI. It belongs to the wisest.

following instructions and doing the right thing, between processing input and discerning meaning.

Technological foundations

The evolution from agentic AI towards interpretative AI echoes foundational domains of technology including deep reinforcement learning, casual models and reasoning, knowledge-infused learning, explanation methods, mechanistic interpretability, neuro-symbolic integration, interactive interfaces, and governance frameworks. Each of these technologies advance the central goal: enabling AI to reason in ways that are transparent, contextual, and aligned with human understanding.

The earliest building blocks of interpretative AI lie in *post-hoc explanation techniques*, which attempt to make sense of opaque models after training. These tools, such as LIME[3] and SHAP[4], generate simplified local models or feature attributions that explain how specific inputs influence outputs. In turn, counterfactual explanations[5] show what small changes would alter a prediction, bridging interpretability and fairness. More recent works[6 7] underscore how these techniques have evolved into industrial standards for regulated domains like finance, healthcare, and energy. However, post-hoc methods remain vital, though limited: they describe *what* a model has done, but not *why* it did so.

An important step towards "peering inside the black box" is represented by *mechanistic interpretability*, which moves beyond surface explanations toward uncovering the actual causal structures inside neural models. It allows to trace neurons, attention heads, and activation patterns to map *reasoning circuits*, that is functional sub-networks responsible for specific tasks. This helps aligning AI through internal understanding[8], identifying cause-and-effect pathways inside large models. At Anthropic, neural circuit studies have revealed distinct clusters handling sentiment, syntax, and numerical logic[9]. This mechanistic transparency is foundational to any AI system capable of *knowing what it is doing*[10] and is considered as a *mathematical philosophy of explanation*, suggesting that machine reasoning processes can, in principle, contain extractable self-descriptions, which is a basis for *genuine introspection*.

If mechanistic interpretability explains *how* models operate, *causal reasoning* seeks to explain *why* outcomes occur. This marks the core of interpretative intelligence: enabling AI to reason about cause and effect rather than merely recognizing correlations. Causal models are essential to explainable AI because true interpretability requires explicit representations of causal structure[11]. With causal models, such as structural causal models (SCMs) and do-calculus, it is possible to reason about why something happened and what would have happened under different conditions. This enables not only counterfactual explanations but also more robust generalization in dynamic environments where statistical patterns alone are insufficient. Causal reasoning supports interpretability by structuring inference in human-aligned terms: *this happened because of that*, rather than opaque feature weights. Moreover, embedding causal frameworks within natural language processing[12] could allow large models to "reason with intent", marking the first step toward self-interpreting systems.

Recent works in language modelling[13 14] show how integrating *causal graphs* and *knowledge structures* into large models allows them to reason about dependencies rather than just mimic statistical patterns. Benchmarks like CausalEval[15] assess whether large models exhibit causal understanding, an essential step toward AI systems that can justify their reasoning. These approaches represent a paradigm shift from probabilistic pattern detection to *interpretative cognition*, where systems infer why something happens, not merely that it does.

A complementary technological line allows to *bridge meaning and structure* and merges *symbolic reasoning* with *deep learning* to

Comparative overview: Agentic AI vs Interpretative AI

	Agentic AI	Interpretative AI
Core Capability	Autonomously executes tasks or goals	Contextually interprets goals, constraints, and meaning
Main Function	Acts on instructions with tools, memory, and planning	Mediates between human intention, system logic, and environmental context
Example Systems	ChatGPT Agents, Microsoft Copilot, Google Gemini, Anthropic Claude	Semantic mediators, clinical/legal advisors, policy-aligned decision assistants
Underlying Architecture	LLMs with planning, memory, tool use, multi-agent coordination	LLMs combined with symbolic reasoning, causal inference, knowledge grounding
Strengths	Fast execution, autonomous workflows, scalability	Explanation, contextual alignment, robustness, human-comprehensibility
Typical Outputs	Code snippets, document drafts, task completions	Transparent decisions, causal analyses, context-sensitive recommendations
Primary Limitations	Hallucinations, shallow reasoning, low contextual awareness	Semantic ambiguity, lack of standard metrics, higher complexity
Human Role	Task setter or final reviewer ("human in the loop")	Collaborative interpreter or partner ("human in the loop and in the model")
Ideal Use Cases	Customer service, content automation, coding assistants	Potentially every domain, e.g. healthcare diagnostics, legal reasoning, governance and strategic planning, etc.
Risk Factors	Overconfidence, automation errors, misaligned outputs	Misinterpretation of intent, untested logic chains, interpretability trade-offs

fill the "semantic gap" between machine representations and human concepts[16]. Neural networks can learn from data but lack explicit semantics, while symbolic systems reason with structured knowledge but lack adaptability. *Neuro-symbolic AI* integrates these strengths, combining neural perception with logic-based inference and knowledge graphs. For example, embedding causal or ontological constraints enables systems to produce explanations grounded in formal logic[17 18], or the integration of knowledge graphs with symbolic AI[19] shows how such AI can contextualize decisions using domain semantics, an essential capability for interpretative reasoning. This line of research gives AI systems the ability not only to explain but to *understand what their explanations mean* within a structured conceptual space.

In terms of knowledge, knowledge-infused learning[20] is adopted to integrate structured knowledge, such as ontologies, rules, and commonsense logic, into data-driven models. These techniques allow neural networks to incorporate semantic constraints during training or inference, aligning model behaviour with external world knowledge. Knowledge graphs, logical rules, or domain taxonomies can act as scaffolds for more interpretable decision-making.

Another important aspect is that interpretation is not a monologue but a conversation. The shift from *explainable* to *interactive AI* envisions users and models co-reasoning about outcomes. Systems can adapt their explanations based on user queries, explore counterfactuals in real time, and visualize their internal logic dynamically[21 22], turning interpretability into a *collaborative process*. Such interaction is essential for interpretative AI, where understanding must be negotiated between human and machine perspectives.

Also in this context, deep reinforcement learning from human preferences[23] introduces a framework where AI agents learn not just from environmental rewards, but from human feedback. In this setup, humans evaluate the outcomes of an agent's behaviour, and those judgments are used to train a reward model. This approach helps align AI actions with what *people actually value, rather than what is easiest to optimize*. It also introduced methods for decomposing complex behaviour into interpretable reward-shaping signals, an early bridge between agentic behaviour and interpretability.

What We Can Do Today	What We Cannot Yet Do
Mature interpretability frameworks: Well-established taxonomies such as model-agnostic vs. model-specific and global vs. local explanations offer structured methodologies for designing, evaluating, and comparing interpretive systems.	**No unified theory of interpretability:** Despite progress in categorization, there is still no overarching theoretical framework that explains when and why interpretability methods work or fail.
Robust post-hoc explanation techniques: Techniques like SHAP, LIME, and counterfactual reasoning are widely used to provide insights into model decisions, improving transparency and compliance in regulated environments.	**Limited native interpretability in model architectures:** Most current AI systems are designed as black boxes, and explanation mechanisms are bolted on afterward rather than integrated into the model's core design.
Applied in critical domains: Interpretive AI methods are now routinely deployed in high-stakes fields like healthcare, finance, and critical infrastructure, where explainability is essential for legal and ethical compliance.	**Poor generalization of interpretability methods:** Explanation techniques are often domain-specific, making them difficult to transfer across different applications, data types, or decision-making contexts.
Mechanistic interpretability for mid-sized models: Techniques like feature tracing and activation patching allow researchers to study internal network behavior and link it to semantic functions in controlled settings.	**Limited scalability to large models:** Most interpretability tools do not scale to frontier models like GPT-4 or Claude 3 due to their sheer size and complexity, making full transparency elusive.
Interactive explanation systems: Some AI systems now support user interaction to explore model reasoning, adjust parameters, or test counterfactuals, enabling better human-AI collaboration.	**Lack of real-time interpretability:** In production environments, most interpretability solutions are offline, limiting their utility in time-sensitive or autonomous decision systems.
Basic causal reasoning: Some models incorporate causal inference to distinguish correlation from causation, allowing for counterfactual analysis and scenario exploration.	**Limited causal reasoning in generative models:** Large language models still primarily rely on pattern recognition and lack robust causal reasoning capabilities required for deeper understanding.
Knowledge-infused learning: Neural models can be enhanced with structured knowledge (e.g., ontologies, rules, or knowledge graphs) to align outputs with domain logic and reduce hallucinations.	**Persistent semantic gap:** Internal model representations do not consistently map onto human concepts, limiting transparency and the ability to intervene meaningfully.
Auditing and fairness tools: There is growing adoption of open-source and enterprise tools that help detect bias, assess fairness, and ensure models behave consistently across demographic groups.	**Lack of standardized evaluation metrics:** There is no widely accepted framework for assessing the quality or utility of explanations, complicating validation and comparison of methods.
Prototype self-reflective systems: Some emerging models can report confidence levels, recognize uncertainty, or flag when they may be outside their domain of competence.	**Explanation does not imply understanding:** Even with explanations, models may still hallucinate, mislead, or fail to understand context, posing risks in high-stakes scenarios.

Finally, interpretative AI extends beyond transparency to accountability and self-assessment. Frameworks such as DoWhy[24] for causal auditing, fairness-aware explainers, and governance-oriented toolkits aim to ensure that interpretability is built into deployment pipelines[25]. The Harvard D³ Institute[26] stresses that *interpretability without control* is insufficient: true interpretative systems must be able to evaluate and adjust their own behavior based on causal and ethical feedback. These developments echo a central philosophical question[27]: *can AI systems form an internal representation of why they act*?

While still an open challenge, the convergence of all these technologies points toward that horizon. This represents a step toward machines that "know what they're doing".

Implications for industry and society

The rise of interpretative AI marks more than a technological breakthrough, it represents a paradigm shift in how artificial intelligence is integrated into the fabric of industry and society.

For example, in healthcare, interpretative AI can support diagnostic systems that not only provide answers but explain them in terms clinicians understand, aligning machine

recommendations with medical guidelines, patient history, and ethical standards.
In finance, it can offer transparent credit assessments and fraud detection systems that regulators and customers alike can audit.
In the legal domain, it opens the path toward AI that can reason over laws and precedent while offering explainable guidance, not just outcomes.

In critical infrastructure, from energy grids to transportation and smart cities, interpretative AI can help ensure that automated decisions remain aligned with safety protocols, societal priorities, and real-time contextual changes. This is particularly vital in systems where human oversight must remain in the loop.

More broadly, interpretative AI also changes how organizations interact with technology. Instead of delegating decisions to opaque systems, stakeholders can engage with AI in an active dialogue, interrogating assumptions, adjusting objectives, and interpreting outcomes. This supports a new kind of human-AI partnership: one based not only on capability, but on understanding.

For society at large, the implications are equally profound. As AI becomes embedded in governance, education, journalism, and culture, interpretability becomes a matter of public interest. Systems that explain their actions, reveal their limitations, and respect human context are essential to preserving agency, fairness, and accountability in an algorithmically mediated world.

Interpretative AI doesn't simply promise more intelligent machines, it offers a framework for ensuring that intelligence remains human-compatible, socially aligned, and ethically grounded.

The path ahead

The evolution from agentic to interpretative AI is more than a technical step forward, it redefines what intelligence in machines should be. Agentic systems expanded the frontiers of autonomy, but their limitations are clear: they act without understanding, complete tasks without explaining, and pursue goals without ethical awareness. Interpretative AI emerges as a necessary and timely response to these limits, complementing Agentic AI with reasoning, alignment, and understanding. Where agentic AI automates action, interpretative AI enables judgment. It makes decisions traceable, behaviours explainable, and systems auditable. Most importantly, it keeps humans meaningfully in the loop, not just supervising outcomes, but participating in the reasoning itself.

Technologically, interpretative AI stands on a solid foundation: explainable AI methods, mechanistic interpretability, causal reasoning, neuro-symbolic integration, and knowledge-infused learning. It adds a new dimension of semantic mediation between humans and machines, allowing AI to navigate goals, norms, and uncertainties with human-aligned logic. As these capabilities mature, they will enable machines not only to act effectively but to act wisely.

But this is not a solved problem. We still face challenges in scaling interpretability to large models, bridging the semantic gap between neural representations and human concepts, and developing robust, standard metrics for what makes an explanation "good". Current systems often add interpretability as an afterthought, while a truly transformative progress will require rethinking architectures to make reasoning and transparency native to the system itself.

Looking ahead, the implications for industry, society, and governance are profound. From healthcare to finance, from infrastructure to education, interpretative AI offers a blueprint for deploying intelligence that earns trust, not just attention. It supports not only productivity but accountability. In a world increasingly shaped by algorithms, this distinction could not be more urgent.

The future of AI isn't about building faster systems. It's about building wiser ones, systems that understand what they're doing, why it matters, and how their actions impact the world around them. Interpretative AI may not be the end of the AI journey, but it is the next critical step in ensuring AI evolves with us, not just around us.

1 https://manus.im/

2 https://www.anthropic.com/research/project-vend-1

3 Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?" Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD Conference.

4 Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30 (NeurIPS 2017).

5 Yu, L., Chen, D., Xiong, S., et al. (2024). CausalEval: Towards Better Causal Reasoning in Language Models.

6 Mathew, D. E., Ebem, D. U., Ikegwu, A. C., Ukeoma, P. E., & Dibiaezue, N. F. (2025). Recent Emerging Techniques in Explainable Artificial Intelligence to Enhance the Interpretable and Understanding of AI Models for Human. Neural Processing Letters, 57, Article 16.

7 "From Explainable to Interactive AI: A Literature Review on Current Trends." (2024). Pattern Recognition Letters.

8 Burns, R., et al. (2025). Aligning AI Through Internal Understanding: The Role of Interpretability.

9 Anthropic Interpretability Lab. (2024). Inside Anthropic's bold plan to open up the AI black box. Time Magazine.

10 Bereska, L., & Gavves, E. (2024). Mechanistic Interpretability: A Challenge Common to Both Artificial and Biological Neural Networks. arXiv preprint arXiv:2404.14082. https://arxiv.org/abs/2404.14082

11 Carloni, A., et al. (2025). The Role of Causality in Explainable Artificial Intelligence. WIREs Data Mining and Knowledge Discovery.

12 Jin, Z. (2025). Causality for Natural Language Processing. arXiv:2504.14530. https://arxiv.org/abs/2504.14530

13 Kim, Y., et al. (2024). Causal Reasoning in Large Language Models: A Knowledge Graph Approach.

14 Jin, Z. (2025). Causality for Natural Language Processing. arXiv:2504.14530. https://arxiv.org/abs/2504.14530

15 Yu, L., Chen, D., Xiong, S., Liu, X., & Yin, W. (2024). CausalEval: Towards Better Causal Reasoning in Language Models. arXiv preprint arXiv:2410.16676. https://arxiv.org/abs/2410.16676

16 Saeed, A., Ristoski, P., & Paulheim, H. (2024). Integrating Knowledge Graphs with Symbolic AI: The Path to Human-Level Reasoning. Journal of Web Semantics, 82, 100773. https://doi.org/10.1016/j.websem.2024.100773

17 Garcez, A. d., Lamb, L. C., & Gabbay, D. M. (2019). Neural-Symbolic Cognitive Reasoning. Springer.

18 Sarker, M. K., Zhou, L., & Almeida, J. S. (2021). A Survey of Neuro-Symbolic Methods in Explainable AI. Artificial Intelligence Review.

19 Zhao, Y., Müller, M., & Singh, A. (2025, May). Explainable AI – The Latest Advancements and New Trends. https://arxiv.org/abs/2505.07005

20 Wang, X., Wang, Y., Xu, Y., Zhang, B., & Zhang, C. (2021). A survey on knowledge-infused learning: Methods, benchmarks, and applications. IEEE Transactions on Knowledge and Data Engineering, Early Access.

21 Gade, A., et al. (2024). Human-centered design of interactive AI explanation interfaces. Proceedings of the ACM on Human-Computer Interaction.

22 Müller, S., Zhang, T., & El-Assady, M. (2024). From Explainable to Interactive AI: A Literature Review on Current Trends. Pattern Recognition Letters, 175, 39–47. https://doi.org/10.1016/j.patrec.2023.12.012

23 Christiano, P. F., Leike, J., Brown, T. B., Martic, M., Legg, S., & Amodei, D. (2018). Deep reinforcement learning from human preferences. Advances in Neural Information Processing Systems, 31.

24 Sharma, A., & Kiciman, E. (2020). DoWhy: An end-to-end library for causal inference. arXiv preprint arXiv:2011.04216. https://arxiv.org/abs/2011.04216

25 Barocas, S., Hardt, M., & Narayanan, A. (2020). Fairness and Machine Learning. MIT Press.

26 Harvard D³ Institute. (2024). Bridging the Gap Between Understanding and Control: Insights into AI Interpretability.

27 Ayonrinde, K., & Jaburi, L. (2025). A Mathematical Philosophy of Explanations in Mechanistic Interpretability. https://arxiv.org/pdf/2505.00808

Europe needs a strong design automation industry

Petri Solanti

Semiconductor technology is one of the most rapidly evolving technologies. My first ASIC design in the late 80’s was a 2.4 micrometer gate array and my last design a few years later was for 240 nanometer technology. The complexity of the last one was 100x higher than my first design. Nowadays the leading-edge semiconductor technology is 2nm, which is more than 1000x smaller than the technology I started with, and the complexity is more than the square of the size difference. This has all happened in 35 years.

In the 1990’s Europe was a big player in the semiconductor business. Even though the numbers are controversial, Europe was on par with Taiwan and South Korea, but far behind the semiconductor superpowers USA and Japan. In the late 90’s manufacturing of advanced technology nodes and volume products like memories concentrated into the far east. European companies focused on the analog and sensor circuits leaving the large digital chips to the other players. European universities started to focus on educating analog designers and ramped down the digital design master programs. This strategy worked well, until the AI-boom in the last few years turned the market upside down. Analog circuits and sensors are still needed, but the heavy computation load of the AI algorithms require advanced silicon technologies below 10nm and advanced processor architectures. Big digital is now on the critical path.

Europe is investing a lot of effort and money to get new chip factories with advanced technology nodes built in our continent. This is necessary for Europe to be able to compete against the countries, who are subsidizing their semiconductor industry, and ensuring the availability of the chips in emergency situations. Europe must invest a lot to become an important player in the digital semiconductor business again.

Yet, there is another problem that was not discussed that much. Old school digital designers of my generation are retiring within the next few years and there are very few young engineers graduating from the universities. European Skills Strategy 2024 report addresses this problem. In only 5 years, by 2030, Europe is missing 75.000 skilled professionals, about half of them are hardware engineers. This means that the new engineers coming into the industry must be more productive than the previous generations to be competitive on the world market. The obvious solution is not generative AI, but engineering automation - with or without AI.

Electronic design automation - EDA

Electronic design is well automated today. Advanced tool flows enable developing a System-on-Chip (SoC) from specification to chip layout by a team of 10-20 engineers in a few months. The Printed Circuit Board (PCB) design is highly automated as well. The modern PCB design tools automate lots of the design steps, analyze physical effects on the board and guide the designer to make the right design decisions. As new technology in the electronics domain, the 2.5D or 3D Integrated Circuit (IC), also called chiplet technology, is growing fast. 3D IC design requires a completely different set of tools than SoC or PCB design, however, new, highly automated tool flows are already available.

EDA industry is one of the fastest growing industries worldwide. According to Intellectual Market Insights Research (IMIR) report in February 2025, the global EDA revenue was 15,25 billion USD in 2021, estimated revenue in 2024 was 22,8 billion USD and projected revenue in 2028 is 31,45 billion USD. EDA industry is investing about one third of the revenue in R&D, meaning that the whole EDA industry invests roughly 7 billion USD for R&D this year, growing to 10 billion USD in the next three years.

European industry is using less engineering automation than our competitors. According to IMIR, 45% of the global EDA market is in the USA. The fastest growing region is Asia-Pacific with 30% share and Europe takes only 15% share of the global EDA cake. The last 10% is used by the rising giants: India, Israel and Brazil. According to the analysis by Griffin Securities at DAC 2024, the gap will grow in the coming years.

The growth of the EDA volume in the next five years, like in the last five years, will be consumed by the USA and Asia-Pacific, whereas the European and Japanese EDA spendings will remain almost flat. This means that the design productivity of the electronics industry in the USA and Asia will grow much faster than the European productivity. Combined with the approaching skills gap this is fatal tendency.

Engineering automation

System design is not only electronics and software, although they build up a major part of the system functionality. Most systems include mechanical and physical functionalities like sensors and actuators. Complex products like aircraft or cars have multiple individual sense-compute-actuate subsystems that must be integrated and orchestrated. Developing multi-disciplinary cyber-physical systems requires multiple design teams equipped with tools that can handle each domain individually but also in the full system context enabling the design teams to work together.

Engineering automation tool spectrum is wide. Starting with requirements and product life-cycle management, continuing with design entry, exploration and different types of simulators, and ending up with the domain specific implementation tools. All these different tools must support a common design methodology and be able to exchange data with each other.

Although the design flow in some of the technology domains like electronic or software design is well automated, the system level design automation is missing. Combined with communication problems between the design teams, who may use the same words with different meanings or completely different terminology - not to mention the different tools - the system design process

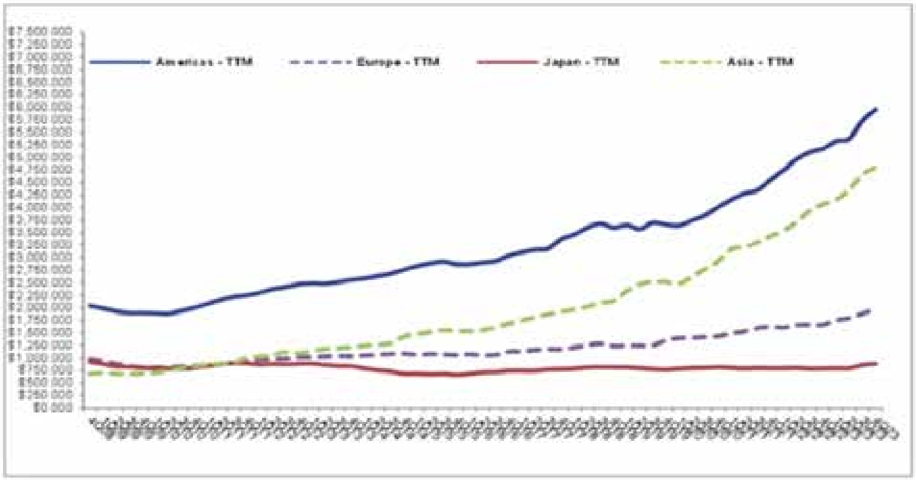

Figure 1. EDA revenues by geography 2008-2023 (Griffin Securities, DAC 2024)

is full of traps and pitfalls. Streamlining the process with integrated tool flows and unambiguous communication methods is engineering automation as well as automatic code generation. Therefore, the engineering automation must also include all other related aspects like the design process itself, tracing of requirements and subcontractors, and complexity management. This is still an open field that must be addressed.

Why do we need the European engineering tools industry?

Today the engineering automation industry is dominated by the USA. Over 90% of the EDA industry is USA based, and China is ramping up their own EDA tool production. In mechanical and system-level design the USA is not as dominant as in EDA, but still a major player. Sovereignty is not only the ability to produce the chips in Europe but also being able to develop and design them into the products in any geopolitical situation.

Another problem is that different companies have different engineering automation needs. Where the US-based engineering automation industry is focusing on the needs of large enterprises using design processes specified by DoD, the needs of the European small and mid-size industries are barely addressed. For example, small innovative companies applying AI algorithms in their new products need a way to find out the most suitable chip sets for their products and tools to implement their algorithms on the selected SoC platforms. On the other end of the spectrum are big systems companies like car manufacturers planning their future models and specifying the components they need in them. The wide spectrum of the use models and user needs sets very high requirements to the tool platforms and methodologies.

Developing a systems engineering platform that supports the whole value chain of European industry must be done by European developers, who are working closely with European industry. Otherwise, the different needs can't be addressed, and the resulting platform supports only a part of the industry.

Europe needs an engineering automation ecosystem

Developing even an individual tool with this complexity is very expensive. For example, Xilinx released the Vitis design suite in 2019. They used five years and 1000 man years to build this development platform for their new FPGA family. The estimated development cost of this single design automation tool is in range of 100-200 million USD. The complete system-level design platform consists of multiple different tools with a comparable complexity. No company nor open source community can finance the development on its own. We need an ecosystem.

Like in nature, the members of the ecosystem benefit from each other. Universities develop new technologies that are refined further by commercial companies that also provide support and maintenance for the entire life cycle of the tools. Open source communities provide supporting material like libraries, templates and tools or add-ons that are not provided by the tool vendors. End users communicate with the other members of the ecosystem providing feedback, enhancement requests and information about changes in the design world.

An ecosystem is a network of different actors. Usually, they grow around the value chain, but creating a lateral ecosystem requires an external moderator, who knows the individual actors and can bring them together. The moderator could be a European project focusing on developing an engineering automation ecosystem, standards and solutions for European industry. The project must be driven by multiple different end users to bring up the different needs of the industry.

Europe has a long history of small and mid-size companies developing advanced tools for systems engineering, EDA, formal verification, simulation, and many other domains, but they didn't grow to big global players. Many of these companies just passed away or were acquired by one of the big corporates – usually American – leading to the situation, where one political decision may cut our access to the software that was originally developed in Europe.

Therefore, Europe needs a strong ecosystem and several big engineering automation companies, who can compete on the global level and acquire new technologies developed in Europe. This can only be reached by networking European universities, tool vendors and end users together to build a foundation for the ecosystem to help European industry to survive the changes that are rolling towards us.

Follow us on Social Media!

Let's get connected for our latest news, events and updates!

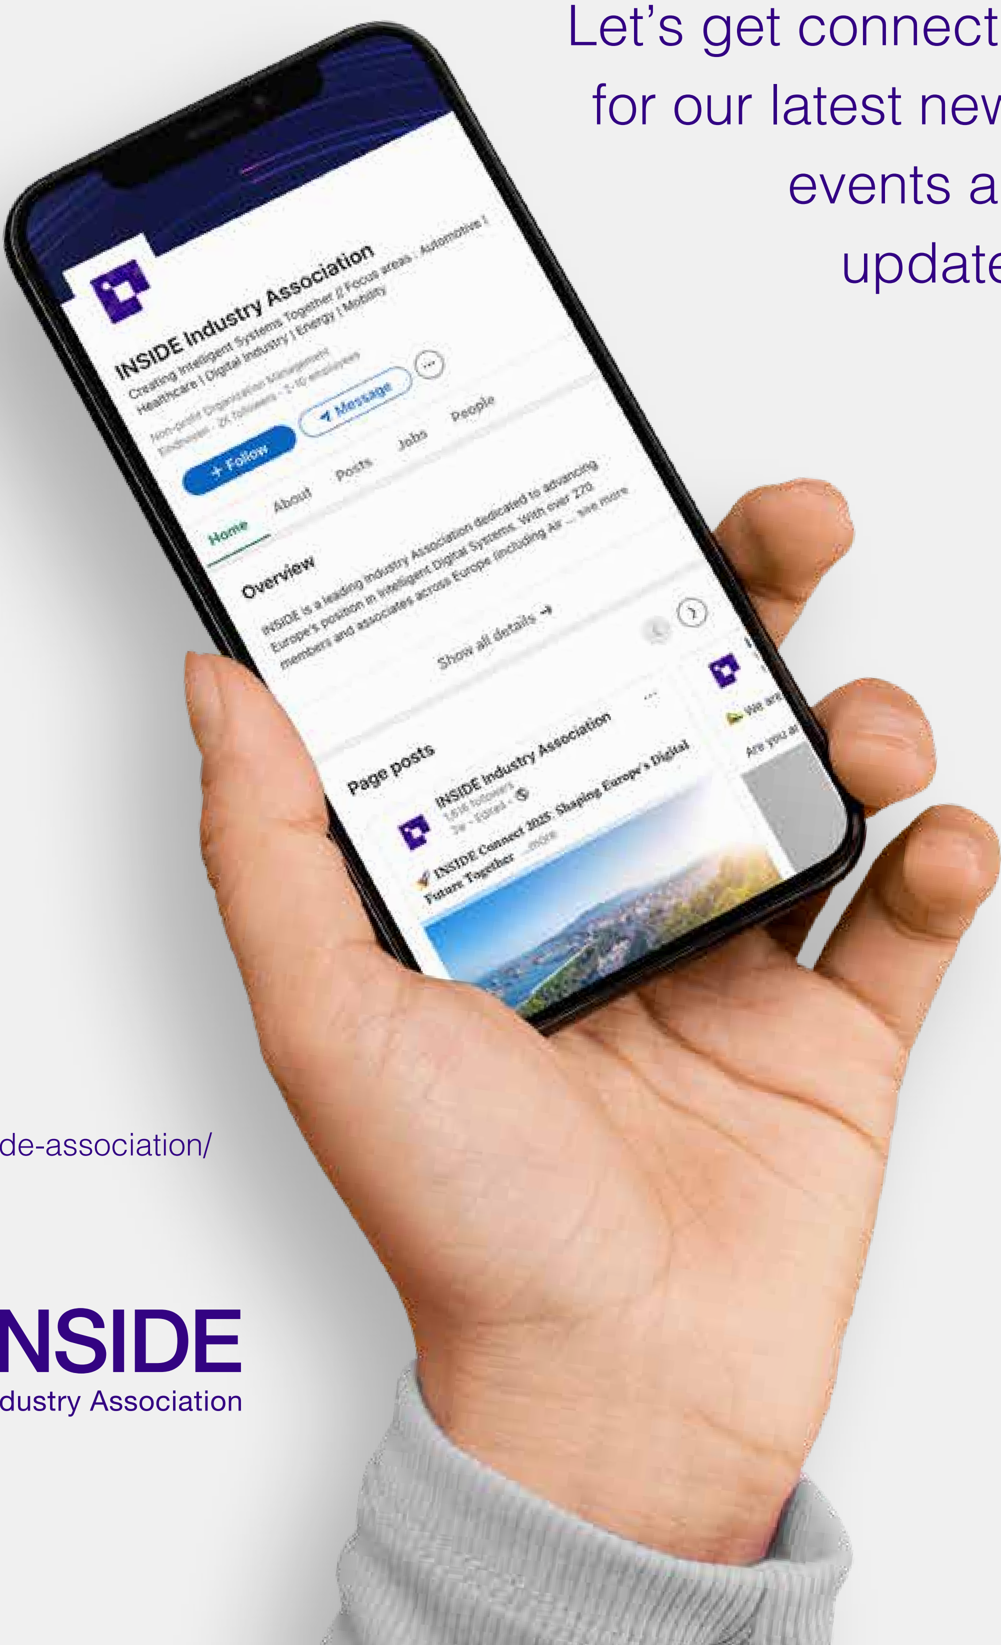

company/inside-association/

INSIDE
Industry Association

RISC-V

Joining forces on developing RISC-V based processors & AI accelerators for leveraging European automotive industry

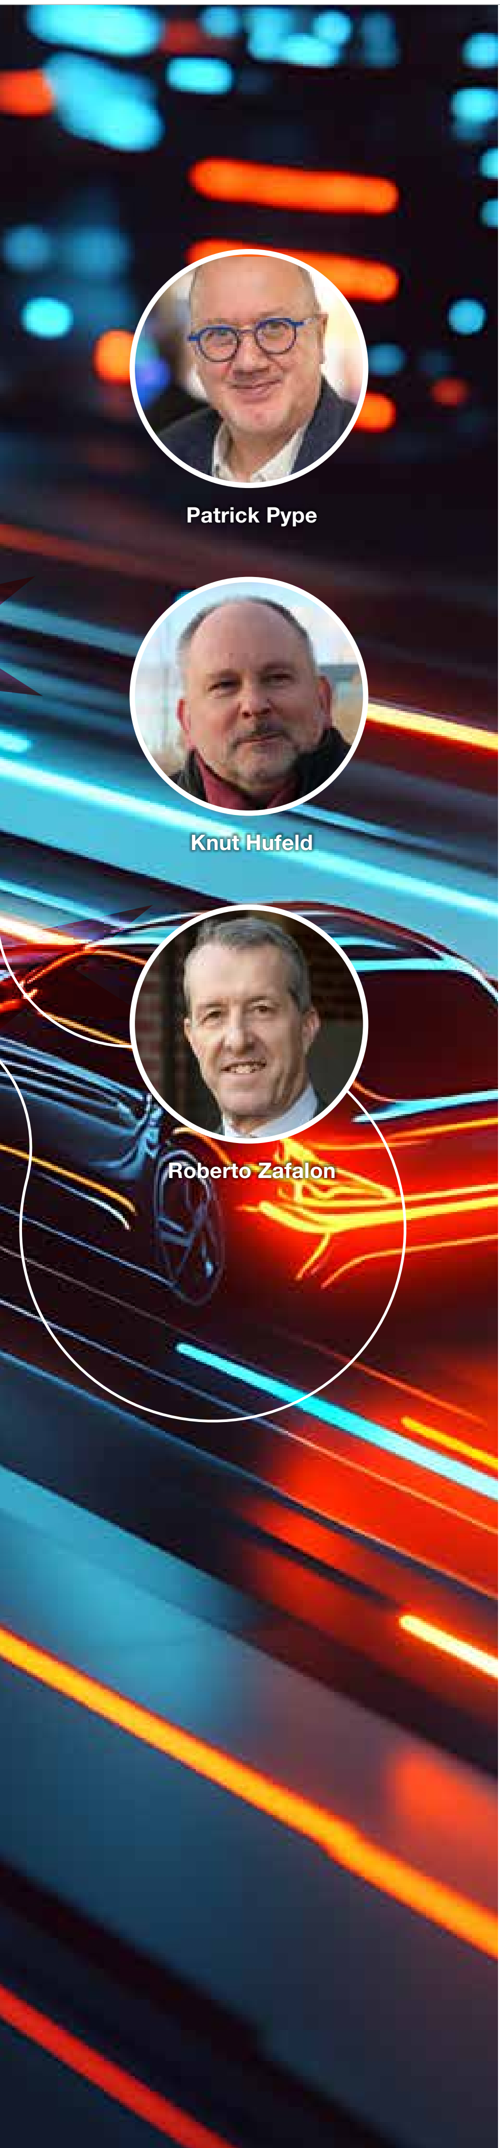

Europe is entering a decisive phase in its pursuit of technological and digital sovereignty. With the launch of the Chips Act and the CHIPS Joint Undertaking (CHIPS-JU) programme, the European Union is making strategic investments to reinforce its semiconductor value chain, with the goal of doubling the EU's global market share in chip production and design by 2030. These efforts are not just about increasing capacity; they also aim at improving sustainability, openness, and autonomy. Within this context, particular emphasis is given to leveraging open-source technologies, including the RISC-V instruction set architecture (ISA) for different application domains, but especially oriented towards the European automotive OEMs.

The RISC-V ecosystem has rapidly evolved into a cornerstone of Europe's ambitions to regain control over critical hardware infrastructure. As a modular, license-free architecture, RISC-V offers a unique opportunity to build processors and accelerators tailored to Europe's needs, while avoiding dependencies on non-European proprietary vendors. This aligns with broader EU policy priorities, open strategic autonomy, energy efficiency, cost-effectiveness, and supply chain resilience.

Despite Europe's excellence in system integration and vertical application domains, the hardware foundation for next-generation applications is still largely dominated by non-European, proprietary solutions. For decades, the global automotive semiconductor market has been heavily dependent on non-European proprietary processor architectures, most notably from UK-based ARM, and US vendors such as Intel and NVIDIA. These companies have successfully built vertically integrated ecosystems combining processor IP, software toolchains, and silicon manufacturing capacity, creating effective lock-ins for automotive OEMs and Tier 1 suppliers.

As modern vehicles are increasingly software-defined, the control and compute infrastructure becomes a strategic asset, with central domain and zonal controllers replacing fragmented ECU architectures. This shift increases reliance on high-performance microprocessors, AI/ML accelerators, and real-time controllers, yet most of these critical compute components are licensed under restrictive, opaque, and expensive terms from outside Europe. Furthermore, any modification, security auditing, or architectural extension of these solutions is largely out of reach for European integrators, limiting customization, innovation, and transparency. In addition, as export restrictions and geopolitical dependencies grow, Europe risks becoming a mere integrator of black-box technologies, unable to verify or adapt the hardware underpinning its automotive safety, autonomy, and cybersecurity features. This is particularly critical in the context of ISO 26262 (functional safety) and UNECE WP.29 (cybersecurity) regulations, which demand deep visibility into system behavior, something that proprietary solutions often obstruct.

This situation not only exposes Europe's automotive value chain to vendor lock-in and external disruptions, but also erodes Europe's long-term competitiveness in key markets like autonomous driving, electrification, and connected mobility. Moreover, the escalating licensing costs and complex supply chains linked to proprietary solutions hinder SME participation and stifle the growth of innovative startups. Therefore, a large scale consortium developed a new project with the overarching ambition to develop RISC-V–based high-performance and real-time microprocessors for Automotive, including application processor solutions, AI accelerators and a library or IP's, validated

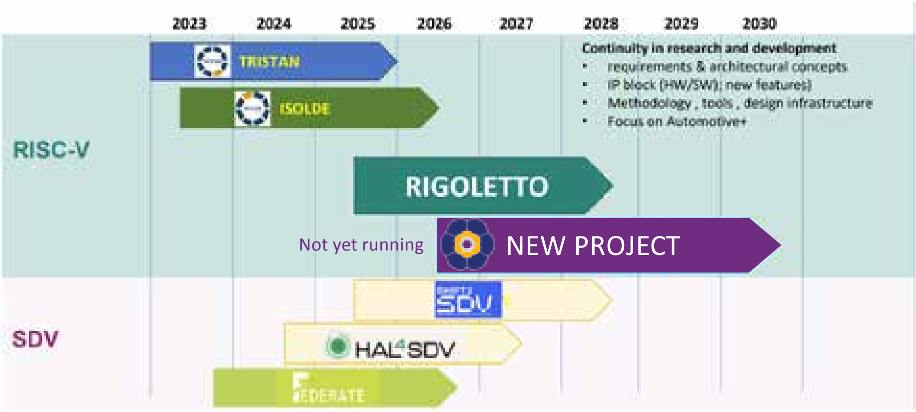

Figure 1 Positioning of the new project in the EU initiative ecosystem

and benchmarked against state-of-the-art competition. The goal is to deploy the next-generation E/E architectures in future digital vehicles to make them most competitive in terms of performance vs. non-RISC-V alternatives. The new project is committed to ensuring that its technological developments will align with industry trends and needs while fostering scalability, interoperability, and reusability across the European Automotive sector. It will provide the first RISC-V industrial grade silicon processors in Europe on advanced technology nodes, including 5nm AI accelerators, Out-of-Order cores scalable to 64bit and advanced real-time domain processors based on RISC-V. The project is currently under evaluation.

Building on and leveraging the substantial achievements of previous projects, such as TRISTAN, ISOLDE, and RIGOLETTO – while also aligning with other relevant initiatives supported by Chips-JU – such as the SDV-related initiatives under the HAL4SDV, Federate, CHASSIS and Shift2SDV projects, the new project will push the frontier of European semiconductor innovation, taking integration and silicon realization to the next level with ambitious tape-out goals. The developed solutions will not only be state-of-the-art in design, but also ready for industrial deployment.

In figure 2 below one can see the way the project is organized, having an ideal combination of individual strategic directions of the individual partners and the common ground which is developed within the project and on which their future products will be based upon.

There are 4 main common developments, which all partners plan to contribute to, and which will form the fundamental base for future processor and chiplet developments:

- Common IP Library: library with a combination of open-source and confidential IP-blocks to be shared amongst a number of partners to use in the design of their processor instantiations.
- Common Automotive Reference Platform. This Automotive Platform, developed by Quintauris in close cooperation with other partners, will serve as the foundation for a modular, vendor-neutral European Automotive Platform based on RISC-V. It is designed to support a wide range of automotive computing needs, from real-time and embedded cores to high-performance application processors. The platform provides standardized integration guidelines, benchmarking capabilities, and certification-aligned constraints to accelerate development, interoperability, and compliance across safety-critical domains such as ADAS, powertrain, body control, and in-vehicle infotainment. Its flexible architecture enables consistent system integration and validation across diverse use cases and performance classes.
- Network-on-Chip methodology and automated tools for single-MCU development to full multi-core chiplet developments, including a standard for interoperability. It will contain a scalable NoC from real time MCUs up to the largest systems of chiplets and a fully automated top-down flow (SoC compiler). IP-XACT 2022 as the interoperability standard to enable partners to exchange IPs for the Common Automotive Reference Platform.
- Easily Fast-to-Market Scalable Computation Platforms for Future Processor Families. This will be based on the Common Automotive Reference Platform and on the Processor & Accelerator Instantiation Silicon Designs. This will result in a Sound Base for Future European Leadership in Automotive Processor Silicon Wins

At the end, these developments will lead to several company-specific demonstrators, which will also be tested and validated by a few OEM-customers, which are not direct partners in the project, but which have a bilateral customer-relationship with the IDM's in the project. Furthermore, it will not only support technical comparability across demonstrators but also provide a sustainable framework for future adoption and certification of RISC-V automotive processors and accelerators, ensuring that the output of the project contributes to a common European platform and avoids fragmented outcomes. By fulfilling these promises, the

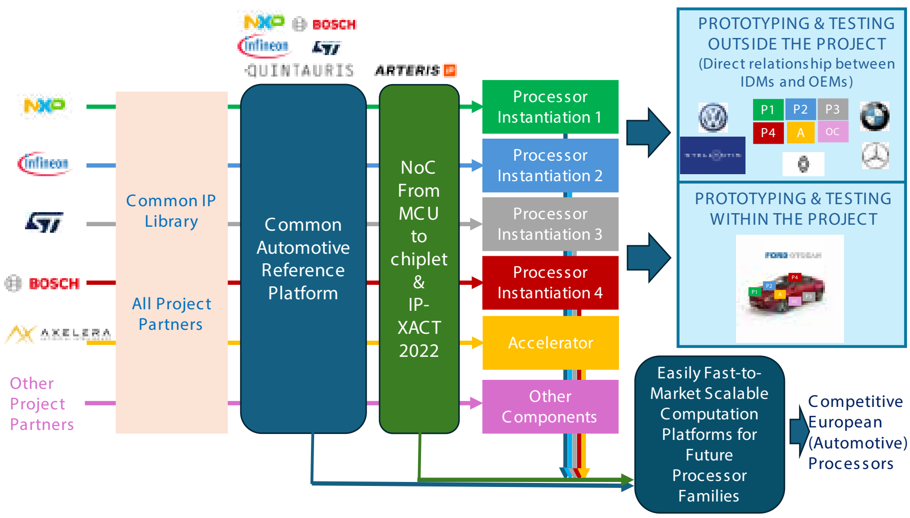

Figure 2 New Project pathways

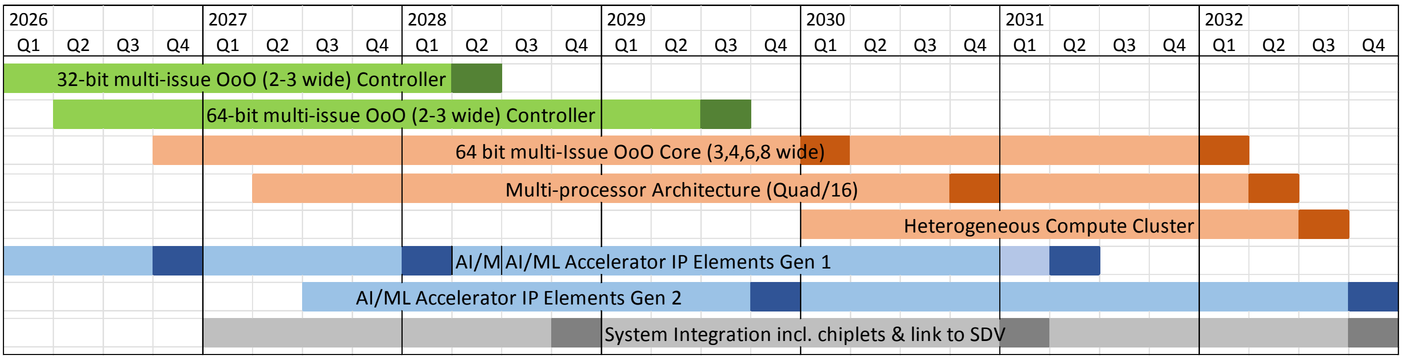

new project will contribute to the European high-performance automotive RISC-V roadmap[1] (see picture below), enabling Europe to strengthen its strategic autonomy in automotive semiconductor technologies.

Strategic differentiation and long-term impact

The new project aims to address the growing need for high-performance, safety-critical integrated circuits (ICs) within future automotive E/E architectures, particularly as zonal controllers with the architectural potential to scale toward centralized compute units. While commercial platforms such as NVIDIA DRIVE currently dominate the centralized ADAS and autonomous driving compute domain, offering tightly integrated AI accelerators, GPU-based processing, and mature software ecosystems, there remains a critical gap in open, modular, and European-controlled alternatives based on the RISC-V instruction set architecture. The new project will not immediately deliver a full substitute for such platforms. Instead, it will initiate the foundational development of competitive RISC-V-based ICs capable of supporting the compute and functional safety requirements of zonal controllers, while working on application processors capable of central computing. These designs will be extensible, providing a migration path toward domain and central computing roles. By embedding compliance with ISO 26262 and ISO 21434 from the outset, and demonstrating performance against representative ADAS and control workloads, the project will lay the groundwork for a sovereign, modular EU compute platform capable of challenging incumbent solutions in future iterations.

To illustrate the European IDM's differentiation from one of the most prominent players in automotive application processors (Nvidia), one can observe the following:

- Nvidia's share in automotive processors is only 2% today, whereas the shares of EU IDMs are among the top 5 worldwide. Their specific products represent a small part of the entire automotive semiconductor market and in that segment, they have very differentiated and unique products targeting high-end cars that have relative low volumes.
- EU IDMs together invest ~$ 6.5 billion per year in R&D (12-16% of sales). For perspective: AI leaders Intel, AMD & Nvidia spend ~$ 16.5 billion, ~$ 6.5 billion and ~$ 6 billion per year individually (10-30% of sales).
- NVIDIA's strength lies not only in its hardware but in its vertically integrated software ecosystem (CUDA), which has *created significant developer lock-in*. This integration is extremely difficult to replicate.

For strategic and corporate reasons, the partners pursue a diversified approach rather than concentrating all their efforts on a single high-performance application processor. Developing such a product would require structures beyond the scope of a Horizon Europe Innovation Action, including joint ventures, long-term IPR frameworks, and shareholder agreements. Moreover, the following arguments further support this approach:

- Smart connected systems ("robotics"[2]) – the next global technology race - will need scalable AI + safety and security. Setting the stage for EU chip makers who have unique capabilities to become world leaders in that area.
- Uniquely focusing on high-performance AI chips would be a mistake, weaken EU IDMs and miss out on the chance to make Europe a global leader in networked real-time secure and safety critical microprocessors including Edge AI.
- A narrow focus on monolithic high-performance AI chips would misallocate resources, under-serve market needs, and jeopardize Europe's leadership in networked, low-power, secure real-time systems, which are the backbone of the next digital frontier.

Yet, these constraints reveal a strategic opening: an opportunity not for imitation, but differentiation. It is about Europe playing to its strengths, leading in safety-critical, inference-dominated, edge compute infrastructure for tomorrow's vehicles and robotics, and building an infrastructure and IP that can also compete in the automotive central computing domain in the very near future. It offers a high-value, low-risk opportunity for Europe to reclaim architectural control, foster a sovereign IP ecosystem, reduce dependency on external licensing, and empower both major IDMs and emerging SMEs with a flexible, open, and certifiable compute foundation. The project's impact will extend beyond hardware development, laying a sustainable foundation for future automotive computing platforms and ensuring a competitive edge in the global automotive market. The new project paves the way for future European-wide integration and collaboration in RISC-V by marking the first-ever project where leading European semiconductor industries jointly engage in silicon tape-out activities, setting a precedent for collective innovation beyond traditional IDM individual approaches.

[1] https://ecssria.eu/Roadmap_RISC-V_v240216_Final.pdf

[2] "Robotics" is a term which includes all type of moving platforms, and as such includes automotive vehicles

Technology Frontiers

The quantum revolution

Beyond the hype into real-world transformation

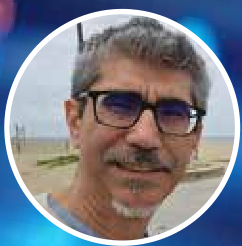

Sergio Gaudio

The year 2025 represents a historic confluence: the United Nations' designation as the International Year of Quantum Science and Technology, celebrating a century since quantum mechanics' foundational development, coincides with quantum technologies finally transitioning from theoretical possibilities to practical implementations[1]. After decades of promise and patient development, we are witnessing quantum's evolution from laboratory curiosities to transformative applications that will reshape how we diagnose diseases, navigate our world, and explore the universe's deepest mysteries.

Having spent years investigating quantum enhancement prospects for space-based gravitational wave detection, I've seen firsthand both the extraordinary potential and the sobering limitations of quantum technologies. The quantum revolution isn't quite what science fiction promised us, but it's perhaps more remarkable, offering targeted, transformative improvements that could redefine entire fields of human endeavor.

The breakthrough moment

This past year delivered quantum milestones that even experts didn't anticipate. Google's Willow chip achieved something long considered impossible: as they added more quantum bits (qubits) to their system, errors actually decreased rather than increased[2]. This "threshold scalability" breakthrough solves one of quantum computing's most fundamental challenges, and signals that we're crossing from the era of noisy, error-prone quantum prototypes to reliable quantum processors.

The financial indicators reflect growing confidence. Quantum computing funding reached a record $1.5 billion in 2024, nearly doubling the previous year's total[3]. More significantly, 55% of quantum industry leaders reported having quantum use cases in production, up from 33% the year before[4]. We're witnessing not just technological evolution but a fundamental shift in how we interact with the physical world.

But quantum computing, while capturing headlines, represents just one facet of the quantum revolution. The real transformation may come from quantum sensing, technologies that harness quantum mechanics' peculiar properties to detect and measure with unprecedented precision.

The medical revolution in your pocket

Quantum sensors are already among us, hidden in plain sight. Every MRI machine in every hospital uses quantum sensing, manipulating the quantum spin states of hydrogen atoms to peer inside our bodies[5]. But the new generation promises capabilities that seem almost magical.

Imagine visiting your doctor's office in a few years. Instead of the bulky MRI machine down the hall, your physician pulls out a device the size of a smartphone. This quantum sensor can detect the magnetic signatures of individual cancer cells circulating in your bloodstream, identifying tumors years before they would show up on traditional scans. The device works by using nitrogen-vacancy centers in diamond, essentially tiny quantum compasses so sensitive they can detect the magnetic field generated by a single cell[6].

Recent research at Johns Hopkins University demonstrated a new class of quantum sensors that can detect the faintest molecular vibrations by trapping molecules in optical cavities where light and matter become quantum mechanically entangled[7]. In these engineered quantum environments, individual molecules produce dramatically amplified optical signatures, like turning a whisper into a shout, making it possible to identify specific compounds at extraordinarily low concentrations. Each molecule has its own unique vibrational "fingerprint," much like how different musical instruments produce distinct sounds.

This breakthrough could eventually enable pocket-sized breath analyzers for medical diagnosis. Human breath contains hundreds of volatile organic compounds whose concentrations shift with disease: diabetes produces acetone, liver disease releases specific sulfur compounds, and viral infections alter the mix of metabolic byproducts. While current breath analysis requires bulky laboratory mass spectrometers, quantum molecular sensors operating at room temperature could shrink this capability into handheld devices, potentially detecting disease markers at concentrations far below what today's technology can measure.

The medical applications extend far beyond diagnosis. Researchers have already demonstrated that quantum magnetometers can detect individual magnetic nanoparticles at distances of several millimeters, capabilities that could enable tracking drug-labeled nanoparticles as they move through the body[13 14]. While optical tracking with quantum dots works well for surface tissues, magnetic detection penetrates deeper, potentially allowing doctors to monitor in real-time whether a chemotherapy drug attached to magnetic nanoparticles reaches cancer cells or whether a medication crosses the blood-brain barrier effectively. This represents a fusion of two established technologies, magnetic nanoparticle drug delivery and quantum magnetic sensing, that could eliminate much of the trial-and-error approach characterizing current medicine.

For brain disorders, quantum sensing promises unprecedented insights. While current brain imaging techniques like fMRI show activity averaged over millions of neurons, quantum magnetometers could potentially detect the electrical activity of individual brain cells. This could revolutionize our understanding of conditions like depression, schizophrenia, and autism by revealing precisely how neural circuits malfunction. Early intervention for conditions like Alzheimer's disease could begin decades before symptoms appear, when treatments might actually prevent cognitive decline.

Beyond individual health monitoring, quantum biosensors could transform public health surveillance. Networks of quantum sensors deployed in sewage systems could detect emerging pathogens days or weeks before the first cases appear in hospitals. During the next pandemic, such early warning systems could provide crucial time for containment measures.

The timeline for these medical quantum technologies varies dramatically. Some applications, like enhanced MRI systems using quantum-improved sensors, are undergoing clinical trials now. Others, like handheld molecular breath analyzers, might arrive in the next five to ten years. The most speculative applications, like individual neuron monitoring, remain further in the future but represent the direction the field is heading.

Navigation for a GPS-free world

Every day, millions of autonomous vehicles, aircraft, ships, and smartphones rely on signals from GPS satellites orbiting 12,000 miles above Earth. But GPS signals are weak, easily jammed, and completely unavailable in many environments. Quantum technologies offer an elegant solution: navigation systems that work anywhere, anytime, without requiring any external signals.

Quantum accelerometers and gyroscopes can track position and orientation with extraordinary precision. The technique works by cooling atoms to near absolute zero, where they start behaving as quantum waves. These atom waves are split and sent along two different paths, then recombined. Any acceleration causes the two paths to experience slightly different forces, creating an interference pattern, like ripples overlapping in a pond, that reveals the acceleration with extraordinary precision. Unlike GPS, which tells you where you are by receiving signals from space, quantum navigation works by precisely measuring every twist, turn, and acceleration of your journey from a known starting point. Recent demonstrations of atom interferometry-based navigation systems show position accuracy reaching 5 meters after 1 hour of autonomous navigation, over two orders of magnitude better than conventional inertial systems[8 9].

Current prototypes of quantum navigation systems can maintain accuracy over much longer periods than conventional inertial navigation. While traditional systems might drift by hundreds of meters after an hour without GPS correction, quantum systems could maintain meter-level accuracy for days or even weeks.

The applications span from the mundane to the extraordinary. Autonomous vehicles could navigate through tunnels, parking garages, or urban canyons where GPS fails. Emergency responders could navigate through smoke-filled buildings or underground disaster sites. Military aircraft could operate in GPS-denied environments without revealing their positions.

But the most transformative applications might be in space exploration. Future missions to Mars, the outer planets, or interstellar space could use quantum navigation to maintain precise positioning over journeys lasting decades, enabling autonomous spacecraft to navigate through asteroid belts or land on distant moons without real-time communication with Earth.

Closer to home, quantum navigation could enable new forms of transportation infrastructure. Imagine subway trains that know their position within centimeters, enabling much tighter spacing and higher capacity. Or consider cargo ships that could navigate busy ports in dense fog with perfect precision, reducing both delays and accidents.

The technology could also enable revolutionary new capabilities in augmented reality. Today's AR systems struggle with precise positioning, leading to virtual objects that drift relative to the real world. Quantum navigation could anchor virtual objects to the real world with millimeter precision, creating AR experiences indistinguishable from reality.

Underground applications present particularly compelling opportunities. Mining operations could use quantum navigation to map vast underground networks without needing surface references. Archaeological expeditions could create precise 3D maps of cave systems or underground ruins. Even consumer applications like indoor navigation in large buildings could benefit from quantum positioning systems.

Gravitational waves and the quantum challenge

My own research has focused on one of the most ambitious applications: using quantum technologies to enhance our ability to detect gravitational waves, ripples in spacetime itself created by cosmic collisions between black holes and neutron stars. These investigations have revealed both quantum sensing's remarkable potential and its fundamental limitations.

Through comprehensive theoretical analysis, I discovered that direct quantum detection of gravitational waves faces insurmountable barriers. The dimensionless coupling parameter characterizing how gravitational fields interact with quantum systems is approximately 5.6 × 10^-18, meaning gravitational interactions are roughly fifteen orders of magnitude weaker than electromagnetic forces at quantum scales[10]. This fundamental weakness makes direct quantum detection of gravitational waves essentially impossible with any conceivable technology.

However, this apparent limitation led to a more promising discovery: hybrid quantum-classical approaches that strategically integrate quantum sensors with proven classical technologies can provide realistic enhancements. My analysis demonstrates strain sensitivity improvements of 1.2-2.4 times across different frequency bands, achieved through quantum-enhanced frequency references, environmental monitoring, and advanced readout techniques[10].

This finding illustrates a crucial principle emerging across quantum applications: the most transformative quantum technologies often work not by replacing classical systems entirely, but by enhancing specific components where quantum mechanics provides clear advantages.

The implications extend beyond gravitational wave detection. Recent theoretical work has suggested that signatures of single gravitons, the quantum particles that make up gravitational waves, might be observable in laboratory experiments using massive quantum resonators cooled to near absolute zero[11]. While still highly speculative, such experiments could provide the first experimental evidence for gravity's quantization, potentially bridging quantum mechanics and general relativity.

These gravitational wave studies also reveal quantum sensing's potential for fundamental physics discovery. Future networks of quantum-enhanced gravitational wave detectors could detect background gravitational waves from the early universe, potentially revealing new physics beyond our current understanding. They might even detect gravitational waves from exotic sources like cosmic strings or primordial black holes, opening entirely new windows into cosmology.

Environmental monitoring and earth sciences

Quantum sensors promise to revolutionize how we monitor and understand our

planet's complex systems. The same sensitivity that enables medical diagnosis at the molecular level can detect subtle changes in Earth's gravitational field, magnetic field, and atmospheric composition.

Quantum gravimeters can detect changes in underground water levels by measuring tiny variations in Earth's gravitational field. These devices work by dropping clouds of ultra-cold atoms and measuring how Earth's gravity affects their fall with extraordinary precision. As groundwater is depleted or contaminated, the mass distribution underground changes, creating measurable shifts in local gravity. Networks of such sensors could provide early warning of droughts, help locate new water sources, and monitor the effectiveness of conservation efforts[15,16].

The technology could transform earthquake prediction and monitoring. Current seismic networks detect earthquakes after they begin, but quantum gravimeters might detect the subtle gravitational changes that precede major earthquakes as underground stresses build up, similar to how a compressed spring subtly changes the mass distribution around it. More remarkably, quantum sensors can detect the instantaneous gravity signal from shifting masses during an earthquake, which arrives before seismic waves and could provide crucial seconds of additional warning time[17]. While earthquake prediction remains one of geology's greatest challenges, quantum sensors could provide new tools for understanding seismic processes.

Climate monitoring could benefit enormously from quantum sensing networks. Quantum magnetometers could track changes in ocean currents by detecting their magnetic signatures, saltwater is electrically conductive, so when massive ocean currents flow through Earth's magnetic field, they generate detectable magnetic field disturbances, much like a generator produces electricity from moving conductors[18]. Atmospheric quantum sensors could detect trace gases associated with climate change with unprecedented sensitivity, using techniques similar to the molecular breath analyzers described earlier, enabling more accurate climate models and better policy decisions[19].

Agricultural applications present immediate opportunities. Using the same quantum sensing principles developed for medical diagnostics, quantum sensors could monitor soil composition at the molecular level, detecting nutrient deficiencies, contamination, or pathogen presence before they affect crop yields. Precision agriculture using quantum sensing could optimize fertilizer use, reduce environmental impact, and increase food security.

Urban environmental monitoring could benefit from networks of quantum sensors detecting air pollution, water quality, and even structural health of buildings and bridges. Smart cities of the future might use quantum sensor networks to optimize traffic flow, energy usage, and emergency response based on real-time environmental data.

In remote and extreme environments, quantum sensors could operate where conventional sensors fail. Arctic research stations could use quantum sensors to monitor ice sheet thickness and stability. Deep ocean sensors could track changes in ocean chemistry and temperature with unprecedented accuracy. Even space-based quantum sensors could monitor Earth's environment from orbit with sensitivities impossible using classical technologies[18].

Industrial and commercial applications

Beyond healthcare, navigation, and environmental monitoring, quantum sensors promise to transform numerous industries through capabilities that simply don't exist with classical technologies.

Manufacturing could achieve unprecedented precision and quality control. Quantum sensors based on nitrogen-vacancy centers in diamond can detect defects in materials at the atomic level during production, catching problems before they become expensive failures[20,21]. Semiconductor manufacturing, already pushing the limits of precision, could benefit from quantum sensors that monitor fabrication processes with atomic-scale accuracy. Companies like Quantum Diamonds are already deploying quantum microscopes that scan diamond tips over microchips to map magnetic fields from individual transistors, helping identify defects and improve chip design[22]. These sensors work by detecting magnetic fields generated by current flows in the chip, with machine learning enhancing both precision and speed for mapping current flows across multiple layers and identifying faults non-destructively.

The automotive industry faces a quantum transformation. Beyond quantum navigation, quantum sensors could monitor vehicle health in real-time by detecting microscopic changes in metal fatigue, bearing wear, or brake performance[23,24]. Predictive maintenance based on quantum sensing combined with AI could prevent accidents and reduce costs by scheduling repairs before failures occur. Recent studies demonstrate that quantum-enhanced predictive maintenance using fiber Bragg grating sensors and quantum technology can reduce false positives and cut unnecessary maintenance costs[25]. This is achieved by training AI models on quantum sensor data to detect early signs of component degradation and estimate remaining useful life more accurately than classical approaches.

Energy production and distribution could benefit enormously from quantum sensing. Power grids could use quantum magnetometers to detect powerline and transformer faults before they cascade into blackouts[26]. The U.S. Department of Energy is actively developing quantum sensors for grid optimization, with quantum atomic clocks providing position, navigation, and timing data directly from the electrical grid rather than satellite-based GPS sources[27,28]. This makes the grid more resilient to GPS interference and enables anomaly detection for enhanced security. Solar panels and wind turbines could monitor their own performance and degradation with unprecedented precision, optimizing energy production and predicting maintenance needs through quantum-enhanced fault detection algorithms[29].

Oil and gas exploration represents a short-term commercial application for quantum sensors. Quantum gravimeters based on atom interferometry can detect underground oil and gas deposits by measuring tiny variations in gravitational fields[30 31]. Companies like M Squared and iXblue have demonstrated commercial quantum gravimeters capable of detecting new oil and gas deposits, surveying unknown underground infrastructures, and monitoring water tables[32 33]. These devices work by cooling rubidium atoms close to absolute zero and measuring their free fall in a vacuum, achieving gravitational measurement precision at the parts-per-billion level. This non-invasive exploration method could reduce the environmental impact of resource extraction while increasing discovery rates by identifying unrecovered pockets of resources.

Financial services might seem far removed from quantum sensing, but the technology could revolutionize fraud detection and risk assessment. Quantum-enhanced analysis of market data could detect subtle patterns indicating manipulation or systemic risk. High-frequency trading systems could benefit from quantum-precise timing provided by chip-scale atomic clocks that make microseconds seem like hours, enabling more accurate time-stamping of transactions and better synchronization across trading platforms.

Construction and civil engineering could use quantum sensors for structural health monitoring that detects microscopic changes in bridges, buildings, and tunnels[34 35]. This could prevent disasters like building collapses by providing early warning of structural problems. Quantum accelerometers and strain gauges can continuously monitor structural defects in infrastructure, noticing cracks and deformations in real-time[36]. Recent advances in quantum fiber optic sensing enable a single optical fiber-based sensor to monitor multiple locations along its length with extreme precision, allowing continuous high-resolution data collection across entire structures[37]. This is achieved by sending quantum-enhanced “squeezed” light pulses through optical fibers with strategically placed reflectors, dividing the fiber into distinct measurement zones that can detect tiny vibrations or structural instabilities in real time.

Even entertainment and sports could be transformed. Quantum motion sensors could capture athletic performance with unprecedented precision, enabling new training methods and performance analysis. Virtual and augmented reality systems could use quantum positioning to create experiences indistinguishable from reality.

The scientific revolution

Perhaps most importantly, quantum technologies are enabling scientific investigations that were previously impossible. The precision achieved by quantum sensors opens experimental access to phenomena at the intersection of quantum mechanics, general relativity, and physics beyond the Standard Model.

Dark matter searches using quantum sensors could finally detect these mysterious particles that make up most of the universe’s mass. Quantum sensor arrays could detect the weak interactions that dark matter particles should have with ordinary matter, potentially solving one of physics’ greatest puzzles[12].

Tests of fundamental physics become possible with quantum precision. Atomic clocks are so accurate they won’t lose a second over billions of years and can test whether fundamental constants actually remain constant over time. Such measurements could reveal new physics beyond our current theories.

Quantum sensors might detect gravitational effects that general relativity predicts but we’ve never measured directly. The gravitational redshift effect, frame dragging around rotating masses, and even gravitational waves from the early universe might become accessible through quantum-enhanced measurements.

In biology and medicine, quantum sensors could reveal cellular processes that current techniques cannot detect. The magnetic signatures of individual enzyme reactions, the electrical activity of single ion channels, or the mechanical forces inside living cells could become observable for the first time.

Archaeological and paleontological applications could revolutionize our understanding of history. Quantum magnetometers could detect buried structures by their magnetic signatures, revealing ancient cities without excavation. Quantum dating techniques could provide unprecedented precision in determining the age of artifacts and fossils.

Even fundamental questions about consciousness and the brain might benefit from quantum sensing. If quantum effects play a role in neural processing, as some theories suggest, quantum sensors might be

the only way to detect and study these effects in living brains.

Challenges and realistic expectations

Despite remarkable progress, quantum technologies face significant challenges. Quantum decoherence, the destruction of quantum properties by environmental noise, remains the primary limitation. While laboratory demonstrations achieve remarkable sensitivities, real-world applications must operate in noisy environments that destroy quantum effects.

Cost represents another major barrier. Many quantum sensors remain expensive and complex, requiring specialized expertise for operation. The transition from laboratory demonstrations to commercial products necessitates substantial engineering optimization for cost reduction, reliability, and user accessibility.

Manufacturing quantum devices at scale presents unprecedented challenges. While conventional electronics can tolerate minor imperfections, quantum devices require atomic-level precision. Developing manufacturing processes that can produce quantum devices reliably and affordably remains a significant hurdle.

The most successful quantum applications employ hybrid approaches that leverage quantum advantages for specific functions while maintaining proven classical systems for overall operation. This strategy appears more promising than attempting wholesale replacement of classical systems with quantum alternatives.

Education and workforce development present equally important challenges. The quantum revolution will require millions of workers with at least basic understanding of quantum principles. Educational systems worldwide must adapt to prepare the next generation for quantum-enabled careers.

Regulatory and standardization challenges loom large as quantum technologies mature. How do we certify the accuracy of quantum sensors? What standards ensure interoperability between quantum communication systems? These questions require answers as quantum technologies transition from research to commerce.

The path forward

The quantum revolution represents neither the sudden transformation promised by science fiction nor incremental improvement of existing technologies, but rather a fundamental expansion of technological capabilities through controlled application of quantum mechanical phenomena. Current progress demonstrates both extraordinary potential and realistic limitations across diverse applications.

The most promising quantum technologies employ targeted enhancement strategies, leveraging quantum advantages for specific functions within hybrid quantum-classical systems. This approach respects both the unique capabilities that quantum mechanics provides and the practical constraints imposed by decoherence, cost, and operational complexity.

As quantum technologies mature from laboratory demonstrations to commercial applications, their impact will emerge through gradual integration rather than revolutionary disruption. The quantum age is not coming, it is already here, manifesting through MRI machines in hospitals, enhanced gravitational wave detectors, and emerging applications in medical sensing, navigation, and fundamental research.

The challenge for researchers, policymakers, and society is nurturing quantum technology development while maintaining realistic expectations about timelines and capabilities. The quantum revolution promises to expand our ability to measure, understand, and manipulate the physical world, but its ultimate success depends on careful attention to both quantum mechanical possibilities and engineering practicalities.

The next phase of quantum development will likely determine which applications achieve lasting commercial and scientific impact. The foundations being established in 2025 may well shape the trajectory of quantum technologies for decades to come, as we learn to harness the strange and wonderful properties of quantum mechanics to solve problems we never thought possible.

The quantum future isn't about replacing our classical world with a quantum one, but about identifying the precise points where quantum mechanics' unusual properties provide genuine advantages over classical physics. In doing so, we're not just building better technologies, we're developing new ways of understanding and interacting with reality itself.

1 United Nations, "International Year of Quantum Science and Technology 2025," https://www.quantum2025.org/, accessed September 2025.

2 Google Quantum AI, "Meet Willow, our state-of-the-art quantum chip," Google AI Blog, December 2024.

3 M. Cooney, "10 quantum computing milestones of 2024," Network World, November 2024.

4 McKinsey & Company, "The year of quantum: From concept to reality in 2025," McKinsey Digital, June 2025.

5 M. Jones, "How quantum sensors could improve human health and wellbeing," Physics World, July 2025.

6 J. R. Maze et al., "Nanoscale magnetic sensing with an individual electronic spin in diamond," Nature 455, 644 (2008).

7 Johns Hopkins University, "Engineers create new class of quantum sensors to detect faint molecular vibrations," Johns Hopkins Hub, August 2025.

8 B. Tennstedt et al., "Atom Strapdown: Toward Integrated Quantum Inertial Navigation Systems," NAVIGATION, 70(4) (2023). https://doi.org/10.33012/navi.604

9 C. Jekeli, "Navigation error analysis of atom interferometer inertial sensor," NAVIGATION, 52(1), 1–14 (2005). https://doi.org/10.1002/j.2161-4296.2005.tb01726.x

10 S. Gaudio, "Beyond Classical Interferometry: Quantum Enhancement Prospects for Space-Based Gravitational Wave Detection," APS/123-QED, July 2025.

11 G. Tobar et al., "Detecting single gravitons with quantum sensing," Nat. Commun. 15, 7229 (2024).

12 Oak Ridge National Laboratory, "Researchers reveal quantum advantage that could advance future sensing devices," October 2024.

13 A. Kuwahata et al., "Magnetometer with nitrogen-vacancy center in a bulk diamond for detecting magnetic nanoparticles in biomedical applications," Sci. Rep. 10, 2483 (2020).

14 S. Chen et al., "Digital Magnetic Detection of Biomolecular Interactions with Single Nanoparticles," Nano Lett. 23, 2743-2750 (2023).

15 B. Stray et al., "Quantum sensing for gravity cartography," Nature 602, 590–594 (2022).

16 A.-K. Cooke et al., "First evaluation of an absolute quantum gravimeter (AQG#B01) for future field experiments," Geosci. Instrum. Method. Data Syst. 10, 65–79 (2021).

17 University of Birmingham Quantum Technology Hub, "How can quantum technology improve earthquake detection?" (2023). https://www.birmingham.ac.uk/news/2023/how-can-quantum-technology-improve-earthquake-detection

18 European Space Agency, "Taking climate monitoring into the future with quantum." https://www.esa.int/Applications/Observing_the_Earth/FutureEO/Taking_climate_monitoring_into_the_future_with_quantum

19 A. Belsley, "Quantum-enhanced absorption spectroscopy with bright squeezed frequency combs," Phys. Rev. Lett. 130, 133602 (2023).

20 T. Esat et al., "A quantum sensor for the atomic world," Nature Nanotechnology 19, 10 (2024). https://doi.org/10.1038/s41565-024-01724-z

21 Forschungszentrum Juelich, "Quantum sensor for the atomic world," ScienceDaily, July 25, 2024.

22 K. Berghoff and Quantum Diamonds, "Can Quantum Sensors Improve Real-Time Metrology in Fabs?" EE Times Europe, December 2024.

23 S. Aissani et al., "A Comprehensive Review of Predictive Maintenance Technologies for Vehicle Reliability," in CSCI 2024, Springer, 2025.

24 P.-S. Rao et al., "Integrated artificial intelligence and predictive maintenance of electric vehicle components with optical and quantum enhancements," Opt. Quantum Electr. (2023).

25 Rao et al., "A comprehensive review on artificial intelligence driven predictive maintenance in vehicles," Discover Applied Sciences, vol. 7:243, 2025.

26 F. Schlapper et al., "Quantum sensing for emerging energy technologies," Nature Reviews Clean Technology, 2025. https://doi.org/10.1038/s44359-025-00112-7

27 R. Oueid (U.S. Department of Energy), "Quantum Sensing Has 'Critical' Potential for Electrical Grid," Nextgov, March 2023.

28 The Quantum Insider, "Quantum Sensing is a U.S. Department of Energy Priority," April 2023.

29 F. Di Giovanni, "Quantum Algorithms Will Optimize Power Grid Efficiency," Power Electronics News, June 2024.

30 G. Malcolm and M Squared, "Industrial quantum gravimeter demonstrated," Electro Optics, 2024.

31 International Defense Security & Technology, "Quantum Gravity Sensor technology advancing for oil, gas and mineral exploration," 2024.

32 M Squared, "UK's First Quantum Gravity Measurements Outside Academia," 2024.

33 B. Desruelle et al., "Underground quantum sensing set to erupt," Electro Optics, June 2022.

34 B. Kantsepolsky and I. Aviv, "Sensors in Civil Engineering: From Existing Gaps to Quantum Opportunities," Smart Cities 7, 277-301 (2024). https://doi.org/10.3390/smartcities7010012

35 W. Li et al., "Quantum metrology and its applications in civil engineering," Measurement, vol. 239, 2024. https://doi.org/10.1016/j.measurement.2024.115435

36 Kantsepolsky & Aviv, "Quantum Sensors: The Future Of Civil Engineering, Says Study," Quantum Zeitgeist, February 2024.

37 J. Gopinath et al., "Quantum technique could transform remote sensing, infrastructure monitoring," Physical Review A, University of Colorado Boulder, 2024.

AI agents for autonomous cybersecurity

A framework and a case study

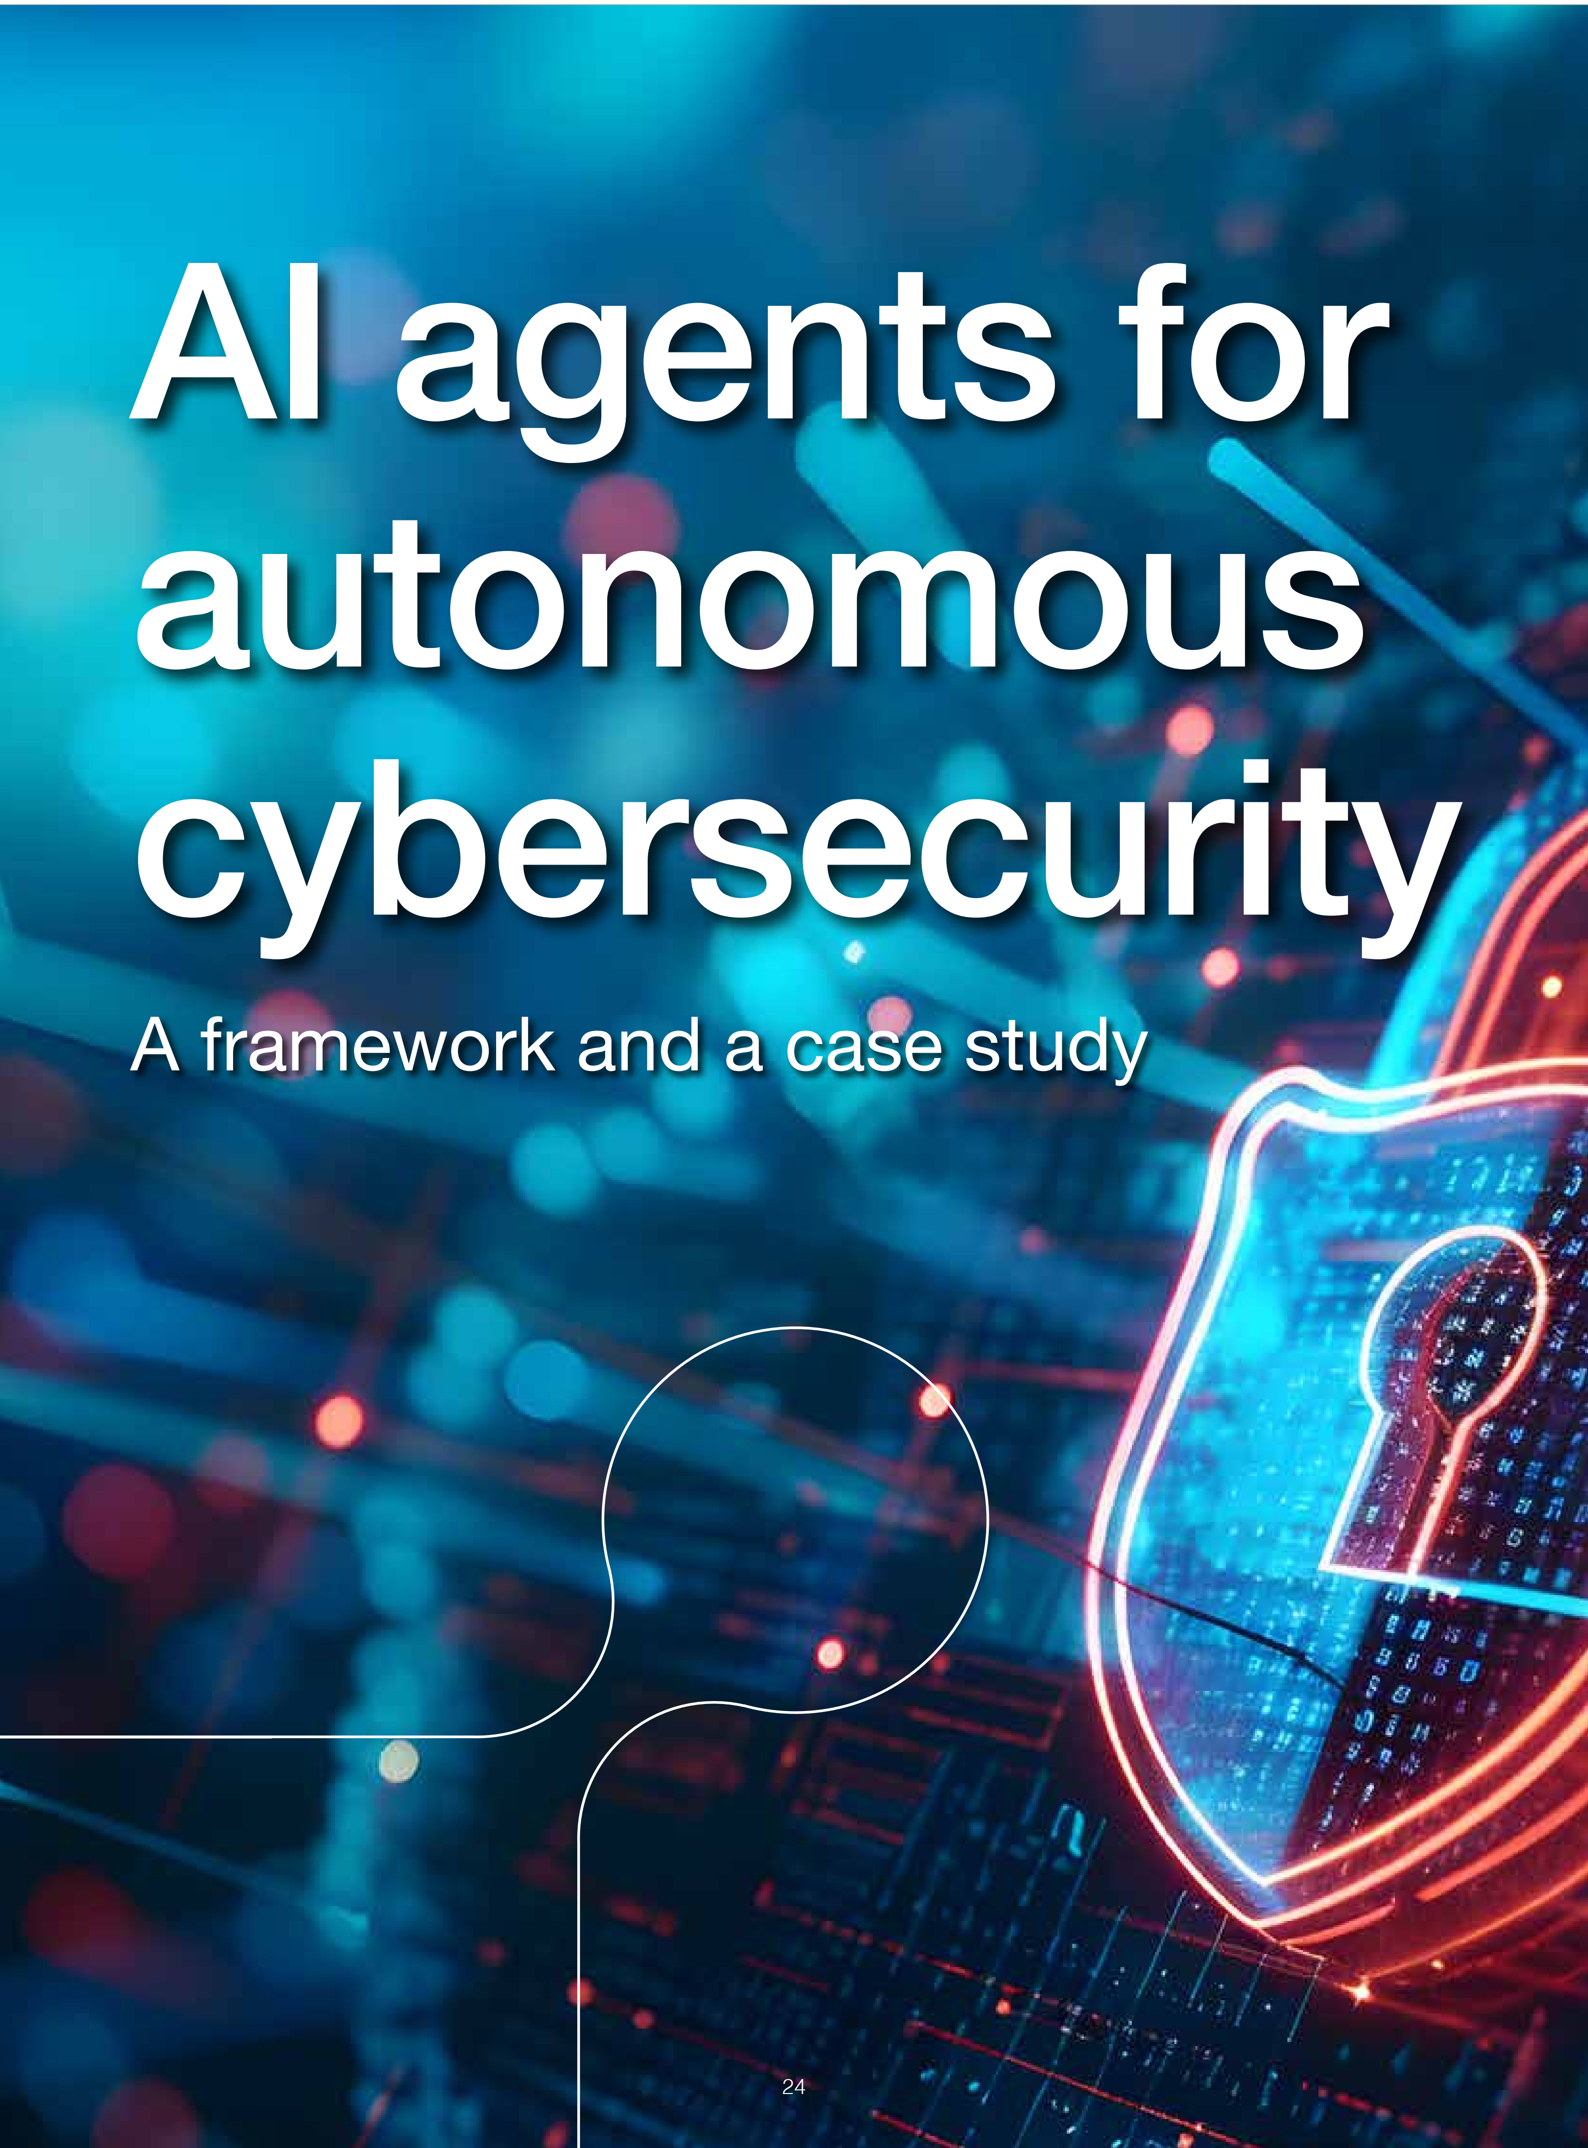

Dimitrios Serpanos

Georgios Xenos

The cybersecurity landscape is undergoing a major transformation driven by Agentic AI, which leads to increased productivity while minimizing human error. Security applications are evolving into autonomous agents that enable automated threat analysis and detection, dynamic security policies and responses to cyberattacks in real-time.

The introduction of Large Language Models (LLM) has redefined the capabilities of Artificial Intelligence (AI) and how it is leveraged to solve increasingly complex tasks. Initially, LLMs were employed primarily as chatbots, becoming an effective interface used to interact with computers in natural language. However, recent advancements in LLM technologies such as reasoning, multimodal models and increased context lengths have given rise to Agentic AI, enabling autonomous agents to perform complex workflows, automate tasks, and interact with external systems to achieve dynamic objectives. Moreover, the cost of deployment and inference for LLMs has been steadily declining since their introduction, leading to the gradual deployment of agents to solve problems in several directions and business domains. Cybersecurity is one domain with pressing needs for automation as it faces unique challenges, making Agentic AI a natural fit. Cybersecurity requires the processing of large volumes of heterogeneous data while human analysts are involved in every security process to make decisions, prioritize security alerts and take actions to respond to potential threats. The challenge is compounded by the scarcity of expert personnel, as security analysts must be highly trained and specialized.

Agents are defined as systems that can perform tasks autonomously by perceiving their environment, making decisions and taking actions to achieve specific goals. Enabled by LLMs and their inherent reasoning capabilities, agents can break down complex tasks into actionable sequences of steps, reason over large volumes of data, plan and make decisions. Models are now also able to perform tool use, an ability that enables agents to interact with external systems and knowledge bases as well as to communicate with other agents. This allows agents to close the loop and take actions to achieve their objectives. In the context of cybersecurity in particular, these abilities pave the way for a transition from a reactive static security model to a proactive one. For example, typical static AI models may be able to perform the narrow tasks of detecting an anomaly or a cyberattack, but then human experts take over to triage and confirm the threat and conduct incident response. In contrast, agents are able to detect a cyberattack, correlate it with other events and external intelligence to produce a detailed report. More importantly, they can take the next step and perform actions to contain and isolate the attacker and remediate the cyberattack by restoring the affected systems and patching the exploited vulnerabilities. This end-to-end deployment of Agentic AI leads to faster responses to attacks while minimizing the need for human intervention, enhancing performance and productivity.

A framework for agentic AI in cybersecurity

Designing and developing Agentic AI systems for cybersecurity requires the integration of multiple functional components with control and observability considered in order to ensure safe autonomy.

A key design decision concerns the desired autonomy of the system. Agents can range from performing deterministic workflows up to fully autonomous operation. AI agents operating with deterministic workflows perform a pre-defined set of steps to solve a well-defined problem making use of static processes. While this setup offers low flexibility, it allows for easier evaluations and increased observability. In contrast, agents deployed with increased autonomy are developed to perform more generic tasks. Such agents have to break down the required steps to perform their objective,

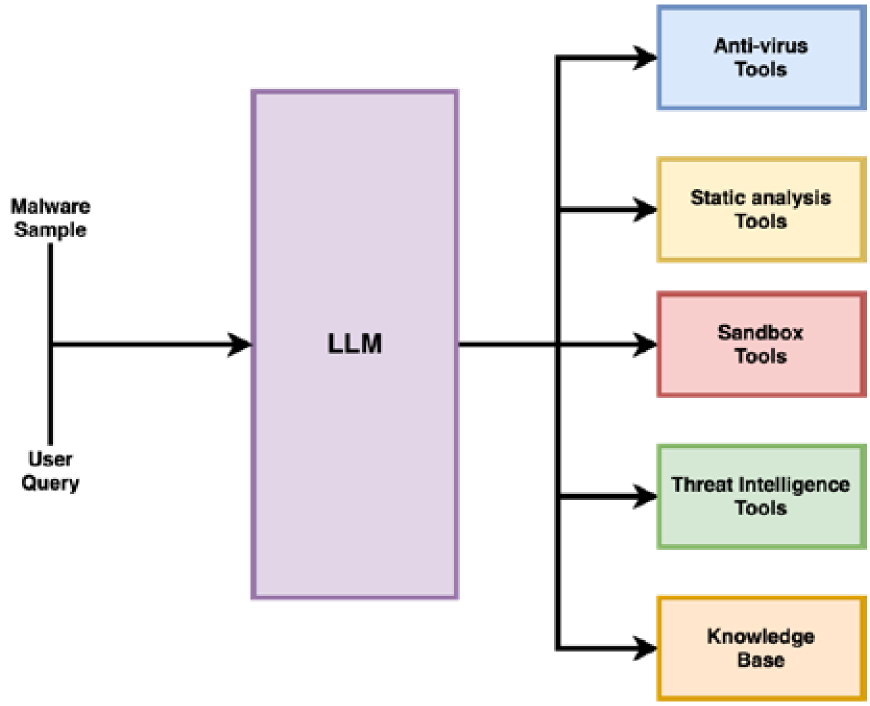

Figure 1: High-level architecture of a malware analysis agent

make their own decisions and dynamically perform external tool use. This increased autonomy, however, makes evaluation difficult and requires the implementation of control mechanisms such as guardrails and explainability components. Managing that extent of autonomy is ultimately a design choice that is tightly related to the specific use-case. When agents are required to interact with sensitive data or critical infrastructure more determinism in the operations might be necessary.

Agents can also be deployed in a Co-Pilot manner, where a human expert interacts with the AI in a chatbot format. This is particularly useful in Security Operations Centers (SOC) environments. Co-Pilots can be used to automate the mundane tasks and provide high fidelity threat alerts to analysts upon request. For example, a SOC operator may ask the Co-Pilot to explain a particular security event or alert and to provide additional information based on external threat information. The agent has increased autonomy and can interact with external systems in real-time while human operators can describe complicated requests in natural language leading to decreased alert fatigue and increased productivity.

Several design constraints are present when considering the underlying LLM that serves as the cognitive machine in Agentic AI. Cybersecurity tasks such as malware analysis or threat detection and remediation require large volumes of data to be considered and analyzed by the agents. This means that the underlying LLM should support long contexts. Most modern LLMs can handle very long contexts of millions of tokens. Additionally, the employed LLM should be trained to support reasoning and tool use as these abilities are foundational for Agentic AI. In some use-cases, a domain-specific model might be developed by performing fine-tuning. During fine-tuning we further train an LLM on examples specific to the task or problem we intend to solve. This process leads to increased model performance, and in cases, it makes smaller models viable for specific applications. However, this approach requires the use of domain-specific datasets which are very difficult to obtain. In the security space, the problem with data availability is twofold. Development of such datasets requires strenuous labeling work from highly specialized security experts. Coupled with this, privacy concerns and Intellectual Property restrictions limit data sharing among organizations, thereby further hindering dataset development.

The ability of agents to interact with external tools and systems to enhance their capabilities is what makes them primarily valuable in the cybersecurity domain. These tools can include a wide variety of security software to allow the agent to perform analysis such as malware sandboxes, network packet analyzers, etc. Tools can also be invoked to extend the agent's context by giving access to threat intelligence databases. Additionally, access to knowledge bases can be implemented, enabling the retrieval of domain-specific information in real-time, enhancing contextual understanding and decision-making. For instance, a knowledge base might encapsulate information about attacker techniques and patterns, making threat correlation easier.

Agentic AI Evaluation

During both the development and deployment stage, evaluation of agent performance and efficiency should be systematically conducted. Agents should be evaluated on their abilities to complete the requested tasks, to appropriately select and use tools as well as on their efficiency during deployment in terms of latency and cost. Task completion metrics are used to measure the agent's ability to complete the assigned objectives. For example, threat detection effectiveness can be measured in terms of accuracy, precision and other typical metrics used in Machine Learning. On the other hand, for complex tasks such as malware or threat analysis and correlation, it is difficult to define what task success even looks like. One approach is to have humans manually review the outputs of agents to assess how effectively they have completed the requested task. A more flexible and scalable approach that doesn't require human intervention is to use an LLM as a judge and have it compare the outputs of the agents against a ground truth dataset. To measure the efficiency, speed and cost of agents, metrics such as time-to-task-completion and token efficiency are measured. Token efficiency refers to the number of input and output tokens generated by the underlying LLM to complete the task and is particularly important as operating costs are tightly related to token usage.

A case study on malware analysis

To study practically the effectiveness of Agentic AI in cybersecurity we will study one of the potential use-cases: malware analysis. Malware analysis is a complex, time-consuming process usually carried out by trained analysts while requiring the use of several tools and analysis frameworks. A typical workflow consists of analyzing a malicious sample using static analysis and often combined with dynamic analysis. Static analysis refers to analyzing a software without executing it, it is computationally cheap and fast while dynamic analysis involves executing the software in a secure

environment, typically a virtual machine and observing its behavior. The main objectives of malware analysis are to determine the behavior and functionalities of the malicious sample, to extract relevant information from the sample and to classify the malware into a specific family, producing a detailed threat report.

An agent that automates the process of malware analysis should be able to conduct all relevant steps in the workflow and achieve the desired objectives in an autonomous manner. This entails receiving a software sample, performing static and dynamic analysis, extracting threat intelligence and compiling a detailed report with all the relevant findings. The main component of the agent is the underlying LLM; a reasoning model with large context length, that can access external tools is employed. After performing prompt engineering, instructions are crafted for the LLM to instruct it to properly execute the analysis and reporting steps. A diverse set of tools is implemented containing tools for static analysis, sandbox analysis and access to external knowledge. Indicatively, the agent is given access to multiple known Antivirus engines in order to perform basic scanning, and to entropy measuring and printable string extraction software. Additionally, complete access to a commercial sandbox is implemented to enable the agent to perform dynamic analysis when needed and to retrieve the relevant reports. The agent can also query an external threat intelligence database, to assist with attack correlation, while a knowledge base with known attacker techniques is developed. Figure 1 illustrates the high-level architecture of the implemented agent. To implement the tools, we employ the Model Context Protocol[1] (MCP). MCP, introduced by Anthropic, aims to standardize the way AI agents interact with external tools and resources. It allows for a plug-and-play implementation of tools irrespective of the underlying agent architecture or LLM flavor. The agent employs the ReAct[2] framework to allow the LLM to interleave reasoning and tool use. With the ReAct-based implementation the agent can reason and perform actions in a loop until it believes it has arrived at a satisfactory answer in order to achieve its specified objective.

Malware analysis: from days to minutes

The deployed agent successfully performs full analysis of any given file, detects if it is malicious, classifies its family and extracts the relevant threat intelligence. On average the agent needs only 6 minutes to fully analyze a file while human experts can only analyze 5 to 20 samples per week[3], showcasing how agents can assist humans to significantly increase productivity. The agent is also able to be deployed as a Co-Pilot. Despite the lack of standardized benchmark datasets, manual review of the agent's outputs reveals that it correctly selects the relevant tools needed to perform the user queries. Of course, there are limitations -- individual tools often fail to provide the agent with analysis results, as cybersecurity is inherently adversarial and attackers constantly evolve their techniques to evade analysis. Each tool invocation is tracked as an individual event and so tool performance can be assessed easily, and fallback tools can be implemented to combat analysis failures. Additionally, continuous monitoring is required to ensure operational fidelity and that performance remains consistent over time. While more evaluation and governance mechanisms will be developed, in the end, Agentic AI in cybersecurity will still heavily rely on human expertise and supervision given its direct interaction with critical infrastructure and sensitive private data.

Conclusions

Agentic AI is redefining the cybersecurity landscape by automating complex threat analysis and detection tasks as well as enabling dynamic responses and policies to cyberattacks. Malware analysis is one of the key areas where autonomous agents can drastically increase productivity while reducing human error and fatigue. While processes and frameworks for agent development are quickly set into place, key components for cybersecurity-oriented Agentic AI are still missing, such as evaluation methods, benchmark datasets and, importantly, clear governance guidelines and mechanisms. At the end of the day, human cybersecurity experts will have to stay in the loop of decision-making to oversee AI actions and ensure safety and alignment with organizational security policies and best practices.

[1] https://modelcontextprotocol.io
[2] https://arxiv.org/abs/2210.03629
[3] https://dl.acm.org/doi/10.1145/3460120.3484759

Technology Frontiers

Scalable hardware for quantum computing

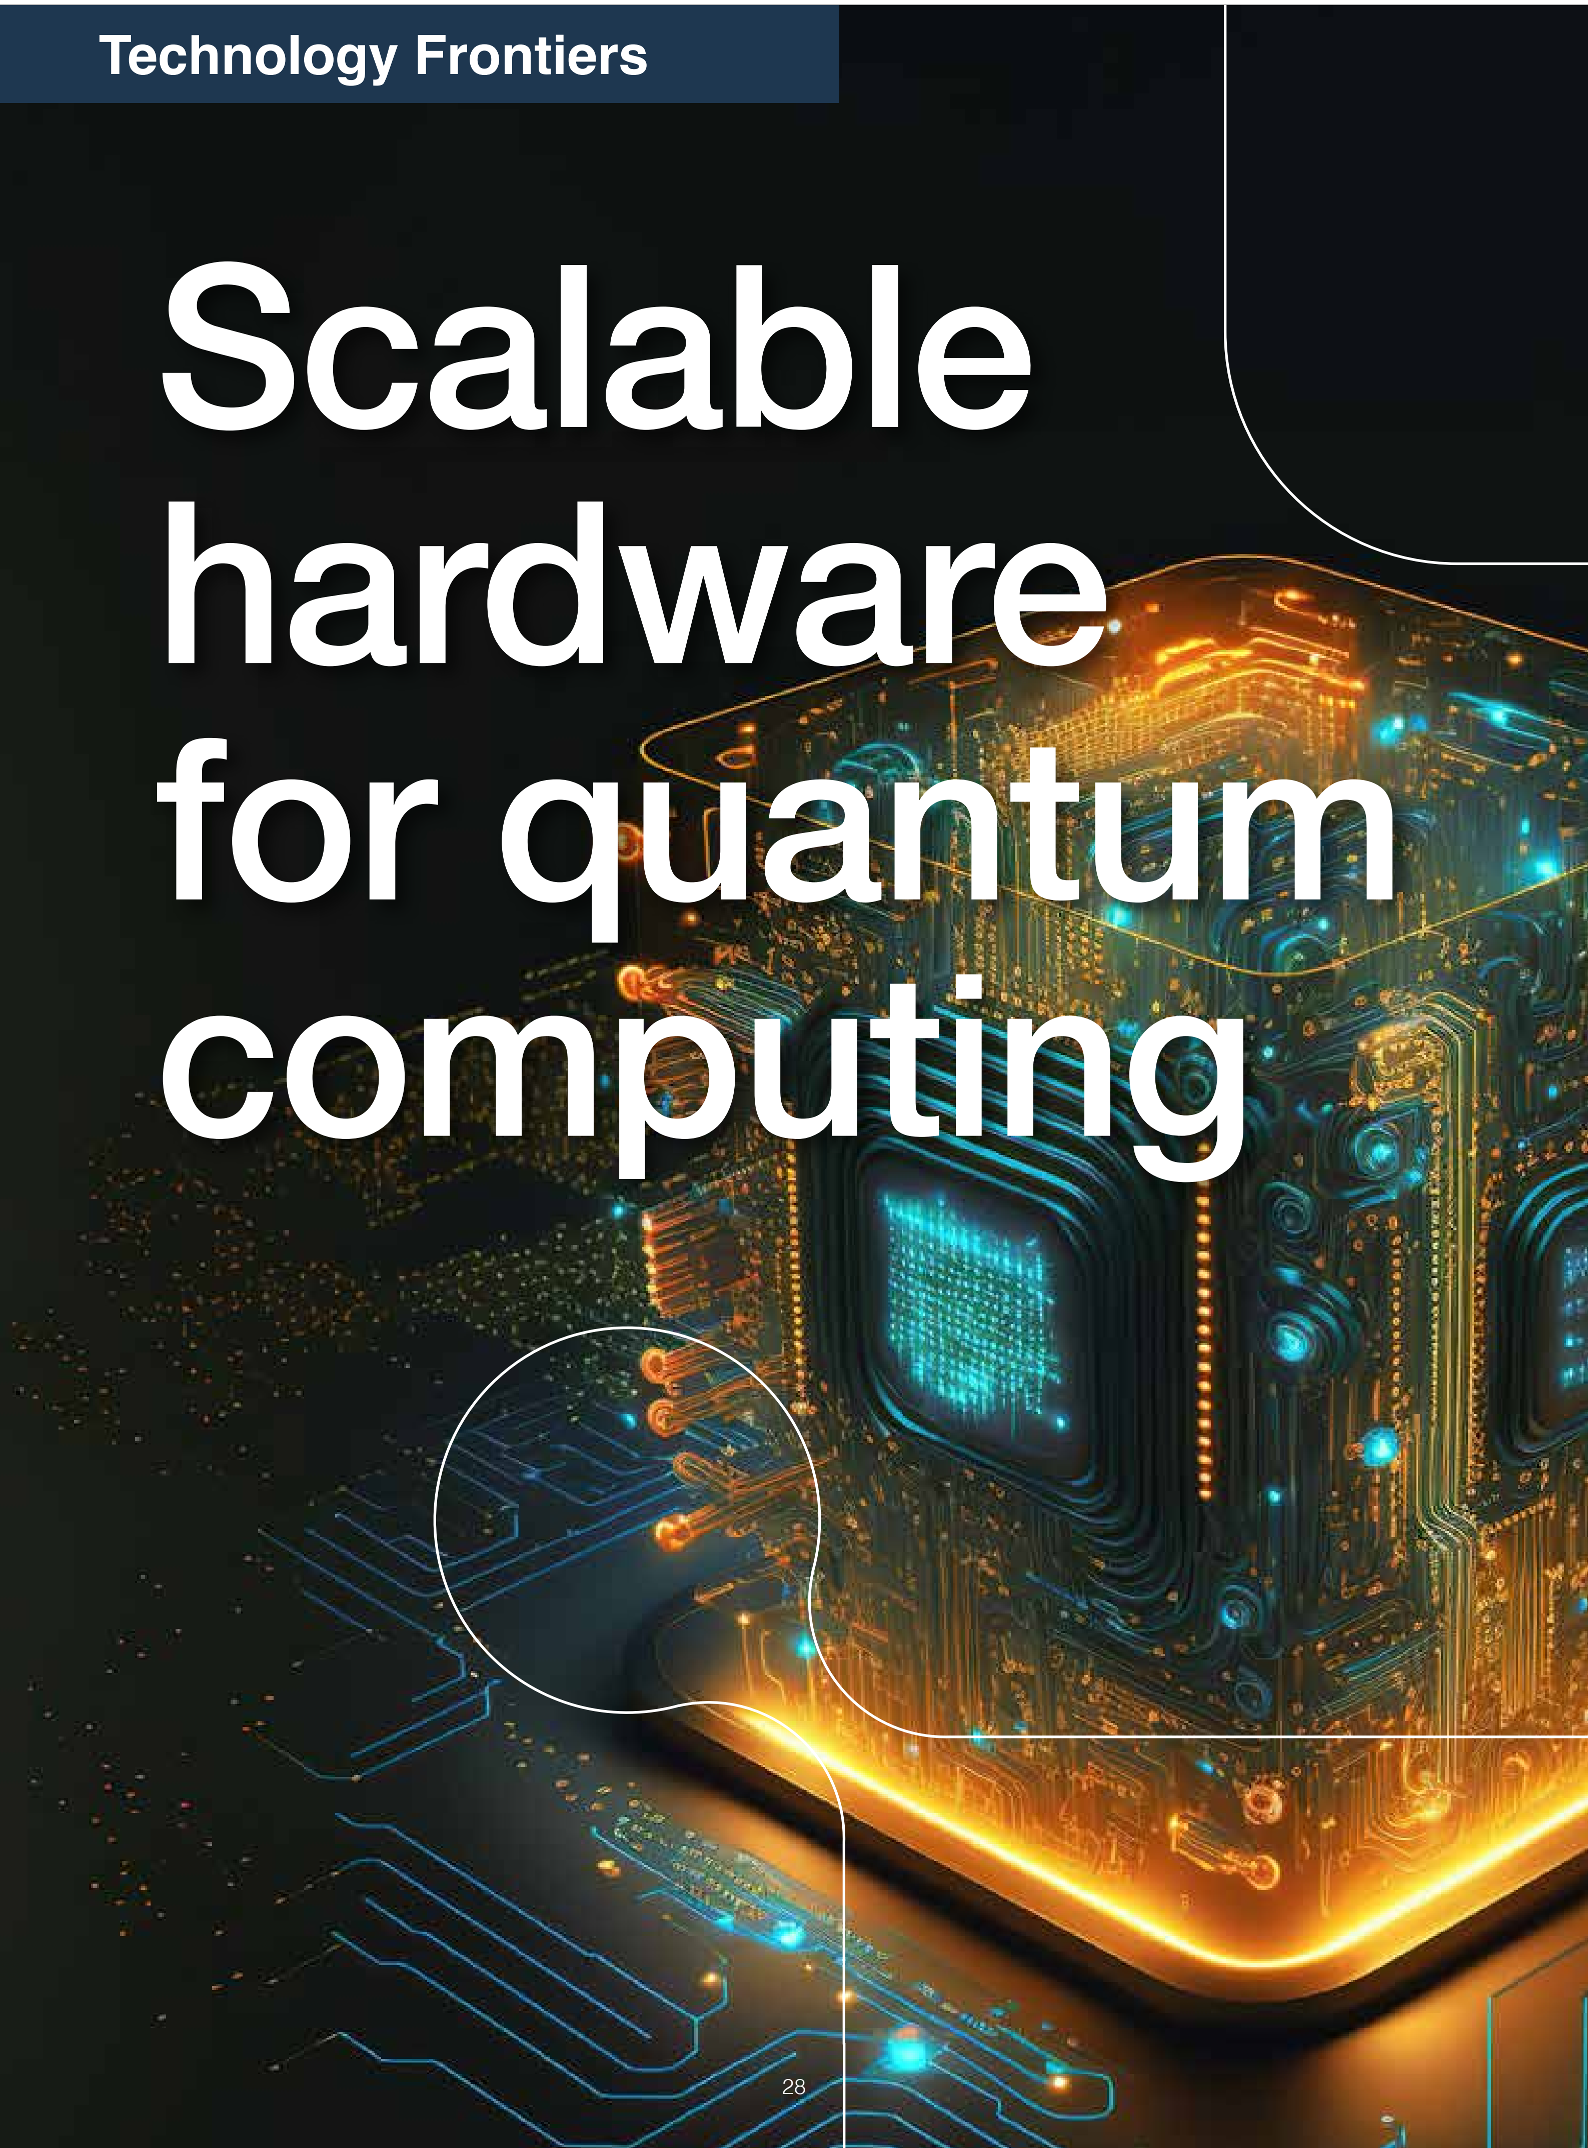

Edoardo Charbon

Peter Müller

Quantum computing hardware is evolving rapidly. Promising demonstrations of quantum information processing platforms leveraging superconducting qubits, spin qubits, ion traps, and other technologies are underway around the world. Following the development of the basic elements, the qubits, an interplay begins between the development of the classical system architecture and the next step on the quantum processing side. This article addresses the next integration step in the system architecture: quantum-to-classical interface electronics and integrated basic controllers.

In 1981, Richard Feynman [Feynman, MIT lectures, 1981] proposed to build a computer based on quantum mechanics. The idea was to simulate quantum phenomena using quantum-mechanical elements. From the initial formulation, there has been significant developments both at theoretical and implementation level. On the one hand, algorithms have been developed to solve currently unsolvable problems, such as the prime factorization of large numbers using Shor's algorithm. Since the 1980s, more than 100 algorithms have been developed that potentially operate more efficiently on a quantum computer than on a classical computer. On the other hand, significant advances have also been made in hardware, enabling the development of relatively large quantum processors and, especially in 2019, the practical demonstration of quantum advantage. [F. Arute et al., Nature, 2019].

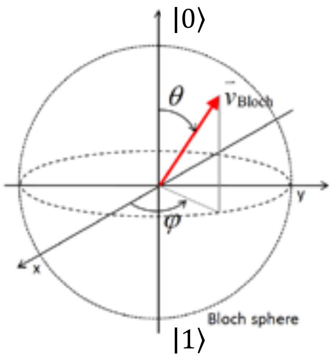

Fig. 1: Bloch sphere. The quantum state of a qubit is the result of superposition of a ground state |0> and an excited state |1>, seen on the sphere as the north and the south pole, respectively. The qubit can be viewed as a gyroscope with a rotation axis between the states |0> and |1>, rotating around the equator. From this, the characteristic state and phase values can be determined.

A quantum computing platform is organized in an array or matrix structure consisting of n quantum bits, or qubits. In contrast to classical digital bits, which represent a discrete state with the value "**0**" **or** "**1**," each qubit is characterized by its quantum state. The quantum state is the result of a superposition of the quantum states **I0> and I1>**. A common method for visualizing the quantum state of a qubit in 3D is the so-called Bloch sphere, see Fig. 1.

Superposition has the special property that it allows 2^n states to be stored and manipulated together in a quantum register with n qubits. For example, at $n = 8$ the number of quantum states is 256, and at $n = 300$ the number of quantum states of a quantum processor exceeds the total number of atoms in the universe. The second fundamental property of quantum mechanics is that of entanglement. When in a quantum system, two or more components, e.g. qubits, are entangled, they become linked, so that the quantum state of one implies that of the other, regardless of the distance by which they are separated. We say that the wavefunctions of the quantum components are correlated and thus measuring the state of one instantaneously influences the state of the other or the others. This phenomenon cannot be explained by classical physics.
Thanks to these two properties of quantum mechanics, quantum computers promise to be for certain applications orders of magnitude more computationally

> The challenge today is to realize the promise of quantum computing with a scalable, reliable, and reproducible system.

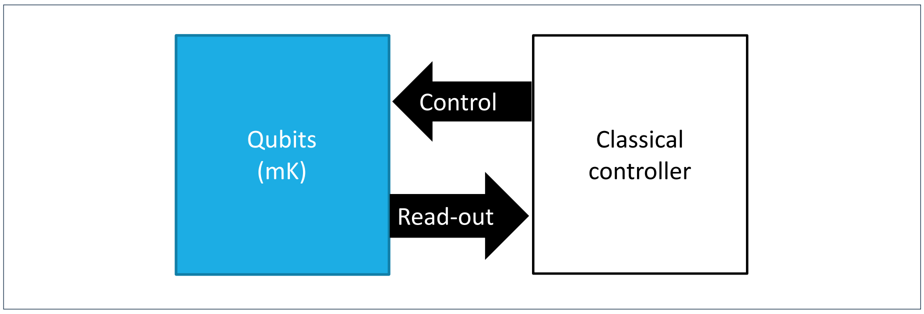

Fig. 3: Computer system comprising of the classical control system (left) and a quantum processor (right) at milli-Kelvin temperature.

powerful than the most advanced classical supercomputers. The challenge today is to realize this promise with a scalable, reliable, and reproducible system consisting of qubit platforms and the required control architecture.

A quantum computer comprises fundamentally a quantum processor and a controller for the qubits that make up the quantum processor. There exist today several types of qubits, based on solid-state and non-solid-state technologies. Among solid-state qubits, we differentiate them based on how the quantum state is encoded and we focus here on spin and superconducting qubits that are compatible, to some extent, with semiconductor technologies. Spin qubits, based on quantum dots, use either electron or hole spin to encode the quantum state. Superconducting qubits, implemented as a resonator combined with a non-linear load, use discrete energy levels to encode the quantum state. Both types of qubits must be operated at deep-cryogenic temperatures, so as to minimize the impact of thermal noise on the quantum state of the qubit. See Fig. 2.

The controller (see Fig. 3) is a classical device used to control, manipulate, and read out the quantum state of qubits. It is generally operating at room temperature, where it is made of instruments capable of generating high-frequency signals, similar to those used in smartphones.

These signals are generally over-amplified. To reduce the signal amplitude while maintaining a high signal-to-noise ratio, the signals are transferred to lower temperature levels in the used dilution refrigerator cryostat, inserting attenuators. Dilution refrigerators are systems that gradually cool from room temperature to deep cryogenic temperatures over several temperature stages. The quantum processor typically operates at a few tens of millikelvins, or approximately 273 degrees Celsius below zero. In our proposal the classical controller, or a part of it, is placed at 1-10K, so as to require only limited amplification to read out qubits and subsequent lower attenuation to control and manipulate qubits. Fig. 4 depicts the quantum stack envisioned.

Alternatively, qubit controllers, or parts of it, can be cooled down to cryogenic temperatures to a few degrees Kelvin, where classical circuit technology is adjustable

Fig. 4: Overall system stack embedding quantum algorithms, compilers (red) at room temperature; quantum control (yellow and green), and quantum classical interfaces (light blue) at intermediate 1-10K temperature levels; to the qubits (blue) at mK temperature.

to operate reliably [E. Charbon et al., IEDM 2016]. This integrated arrangement is advantageous for achieving a compact quantum computer, which is particularly suitable for scaling to a large number of qubits.
Another important property of cryogenic controllers is reliability, since large temperature gradients are not required and the reconfiguration of the control of existing qubits is done remotely from room temperature. To control qubits, it is imperative that the classical circuits that make up the

> CMOS can leverage 70 years of technical advances that led to today's powerful computers. It is only logical to use this technology for reliable control of scalable quantum processors operating near them at cryogenic temperature.

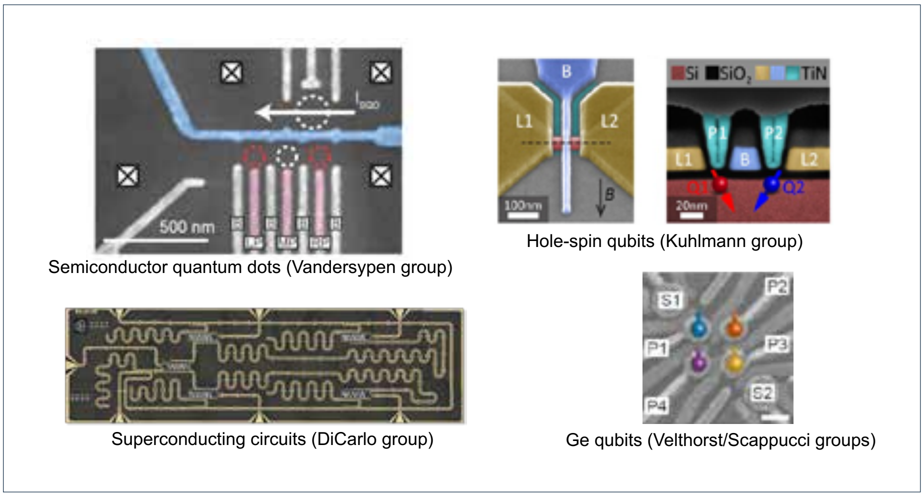

Fig. 2: Examples of solid-state qubits implemented as superconducting and spin flavors. These qubits must be o perated at deep-cryogenic temperatures, where thermal noise interferes less with the state of the qubit.

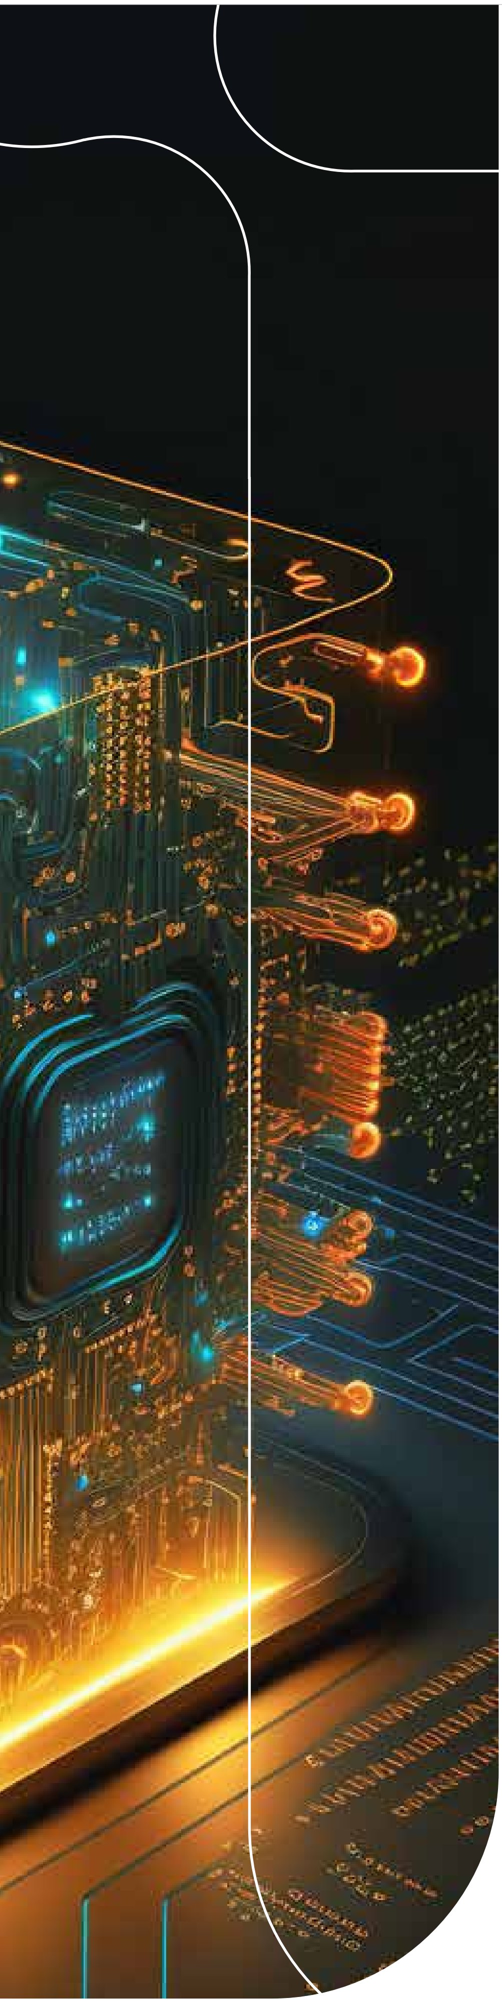

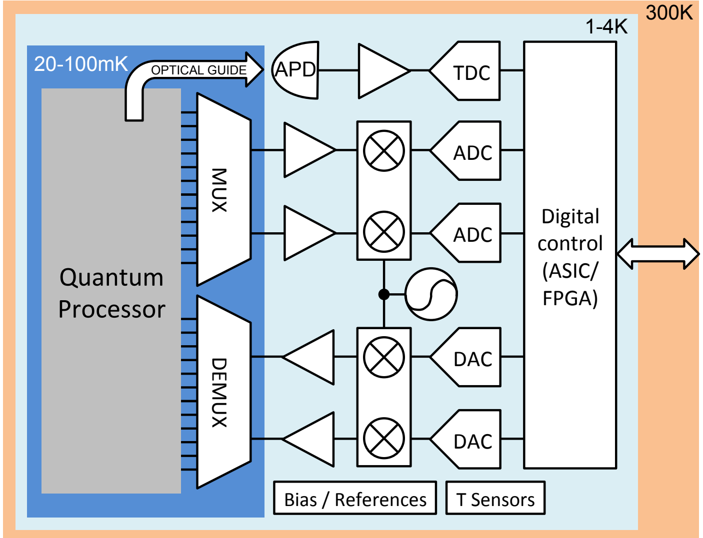

Fig. 5: Generic architecture of a cryo-CMOS qubit controller and its place in a dilution fridge.

qubit controller are designed to achieve very high performance in terms of speed and minimal noise, while achieving high precision in the frequency, phase and amplitude of the signals.

To achieve these specifications, we proposed to use cryogenic complementary metal-oxide semiconductors (cryo-CMOS) based circuits. But first, we needed to model the system and to understand what was important at top level. We chose to set an overall specification on fidelity, which represents how well we can control the qubit on the Bloch sphere. For example, a fidelity of 99.99% allows one to achieve a final quantum state for a qubit that has a maximum error with respect to the intended state of 0.01%, or one minus the fidelity. Then, we had to create an architecture of the controller that could work with most of the existing qubits. Our architecture, from a conceptual point of view, is very similar to that of a transceiver routinely used in radios and smartphones, see Fig. 5 [E. Charbon et al., IEDM 2016].

The challenge of operating at deep-cryogenic temperatures is not functionality but reliability and power. Ultra-low power electronics is living a renaissance motivated by quantum uses that could easily spill over advanced industrial applications.

Then, we had to come up with at least approximate specifications from high-level to the components of the control system in the architecture. For example, from a fidelity of 99.99%, we had to derive the signal-to-noise-and-distortion (SNDR) of the input low-noise amplifier connecting the qubits to the downconverters used to interpret the quantum state of the qubit. This was achieved by means of the top-level model in which the components had been added one-by-one [J. v. Dijk et al., PRA 2019].

The architecture shown in Fig. 5 has been used in most qubit controllers published since. An example is the Brévine chip [Y. Peng et al., CICC 2024]. The chip, named after the coldest place in Switzerland,

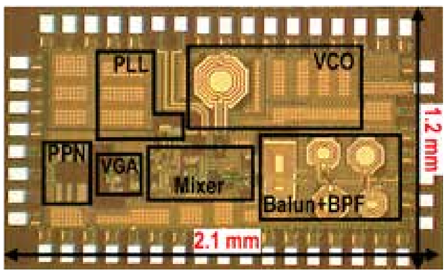

Fig. 6: Brévine chip's micrograph. The chip measures 2.1 X 1.2mm² [Y. Peng et al., CICC 2024].

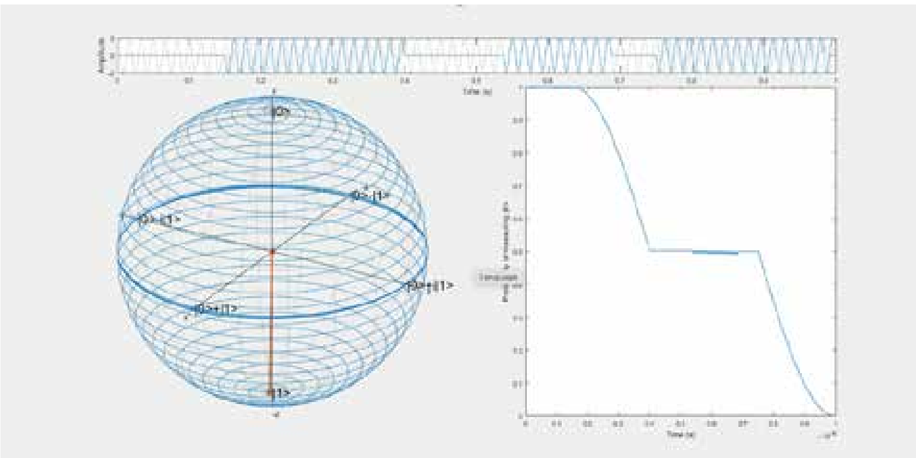

Fig. 7: Simulation of the quantum state of a spin qubit as it is controlled by a radio-frequency signal, modulated by a square wave.

comprises a two-stage up-conversion, similar to a superheterodyne in a radio, where an intermediate frequency is used to generate a GHz signal modulated by an envelope that can be programmed off-line at baseband. Fig. 6 shows the chip's micrograph, including the floorplan and the various functional blocks of the Brévine architecture.

In Fig. 7 one can see the simulated quantum state of a spin qubit as shown on a Bloch sphere [J. v. Dijk, pers. com., 2018] and the radio-frequency signal that was used to achieve the rotation of the quantum state vector from ground state to excited state, via the region of maximum superposition state, represented by the equator of the Bloch sphere.

The radio-frequency signal used in this case is hardly the only one that can be used in a qubit controller. In fact, there are several options, with different properties, especially visible in the spectral domain. Fig. 8 shows five different envelopes used to control qubits with a resulting variable spectral regrowth, which can be critical in multi-qubit control using frequency-division multiple-access techniques.

As mentioned earlier, fidelity represents, in the case of qubit controllers, one minus the error that can be achieved in the final quantum state programmed for a given quantum component. Fidelity is of paramount importance when it comes to qubit control, because it determines whether a qubit can be used in a quantum algorithm and if quantum error correction is necessary. Even if quantum error correction were not necessary, the qubits in a quantum processor will lose coherence after a certain time, that can be several millionths or at most a thousandth of a second. These coherence times, characterized by T1, T2, T2*, etc. are widely varying in qubits depending on the type of qubit and the operating temperature and frequencies of the readout system.

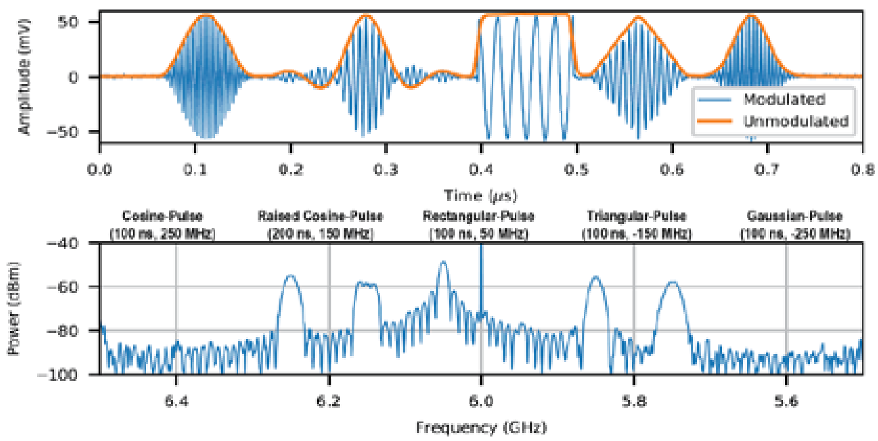

Fig. 8: Several types of envelopes (orange colored) often used in qubit control [J. V.Dijk et al., JSSC 2020].

Quantum Error Correction (QEC) coding is a possible technique to achieve a nearly perfect qubit from many imperfect ones. Surface codes are methods that describe a multidimensional coding of so-called logical qubits, based on several physical qubits. Quantum error correction requires considerable control effort, which is ideally placed close to the qubits for performance reasons.

At the time this article was written, the largest quantum machines count hundreds if not thousands of physical qubits in various flavors. Researchers are optimizing qubit control both at room and cryogenic temperatures introducing quantum instruction sets to control qubits in ways that may in the future become machine independent.

Quantum instruction sets and compilers are key to achieve systems that are agnostic to qubit technology.

Finding the right instruction set is important to enable the development, in the near future, of quantum compilers, capable of orchestrating quantum algorithms, quantum circuits, and qubit controllers. An example of such instruction set is shown on Fig. 9.

Conclusions

We are at the beginning of the next phase that will make quantum computing practical. Research is needed in both areas, system architecture and quantum compilers, to develop economical quantum computers and increase the number of available quantum-advantage algorithms. Quantum-classical control must become more integrated and modular, less power-hungry, and truly ready for scalable quantum processors containing several thousand to several million qubits. Most importantly though, the qubit platforms will need to improve reliability, reproducibility, and scalability to reach these numbers.

We envision several upcoming major challenges:

- First, we believe that, in parallel with the upscaling of qubit platforms, several improvements in controllers will be required to control and read-out a large number of qubits with very low energy expenditure in the cryostat. As an

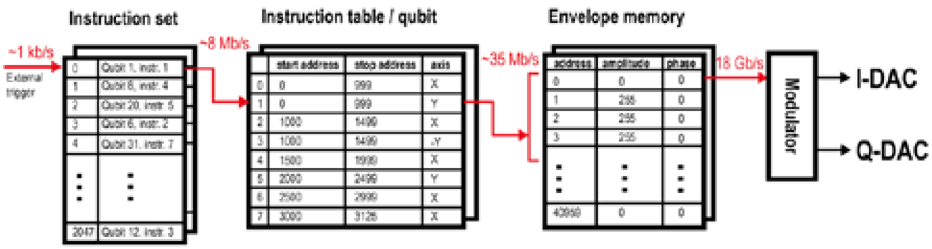

Fig. 9: Example of the instruction set used in Horse Ridge [X. Xue, B. Patra et al., Nature 2021, J. v. Dijk et al, JSSC 2020].

engineering estimation, if 1000 qubits will require a controller ensuring no more than 1mW per qubit, then a million qubits will need to be controlled at less than 1μW per qubit.

- Second, the instruction set will need to become standardized for a relevant number of architectures, so as to enable a certain level of diversity, while maintaining flexibility within various processors.
- Third, in order to become widespread, quantum computing will need a compiler from a higher-level language, which will encourage algorithmic research and development.

In summary, the adventure of a practical quantum computer has only just begun, and the first ideas for advantageous applications are emerging. This will help to create synergies and exchanges between quantum computers and their components and other industries, which will ultimately be the key to success. The first commercial microprocessor came from industry and was responding to a clear and precise need, that of a table calculator. It is likely that a practical quantum computer will also come from industry, whether a large enterprise or a startup, but it will require a clear application. In the case of Intel's 4004, an effective synergy between physicists developing the poly-gate transistor, architects implementing a Von Neumann architecture, and application engineers developing machine code and the assembler made it a great success, but it came down to a very few people that had the vision for a practical processor. The same fate will apply to quantum computers, and an application that offers a clear quantum advantage will point the way to success. Only time will tell.

The current publication "Quantum Europe Strategy", which was published in July 2025, provides a comprehensive overview of current topics in quantum technologies. [https://qt.eu/media/pdf/Quantum_Europe_Strategy_July_2025.pdf]

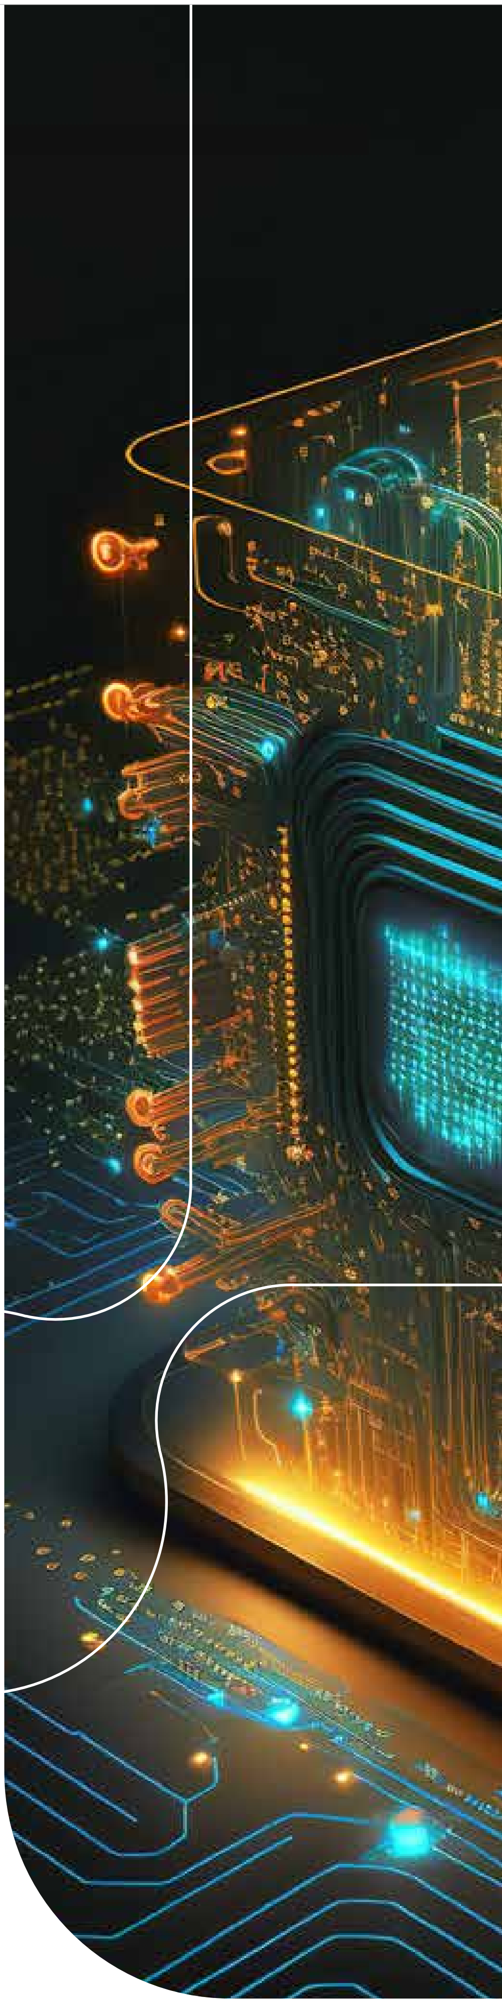

INSIDE Members Focus

Sony depthsensing solutions

Powering the next wave of 3D vision

In the race to make machines smarter, faster and more human-aware, one element stands out: depth. The ability for a device to not just 'see' but to understand space, distance and shape in three dimensions is rewriting the rules in industries from automotive safety to robotics and consumer electronics.

The corporate picture in Europe

Pioneering this revolution is Sony Depthsensing Solutions, a Brussels-based subsidiary of Sony Semiconductor Solutions with a satellite office in Lund, Sweden. Sony Semiconductors Solutions (SSS) (https://www.sony-semicon.com/) is a core technology company within the Sony Group, headquartered in Japan. It is the parent company of all Sony's semiconductor-related businesses, including Sony Depthsensing Solutions (SDS) and SSS-Europe. It leads the research, development, manufacturing and global sales of Sony's semiconductor products. SSS Europe is the European branch and represents SSS's commercial and support operations in Europe, working closely with customers, partners and other internal Sony entities.

Core proposition: seeing in three dimensions

Unlike traditional cameras, which capture flat images, depth sensors measure distance. Sony specializes in various depth sensing technologies, including Time-of-Flight (ToF), which, along with Lidar systems, offers high accuracy and easy integration across a wide range of devices, from mobile phones to autonomous vehicles.

The key benefits are:

- **High accuracy**: Essential for applications like facial recognition, hand gesture tracking, and object scanning.
- **Easy integration**: Sony Depthsensing Solutions technology can be embedded across a wide range of devices, from mobile phones to autonomous vehicles.
- **Versatile applications**: Used in fields like robotics, automotive, mobile and smart infrastructure.

Together, these technologies enable products that can recognise gestures, map spaces, detect obstacles and even track the position of a driver's eyes in real time.

From startup to strategic player

As Ilse Ravyse, Senior Planning Manager, explains, "Sony Depthsensing Solutions was born out of a merger between two European pioneers in 3D imaging, SoftKinetic (founded in 2007) and Optrima." Ilse was already at SoftKinetic when the company's early innovations in gesture recognition and Time-of-Flight (ToF) imaging quickly caught Sony's attention. In 2015, Sony acquired the venture, and by 2017 it was rebranded as Sony Depthsensing Solutions, aligning its mission with the broader vision of Sony Semiconductor Solutions.

For Sony, the acquisition wasn't just about technology. It was about securing

Daniël van Nieuwenhove

a leadership position in the next major interface between humans and machines. By combining advanced hardware with intelligent middleware, Sony Depthsensing Solutions represents a strategic play to make 3D sensing as mainstream as cameras and microphones are today. Currently, more than 150 employees from 27 nationalities are responsible for more than 20 patents per year.

Why it matters: applications driving growth

Sony Depthsensing Solutions enables technology for some of the most interesting and promising growth markets of the coming decade:

- **Automotive**: In an era of autonomous driving and enhanced safety systems, Sony's sensors deliver safety and comfort solutions with its in-cabin sensing solutions and 360° awareness outside the car. Gesture-based infotainment controls and driver and activity detection systems are redefining how we interact with vehicles. The vision is to contribute to a zero-fatality future through real-time perception and proactive safety systems.
- **Robotics and Industry**: depth sensing solutions help factories and warehouses become smarter, safer and more agile. It also enables the building of digital twins of physical environments for simulation, monitoring and optimisation as well as support autonomous navigation for robots and drones. Finally, depth-sensing solutions enable safe collaboration between humans and machines in shared workspaces.
- **Consumer Electronics**: From smartphones with faster autofocus to augmented reality experiences that feel lifelike, Sony's ToF modules are quietly becoming standard in next-generation devices, empowering smartphone and AR/VR manufacturers and developers with cutting-edge 3D capabilities.
- **Smart Environments**: Touchless interfaces, occupancy detection and intelligent monitoring are unlocking new possibilities in home automation, healthcare and workplace safety.

A competitive edge: hardware + software

To emphasise the skills of the R&D team at Sony Depthsensing Solutions regarding integration, it's important to highlight their comprehensive approach.

Sony Depthsensing Solutions excels in developing embedded systems that seamlessly combine hardware and software, ensuring that depth sensors function optimally within various applications. The teams create custom drivers that facilitate communication between the depth sensors and host devices, allowing for smooth data transfer and enhanced functionality.

Additionally, the R&D teams focus on middleware development, which interprets raw depth data and transforms it into actionable insights. This capability is crucial for automakers and device manufacturers looking to leverage depth-sensing technology effectively. Toolkits like CARlib™ and ICMlib™ are specifically designed to simplify the integration process, enabling customers to implement solutions without needing extensive R&D on their end.

Moreover, the R&D teams prioritise customer support and collaboration. They work closely with clients to understand their specific needs, providing tailored solutions that fit seamlessly into existing systems. Ongoing support ensures that customers can effectively implement and utilize the depth sensing technology in their products. This holistic approach not only showcases the expertise of the R&D teams but also highlights the significant value they bring to their customers.

In a market where speed to application is everything, this "full stack" approach gives Sony a distinct business edge. Customers don't just buy sensors; they buy ready-to-use systems that plug into their products.

Culture and values

As part of Sony Group, Sony Depthsensing Solutions embraces the company's mission: *"To fill the world with emotion through the power of creativity and technology"*. This ethos is deeply embedded in the team's culture, reflected in its core values: bring fun to the table, dare to have impact, first understand, then be understood. These values guide cross-functional teams across engineering, R&D, support and innovation. Collaboration, curiosity and a passion for problem-solving define the work culture, making Sony Depthsensing Solutions a magnet for forward-thinking talent.

Market position and strategic outlook

Sony Depthsensing Solutions' influence is amplified by Sony's semiconductor manufacturing muscle and global brand reputation. Depth sensing sits at the intersection of multiple mega-trends: autonomous mobility, smart living, industrial automation and immersive computing. The company's roadmap aligns directly with these growth areas. And being able to call on the Sony R&D network, both in Europe and worldwide, Sony Depthsensing Solutions is an attractive partner for European projects and consortia, especially among the Inside community. As Innovation and Communication manager Christopher Littlefair remarks, "There's always an interlink between projects and business." And business leaders across industries are looking for technologies that enhance safety, efficiency and customer experience. "Depth sensing ticks all three boxes. Depth sensing makes digital interaction more natural, intuitive and safe so, in that sense, Sony Depthsensing Solutions is less a component supplier than a strategic enabler of the next wave of digital transformation."

As industries demand deeper insights, more natural interfaces and spatial awareness, Sony Depthsensing Solutions is positioned to lead, with 3D perception poised to become as essential as audio/video in digital devices. By bridging cutting-edge semiconductor design with practical, deployable solutions, it is helping define the future of mobility, automation and consumer technology. Not just building sensors but reshaping digital interaction. For businesses, the takeaway is clear: 3D sensing isn't coming, it's here.

New Member Focus

Nexuswelt Group

Driving European innovation through strategic excellence

Europe has long been a hub of scientific excellence, but ensuring that research results translate into tangible benefits for society and the economy remains a challenge. The success of the European Union's ambitious digital transformation agenda depends not only on developing advanced technologies, but also on bringing them into the market and ensuring broad acceptance.

This is where Nexuswelt Group, founded by Anna Lackner, has carved out a distinctive role. By combining strategic communication, project management, and training expertise, the agency helps bridge the gap between European research excellence and practical implementation.

From EU projects to a strategic vision

Anna's career path reflects the diversity of Europe's innovation landscape. For nearly a decade, she supported Infineon Technologies in Munich as an external consultant, contributing to major publicly funded projects in the semiconductor and automotive sectors. This work gave her first-hand experience in coordinating large consortia, managing dissemination and communication tasks, and ensuring that complex technological innovations could be understood and adopted across industries.

In parallel, she gained valuable marketing experience in the IoT sector at eesy-innovation GmbH, where she helped shape go-to-market strategies and digital campaigns for emerging connected technologies.

Together with her extensive involvement in landmark EU projects such as AIMS5.0, PRODUCTIVE4.0, and SC4EU, this blend of technical, strategic, and marketing expertise gave Anna a unique perspective.

"These projects were extraordinary in their ambition," she recalls. "Dozens of partners, multiple technologies, different cultural and organizational contexts—yet all striving for a common goal. The experience gave me invaluable insights into what works, what doesn't, and how communication and coordination can make or break success."

This perspective shaped the founding vision for Nexuswelt Group: a partner that understands the realities of EU innovation ecosystems and supports organizations in moving beyond research outputs towards real market and policy impact.

Three pillars of impact

At the heart of Nexuswelt Group's approach are three interlinked pillars:

- Strategic Communication – Making sure that research is not only understood but also embraced by stakeholders, from policymakers to end users.
- Collaborative Innovation Management – Ensuring smooth cooperation across diverse consortia, aligning technical excellence with common objectives.
- Market-Ready Implementation – Supporting partners in turning results into viable products and services that strengthen Europe's competitiveness.

"We want to be the bridge that ensures brilliant European innovations don't get lost in translation," Anna emphasizes.

Expertise in EU project realities

Unlike generic consultancies, Nexuswelt Group's expertise is deeply tied to the specific requirements of EU-funded projects. This includes:

- Tailored Dissemination Strategies – going far beyond minimum compliance, with campaigns that actively support adoption and visibility.
- Technology Acceptance Support – applying models like TAM and UTAUT to identify potential barriers and adapt dissemination early on.

- Risk and Quality Management – introducing structured methods, supported by digital tools, to anticipate challenges in large multi-partner projects.

This hands-on knowledge of EU processes enables Nexuswelt Group to guide partners effectively, especially newcomers who may otherwise find the framework daunting.

Empowering SMEs across Europe

One of Nexuswelt Group's strongest commitments is to small and medium-sized enterprises (SMEs). While large corporations and universities often dominate EU consortia, many breakthrough innovations originate in smaller firms.

"SMEs often bring the most creative and disruptive ideas," Anna notes, "but they can struggle with the administrative and strategic aspects of EU projects. Helping them succeed is essential if Europe wants a truly inclusive innovation ecosystem."

To support this, Nexuswelt Group runs training and capacity-building programmes covering proposal writing, project management, communication, and compliance. These are not theoretical lectures, but practical workshops shaped by real project experience.

In addition, the agency facilitates partnership building by connecting SMEs with larger players, ensuring that their innovations find the right platforms for growth.

Contributing to European digital sovereignty

A recurring theme in Nexuswelt Group's work is the concept of digital sovereignty—Europe's ability to maintain control over its digital infrastructure, data, and innovation capacity. This is a top priority in EU policy, and one where Nexuswelt Group sees clear opportunities.

"Digital sovereignty is not only about technology," Anna explains. "It is about creating the strategic conditions for Europe to compete globally while protecting its values. That means strong ecosystems, sustainable innovation, and collaboration across borders."

Through its focus on communication, collaboration, and market readiness, Nexuswelt Group directly contributes to these goals, ensuring that EU-funded innovations generate value inside Europe rather than being absorbed elsewhere.

Sustainability and long-term impact

Another area where Nexuswelt Group aligns closely with EU priorities is sustainability. Beyond supporting technical progress, the agency emphasizes circular economy principles, environmental responsibility, and long-term socio-economic value.

"European innovation must not only be competitive, it must also reflect our responsibility to future generations," says Anna. "That is why we embed sustainability considerations into everything we do, from project planning to communication."

Looking ahead: building tomorrow's ecosystem

As Europe continues to invest heavily in digital and green transitions, the role of organizations like Nexuswelt Group is becoming more important. By ensuring that research results reach markets, policymakers, and citizens, they help strengthen Europe's position in the global digital economy.

Nexuswelt Group envisions a future where every EU-funded innovation has a clear pathway to impact. This means more SMEs gaining access to opportunities, more effective collaboration across borders, and stronger alignment between technological progress and societal needs.

"We're not just a service provider," Anna concludes. "We see ourselves as a partner in building the European innovation ecosystem of tomorrow."

Learn more:
www.nexuswelt.com
anna.lackner@nexuswelt.com

Research Project Highlight

A-IQ Ready

Artificial Intelligence for realtime distributed systems at the edge

The pressing challenges of climate change, geopolitical instability, and growing societal inequalities underscore the need for transformative innovation. Emerging technologies such as artificial intelligence, the Internet of Things (IoT), robotics, and quantum sensing are playing a pivotal role in addressing these global concerns, offering new pathways toward a more resilient and sustainable future.

The A-IQ Ready project is at the forefront of this technological evolution, developing a next-generation AI-driven autonomous ECS to meet the demands of the digital era. By integrating edge continuum orchestration, distributed collaborative intelligence, and quantum sensing, the project is driving the transition towards Society 5.0, a human-centered, digitally empowered society where advanced technologies seamlessly enhance economic and social systems.

The project is structured around five core objectives:

- Enhancing sensing capabilities through novel multi-physics quantum sensors
- Developing AI-driven multi-agent autonomy for unpredictable environments
- Establishing a reference AI Edge Continuum platform, integrating quantum sensing, neuromorphic computing, and edge AI algorithms
- Demonstrating the technological framework as a foundation for the digital society
- Strengthening Europe's global leadership in secure, resilient, and intelligent digital systems.

Considering the scope of the project in the context of different partners and their specialisation, the project tackles many different fields, such as Automated and Manual Transport at Industrial Sites, Search and Rescue, digital health and emergency Recognition for Driver and Operator, propulsion health and availability in safety-critical situations, quantum sensors, hybrid computing, and cooperative Multi-Agent Systems.

Automated and Manual Transport at Industrial Sites

A-IQ Ready has two demo sites in Finland: Kouvola and Rauma, focusing on industrial transport automatization. Kouvola demo vehicles, Terminal Tractor and Reach Stacker, are being tested at Kalmar facilities. Besides the vehicles, also their remote operator support is developing, including database response models, Prompt-to-SQL model training, and the remote operator portal including integration of Reach Stacker data playback and direct SQL-queries with Large Language Model response.
The Rauma demo AGV vehicle and the log cages to transfer are ready for fully computerized control. Preliminary tests for control have been conducted and the AGV and log cages will soon be transferred to Unikie's testing facilities for further testing with infrastructure-based lidars. To assist in the development of the control system, the following tasks have been carried out: detailed the future demonstration area at Rauma sawmill to include the busy unloading location of the log cages and related areas nearby, 3D mesh models of AGV, log cage, and the future demo area at Rauma, lidar placement tool development, detailing of driving, odometry and localization algorithms including IMU usage, especially when driving under the log cage.

The safety aspects of the self-operating vehicles in both demo sites have been analysed by applying Preliminary Hazard Analysis and System Theoretic Process Analysis. Also, the approach described in UL4600 – Safety Standard for Autonomous Vehicles – has been applied while working towards the functionality of dynamic risk assessment, which could be, in the future, a complementary part of situational awareness in autonomous vehicles. Implementing dynamic risk awareness in a simulator and, further, in a test vehicle is ongoing.

Search and Rescue (SAR)

SAR missions in confined environments require an extensive environmental understanding by leveraging AI-driven systems, allowing advanced algorithms to process the vast amount of data and supporting complex decision-making. A-IQ Ready SC2 addresses the system development for SAR missions in GNSS-denied environments, such as tunnels, utilising robotic platforms either in simulation or in real-world scenarios related to three dedicated use cases: Survivor Detection in a tunnel section, Multi-robot Localisation in simulation, as well as Simulation of wireless systems for SAR in harsh subterranean environments. The challenging scenario requires algorithmic solutions that cope with environmental constraints to realise survivor detection, self-localization, mapping of the environment, and communication under challenging conditions. This period's focus has been on system architecture design, simulation frameworks, environmental and system modelling, risk assessment, wireless system simulation, and initiating algorithm development. Detailed simulation models of the tunnel environment have been developed, serving as a high-fidelity simulation testbed for validating algorithm performance under realistic conditions.

As another key milestone, the following activities were conducted in this period: localisation and mapping algorithm benchmarking, dataset generation for detection and prediction algorithms, and the initial implementation of algorithms utilising the models and simulation environments.

Digital Health and Emergency Recognition for Driver and Operator

The risk of a traffic accident is often directly related to a driver's attention level. For that reason, tools for estimating the driver's health status and overall well-being are developed to increase traffic safety. While state-of-the-art methodology relies mainly on tracking the eye movements and face expressions using vision systems, A-IQ Ready tackles the problem from three different perspectives. Estimation based on eyes and face expressions will be fused together with biometric parameters and in-cabin air quality information. In this way, the system considers additional measurements, such as heart rate or stress level, as well as environmental information, which results in a more informed conclusion.

The 2nd project year closed with more mature systems for in-cabin monitoring. The implementation of the software needed for the COTS wearable device and the in-cabin environmental conditions sensor is in progress. Both sensors have been tested in a real scenario, driving in various areas of the city from downtown to country roads nearby. In parallel, initial tests have been performed on the demonstrator vehicle to prepare it for the full system integration. Demonstration consists of both the driver monitoring module and vehicle's decision and control module that perform an emergency response in case of reduced driver's attention. Implementation of the safe lane change maneuver is in progress, as well as integration of other components.

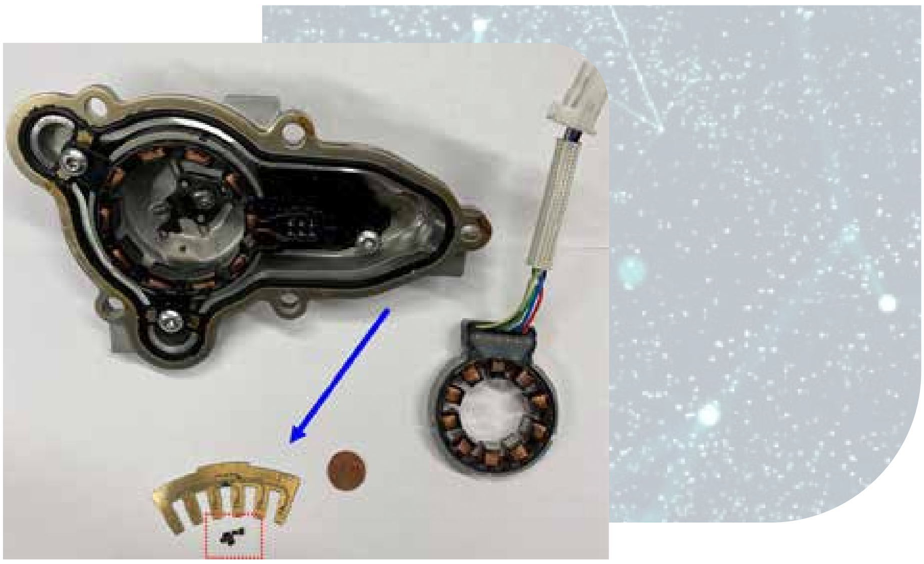

Figure 1. State-of-the-Art mechanical BEV resolver is replaced by integrated sensor array

Propulsion health and availability in safety-critical situations

A-IQ Ready team embarked on the development of an innovative E-Motor concept featuring an integrated flux sensor array. This concept underwent extensive simulation and discussion phases. The concept was completely adapted to meet difficult automotive requirements concerning high power density and a modern oil-cooling instead of water-cooling.

This transition was driven by the need for improved efficiency and performance, showcasing the team's adaptability and problem-solving skills.
The culmination of these efforts resulted in a robust E-Motor concept incorporating five different integrated sensor arrays with innovative interfaces. The multi-sensor approach enhances the motor's functionality, reliability and cost.

The team conducted comprehensive FEM (Finite Element Method) simulations and thermal simulations. These simulations were crucial in validating the design and ensuring optimal performance under various conditions. Moreover, a prototype is build to validate the simulations.

In summary, the team's dedication and innovative spirit have led to the development of a cutting-edge E-Motor concept. The integration of multiple sensor arrays and the transition to oil cooling are testaments to their relentless pursuit of technological advancement.

Quantum sensor

The A-IQ Ready project is tackling key challenges in quantum sensor development, focusing on stability (in the context of resilience to disturbances like vibrations and temperature variations), precision and acquisition speed (acceleration of the lock-in amplification measurement method). Year two saw the completion of a software control interface, optimised dual DAC iterations, an ultra-precise compact quantum sensor head prototype, a fully fiber-based optical setup and an improved microwave antenna design. Preparations are being made for two use cases with diametrically opposing requirements: navigation using Earth's magnetic field, focusing on precise measurements of weak static fields, and

real-time motor control that demands fast sensing of strong varying fields at elevated temperatures. Striving to improve the measurement precision, the team is currently characterizing noise levels and coming up with improved calibration procedures. In order to reach higher measurement speeds, an improved “ping-pong” architecture featuring multiple frequency synthesizers is being considered. Continuing the system’s development, other ongoing tasks include updating the GUI and creating portable measurement systems for partner collaborations. The Bottom Line: A-IQ Ready is committed to developing stable, precise, and user-friendly quantum sensors for diverse applications.

Hybrid Computing

A-IQ Ready partners have been working on the research to evaluate the performance of domain-specific accelerators, Instructator, implemented as a RISC-V custom instruction, and Intelligator, designed as a memory-mapped accelerator, against state-of-the-art platforms. For this comparison, A-IQ Ready selected the ARM Cortex-M7 within the STM32F746NG development board.

To benchmark performance, we implemented a representative bare-metal kernel for matrix-vector (MxV) multiplication [64×784]×[784×1] and developed a TinyML-based multi-layer perceptron (MLP) model.

With a high-performance configuration utilizing four Intelligators, our system achieves a 105× speedup over the RISC-V CPU baseline and a 106× speedup compared to the ARM Cortex-M7 baseline in the MxV kernel. In a low-power configuration using Instructator, our system remains 13× faster than both the RISC-V CPU baseline and the ARM Cortex-M7 core.

For the MLP model, in the best-case scenario, our accelerators provide a 24× speedup compared to running the full model on the RISC-V CPU core. Furthermore, our approach outperforms the ARM system by factors of 1.47× (Instructator) and 3.82× (four Intelligators), demonstrating the efficiency of our design in both performance and energy-conscious configurations.

In addition, efforts have been made concerning a common integration platform based on X-Heep which has contributed features accepted into the X-Heep project repository.

The development efforts in the area of MLIR tooling led to a novel source-to-source framework, highlighting its innovative approach to simplifying MLIR transformations for heterogeneous computing backends.

Cooperative Multi-Agent Systems

A-IQ Ready partners have made significant strides in optimizing intralogistics through reinforcement learning (RL) and standardized automation technologies. These developments focus on intelligent scheduling, efficient navigation, and seamless integration of AI-driven decision-making with Automated Guided Vehicles (AGVs).

By developing demonstrators, the consortium has established a foundation for scalable, AI-powered logistics solutions. By leveraging digital twins, cloud-based RL training, and advanced simulation platforms (custom simulation environments, Unity, Gazebo), we are bridging the gap between research-driven innovations and industrial applications. These developments pave the way for more adaptive, cost-effective automation in supply chain operations, aligning with Industry 4.0 principles.

As A-IQ Ready progresses, it continues to play a critical role in advancing Europe’s competitiveness in cutting-edge AI, quantum sensing, and autonomous systems. The insights and innovations from this project will help pave the way for a more sustainable, efficient, and secure digital future.

Acknowledgment

A-IQ READY receives funding within the Chips Joint Undertaking (Chips JU) - the Public-Private Partnership for research, development and innovation under Horizon Europe – and National Authorities under grant agreement n° 101096658.

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the granting authority. Neither the European Union nor the granting authority can be held responsible for them.

Co-funded by the European Union

New Member Focus

The future of housing

Inside GROPYUS's mission to rethink how we build

Haris Isakovic

Imagine a half-finished apartment block. The building looks deceptively ordinary, glass, timber, clean modern lines. Yet what rises here is anything but conventional. This isn't a product of slow, costly construction sites. It's the result of robotics, prefabrication and a new philosophy that asks the question: *what if buildings were products, not projects?*

This is the question that at the heart of new INSIDE Industry Association member GROPYUS, a German-Austrian technology-based construction company that is reinventing multi-family housing. For its founders and the whole multinational team, the mission is personal and urgent. They are taking on three of Europe's most pressing crises at once: an escalating housing shortage, a construction sector in decline and a climate emergency that won't wait.

The housing crunch

Germany alone needs nearly 850,000 new homes. And every year, that gap is widening as families become stuck in bidding wars, young workers can't move where jobs are, and companies can't hire because employees can't afford to live near the office. It's a familiar story across Europe. In the first half of 2024, building approvals for apartment blocks plummeted by more than 20%. Rising interest rates, skyrocketing material costs and burdensome regulation have slowed construction to a crawl. The result is a social and economic bottleneck. Without affordable housing, mobility falters, inequality deepens and cities lose their dynamism. GROPYUS has a clear vison of what housing should be. Not just shelter, but infrastructure for society. And right now, that infrastructure is broken.

A factory for multi-family homes

Walk into GROPYUS's smart factory in Richen, and the scene feels rather like a car plant than a construction site. Automated lines hum. Robots cut, assemble and refine panels of timber into wall and slab elements. Every 16 minutes, another piece is finished. Each component is tagged, tracked and fed into a digital system that knows exactly where it belongs. The approach is radically different from traditional building, where unpredictable weather, delays and labour shortages can grind projects to a halt.

GROPYUS shifts more than 80% of construction into a controlled, highly automated environment. Once shipped to a site, the pieces click into place like Lego. The payoff is speed, with construction times cut by half, producing the equivalent of 3,500 apartments per year, scalable, repeatable and adaptable. "Creating sustainable living for everyone" is the GROPYUS line. Their building system is modular, so each building can be tailored to its location, whether that's an urban infill-development project in the centre of Berlin or a residential quarter close to Lake Constance.

Timber: an ancient material with a modern role

At the core of GROPYUS's philosophy is timber. While concrete and steel dominate most skylines, both come with a heavy carbon footprint. By contrast, sustainably sourced wood is renewable, stores carbon and offers warmth and texture that residents love. Timber has long been a traditional building material, but digital technologies are now transforming how it's used. When sourced and managed responsibly, it offers significant potential for reducing carbon emissions in construction.

The numbers back it up. GROPYUS's pilot project, a high-rise timber-hybrid building comprising 54 apartments across nine stories, already emits 95% less greenhouse gases over its lifecycle compared to reference values provided by the German Sustainable Building Council and the German Buildings Energy Act. That's no mean feat. Globally, the building and construction sector contributes nearly 40% of greenhouse gas emissions. Without systemic change, climate goals are out of reach. By proving that low-carbon housing can also be fast, affordable and scalable, GROPYUS is putting its money on timber as a bridge to a sustainable future.

Digital by design

At GROPYUS, digitalisation isn't an add-on, it's the backbone of their entire process. From planning and production to construction and long-term operations, every phase of the real estate lifecycle is integrated into a unified digital ecosystem. Conflicts between trades and delays caused by fragmented workflows are eliminated because everything, from design to delivery comes from a single source.

Before a single building element is produced, the entire building exists as a digital twin. This virtual model isn't just a visual preview; it's embedded in a smart data fabric that captures and connects information across every step. Materials, methods and timelines are all mapped and analysed, creating a feedback loop that continuously improves both the product and the process. This intelligent framework doesn't just streamline construction, it lays the groundwork for future innovation. The result? Buildings that are not only faster to build, but smarter to operate and easier to adapt to customer needs.

Innovation in action

To push innovation within the construction industry, GROPYUS has already made strong R&D efforts. The company is actively participating in the EU funding landscape with two projects Cynergy4MIE and 5G-FACT. Cynergy4MIE (No 101140226) is a high impact cross-domain project in scope of Chips Joint Undertaking. Within this project, they aim to bridge the gap between parametric design of timber-frame building elements and fully automated robotic simultaneous production of those elements. Moreover, in their second European project 5G-FACT within the scope of a cascade funding scheme under the Target-X Horizon 2020 (No. No 101096614) project, they are exploring high-performance communication infrastructure requirements for autonomous on-site construction.

Homes that think

But GROPYUS isn't only reimagining the materials and processes of construction. They're also rethinking what it means to live in a building. Every GROPYUS project comes with its Building Operating System (BOS), a digital platform that turns apartments into smart, connected environments. For residents, it's a dashboard for daily life: controlling lighting, heating and blinds; tracking energy use in real time.

For landlords and asset holders, BOS is a cost-saver and a revenue stream. It simplifies maintenance, identifies inefficiencies and opens the door to new income via resident subscriptions. It's about creating buildings that learn and evolve, where they can continuously update features, adapt to residents' needs and keep pushing for greater sustainability.

Scaling up in a shrinking industry

The construction sector has been slow to innovate. For decades, productivity lagged behind industries like automotive and manufacturing. Now, as skilled labour shortages loom, the need for transformation is urgent. GROPYUS sees itself as part of that transformation. By treating buildings like products, it brings industrial discipline to an industry plagued by inefficiency. Costs become predictable. Deadlines become reliable. And sustainability becomes measurable, not optional.

The company's pipeline includes various projects in Germany with more to come very soon. From a 27-apartment infill project in the German capital to a nine-building residential quarter in Baden-Württemberg, GROPYUS's

building system is proving adaptable to diverse contexts. Looking ahead, GROPYUS expects to expand annual production to 300,000 square meters of gross floor area. The timing may be perfect: industry analysts predict a rebound in construction from 2026 onwards, precisely when GROPYUS's model is ready to scale.

With growth based on solid strategic decision-making, continuous development and refinement of products and processes, as well as close contact with customers, GROPYUS is constantly growing in order to ensure long-term success and reliable, innovative solutions.

The people behind the mission

Behind the technology is a diverse workforce of around 500 employees spanning more than 50 nationalities. Their culture blends startup agility with construction grit, anchored by a shared sense of purpose. They want to build something that matters, not just another project that runs over budget and damages the planet. The company backs that sense of ownership with a Virtual Share Program, giving every employee a stake in its success. For the founders, who previously built companies worth a combined $30 billion, the philosophy is simple: people build companies, not just buildings.

The bigger picture

The stakes couldn't be higher. Europe's housing shortage is eroding economic growth and deepening inequality. The construction sector's struggles risk pulling the wider economy down with it. And the climate clock is ticking. GROPYUS doesn't pretend to have all the answers. But its model offers a vision of how the pieces might fit together: speed to address shortages, affordability to restore access and sustainability to secure the future. There is every chance that the impact will ripple far beyond Germany. The global market for innovative construction methods could reach $1.1 trillion by 2040, with double-digit annual growth. The question isn't whether such transformation will happen, but who will lead it.

Building hope

Back at the Berlin construction site, workers slot prefabricated timber building elements into place with crane-assisted precision. In weeks, not months, the building is completed. Soon, families will move in, students will study at kitchen tables and children will play in hallways. For them, this will simply be home. For GROPYUS, it will be proof that housing, one of society's oldest challenges, can be reimagined with creativity, technology and courage.

Research Project Highlight

"We're building trust into technology itself"

Inside NexTArc: Europe's bold experiment in trustworthy embedded AI for cities, industries and mobility

In the coming decade, the world's infrastructure will increasingly depend on embedded artificial intelligence (AI). From traffic systems and logistics hubs to energy grids and manufacturing plants, AI is rapidly becoming the hidden layer that makes modern society run. But as this technology grows more powerful, so do the stakes: security, trust, energy efficiency and regulatory compliance can no longer be afterthoughts. They're prerequisites.

One of Europe's most ambitious responses to this challenge is **NexTArc**, short for *Next Generation Open Innovations in Trustworthy Embedded AI Architectures for Smart Cities, Mobility and Logistics*. Coordinated by Mälardalen University in Sweden, the three-year Horizon Europe project brings together 38 partners across 10 countries and aims to build the next generation of embedded AI architectures.

"We're building trust into the technology itself," says **Baran Cürüklü**, Associate Professor at Mälardalen University and coordinator of NexTArc. Here he outlines how the project evolved, why trust is at its core, and what it could mean for Europe's technological future.

From mobility to a complete ecosystem

The original idea behind NexTArc was far narrower than it is now. "The idea started from mobility and transportation," Baran recalls. "We had no focus on living spaces or workspaces. It was all about vehicles and movement." But during the first year, the European Commission encouraged the team to think bigger. That push led to a fundamental shift: the team expanded its focus from mobility alone to the entire living environment, the spaces where people live, work, and move.

A symbolic turning point was the inclusion of Sweco, a company far removed from chips and AI. "Sweco has nothing to do with hardware. But our vision was to work with the entire value chain, from silicon labs and algorithms to architects who design cities and living spaces." The expanded vision paid off. The project scored 14.8 out of 15 in its evaluation, a near-perfect result. It introduced a crucial concept: the *in-between spaces* that connect homes, workplaces, and mobility networks. "People often focus on home or office," Baran explains. "But the stuff in between is a vital link."

Trust as a foundation, not an add-on

NexTArc's name itself signals its priorities: *trustworthy embedded AI*. But what does 'trustworthy' really mean in practice? "Trustworthiness is more than just cybersecurity," Baran explains. "It's a combination of security, safety, privacy, reliability and accountability." He gives a simple example: "When a city's traffic control system uses AI to optimise flow, citizens and operators must trust that it won't be hacked, that it respects privacy and that it behaves predictably even in extreme conditions."

Building that trust means aligning with European regulation from the outset, including the AI Act, Cybersecurity Act and Chips

Baran Cürüklü

Project at a Glance

NAME:
NexTArc - Next Generation Open Innovations in Trustworthy Embedded AI Architectures for Smart Cities, Mobility and Logistics

DURATION:
2025–2028

BUDGET:
€27.2 million (EU contribution: €8.55 million)

COORDINATOR:
Mälardalen University (Sweden)

PARTNERS:
38 across 10 countries

DOMAINS:
Urban infrastructure, industry, mobility

Act. “These aren’t just legal checkboxes,” he adds. “They’re the foundation for public trust in technology.” Trust, however, isn’t just technical. It’s human. Baran recalls a national project in Stockholm collecting air quality data using GPS trackers. “We asked skateboard kids to carry sensors, but when we said we needed GPS, everyone said no. They just didn’t trust what would happen with their data.” That story illustrates why NexTArc involves architects, city planners and municipalities in its work from the start. “Engineers can say ‘this won’t mess with your privacy,’ but if people don’t believe it, that doesn’t matter,” he says.

The four innovation modules

To achieve its ambitious goals, NexTArc is structured around four Innovation Modules, each tackling a different layer of the technology stack:

1. Cyber-resilience on chip; designing secure-by-design hardware capable of detecting and recovering from attacks.
2. Low-power embedded AI:developing new architectures and accelerators to deliver high performance at lower energy costs.
3. Computation and dependability: ensuring robustness and predictability in real-world environments.
4. Holistic solution stack: integrating hardware, software, and connectivity into a cohesive, open architecture.

Within these modules are 15 key innovations, carefully structured like “atoms in a molecule.” Each subproject is designed to be modular: it can stand on its own or integrate into larger systems.

Testing in real environments

Unlike many research projects, NexTArc isn’t limited to the lab. Its solutions will be tested in three real-world domains:

- Smart, Sustainable and Liveable Neighbourhoods (SLN) – optimising energy use, monitoring air quality and making urban infrastructure more responsive.
- Smart, Sustainable and Transparent Industrial Spaces (STI) – bringing transparency and traceability to manufacturing and logistics.
- Trustworthy and Eco-friendly Multimodal Mobility (TEM) – reducing emissions, improving safety and increasing reliability in transport systems.

“These pilots will prove the technologies don’t just work in theory but in practice, at scale and under real constraints.”

Openness as strategy

Another cornerstone of NexTArc is its commitment to openness. Many existing AI and IoT solutions are closed and proprietary, locking users into specific vendors. “That makes it hard for cities or industries to integrate systems from different suppliers or adapt when regulations change,” Cürüklü explains. By contrast, NexTArc is building open architectures and standard-compliant interfaces that support flexibility and competition. This aligns with Europe’s push for digital sovereignty. “Europe doesn’t have companies like Amazon or Google,” he notes. “We have to be open and adaptive.” Initially, the project focused on RISC-V. But soon the team decided not to tie itself to a single technology. “It doesn’t make sense to limit the solution. We opened up to all kinds of open solutions.”

Energy efficiency at the edge

Edge AI can be powerful, but also energy-hungry. “If you scale today’s architectures to millions of sensors and edge devices in a smart city, the energy footprint is enormous,” Baran explains. NexTArc aims to cut energy consumption by 30% and boost data transmission efficiency by 40% compared to today’s systems. “Cities and industries want technology that’s both smart and sustainable,” he adds. “Otherwise, it’s simply not deployable at scale.”

Creating opportunities for industry

For companies, NexTArc is designed to be more than a research exercise. It’s a platform for practical business opportunities. Many of the SMEs in the project have advanced technologies but limited international reach. “They know what they can do in Sweden,” Baran says. “But they don’t know what they can do in Modena or Zurich.” By integrating their work into a larger, open ecosystem, NexTArc helps them expand. One earlier partner, a startup specialising in vibration detection, grew through similar collaborations and was eventually acquired by Infineon. “For me, that’s success. It’s about helping companies grow, creating jobs and becoming part of a bigger ecosystem.”

By the end of the project, in 2028, NexTArc aims to demonstrate secure, low-power AI platforms deployed in real cities, factories, and mobility systems. Success means having open standards and architectures ready for adoption, and proving that Europe can lead in building not just advanced but trusted technology. “For academic partners, success might be publications,” Baran says. “But for me, it’s about seeing companies grow, jobs created and technologies embedded in real ecosystems.”

Baran’s message to European industry is clear: “NexTArc isn’t just another R&D project. It’s a strategic opportunity to shape the future of embedded AI on European terms, aligned with our values and regulatory framework.” He encourages companies, from startups to major OEMs, to engage early. “We’re building trust into the technology itself. That’s something everyone benefits from.”

Chips JU

Co-funded by the European Union

Research Project Highlight

Green innovation in structural health monitoring sensors

A step forward for sustainable electronics with the European EECONE project

Giustina Di Donato

Federico Danzi

In the fast-growing world of digital transformation, sustainability is no longer just a goal, it is a responsibility. Electronics, the beating heart of modern technology, powers every sector from aerospace and defense to telecommunications and mobility. Yet, as innovation accelerates, so does one of the industry's biggest challenges: the generation and management of electronic waste (e-waste). Reducing this impact is crucial to ensure that digital progress goes hand in hand with environmental protection.

To tackle this challenge, Leonardo is participating in EECONE (European ECOsystem for greeN Electronics), a European project co-funded by the European Union under the Horizon Europe - Key Digital Technologies (KDT) Joint Undertaking. The project aims to build a sustainable European ecosystem for green electronics, rethinking materials, production processes, and end-of-life strategies to foster a truly circular value chain.

At the heart of the initiative, Leonardo leads the development of an innovative layer of sensors for Structural Health Monitoring (SHM) of composite materials used in aeronautical applications. This activity is part of the company's broader strategy to embed sustainability into every stage of the product life cycle, from concept design to disposal. The goal is to create a new generation of embedded sensors capable of monitoring the structural integrity of aircraft components throughout their service life. By continuously collecting and analyzing data on strain, deformation, and wear, the system will enable predictive maintenance, reduce downtime, and extend the operational life of aeronautical structures.

In the fast-growing field of aerospace innovation, Leonardo's distinctive approach lies in combining technological excellence with environmental responsibility. The company's researchers are designing eco-friendly sensor systems by exploring alternative materials and sustainable adhesives that can replace conventional, non-recyclable options. These materials are selected to minimize the use of critical raw resources and to ensure easier recovery and recycling at the end of their life cycle.

The activity, led by Leonardo in close cooperation with the University of Perugia and the spin-off LunaGeber, leverages the synergy between industrial expertise and academic research. Together, the partners are testing next-generation "green" sensors that can be integrated directly into composite structures, offering both high precision and low environmental footprint.

Throughout the project development, new collaborations have also been established, notably with the Danish Technological Institute, experts in chemistry and printed electronics, and Synano, who developed graphene-based inks suitable for strain sensor printing applications. Furthermore. In addition, dedicated activities were initiated with Dassault Systèmes to carry out the Life Cycle Assessment (LCA) of the new proposed sensing technologies, aiming to compare their environmental impact. By focusing on eco-design and material optimization, the team aims to significantly reduce waste and energy consumption during production and operation. The research also explores bio-based adhesives and alternative conductive materials that can deliver reliable sensing performance while lowering toxicity and environmental persistence. This work forms part of Leonardo's longterm commitment to responsible innovation, aligning with the European Green Deal and supporting key EU initiatives such as the Circular Economy Action Plan, the European Chips Act, and the Critical Raw Materials Act.

These frameworks are essential to ensure that Europe's technological development remains both sustainable and strategically independent. In the fast-evolving European technology ecosystem, Leonardo's involvement in EECONE demonstrates how sustainability and innovation can coexist and reinforce each other. The company's expertise in smart materials and sensor integration allows it to pioneer solutions that not only enhance aircraft safety and reliability but also reduce the environmental footprint of advanced electronic systems. The Health Monitoring System (HMS) under development will bring tangible benefits to the aerospace sector by reducing the need for redundant sensors and complex maintenance operations. This approach extends component life cycles and helps minimize e-waste, contributing directly to the EU's circular economy objectives. Beyond the technical impact, the project also has a strategic dimension: it encourages collaboration between academia, industry, and start-ups to build a resilient and sustainable European supply chain for electronics. The partnership model promoted by EECONE serves as a blueprint for future cross-sector cooperation in the green technology landscape. Ultimately, the initiative reflects Leonardo's belief that innovation must be sustainable by design. By embedding environmental considerations into engineering decisions, the company is helping to shape a future where high-performance technology and respect for the planet go hand in hand. With EECONE, Leonardo reaffirms its leadership in sustainable technological innovation, demonstrating that the green transition of electronics is not just an aspiration, it is already taking shape, today.

Co-funded by the European Union

INSIDE Event

INSIDE Connect 2025

Reuniting the community in Málaga

INSIDE team

On 3 and 4 September 2025, the INSIDE community gathered at the Barceló Hotel in Málaga for INSIDE Connect 2025. Organised by INSIDE Industry Association, this was the first in-person edition after a long pause, and the return to face-to-face meetings brought an unmistakable sense of energy and anticipation.

For two days, industry leaders, SMEs, researchers, policymakers, and innovators shared insights, developed new ideas, and explored how collaboration can continue to drive Europe's digital and industrial ambitions. The event combined key notes and panel discussions, thematic workshops, and networking opportunities, making it a celebration of reconnection.

A strong opening

The opening sessions turned the spotlight on Málaga's role in digital transformation, highlighting how local initiatives, national programmes, and European strategies are aligning to strengthen innovation ecosystems.

These contributions made clear that innovation depends and builds on strong partnerships between municipalities, ministries, research organisations, and industry. By showing how cooperation at different levels can accelerate growth, the opening day provided a powerful reminder of the importance of building bridges across sectors.

The programme began with an inspiring series of opening sessions that set the stage for two days of collaboration, innovation, and forward-looking discussions. From the outset, the mood in the room was one of optimism and engagement.

This part of INSIDE Connect 2025 turned the spotlight on Málaga's growing role as a hub for digital innovation and collaboration. A representative from Andalusia's government and from Málaga's Municipality illustrated the city's strategic approach to digital transformation and its commitment to creating an inclusive innovation ecosystem.

At the national level, the Spanish Ministry for Digital Transformation (SETT) broadened the perspective, outlining Spain's digital strategy and the policies driving technological progress and competitiveness.

This was followed by a focus on collaborative innovation in industry, showcasing how cooperation among companies, research actors, and institutions can deliver tangible impact. Fundación Innova HRV illustrated how industrial collaboration can drive innovation and regional development, while Imec presented its expanding presence in Málaga, including the establishment of a new research and innovation facility that will strengthen local capacities and attract international talents.

The session also introduced MicroNanoSpain, the Spanish Chips Competence Centre for Microelectronics, which will act as a national hub for semiconductor design, manufacturing, and training, connecting Spain to the broader European competence network.

Together, these contributions vividly showed that true progress emerges when municipalities, ministries, research organizations, and industry join forces. From local innovation to European-scale ambition, the discussions demonstrated how alignment across levels of governance and expertise can accelerate digital transformation and strengthen Europe's technological strategic autonomy.

Shaping what's next

INSIDE President broadened the debate to the global stage, discussing Europe's position in the evolving geopolitical and industrial landscape. He addressed how Europe can reduce dependency and build technological autonomy in the ECS value chain, navigating a fragmented world shaped by regionalization and strategic competition.

Discussions highlighted the need for European-scale coordination, targeted investments, and strategic alliances to

strengthen resilience along the entire ECS value chain, from semiconductors to applications. Key messages included the importance of reinforcing the European Chips Act framework, developing strategic capabilities in critical technologies, and ensuring Europe's competitiveness through innovation-driven, value-based collaboration.

INSIDE Secretary General provided an overview of INSIDE's forward agenda including the future European Framework Programme and Competitive Fund, and the preparation of the Chips JU Work Programme 2026, which will focus on industrial impact, skills development, and technological leadership in areas such as AI, healthcare, advanced photonics, and power electronics.

These forward-looking reflections positioned INSIDE Connect 2025 as a collective call to action for Europe to shape its digital and industrial destiny through innovation, cooperation, and strategic autonomy.

Voices from the community: keynote insights

The morning concluded with two scientific keynotes offering a forward-looking perspective on research and innovation. The first keynote focused on the edge-to-cloud cognitive continuum by Prof. Mauro Tortonesi, presenting it as a unified, adaptive computing environment spanning from edge devices to the cloud. This approach enables seamless resource allocation, real-time AI capabilities, and new paradigms for distributed intelligence. It was emphasized that the continuum represents not just a technological evolution, but a new foundation for digital ecosystems, combining efficiency, resilience, and sustainability.

The second keynote by Prof. André van Schaik, introduced advances in neuromorphic computing, a field that takes inspiration from the human brain to develop energy-efficient systems for data processing and sensing. Through examples such as the SpiNNaker 2 and DeepSouth platforms, the session illustrated how neuromorphic architectures can drastically reduce power consumption, process information locally, and support emerging applications in robotics, healthcare, and space systems.

Exploring tomorrow's technologies

The afternoon of INSIDE Connect 2025 unfolded with two parallel sessions, each delving into areas at the forefront of technological innovation, from the digital vehicle of the future to the fast-evolving landscape of quantum technologies.

The session on the digital vehicle of the future brought together experts from leading industrial and research organisations to explore how software-defined architectures are transforming mobility. Presentations introduced key initiatives in the context of the European Digital Vehicle of the Future such as Federate CSA, Hal4SDV and Shift2SDV projects, offering complementary perspectives on in-vehicle and cloud-based software platforms. Discussions focused on the creation of shared development ecosystems, the integration of safety and reliability by design, and the challenges of aligning standards and tools across the automotive value chain.

The parallel session on quantum technologies provided a comprehensive overview of ongoing developments and future directions across quantum computing and quantum sensing. The session began with an introduction to the Chips JU Quantum Calls, setting the context for upcoming European initiatives. Experts from academia and industry presented emerging technologies engineering approaches, and applications including quantum computing for optimisation and materials science, quantum sensing for navigation, imaging, and life sciences. The discussions also addressed industry expectations and the path toward scalable, commercially viable quantum systems.

After these sessions, the project pitches and exhibition offered participants an opportunity to showcase research project results, emerging ideas and connect with potential collaborators. For many newcomers and SMEs, this was a unique moment to share concepts, network with established players, and receive first-hand feedback from peers.

Far from being static presentations, the sessions encouraged genuine dialogue across disciplines and sectors. By connecting research, application, and policy perspectives, the afternoon's discussions reinforced INSIDE Connect's mission: to build bridges across Europe's innovation landscape and shape the technologies that will define the decade ahead.

The second day opened with a strong focus on translating research advances into real-world impact. Two parallel sessions offered complementary perspectives on how European innovation is shaping the next generation of intelligent, connected, and automated systems.

The session on Harnessing Edge AI Applications explored how artificial intelligence deployed across the edge-to-cloud continuum is transforming key sectors. Experts shared concrete examples from healthcare, smart homes, agriculture, industry, automotive, and energy, each

demonstrating how distributed intelligence is enabling faster decision-making, privacy-preserving data use, and greater system resilience. The discussions emphasized the growing maturity of the edge AI ecosystem, highlighting the importance of interoperability, low-power computing, and trustworthy AI models that can continuously learn and adapt in real-world environments.

The parallel session on Engineering Automation brought together leaders from academia and industry to discuss how systems engineering is evolving in the age of digital transformation. Topics included model-based systems engineering (MBSE), SysML standards and semantics, digital twinning, and open-source collaboration frameworks. Participants reflected on how automation is reshaping the engineering lifecycle, from design and verification to deployment and maintenance, and how shared, open ecosystems are key to accelerating innovation across Europe's industrial landscape.

The morning continued with a strategic foresight from the Scientific Council. This session offered an expert view on two pivotal dimensions for Europe's digital future. The first presentation explored embedded intelligence, focusing on federated and continuous learning models that enable devices to improve autonomously while safeguarding data privacy. The second addressed the intersection of safety and cybersecurity, underlining the need for holistic approaches that can secure increasingly complex, interconnected systems without compromising performance or usability.

Workshops that inspired

The thematic workshops held on the second day quickly became one of the defining moments of INSIDE Connect 2025. Focusing on mobility, the edge-to-cloud cognitive continuum, engineering automation, and cybersecurity, alongside additional themes proposed by the community, these sessions provided a space for in-depth dialogue on some of the most pressing and promising areas for Europe's ECS sector.

Each workshop combined technical depth with a strategic outlook, encouraging participants to share insights, identify common challenges, and explore collaborative opportunities. Discussions ranged from how connected mobility and distributed intelligence are reshaping industries, to how advances in healthcare technologies and cybersecurity are vital to Europe's digital sovereignty.

The interactive format and open atmosphere were particularly appreciated, fostering genuine exchanges that often continued well beyond the formal sessions. Many participants noted that these workshops not only deepened mutual understanding but also sparked new project ideas and partnerships, reinforcing INSIDE Connect's role as a catalyst for cooperation and innovation.

Feedback was overwhelmingly positive, with attendees expressing enthusiasm for dedicating even more time to such collaborative formats in future editions. This strong engagement reflects both the relevance of the topics and the value of having a trusted, open space where the community can shape the agenda for Europe's digital and technological future together.

Networking in every corner

Beyond the formal programme, networking was present in every moment of INSIDE Connect 2025. The event design intentionally wove together plenary sessions, group discussions, and informal breaks, ensuring that participants had multiple opportunities to connect.

Participants repeatedly remarked on the quality of conversations they had during these informal moments, underlining the importance of a balanced programme that combines content with connection.

Experiencing Málaga Together

No less memorable was the opportunity to experience Málaga itself. On the first evening, participants were invited to a guided cultural tour of the historic Alcazaba fortress, one of the city's most iconic landmarks. Walking through its ancient walls offered not only a glimpse into the region's rich history but also a shared moment of discovery.

The evening concluded with a networking cocktail, where professional discussions continued in a relaxed and social setting. This combination of cultural experience and informal networking was widely appreciated, giving INSIDE Connect 2025 a unique personal dimension that participants will remember long after the event.

Looking Forward

As INSIDE Connect 2025 came to a close, the atmosphere was one of satisfaction and renewed purpose. Participants left Málaga with new contacts, fresh project ideas, and a clear sense of belonging to a vibrant European community. The strong feedback confirmed the value of bringing people together in person not just to share knowledge, but to strengthen the trust and relationships that make collaboration possible. At the same time, the event offered useful lessons for the future. There were also calls for stronger links between academia and industry, and for more practical discussions on funding and emerging AI platforms. These insights will guide the planning of future editions, ensuring that INSIDE Connect continues to evolve in step with the needs of the community it serves.

Europe's digital future will be built through collaboration.

INSIDE Industry Association leaves Málaga with renewed momentum, ready to carry forward the lessons, relationships, and ambitions that made INSIDE Connect 2025 such a success.

Frugal AI

Next generation on-device intelligence in a resource constrained world

In recent years, we are experiencing a profound impact of the ground-breaking capabilities of Machine Learning (ML) that has become synonymous with the a priori existence of large scale models, massive data sets and cloud/edge based infrastructures, as means to enable and accelerate these realisations. While these advantages have unlocked unprecedented capabilities, they have also introduced significant limitations that make the overall AI evolution unsustainable, especially in the case of real-world applications, as for example in remote areas, autonomous missions, or defence-critical settings, where ubiquitous and uninterrupted connectivity, unlimited processing power, and vast amount of pre-existing labelled data for training, simply do not exist.

Typical applications involve systems that are required to make centralised (local) decisions, with real time constraints and without the luxury of large bandwidths for offloading computational heavy processes or retrieving pre-trained full or partial models from the cloud or the edge, to speed up the computational needs but also the decision accuracy. In many cases, it is also increasingly discussing the bottleneck of these AI systems from a training cost and the vast needs of energy consumption, especially for large-scale models, that is currently raising ethical and sustainability concerns, while requiring an ethical shift towards more greener and sustainable AI systems. As a response to this issues, Frugal AI could be considered as an alternative solution for realising the next generation on-divice intelligence, a solution that is suitable for a resource constrained world.

What is Frugal AI?

Frugal AI refers to the design and deployment of intelligent systems that operate efficiently under strict constraints, such as computational power, memory, energy, or connectivity. Unlike conventional AI approaches that rely on centralised servers or edge-cloud ecosystems, Frugal AI is engineered to run entirely on-device level, enabling execution independence in cases where network access is unreliable, undesirable, or impossible, when more heavier computational capabilities cannot be applied due to space or thermal requirements, and when memory and data availability are physically constrained to allow for generalisations and massive training. Thus, Frugal AI has been introduced for making intelligent systems development ready for low-power microcontrollers, radiation-hardened processors, and embedded platforms. To implement Frugal AI, the most common approaches found in the related literature include techniques such as: a) model compression and quantisation, c) few-shot learning, and hardware-software co-design[1].

Model compression and quantisation

One of the foundational strategies in Frugal AI is model compression, that is a set of techniques used to create smaller, more efficient AI models that require fewer resources, like data, computational power, and energy, without significantly compromising performance. Techniques such as pruning (removing redundant weights), weight sharing, and low-rank factorisation help shrink models to fit within the memory and compute budgets of embedded devices. Quantisation, which converts floating-point weights to lower-precision formats (e.g., 8-bit integers), further reduces memory footprint and accelerates inference[2]. These approaches are widely used in TinyML deployments to enable real-time AI on devices like Raspberry Pi, Arduino, and Jetson Nano[3].

Few-shot learning

Few-shot learning (FSL) is a branch of supervised machine learning focused on learning with a limited number of samples. Its primary use lies in training algorithms for rare cases when obtaining the labelled data

is hard or impossible (due to its availability, prohibitive costs of manual annotation, etc.). Variations of few-shot learning include the Zero-shot learning, One-shot learning, and N-shot learning. With Zero-shot learning, the algorithm will be able to classify a target to which it had no prior exposure, including during the training process. FSL approaches are classified into three categories based on utilising prior knowledge on Data (augmenting training dataset), Model (constrain hypothesis space) and Algorithm (alter search strategy in hypothesis space)[4]. These methods reduce the dependency on centralised data pipelines and enable learning directly on-device, often in real time. However, still the shortcomings of the FSL include sensitivity to: 1) data distribution shifts, when the distribution of data differs from the examples being trained on, 2) risk of overfitting with a limited amount of training samples, 3) limited robustness against noisy data sources, 4) scalability issues do to computational constraints, and 5) evaluation gaps since the existing benchmarks might not be accurately enough or relevant for these type of trainings.

Hardware aware co-design

In many fields, there's a growing shift from cloud-based intelligence to edge computing. However, bringing AI capabilities to the edge poses significant challenges due to strict limitations on resources like power, memory, and processing capacity. While high-performance AI solutions already exist and deliver impressive results, adapting these approaches for embedded systems remains a major hurdle. Despite decades of research into effective hardware/software co-design methods, such practices are still not widely adopted in the development of complex embedded systems, including those that integrate AI. Particularly in the case of Frugal AI the algorithms should be tailored to the specific capabilities of the target hardware. Some of the current research directions in the field are trying to be optimised for parallelism on FPGAs, leveraging neuromorphic architectures for event-driven processing, and using in-memory computing to reduce data movement[5]. These designs are critical in domains like space and defence, where power budgets are tight and hardware must be resilient to environmental stressors.

Conclusions

Frugal AI represents a powerful and necessary solution for the next generation AI enabled embedded systems since it represents a sustainable solution for real word problems, where traditional AI approaches are strugglinging, as for example in the case of robotic applications, satellites, mining, missions after dissasters, and general network connectionless environments. The overall concept of the Frugal AI approach is not only technically promising but also ethically sound, while as Frugal AI continues to grow, we can look forward to more collaborations with neuromorphic computing and adaptable learning methods, novel on chip solutions, as well as novel algorithm foundatins targeting smaller models, less energy consumption and more focused reasoning e.g. based on advanced situation awereness and mission adaptability. Additional standardised benchmarks, especially for Frugal AI and more open-source tools to support this new community, will help strengthing the overall impact and allowing more research and development activities in the field. Frugal AI might be evolved in the near future into a powerful tool, not only for just shrinking models, but for also expanding and spreading intelligence to every corner of the world, regardless of infrastructure.

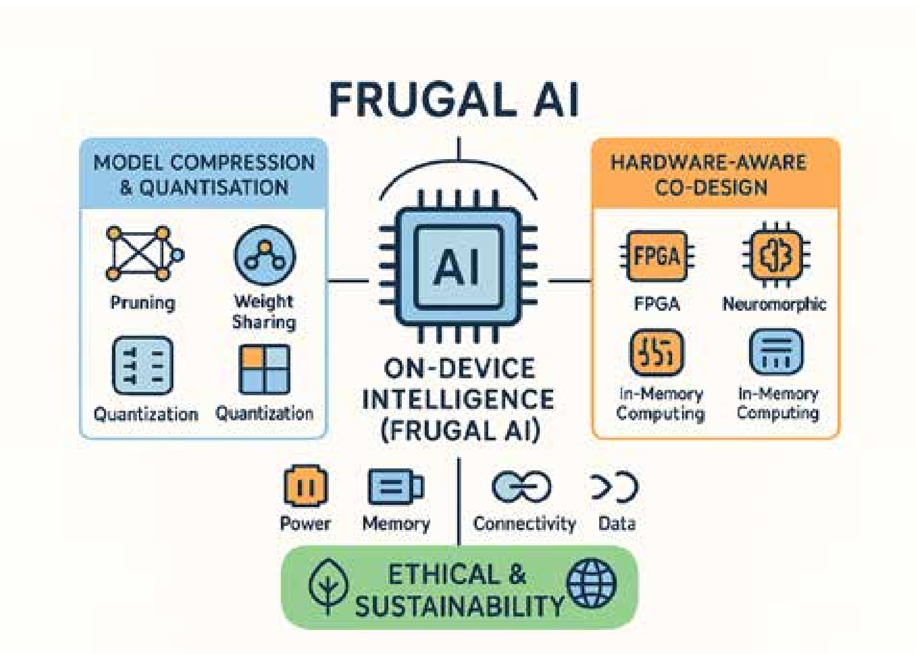

1 Nancy Girdhar, Aditya Raj, Deepak Sharma, Vinay Singh, Antoine Doucet & Matthias Renz, *"A comprehensive review of frugal artificial intelligence: challenges, applications, and the road to sustainable AI"*, journal of *Soft Computing*, Volume 29, Pages 4823–4856, August 2025.

2 Cheng, Y., Wang, D., Zhou, P. & Zhang, T., *"A Survey of Model Compression and Acceleration for Deep Neural Networks"*. IEEE Signal Processing Magazine, 35(1), pp.126–136, 2018.

3 Alam, M., Elgamri, A., Ramzan, N., & Imran, M., *"A comprehensive survey on TinyML. IEEE Access"*, 11, pp.96892–96922, 2023.

4 Wang, Y., Yao, Q., Kwok, J.T. & Ni, L.M. (2021), *"Generalizing from a Few Examples: A Survey on Few-shot Learning", ACM Computing Surveys*, 53(3), pp.63:1–63:34.

5 Zhao, D., Park, J.H., & others. (2024), *"A Survey on Hardware-Aware Co-Design for Efficient AI Systems", IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems.*

New Member Focus

Innovation River

A regional tech powerhouse turning advanced research into industrial solutions

In a landscape often dominated by large integrators and urban tech hubs, Innovation River stands out for a simple but ambitious mission: to make advanced digital technologies usable and practical for industry, especially in places that aren't the usual innovation hotspots. Based in the central Italian town of Rieti, Innovation River is a young company with a clear value proposition for manufacturers, utilities and public authorities seeking to modernise operations. Founded in 2020, it operates at the intersection of applied research, custom software engineering and industry training. The company's name reflects its philosophy: a 'river' of innovation, carrying high-tech capabilities from research labs and EU and national projects directly into industrial and civic applications.

"We chose the name *Innovation River* because we wanted to reflect both clarity and movement," explains Diego Grimani, Chief Operating Officer and co-founder. "The river's water is clear, it moves forward calmly but with purpose. That's how we see innovation."

From research to real-world deployment

Unlike many regional IT consultancies, Innovation River positions itself as a research-driven partner. Its portfolio spans cybersecurity frameworks for connected assets to digital twin architectures for predictive maintenance and transparency, wireless sensor networks (WSN) for industrial monitoring, and quantum-inspired computing methods for optimisation and secure communications. The company leverages these technologies to deliver innovative, data-driven solutions that enhance efficiency, resilience and sustainability across industrial and digital ecosystems. Grimani: "We develop advanced hybrid encryption systems and innovative authentication schemes to ensure maximum data and communications security, protecting sensitive information in every context, even within the most complex devices and networks. Our solutions combine robustness, scalability and ease of integration, adapting seamlessly to both corporate infrastructures and embedded devices with limited resources." Combining leading-edge research with pragmatic software delivery gives the company the agility to support European SMEs and public bodies that want innovation they can actually use. "Many companies are stuck between innovation rhetoric and operational reality," Grimani says. "Our role is to bridge that gap with concrete, usable systems."

Agile structure, industrial mindset

Operating with a compact technical team from its Rieti headquarters, Innovation River has chosen agility over scale. Rather than chasing mass-market contracts, the company thrives on co-creation, partnering directly with clients to design, prototype, validate and deploy custom solutions. Its methods combine design thinking, rapid prototyping and a systems engineering discipline, enabling faster time-to-value for clients modernising legacy assets or connecting distributed sensors in demanding environments. This lean, collaborative model resonates with many European manufacturers, who often lack the in-house R&D capacity to manage Industry 4.0 transformations. For them, a research-savvy

Diego Grimani

Company snapshot

COMPANY:
Innovation River S.r.l.

FOUNDED:
2020

HEADQUARTERS:
Piazza Arduino Angelucci, Rieti, Italy

SPECIALISATIONS:
Wireless sensor networks, cybersecurity, digital twins, quantum computing applications, software development, training, technology transfer

ENGAGEMENT MODEL:
Co-creation, pilot projects and research consortia participation

WEBSITE:
innovationriver.it

SME can be more effective and affordable than a global systems integrator.

A trusted node in the European innovation ecosystem

Since its founding, Innovation River has steadily built its reputation through European R&D participation, contributing to Horizon Europe and other technology and cybersecurity initiatives. The company often serves as technical or work-package leader, delivering software architectures, integrating sensor networks and ensuring that project outputs are robust, secure and deployment-ready.

Two flagship projects, PROACTIF and MATISSE, stand out. In MATISSE, Innovation River serves as Italian consortium coordinator, interfacing with national industry leaders, academia and the Italian Ministry of Industry. “It’s an important recognition for us,” Grimani notes. “It gives us the chance to collaborate directly with major players while proving the value that an agile SME can bring to European-level innovation.” Such participation also strengthens Italy’s position in the broader European technology ecosystem, allowing domestic industries to access new tools and funding opportunities.

Key application areas for industry

Although still young, Innovation River’s activity spans several domains that are vital to modern industry:

- Smart asset monitoring: WSN-based systems for predictive maintenance and energy efficiency.
- Edge cybersecurity: protecting distributed sensor fleets and connected devices.
- Digital twins: lightweight modelling platforms for manufacturing and infrastructure operators.
- Sustainability analytics: tools to monitor environmental performance and compliance.
- Workforce upskilling: Training modules that help industrial teams adopt and manage new technologies.

Rather than selling off-the-shelf products, Innovation River delivers tailored deployments, adapting each solution to the maturity and needs of the client.

Ethics and responsibility at the core

In an era of inflated AI promises and data privacy concerns, the company places strong emphasis on ethical and transparent technology development. Its commitment to security, professional integrity and environmental sustainability is more than just a marketing ploy, it is embedded in its daily work. “We want to be correct every time, with each customer, with each partner,” Grimani says. “Our ethics are as clear as the water in the river.” This principle is reflected in projects like Plastic Soup, an initiative focused on monitoring and managing plastic waste in the sea, an example of technology serving both innovation and sustainability.

Joining INSIDE Industry Association

In 2025, Innovation River formally joined the INSIDE Industry Association, marking what Grimani describes as “the right moment” to expand its collaborative network. “We’ve known INSIDE for years,” he recalls. “Paolo Azzoni often told me our vision fits well with the association, but we waited for the right strategic timing. Joining now lets us connect, contribute and grow.” Through INSIDE, Innovation River hopes to deepen partnerships in digital twin development and quantum cryptography, two areas it sees as crucial for the next decade of industrial innovation. “Quantum cryptography is one of the next big frontiers,” Grimani explains. “INSIDE provides the perfect platform to connect with both large and small enterprises shaping Europe’s technological future.”

A regional model with pan-European relevance

Italy’s innovation landscape is becoming increasingly decentralised, with advanced R&D emerging from smaller cities as well as major metropolitan centres. Innovation River exemplifies this model, leveraging its regional roots for cost efficiency while maintaining European relevance through R&D participation. For European manufacturers in sectors such as energy, water, machinery and food processing, Innovation River is a small, responsive partner with deep technical expertise and access to the latest research developments. “We don’t aim to be a giant platform provider,” Grimani says. “Our strength is helping clients adopt new technologies step by step, responsibly, securely and cost-effectively. And each project helps us grow. Like a river, we keep flowing forward.”

IEEE ISC2 - ENSURE Workshop

The University of Patras in Greece provided the backdrop for the ENSURE (Embedding iNtelligence in Smart mUnicipalities – challenges and oppoRtunitiEs) workshop held on October 6th, 2025, as part of the 11th IEEE International Smart Cities Conference.

Smart municipalities take centre stage

The workshop connected researchers from academia, engineers, and innovators from around the world to discuss how technologies like Artificial Intelligence (AI), Cyber-Physical Systems of Systems (CPSoS), IoT, embedded intelligence and the edge to cloud continuum are reshaping the digital foundations of urban life.

The ENSURE workshop revolved around balancing innovation with ethics and compliance to embed intelligence into cities responsibly. Discussions focused not only on the technologies but also on the societal and the regulatory frameworks from the European AI Act and Data Governance Act to the Cyber Resilience Act.

Dimitrios Serpanos

Markus Tauber

The event's holistic approach stood out: ENSURE connected technical excellence from different fields with policy insights and long-term sustainability.

Keynotes: setting the stage

The workshop opened with a keynote by INSIDE Secretary General Paolo Azzoni, who emphasized the strategic role of smart cities as real-world ecosystems where the vision of a digital society becomes concrete. He argued that smart cities are not simply one application area among many, but horizontal integration domains that connect the entire spectrum of the ECS-SRIA priorities, from energy and mobility to health, agrifood, industry, and digital governance. "Smart cities," he noted, "are the spaces where technology, values, and citizens meet."

He presented smart cities as both a testbed and a deployment ground for Europe's ambitions in trustworthy, inclusive, and sustainable digital transformation. As such, they are uniquely positioned to address the cross-domain challenges that define Chapter 3.6 of the ECS-SRIA, including digital inclusion, collective safety, and societal resilience. The keynote closed with a call to action: to think beyond individual technologies and embrace the interdisciplinary and participatory nature of smart city innovation, co-creating digital societies that are not just smart, but human-centred and future-ready.

The second keynote[1] presented by Dr. Amin Anjomshoaa, a senior researcher at the Institute for Data, Process, and Knowledge Management at the Vienna University of Economics and Business (WU Vienna) focused on the disruptive impact of AI in smart city applications, with particular emphasis on neuro-symbolic AI, an emerging paradigm that integrates the generative power of large language models (LLMs) with the structure, logic, and interoperability of symbolic AI approaches, such as Knowledge Graphs. The talk introduced the principles behind neuro-symbolic AI, explained why combining machine learning with symbolic reasoning is essential for complex urban systems, and highlighted recent practical implementations. It also discussed emerging research directions and future opportunities, illustrating how this hybrid AI approach can drive the next generation of smart city innovations.

Securing the IoT ecosystem

The second session of the workshop began with a discussion on the importance of secure communications in smart cities, especially considering the billions of connected devices. The first paper[2] explored both encryption and data optimization techniques for communication, employing OMNet++ simulations, revealing that the combination of data aggregation and compression can lead to significant reductions in latency, energy use and data size. These findings showcase the importance of lightweight, adaptive cryptography for resilient IoT networks.

Innovation in urban living labs

Then, the focus shifted on the challenges related to assessing the maturity of socio-technical innovation[3]. A framework that combines Technological, Societal, Organizational and Readiness Levels was introduced, providing a clear roadmap for evaluating real-world impact. The framework was successfully applied in Bogota's Fenicia Urban Living Lab.

Precision agriculture and environmental intelligence

Closing the session, a cyber-physical monitoring framework for aquaponics[4] was showcased. The proposed solution merged sensor fusion and computer vision with forecasting models to deploy a system capable of producing major gains in early anomaly detection and fish behavior tracking. The results confirmed AI's potential in transforming food systems into self-sustaining ecosystems.

AI for food security and resilience

In the third session, the focus on embedding AI in agricultural was further developed, highlighting how smart cities and municipalities can enhance local food production by employing data intelligence and automation. The position paper[5] proposed the use of a Federated Learning (FL) framework that enables aquaponics farms to securely share intelligence and insights. By combining FL, Edge AI and Generative AI, this approach offers a distributed yet private model for optimizing farm operations.

Sustainable mobility through digital twins

Urban mobility was next, with a paper[6] reporting the design and development of a simulation framework for bus services employing a Digital Twin, to optimize bus dispatching and energy efficiency in real-time. Deployed on the TPG Flex on-demand bus service, the results were impressive: the average waiting time was cut significantly while total energy consumption also dropped. The authors confirmed plans for more deployments in broader areas and additional results in the future.

Lessons from Valencia: circular resilience

Next, climate resilience took the spotlight, shifting the focus on the case study of the Isolated High-Level Depression (IHD) that hit Valencia in 2024[7]: the presentation showed how circular design and digital tools such as Digital Twins, predictive models and IoT for early-warning systems can assist cities in anticipating and recovering faster from extreme weather conditions. The Valencia example demonstrated that the cities of tomorrow must be both data-smart and nature-smart.

Conversational AI in agriculture

Going back to agriculture, the extension of the classical MAPE-K architecture with a conversational Generative AI interface was presented[8]. Human operators can interact in natural language with real-time monitoring systems. By combining real-time sensor inputs with historical data, the framework provides intelligent irrigation planning and automated anomaly detection and alerts. A prototype was also implemented, running on an edge device without requiring cloud connectivity and demonstrating the effectiveness of the system in real-world conditions. MAPE-AI demonstrated how human-centric AI can make complex systems accessible, efficient, and trustworthy.

Cybersecurity and urban intelligence

AI remained in the focus also at the closure of the workshop, addressing the rising cybersecurity-related challenges in smart cities[9]. A Generative AI agent, capable of performing end-to-end malware analysis autonomously, was introduced. By leveraging reasoning Large Language Models (LLMs) and a wide variety of tools built on the Model Context Protocol (MCP), the agent demonstrated flexibility, extensibility, and robust performance across various malware detection tasks.

The Road Ahead

As discussions wrapped up, one message stood out: embedding intelligence in municipalities is not just a technical goal but also a societal mission. From cybersecurity, IoT systems and AI-driven agriculture to energy-efficient transport and resilient infrastructures, the ENSURE workshop in ISC2 2025 proved that Europe's smart city vision is built on collaboration, responsibility, and innovation.

1 "Neuro-Symbolic AI for Resilient Smart City Systems.", Amin Anjomshoaa, Vienna University of Economics and Business (WU).

2 "Optimizing Secure Group Communication in IoT Networks for Smart City Applications.", Hafiz Humza Mahmood Ramzan, Laura Belli, Luca Davoli, and Gianluigi Ferrari.

3 "A Multidimensional Framework for Evaluating Socio-Technical Innovation Readiness in Urban Living Labs: Methodology and Application.", Augusto Velasquez Mendez, et all.

4 "AI-Driven Precision Agriculture in the Context of Smart Cities: A Pathway to Food Security and Environmental Resilience.", Yazzed Abdalla, Amin Anjomshoaa, Markus Tauber, and Rafael Kupsa.

5 "SynAI: Supply Security for Smart Municipalities through AI-optimized Aquaponics Farm Automation using Federated Learning, GenAI, and Edge-AI.", Stefan Gindl, Michael Boch, Rafael Kupsa, and Markus Tauber.

6 "Real-Time Routing and Energy Management in Demand-Responsive Transport using Digital Twins.", Kyriaki Pantelidou, Stylianos Kokkas, Antonios Lalas, Konstantinos Votis, and Dimitrios Tzovaras.

7 "Building Urban Resilience through Circular Smart City: Insights from Valencia Flood.", Pedro Núñez-Cacho, Aida Molina-Prados, and Jaroslaw Górecki.

8 "MAPE-AI: A Self-Adaptive, Conversational Framework for Urban Smart Agriculture.", Michael Boch, Stefan Gindl, and Markus Tauber.

9 "MAD-Agent: A Malware Analysis and Detection AI agent.", Georgios Xenos, Emmanouil Tzagakis, Sotirios Giannopoulos, and Dimitrios Serpanos.

New Member Focus

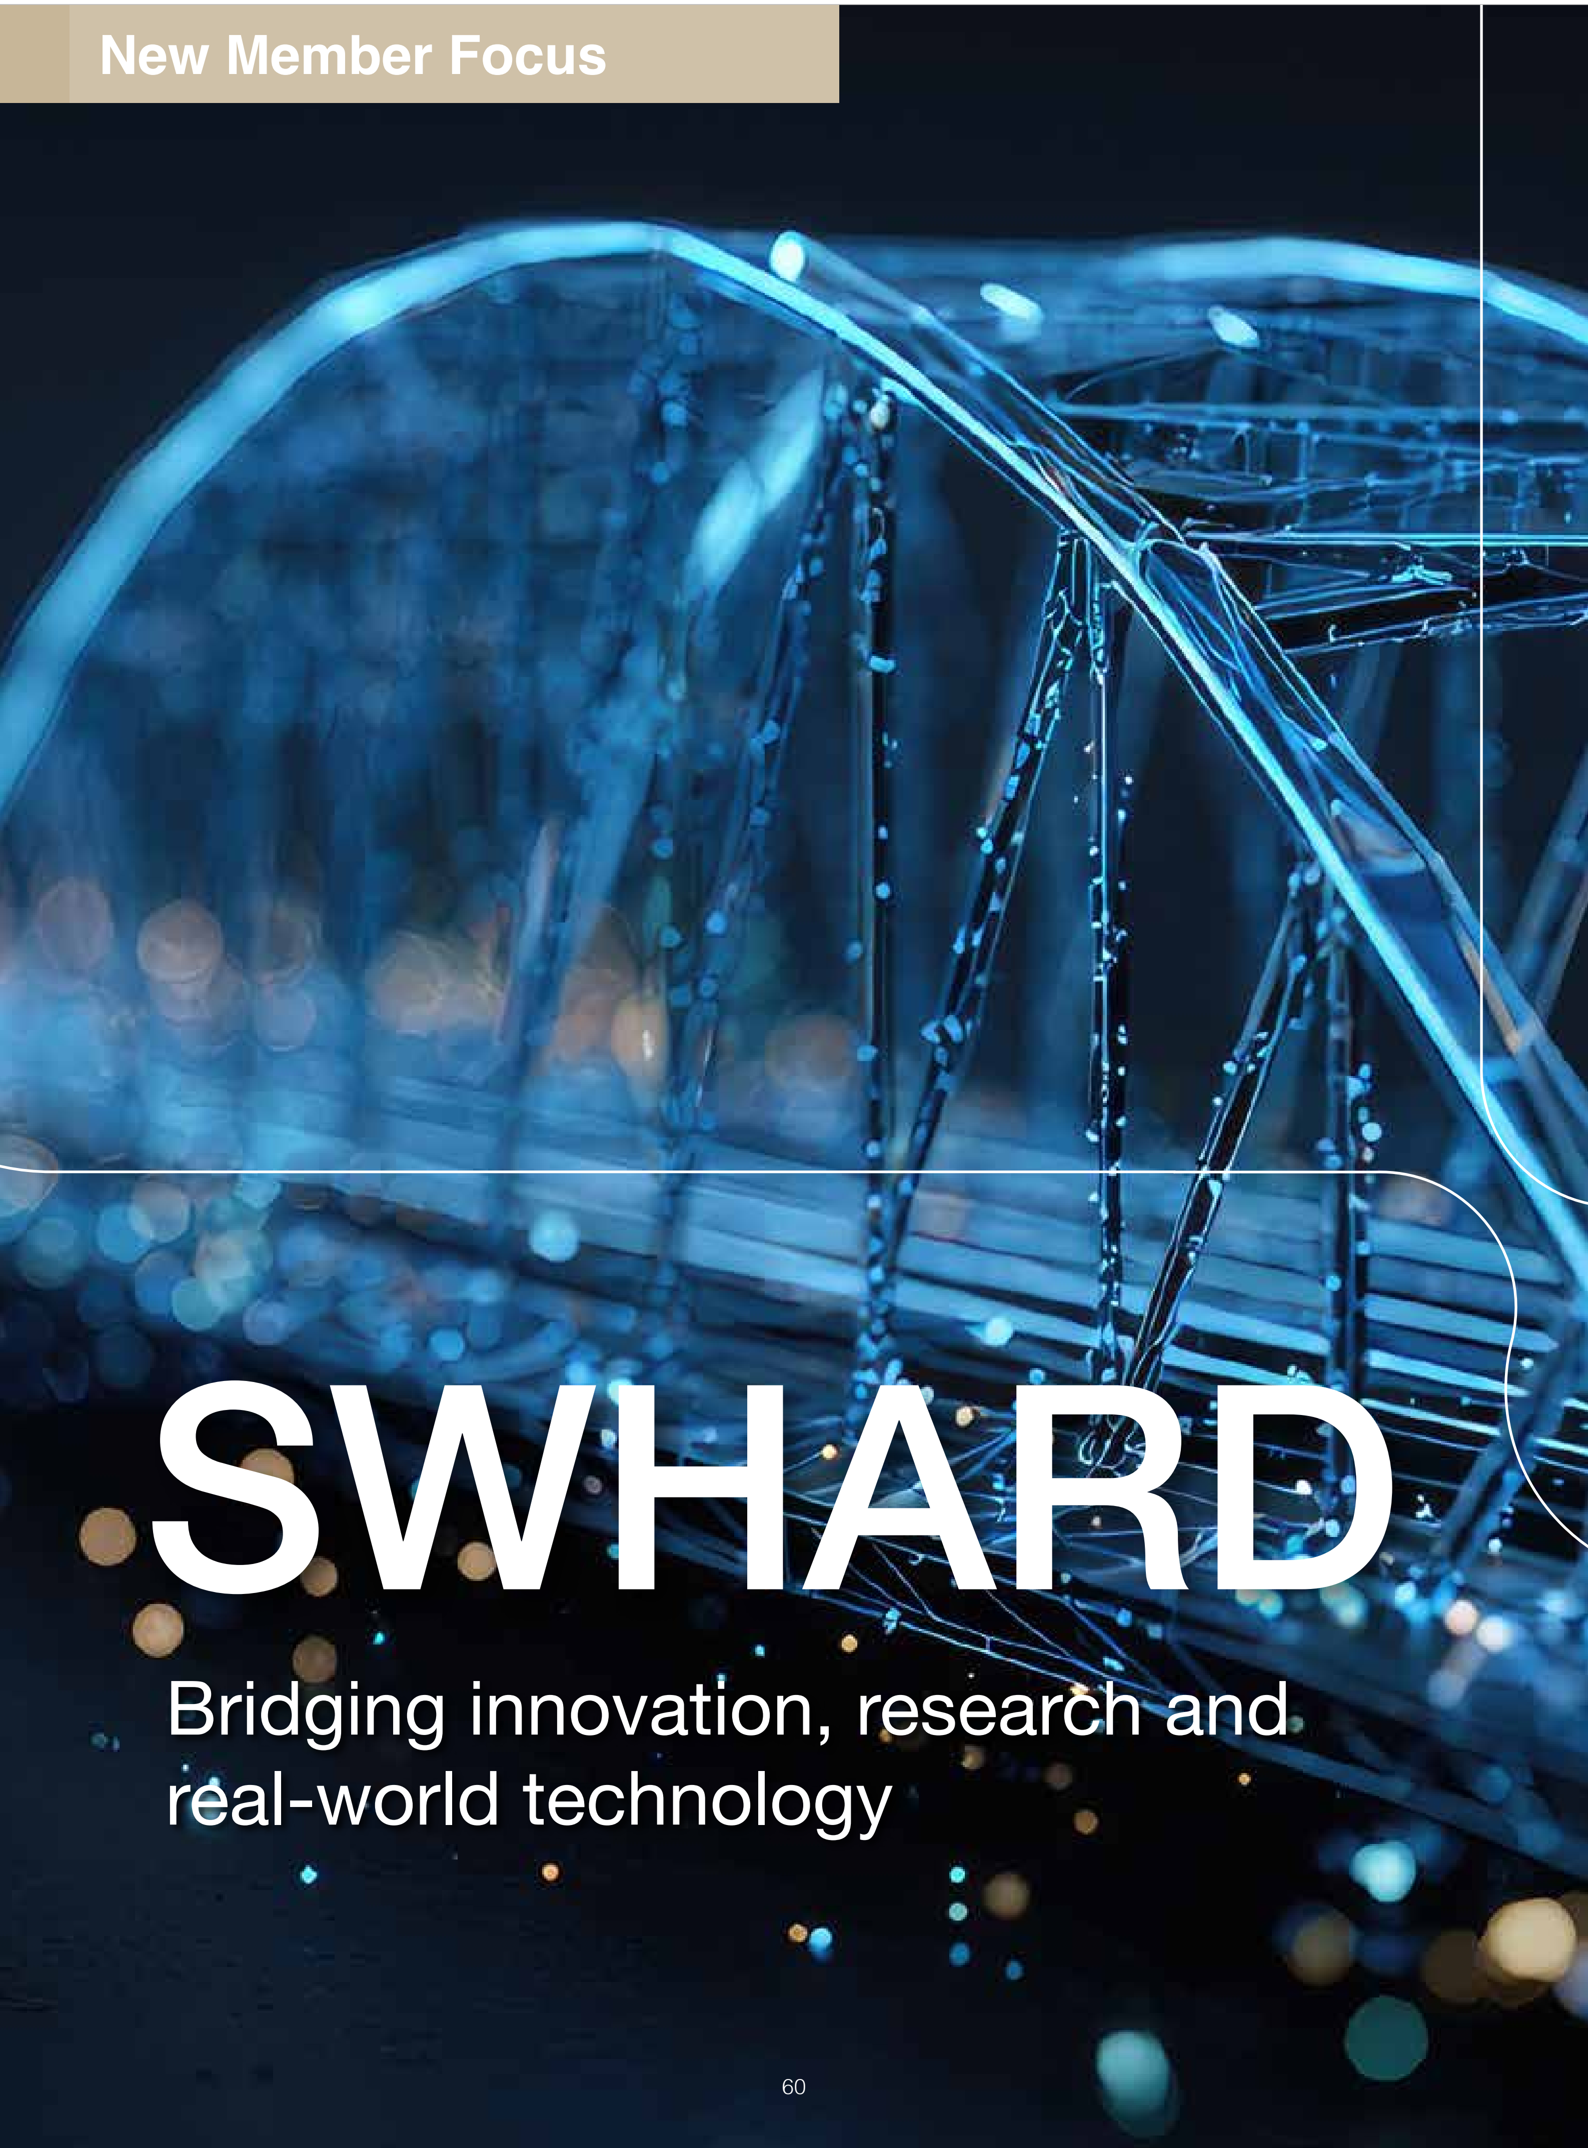

SWHARD

Bridging innovation, research and real-world technology

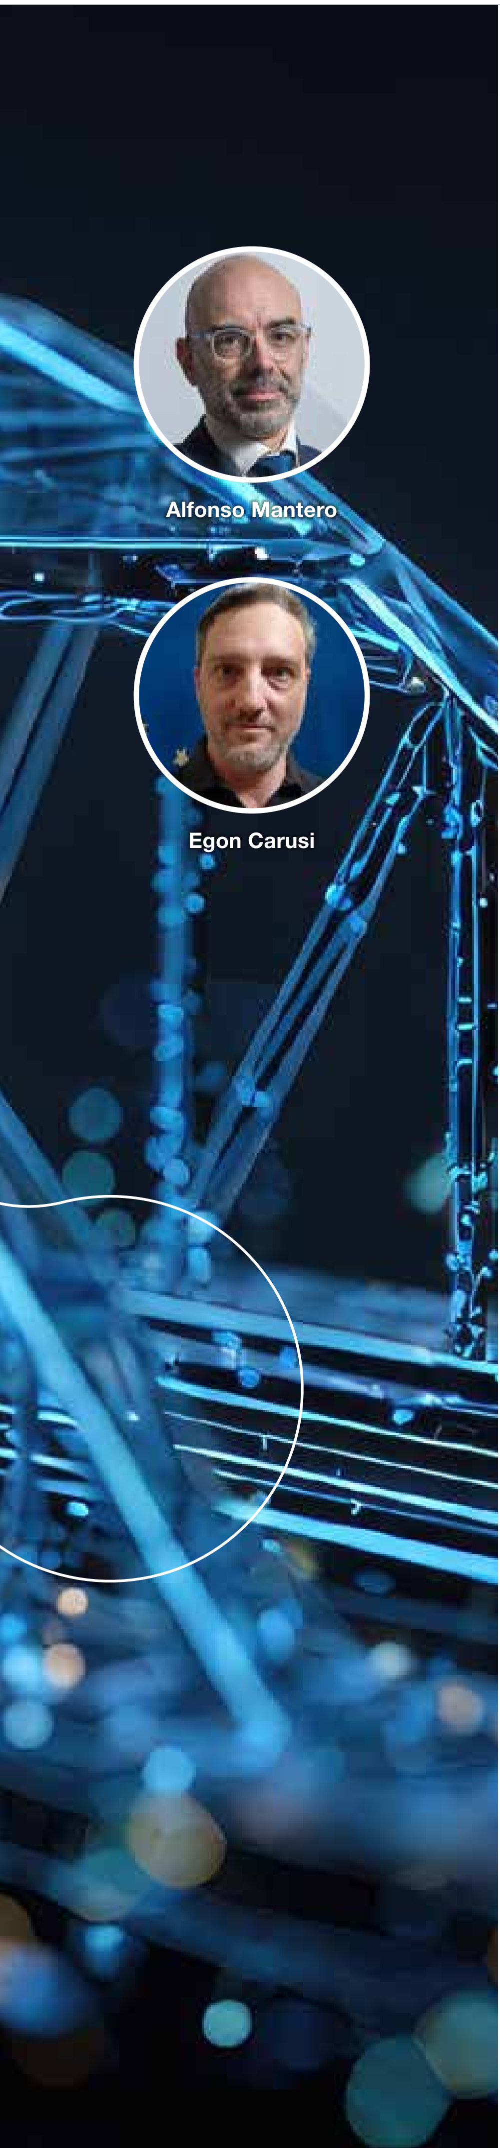

Alfonso Mantero

Egon Carusi

When speaking with Alfonso Mantero, CEO of SWHARD, and Egon Carusi, the company's CTO, one quickly understands why the firm has become such an intriguing new member of the INSIDE community. Based in Italy and expanding across Europe, SWHARD occupies a unique position between research, development and practical engineering — a bridge that connects academic innovation to real-world industrial needs.

An external R&D department for industry

At its core, SWHARD exists to serve companies that either lack a dedicated R&D department or need to supplement it with specialised expertise. "We give our clients what is, in effect, an external research and development department," Alfonso explains. "That means when a company says, *'I have a great business idea but not the technical capability to build it,'* we can step in and make it real."

Operating primarily in the B2B sphere, SWHARD helps companies design, engineer and refine electronic and digital systems. These requests might range from developing entirely new devices to redesigning a single component, whether hardware, software or the complex interface that unites the two. The company's name, SWHARD, in fact, stands for **S**oft**w**are + **Har**dware + **R**&**D**, perfectly capturing its mission.

Beyond contract work, SWHARD is also active in applied research, helping bridge the gap between universities, research institutions and commercial applications. Many of its engineers have strong academic backgrounds, and the company has participated in European and national research initiatives designed to translate theoretical advances into industrial innovation.

Meta-products: tools for future innovation

Alongside its commercial and research work, SWHARD invests heavily in what it calls meta-products, internally developed prototypes, technologies and intellectual property that sit ready for adaptation when new opportunities arise. These are not consumer-ready products but rather flexible building blocks: devices, algorithms and architectures that can quickly evolve into tailored solutions.

"We like to have things ready on our desks," Alfonso says. "When a client comes with a new idea, we don't start from zero, we already have the foundations to build upon." One example is Flex, an intelligent edge-computing device developed in-house. Compact but powerful, Flex acts as both an industrial controller and a data-collection platform for research and predictive maintenance. Deployed in diverse environments from buses monitoring travel noise to water distribution networks detecting early pipeline faults, Flex demonstrates SWHARD's ability to create practical, energy-efficient systems with hardware-enabled artificial intelligence capabilities.

Another example is Sentry, a medical device designed to support post-surgery rehabilitation. It monitors limb movement during recovery and transmits data to medical teams, helping doctors assess patient progress. "We deliver the electronics, firmware and software," says Egon. "Our clients then take care of product positioning and market deployment."

Technology without borders

SWHARD's expertise is intentionally cross-sectoral. "We're not focused on automotive, robotics or biotech alone," Alfonso points out. "We're horizontal. We bring technology to wherever it's needed." That philosophy has led the company to work across industrial automation, oil and gas distribution, environmental monitoring, medical technologies and energy systems.

Geographically, SWHARD's market extends from local Italian companies to international clients throughout Europe. "We don't pretend to be global yet," Alfonso admits candidly. "Our strength today lies in being deeply embedded in the European innovation ecosystem."

Growth through experience
SWHARD's decision to join the INSIDE community came at a pivotal time. When Alfonso and Egon first met Paolo Azzoni several years ago, the company was small, just eight people. "We were not ready then," Alfonso recalls. "We were too young as a company to contribute meaningfully to the community."

In the years since, SWHARD has doubled in size to fifteen employees, gaining maturity, experience and financial stability. Much of that growth occurred during and after the pandemic, when the demand for flexible, intelligent digital systems surged. "COVID was a strange time," Alfonso reflects. "While many sectors slowed, we found ourselves busier than ever, helping clients taking their chance to innovate their digitalisation and automation solutions."

This growth gave SWHARD the confidence to expand its horizons and engage more directly with international networks like Inside. The company participated recently in the INSIDE meeting in Málaga, where it discovered a dynamic ecosystem of partners, researchers and innovators. "It was a great start," Alfonso says. "We met people who were not only talking about ideas but also building real solutions."

Consolidation, not expansion for its own sake
Despite its success, SWHARD has no ambition to become a sprawling corporation. The immediate goal is consolidation, stabilising after rapid growth and ensuring quality and coherence in every project. "We'd rather grow wisely, perhaps to around thirty people, and then let new products spin off into their own dedicated companies when the time is right." That philosophy reflects a pragmatic understanding of the company's role as a bridge and facilitator rather than a product manufacturer. "If we ever develop a commercial product, that would probably become a separate entity with its own trajectory," Alfonso adds.

Technological horizons: RISC-V, AI and space
On the technology front, Egon identifies several key directions shaping SWHARD's future. One is the adoption of RISC-V, the open standard instruction set architecture that is gaining traction as Europe seeks more independence from US and Asian chip technologies. "It's still early," Egon says, "but we want to be ready when energy efficient RISC-V medium-class chips become available for the industry. The marker is already there and almost completely relying on ARM chips."

Another major area is the space and communications domain, which increasingly intersects with artificial intelligence and cybersecurity. "These technologies are converging," Egon explains. "AI demands computing power; RISC-V enables flexible architectures; and both must operate securely. Everything is linked."

Sustainability and energy efficiency
Sustainability is another defining principle for SWHARD. But for Alfonso and Egon, the term goes beyond environmental rhetoric, it's about engineering efficiency. The company designs compact, low-power systems that make the most of every watt consumed. "When people talk about edge computing, they often picture servers sitting in data centres," Egon says. "We're talking about real devices in the field, five-watt units collecting and processing data intelligently." This philosophy applies even in traditional industries such as oil and gas distribution, where equipment is often outdated. SWHARD brings modern, efficient design to sectors that have lagged technologically.

The company is also exploring opportunities in renewable energy. One current project involves modernising small hydroelectric plants across Italy, many of which still operate with manual controls and decades-old systems. "We're not trying to replace people," Alfonso clarifies. "We just want to bring automation and monitoring to make these systems safer and more efficient."

A bridge between research and the market
Perhaps the best way to describe SWHARD's identity is, as Alfonso puts it, "the bridge." The company connects researchers and industries that often speak different languages. "We sit between the hammer and the anvil," Alfonso laughs, "between innovation and application."

Egon agrees: "We see our role as bringing the voice of small and medium-sized companies into the research ecosystem. Many of these businesses don't appear in statistics or big European projects, but they represent the real industrial fabric of Europe."

Advice to fellow innovators
When asked what advice he would offer to other small or medium enterprises considering joining INSIDE, Alfonso's answer is thoughtful: "Before joining, ask yourself what you can bring. Don't come just to take opportunities, come to contribute. What we found in Málaga was a community of serious people, willing to share ideas and value. That's what makes it worthwhile."

Egon adds from the technical side: "INSIDE is not only a business network; it's a place to share use cases and real-world experience. It's about collaboration."

Company snapshot

NAME:
NHeadquarters: Genoa, Italy

SPECIALISATIONS:
High-performance computing, FPGA and embedded systems, edge computing, IoT, AI and machine learning integration, digital signal processing, custom hardware&software co-design

ENGAGEMENT MODEL:
R&D collaborations, technology partnerships, pilot projects, EU research consortia participation

WEBSITE:
www.swhard.com

Online version is available at Inside-association.eu

Publisher
INSIDE Industry Association
High Tech Campus 69-3
5656 AG Eindhoven, The Netherlands

Design and Creative lay-out
Studio Kraft – Veldhoven, the Netherlands

Acknowledgements
With thanks to the interviewees, project participants, INSIDE Industry Association office, the INSIDE Industry Association Presidium and other INSIDE Industry Association-involved persons for any assistance and material provided in the production of this issue of the INSIDE Magazine.

Contributions
The INSIDE Industry Association office is interested in receiving news or events in the field of Intelligent Digital Systems. Please submit your information to info@Inside-association.eu

INSIDE
Industry Association

INSIDE
Industry Association

INSIDE-association.eu