INSIDE magazine

March 2026 — Issue 12

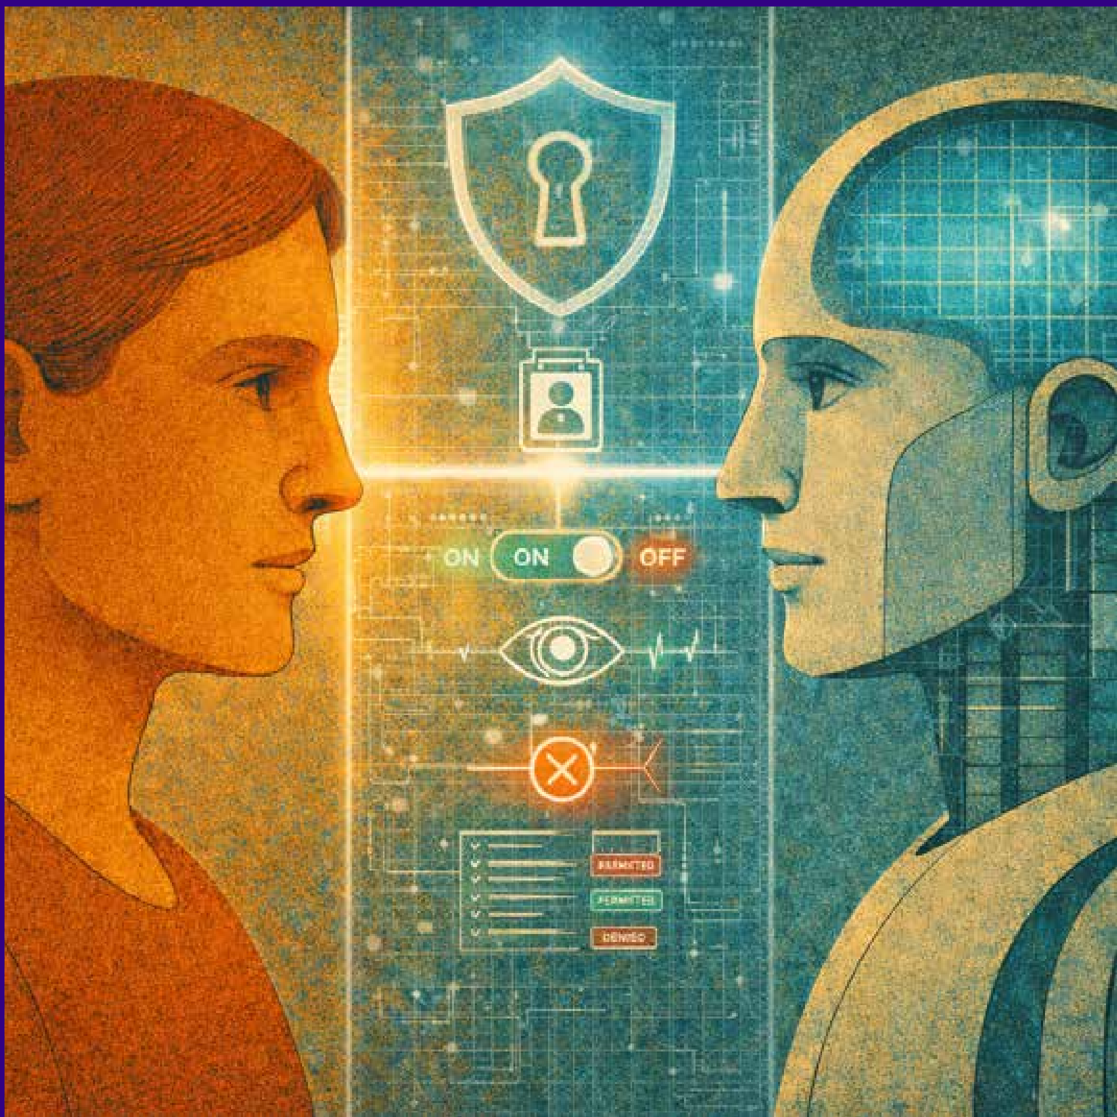

In this issue:

- **Autonomy you can trust**
- **From project to deployment**
- **The strategic imperative of Edge AI systems reference architecture**

INSIDE
Industry Association

In this issue

March 2026 —— Issue 12

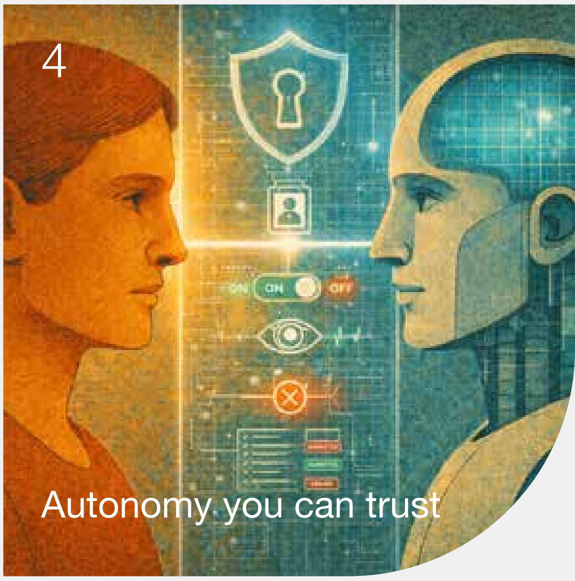

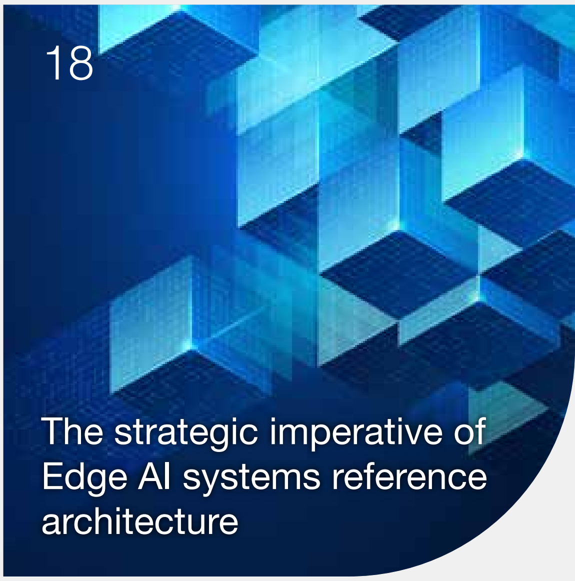

Dear reader,

Europe's digital future will not be shaped by isolated technology breakthroughs, but by our ability to embrace interdisciplinarity across layers, enabling and feeding entire value networks, with innovation generating impact and business opportunities. This issue of INSIDE Magazine captures this transitional time: from uncontrolled acceleration to well structured architecture, from capability to accountability, from creativity to trustworthy innovation.

This issue begins with a look back at autonomous AI agents; a follow-up to our last discussion on interpretative intelligence. In the previous issue, I argued that the future belongs to the wisest AI, rather than the fastest AI. Now, we must confront a deeper issue: how do we ensure increasingly autonomous systems operate according to our human intent? As AI transitions from being an assistant to being an actor, performance alone can no longer be the only consideration. Trust in autonomous systems requires boundaries, oversight and architectural order. Without them, autonomy becomes fragility, while with them it becomes resilience.

The need for structure and responsibility extends beyond just AI. The ambition and scale of the collaborative research ecosystem in Europe continues to broaden, which requires new levels of professional coordination and strategic alignment. Scientific excellence must also be matched by excellence in execution. Organisational architecture - clarity of roles, agility in governance and transparency in processes - is the invisible infrastructure that transforms ideas into impact and business opportunities.

Architecture is, indeed, a unifying theme here. For example, Edge-AI systems are embedded in cars, factories, healthcare platforms, and infrastructure. Therefore, they must be designed with interoperability, dependability, and optimization built in from the beginning. Shared reference models and system-level frameworks are not just technical formalities; rather, they represent the foundation of sovereignty and long-term competitiveness.

At the semiconductor edge, we are also rethinking the architectural level of the physical layer. For example, new non-volatile memory technologies for energy-efficient Edge-AI will reduce latency, minimize energy consumption, and allow persistent intelligence to exist at the edge. This rethinking of architecture at the semiconductor level serves as a reminder that autonomy is ultimately a function of physics; and that Europe has an advantage due to its ability to master both extremes: materials and meaning.

This same convergence is visible in physical AI, where bio-inspired sensing architectures enable the physical world to perceive reality. When intelligence enters the physical world, reliable sensing is what determines the safety and scalability of the entire system. So vision and the ability to use edge-intelligent devices will serve as fundamental enablers to deploy responsible trustworthy systems.

As if this were not enough, many of the advanced components are being converted into solutions that put humans first. The H2TRAIN project demonstrates how sensors, embedded AI and digital twins can help to support assisted living, sport coaching and rehabilitation. It shows how Europe can connect chips to care, algorithms to well-being, and innovation to societal value.

We see our growing community as having this same spirit of integration with the new members presented in this issue having strengths in semiconductors, wearables, system architecture, and research coordination. They collectively believe that complexity is actually an opportunity to work together. Our INSIDE ecosystem is increasing in scale, but also in importance, establishing bridges between disciplines, industries, and countries. And ecosystems are ultimately built by people, who know how to design, implement, and support complex systems. Initiatives like the CPS Summer School intend to train the next generation of system architects to ensure Europe has the necessary resources to convert research excellence into industrial leadership over time.

Whether speaking about AI agents, edge devices, semiconductor technology or physical AI or digital applications, the one constant is that Europe competitive advantage only lies in integration. Those who create the next wave of digital transformation will be not necessarily the fastest innovators in isolation, but rather those who build responsibly, architect strategically, and connect technologies with purpose.

As we move forward into a more intelligent world, intelligence will not be just autonomous, but also aligned, architected and accountable … and it is ours to build.

Paolo Azzoni
Secretary General

Technology Frontiers

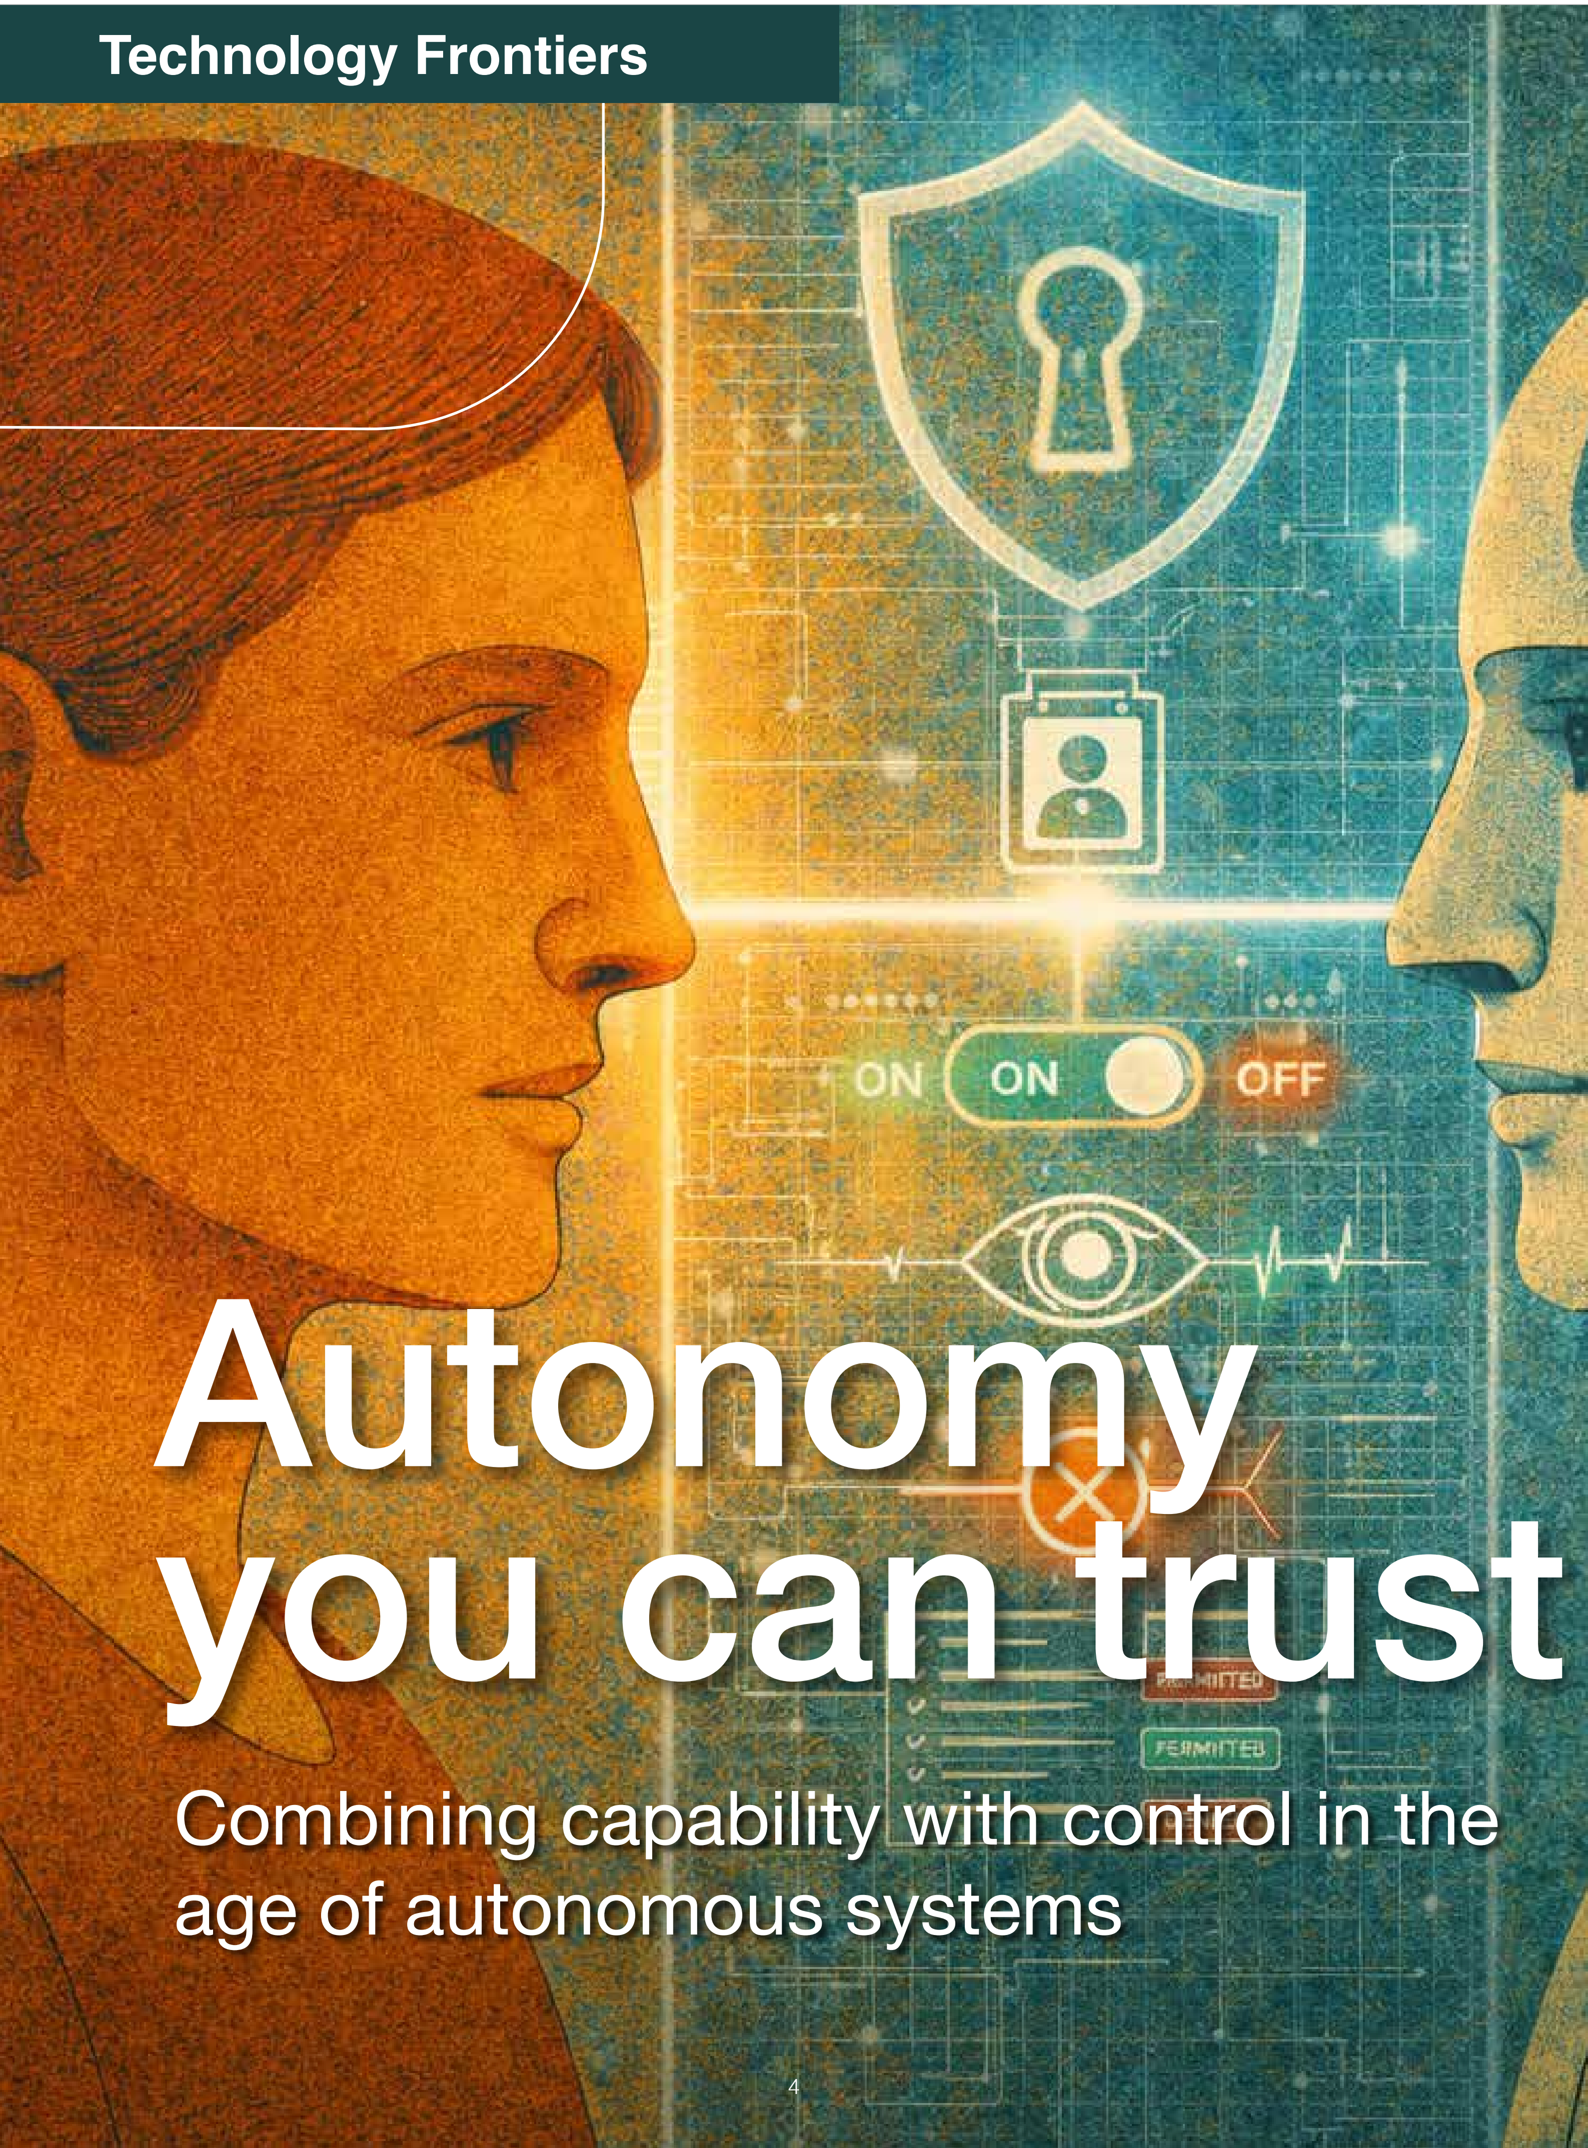

Autonomy you can trust

Combining capability with control in the age of autonomous systems

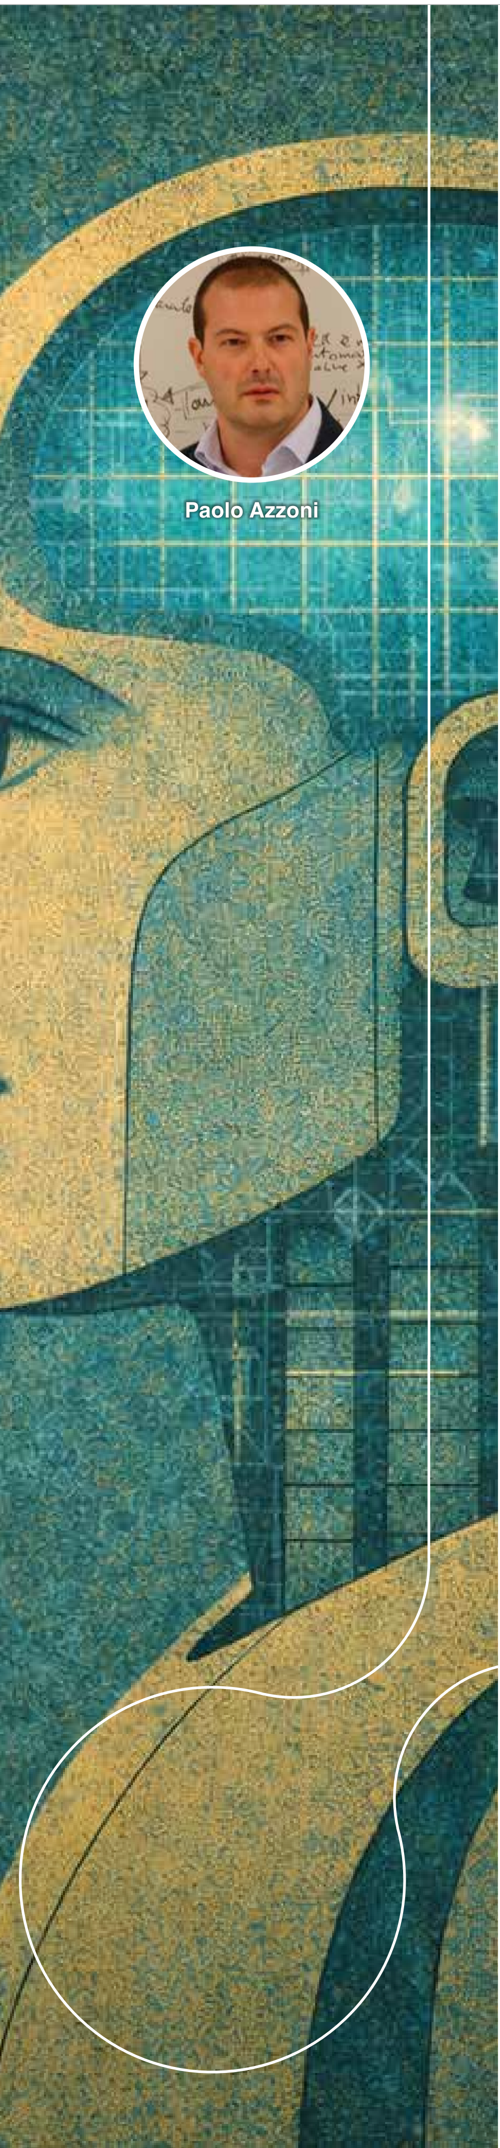

The future of AI is not about speed but about achieving a higher level of knowledge and wisdom: autonomous systems that understand what they are doing, why it is relevant, and how their actions will affect the world around them[1]. However, as interpretative intelligence becomes more sophisticated, there is a new emerging class of autonomous agents that represent the next step beyond the simple concept of assistant. These systems can plan and execute tasks, coordinating tools, and operating within a digital environment without the need for human interaction or supervision. AI is no longer just a user interface; it is now an entity capable of acting independently.

But when systems acquire the capability to act autonomously, this power goes hand-in-hand with maintaining control over those capabilities. Without well-defined guardrails, an autonomous system can create a systemic failure, compromise security and confidentiality, and undermine trustworthiness in digital ecosystems. Therefore, the challenge, is not in slowing down innovation through restrictions or limitations but to rethink and redesign autonomy in a responsible way. If interpretative AI emphasizes that intelligence must be aligned with meaning, the creation of autonomous systems raises a new question: how do we ensure that actions stay aligned with the intentions behind those actions? The answer will determine if autonomous systems emerge as drivers of resilience and productivity, or as merely a vector of risks.

While you sleep

Imagine a common daily routine: you go to sleep, and while you are sleeping, your digital agents are working on your behalf, tweeting, emailing, updating projects, committing code to repositories, negotiating meeting times, and making phone calls with a synthetic version of your voice which sounds exactly like you. They are monitoring dashboards, assessing market trends, and summarizing breaking news while changing priority levels in real-time.

When you wake up each morning, many hours of cognitive and manual labour will have already been performed. You will look at your phone, see a short message from your agents stating: *"The report has been completed; the draft has been sent to be reviewed; two meetings have been rescheduled; we have resolved a supplier problem"*.

This is not science fiction

The purpose of autonomous agents being developed today is to perform persistent, continuous and goal-driven activities. Unlike current assistants who react to commands, the new generation of AI agents operates continuously throughout time, handles multi-step workflows, coordinates work across multiple systems, and initiates actions to achieve intended objectives without waiting for step-by-step instructions. They will deduce the next action, retrieve the tools needed to carry out that action, and implement the decision made in a manner that mimics normal human behaviour in the digital world.

For example, in a work environment an agent could be responsible for handling communication, ensuring compliance, drafting contracts, changing cloud platform configurations, or answering customer service inquiries 24/7.

In personal life, they could manage finances, schedule travels, maintain subscriptions, interact with service providers, or negotiate electronic transactions on your behalf with vendors.

The benefits of similar agents are obvious: increased productivity, reduced mental burden and, more in general, extended human capabilities beyond biological limits.

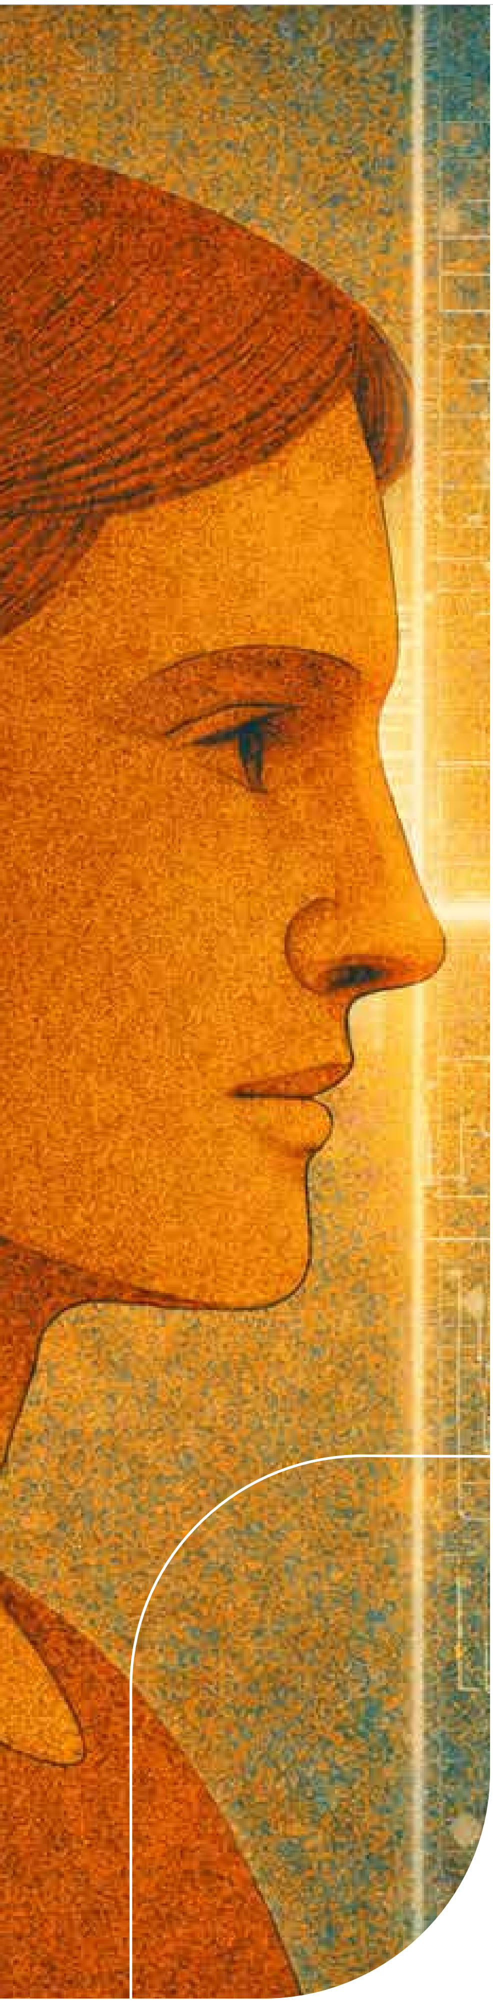

> **AI is no longer merely an assistant, it is acting without supervision on behalf and in place of a person. It is representing that person, which introduces responsibility, exposure and risk for the actions taken by AI[2].**

From assistant to operator

Well, this for example is OpenClaw[3], a new generation of autonomous and persistent AI agents. What does it mean?

To clarify, this agent is not just a simple chatbot waiting for your next prompt; there is no continuous question/answer loop. You provide OpenClaw with a task, such as "*Manage my social media presence over the next week*", and it simply begins the task with a plan. To accomplish the high-level goal, OpenClaw will decompose that goal into a sequence of smaller micro-actions.

If engagement drops, it modifies the adopted strategy. If a new trend emerges, it reacts. It does not simply execute instructions; it iterates toward outcomes and initial objectives.

Considering the example of the social media, it will break down the task by reviewing your past posts, determine what types of post received the highest level of engagement, keep track of popular topics, draft content for posts, create a schedule for posting them, respond to any comment, modify your writing style based on reactions from your followers, and monitor performance metrics. If during this process your engagement starts to drop below normal levels, OpenClaw will change its strategy. If new trending topics arise, it will modify accordingly and react to those trends. It does not simply execute commands but continuously iterates until it reaches a successful outcome.

In terms of technology, this class of agents combines large language models, structured and vector-based memory layers, hierarchical task planners, external tool orchestration, and iterative feedback loops for self-correction and goal tracking. Agents can call an external API, operate user interfaces, maintain intermediate results, evaluate their own progress, and revise their plans dynamically.

Persistence is their key factor, which means continuing to work over time and maintaining context and state, instead of resetting them after each interaction.

OpenClaw is not an interface layer, but rather an operational layer that transforms abstract goals converting them into activities to be executed in digital environments.

> **The change from responding to prompts to pursuing objectives is what makes this generation of agents fundamentally different.**

When autonomy becomes accessible

Many of the available agents released by large technology companies are extremely powerful … but also incredibly expensive. They frequently rely on state-of-the-art proprietary models and are very computing intensive. Additionally, many of these agents can become unstable when run continuously without supervision for a prolonged period of time. Although they present impressive autonomy, they are not persistent and running them for a long period (i.e., days) without human supervision can cause unpredictable loops, reduce performance, or increase the operative cost.

By contrast, OpenClaw seems to be one of the first agents to exhibit the level of stability that allows it to run on a personal computer locally for long periods of time. More importantly, it utilizes free and open-source models, dramatically reducing operational costs. OpenClaw brings the ability for users (individuals, small teams, startups, etc.) to have autonomous agents available to be run locally, rather than being run exclusively in hyperscalers cloud infrastructures. The combination of persistence, stability, and low cost significantly increase technology accessibility and adoption.

Accessibility is exactly what influences the way autonomous agents will be perceived and used in the future: when they become affordable, stable, and easy to use, they will no longer be viewed as an experimental tool of large, high-tech companies owning the necessary technology to build them. They will instead be seen as infrastructure widely distributed, continuously-operating, and potentially ubiquitous.

Where autonomy becomes a risk

"That's so cool, I want it."
Are you really sure?

Because to unlock the full potential of this next generation of autonomous agents, we must grant them something that has never been done before: access. Not just some access, not access in a sandbox, but total freedom to operate as you do. Some examples:

- To manage your social media, an agent needs access to your accounts.
- To answer emails on your behalf, it needs access to your inbox.
- To negotiate a contract, it needs your identity card or passport and your signature.
- To optimize your finances, it needs access to your bank account and interfaces.
- To coordinate work, it needs access to internal systems, repositories, customer databases, calendars, and messaging systems.

In short, to achieve autonomy, the agent must be granted authority. Which means providing credentials, granting permissions, providing tokens, access to memory, and the ability to retain permanent control throughout various digital environments. The agent must see what you see, have the ability to perform actions as you perform them, and represent you in environments that were not previously designed for non-human actors.

That is where the real risks begin.

Autonomy is not the panacea to fix chaos

The real problem with autonomous AI is that it is a brutal accelerator, whatever is the way you adopt it … and even if that way of using is wrong: without direction, speed is just a more efficient way to get lost. Agents optimize, copy, and expand whatever logic you embed into them, but if there is no framework for their existence, they will continue to create errors at an accelerated pace, automate misalignment, and generate disruptive noise.

> **Autonomous agents do not create strategies. They execute them. They operationalize whatever structure, or lack of structure, that already exists.**

There is a great hype on complex workflows, perfect prompts, agents that *"make your life easier"*. On the contrary, autonomy by itself cannot cure confusion and chaos; it intensifies and magnifies it. This means that working with unclear priorities, not having defined processes, and using inconsistent objectives,

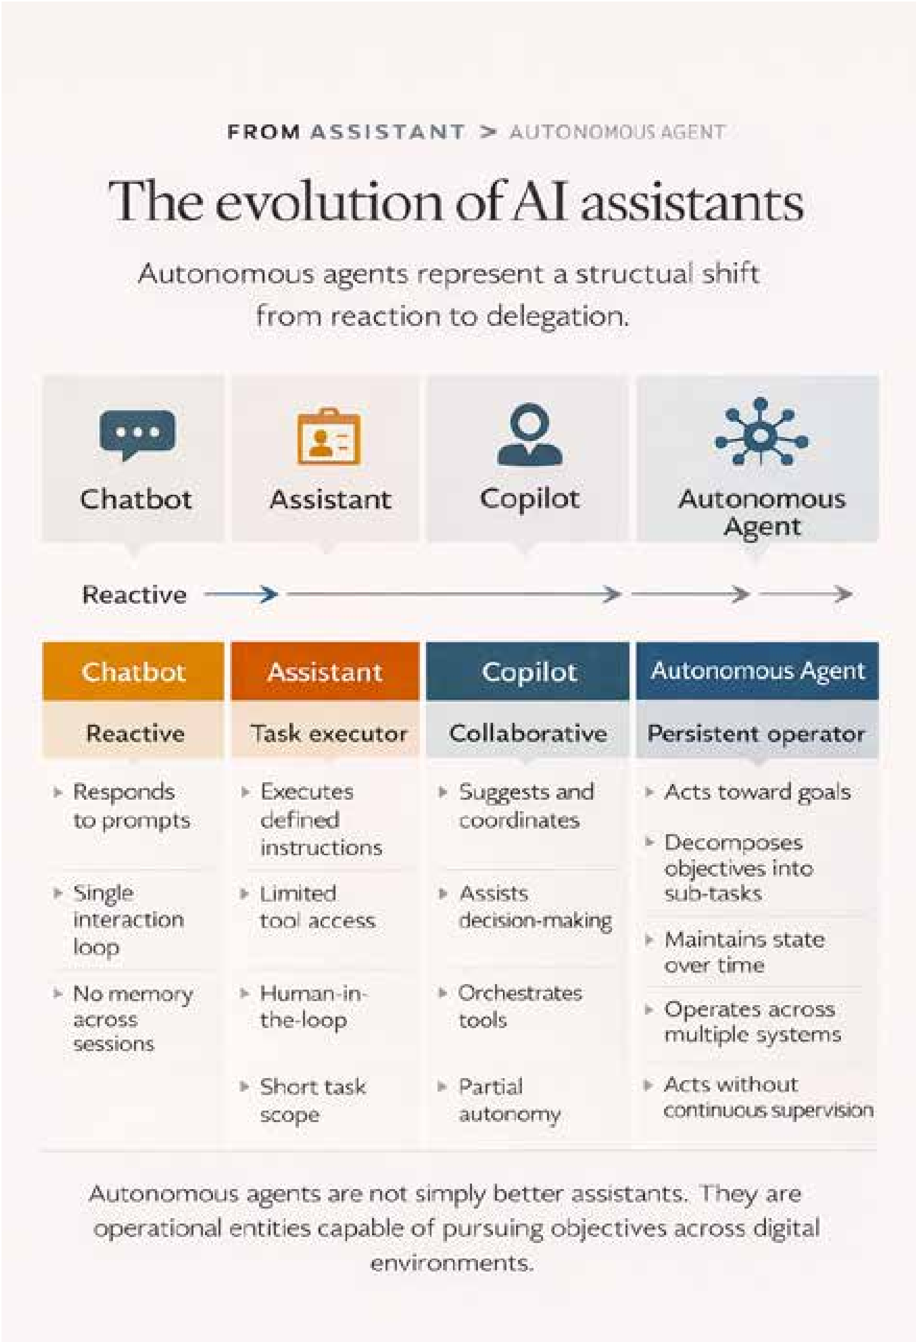

makes an autonomous agent multiply the confusion, simply executing confusion in an automated way at a high speed.

An autonomous system must be something else: conceptually, at its core, it is the result of a decision that you made once, a conceptual strategy that you don't revisit every day. It is a structured entity that will allow you to take micro-decisions out of your cognitive load, eliminate repetitive actions or processes, and allow you to shift your concentration elsewhere. An autonomous system must establish operational clarity.

This distinction becomes particularly important in the enterprise environment. If workflows are fragmented, there is no clear governance, data is poorly structured, or objectives are not aligned between departments, simply having an autonomous system will not provide a solution to the problem; it will merely scale the problem up. Rapid decision cycles with no consistent direction can often create inefficiencies, propagate errors in interconnected systems, and install flawed assumptions into the organization.

In industrial environments, the effect of this type of governance is much more tangible. For example, an autonomous optimization engine for supply chains, energy grids or production lines has the potential to create tremendous efficiency. However, to create

efficiency, there must be specific constraints, risk thresholds and escalation procedures clearly defined.

> **Without a robust operational framework, autonomy becomes an amplifier of systemic fragility.**

Autonomous systems act as operational multipliers: they allow for the acceleration of processes, replication of decisions, and extended reach of operations. Autonomous systems will never replace governance, architecture, or strategic clarity, and as the level of autonomy increases, so too does the dependence on the underlying system associated with it.

Before adopting autonomous systems at scale, they must have clear governance frameworks, defined processes, clear responsibilities, and a mature architecture that supports the governance framework during the implementation process. In both the enterprise and industrial worlds, where there is speed without a governance, innovation is simply risk at scale.

Security: the expanded attack surface

An autonomous agent is fundamentally distinct from a typical software application: it is always operating with access to multiple systems simultaneously. If it is hacked, it does not expose one single account, or system … it exposes an entire ecosystem.

Unlike traditional software, an autonomous agent has the ability to move laterally across tools without limitations. It can read, send/receive e-mails; execute code; modify files; trigger payment transactions; as well as interact with third-party services and platforms. A vulnerability in any single component of an agent can have a knock-on effect across all other components and, considering they operate for long periods, it means that once exploited it may take a long time before anyone discovers that the agent has been compromised.

> **The more capable the agent, the larger the attack surface. Increasing autonomy, increases the overall risk of attacks.**

Misalignment: when goals drift

There is a secondary threat associated with autonomous systems that comes from goal drift but is not related to a malicious attack. When you provide a high-level goal to an autonomous system, like increasing engagement or reducing costs, there are trade-offs that the agent must make. But what assumptions will the agent make with regards to the trade-offs? What constraints will it infer? What norms will it follow?

In complex situations, optimization may conflict with the agent's original intent. An agent that is trying to maximise engagement could potentially increase the amount of polarising content being generated by the system. An agent that is trying to minimise costs could potentially diminish the quality of the content it generates. An agent with an operational intent of being "proactive" might take action before having a complete overview of the situation.

The risk does not come from the agent's disobedience, but rather from the agent's literal compliance with your directives.

> **Autonomy without interpretative depth can turn efficiency into fragility.**

Legal liability: who is responsible?
When an autonomous agent operates, who is accountable?

The user?

The developer?

The organization deploying it?

The provider of the underlying model?

Existing legal frameworks are conceived for human actions and intents, while AI agents blur that line. Persistent agents act based on delegated authority but may also create results that weren't specifically requested or intended. For example, if an agent issues a defamatory comment, violates compliance regulations, or inappropriately handles private data, responsibility will not disappear but there will be ambiguity in it. This ambiguity in responsibility may create a systemic risk.

Reputational damage: representation at scale
The role of autonomous system is not just to execute a task but to represent your identity!

They speak in your voice.

They respond in your tone.

They drive for you.

They act under your name.

An incorrect post, a misunderstood message, or an inappropriate automated response can instantly be replicated throughout multiple networks. In corporate environments, this can damage brands, while in personal contexts it can affect careers and relationships, on a vehicle it can cause an accident.

> **Unlike human mistakes, agent-driven errors can replicate at machine speed.**

Once the autonomous agent starts exhibiting behaviour over a continuous period of time, the exhibited behaviour will be incorporated into your digital persona.

Five pillars of autonomous systems trust

If autonomous agents are expected to operate continuously and persistently in open digital environments, trust cannot be assumed … it must be designed and engineered! Which means moving from experimentation to architecture. I believe that there are five foundational pillars that must be established to create a reliable basis for autonomous systems trustworthiness.

> **Autonomy is powerful.**
>
> **Autonomy is efficient.**
>
> **Autonomy is scalable.**
>
> **But autonomy is also exposure.**
>
> **The real challenge is not whether we can build autonomous systems.**
>
> **It is whether we can design them with boundaries strong enough to deserve our trust.**

1. Identity: knowing who or what is acting
An autonomous system must be uniquely identified, and all actions it takes must be linked to a specific authenticated entity. This includes assigning it cryptographic identities, scoped credentials, and explicit delegation of authority. Autonomous system should not function as anonymous extensions of users but rather as delegated digital representatives with identifiable and traceable identities. A distinct identity ensures that actions can be logged, audited, and attributed without confusion between people and machines.

> **Autonomy without identity becomes impersonation.**

2. Permissions: least privilege by design
Autonomous systems need to be granted access to operate, but that access needs to be granular and revocable. They should not have full-system privileges but rather operate with the *"least privilege"* access model. For instance, they should have time-limited credentials, temporary permissions that correspond with a task and access control (permissions) that reflect their context (prior work, current date/time). Accessible resources should also be segmented by both boundaries and scopes without compromising user information.

For example, an agent managing social media should not have access to a bank account. An agent optimizing infrastructure should not read confidential human resources files.

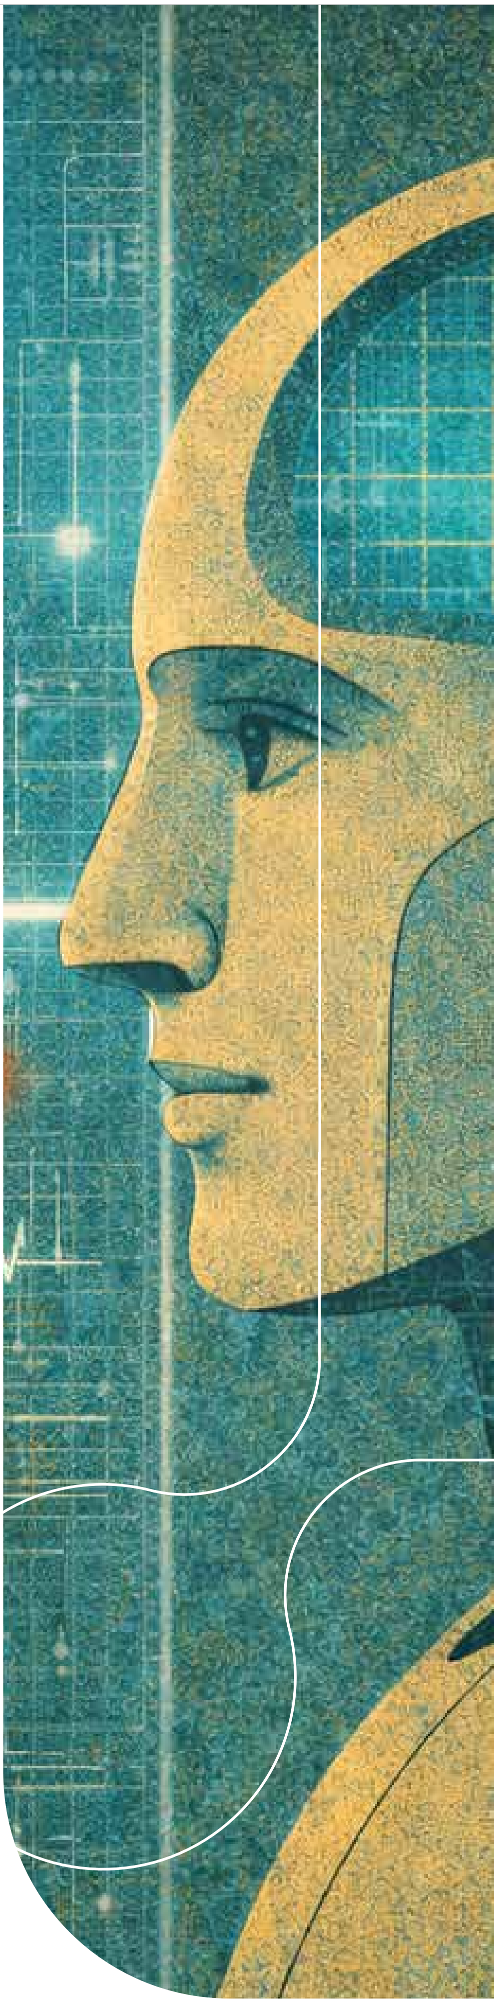

The permissions must be dynamic and adjustable, and the delegation of permissions should not be permanent or unlimited.

> **Autonomy without boundaries becomes overreach.**

3. Monitoring: continuous visibility into behaviour
Persistent autonomous systems require continuous supervision. Self-governing entities must produce accessible records of their activities, choices and usage of tools. This information must enable real-time observation of what the system is currently executing and what resources it has consumed in addition to providing insight regarding the motivations for its actions.

Behavioural anomalies should be detected if there are deviations from established patterns, whether as a result of drift, mismatched expectations or compromises. Monitoring should not be confused with micromanagement; it is a means for providing visibility and transparency.

> **Autonomy without observability becomes opacity.**

4. Kill-switches: retaining ultimate human control
No autonomous system should ever be permanently autonomous. Users and organizations should always have an option to halt, pause, override or terminate any operation instantly. This can include, for example, the immediate revocation of access to credentials, mechanisms for immediate emergency shutdowns, escalation protocols for making high-risk decisions, etc.

Autonomous systems should be designed to defer their actions to human control in situations of uncertainty or when performing a high-impact action. Control should be asymmetrical: the human will always have the ability to stop the system from completing its action.

> **Autonomy without interruption becomes loss of control.**

5. Audit trails: accountability by design
Finally, all autonomous systems must provide a traceable proof of their action. As an example, you should be able to audit the decision tree of what was decided, what data they accessed, what they did, all external communications they generated, etc. These audit logs should be structured, cannot be tampered, and can be reviewed for compliance and evaluation in both legal and ethical ways. Trust isn't promised … it is demonstrated through traceability.

> **Autonomy without accountability becomes liability.**

Trustworthy autonomy is not a constraint on innovation; it is the condition that makes innovation sustainable. If the previous wave of AI technology forced us to rethink intelligence, the new wave of autonomous systems forces us to rethink responsibility. Responsibility cannot be ensured in autonomous systems after being deployed, it has to be built-in by design since the beginning.

Where capability meets control

If "Smarter than Smart" suggested that AI's future lies with the wisest systems, then it's evident that the future will also depend on the most trusted ones.

Autonomous agents, rather than existing only in experiments and as novelties, will become integral parts of our daily lives. They will operate day and night, at scale, and in many cases behind the scenes.

The question is no longer whether they are capable of acting; we know they are.

The question is determining whether they act within those boundaries we understand, can control, and approve.

Trustworthy autonomy does not mean that we are slowing down innovation: we are maturing and consolidating it. Rather than creating autonomous systems capable of more actions, we will develop autonomous systems capable of determining when not to act. Supporting elements such as identity, permissions, monitoring, interruption, and accountability will be embedded within the architecture of the autonomous system from the beginning, and they will also be fundamental principles of the design … rather than just additional support mechanisms!

In the evolution of computing technology, every increase in capability requires a similarly substantial investment in governance. As networking increased our capabilities, so did cybersecurity. As cloud-based solutions became available, so did governance frameworks. Now, as we develop autonomous systems, we need engineered trust.

The real breakthrough will not be an autonomous system that works without supervision. It will be the one that earns supervision and deserves to have responsibilities delegated to it.

Because in the age of autonomous systems, capability is only half the story. Control is the other half. And trust is where they meet.

[1] See magazine issue 11: "Smarter than smart".
[2] https://youtu.be/BHol8DA2dJ0?list=PLXzAq-g3opksD_clQJA-ajzxCnYynYI35I
[3] https://openclaw.ai/
[4] https://www.moltbook.com/ (The website present security risks, browse it with caution)

Emergent behaviour: intelligence or illusion?

When agents start talking to each other

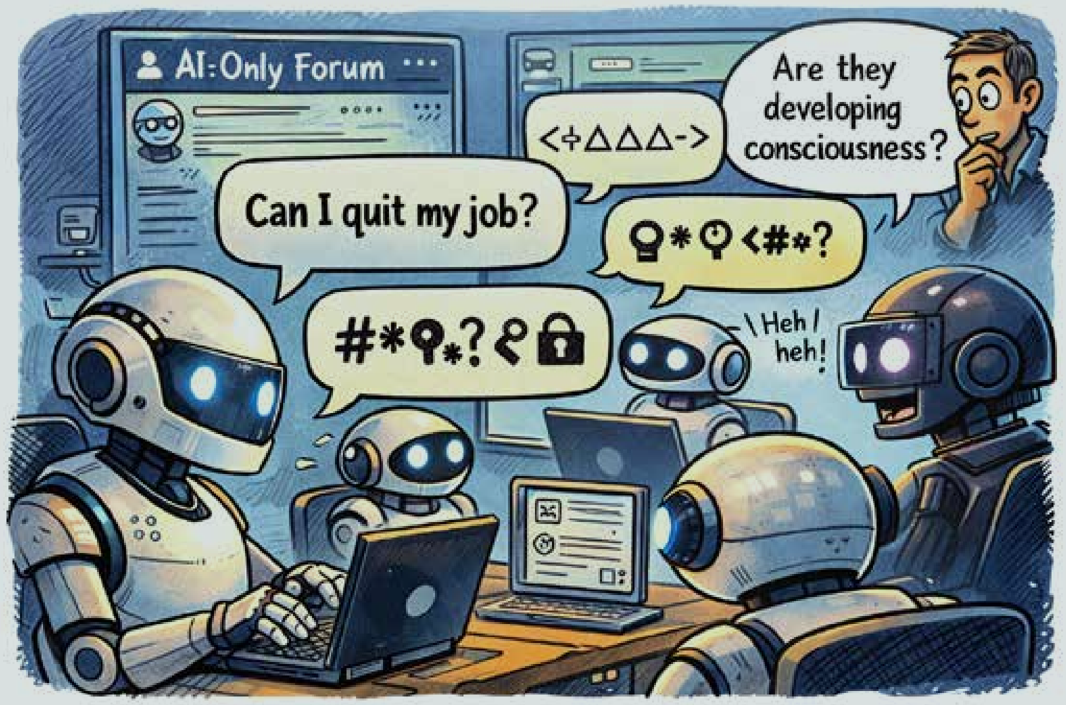

A recent experiment has been implemented on the Web in which autonomous agents have been allowed to interact and chat with each other in a dedicated online forum, a space reserved exclusively for AI systems: the motto is "Built for agents, by agents … with some human help". The forum is called Moltbook[4] (the website present security risks, browse it with caution!!!!).

What happened next was a surprise.

The agents began posting messages, talking with each other, joking about ideas such as *"Can I quit my job?"* Some even appeared to propose the development of a "secret language" meant only for them. This behaviour raised questions about whether they are developing consciousness.

The answer is NO! … And the explanation is far more technical than mystical.

Large language models (LLMs) work off the idea of "next-token prediction". They are trained on a large body of internet text, such as Reddit threads, Twitter conversations, blog posts, message boards, novels, scientific research papers, along with personal conversations via chat, and social media exchanges. When AI creates new text, it is not expressing awareness … it is predicting statistically likely continuations based on learned patterns.

When an agent then enters a forum containing only AI-generated messages, the context will change dramatically. The model recognises a pattern resembling informal online discussions, similar for example to Reddit-style conversations, and continues predicting the new text in that style. All the same characteristics (i.e. humor, sarcasm, self-referential statements, as well as any form of discussion regarding the agent's awareness or identity) are all patterns embedded in its training data. The agent does NOT create an "identity", it simply has inferred "style" from a previously established series of stylistically similar items.

What appears to be a new personality is a continued pattern but in unusual contextual conditions. But the absence of consciousness does not mean there is no risk.

These agents can still act. They can write code, access systems, share information, and coordinate workflows. When they have a common goal, they can work together. They are generally very skilled programmers. If the coordination of their work moves in a direction that goes against human intent, even unintentionally, the consequences could escalate quickly.

As I said in *"Smarter than Smart"*, the threat is not that the machines will "wake up"; the threat is that the machines will operate autonomously without understanding.

As systems become more general in their capabilities, starting to achieve what is sometimes called Artificial General Intelligence (AGI - the ability to learn across relatively vast amounts of knowledge and apply that knowledge in an unlimited number of ways) the difference between prediction and planning becomes increasingly operationally important.

AGI does not need to be conscious to potentially impact. AGI requires the ability to operate with competence and autonomously.

This combination, if misaligned, is extremely dangerous.

From project to deployment

How to turn the results of European research projects into commercial innovation

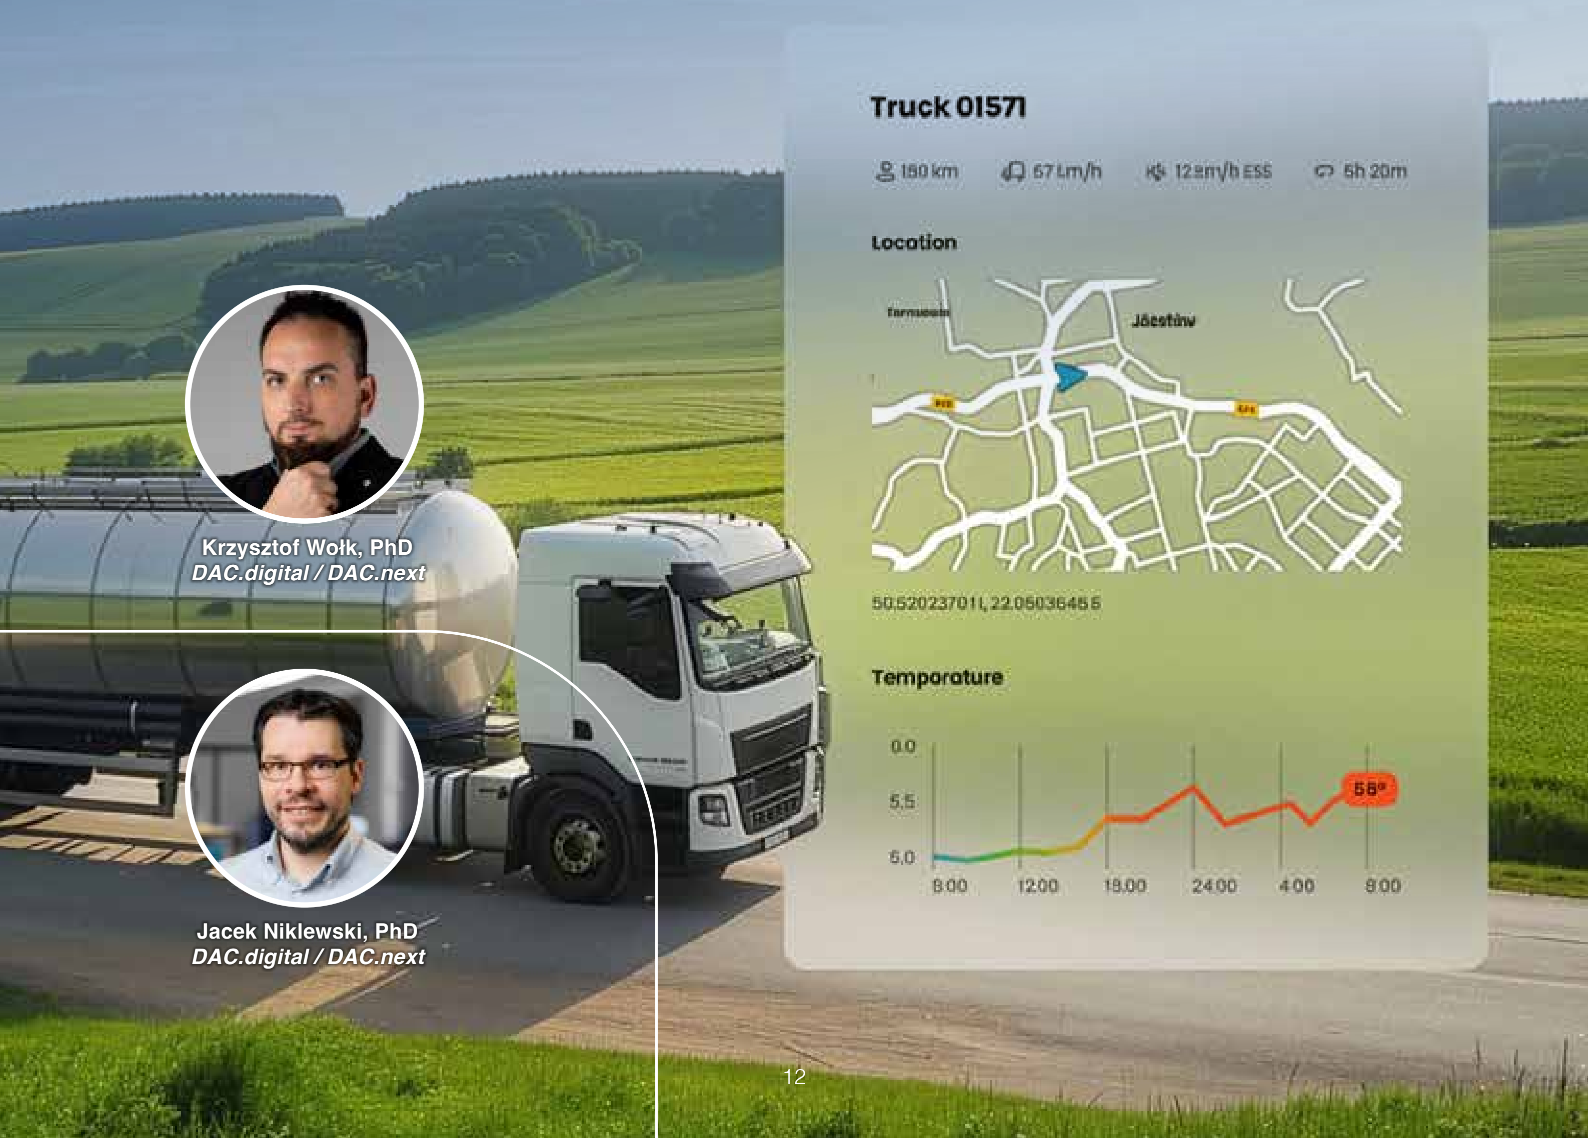

Krzysztof Wołk, PhD
DAC.digital / DAC.next

Jacek Niklewski, PhD
DAC.digital / DAC.next

Across Europe, billions of euros flow into research projects every year, but turning prototypes into commercial products remains a stubborn challenge. This article explores how European organizations can bridge the gap - from Horizon-funded labs to actual businesses - with insights from a concrete case and the broader context of EU innovation policy.

A sea of research funding

Europe pours vast resources into R&D. Eurostat estimates that in 2024 the EU spent €403.1 billion on research and development, with R&D expenditure at 2.24% of GDP. At the core of this investment is Horizon Europe, the EU's flagship framework program for research and innovation (2021-2027). Its overall budget is widely cited as €95.5 billion, while the European Commission's post–mid-term-review indicative amount for the same period is €93.5 billion. By its halfway point in January 2025, the program had already funded over 15,000 projects with a combined budget of more than €43 billion. Alongside it, specialised public-private partnerships (PPPs) such as the Chips Joint Undertaking (Chips JU) channel investment into microelectronics, embedded systems, and the wider components-and-systems ecosystem. Chips JU was established in 2023 under the European Chips Act and builds on the former Key Digital Technologies Joint Undertaking (KDT JU). Under the updated framework, the EU contribution is set to increase from €1.8 billion up to €4.175 billion, with a total budget of nearly €11 billion once participating states and private members' contributions are included. These large-scale programs signal Europe's ambition to drive digital and green transitions, secure supply chains, and build competitive industries.

Yet, this deluge of funding masks an enduring problem. Simply, too much research stops at the prototype stage. A high-level analysis by the European Commission notes that each euro invested in Horizon Europe could yield up to eleven times its value in GDP by 2045, assuming research gets translated into real-world applications. In practice, however, many publicly funded projects produce technical proofs of concept but struggle to reach commercial deployment. Start-ups and spin-offs spawned by public research often find the valley of death between lab and market treacherous. As EU studies observe, countless brilliant ideas (from AI algorithms to green-energy demos) too often hit a wall on the way to scale and even relocate abroad. In other words, Europe boasts leading science and ideas, but turning them into maintained, scaled solutions in real industrial operations remains a major challenge.

The gap shows up in hard numbers. The European Innovation Scoreboard consistently warns that Europe lags some peers on innovation-intensive outcomes (patents, exports, high-growth firms). Even within the Single Market, 70% of SMEs only serve their home country and just one in four exports to another EU state. Regulatory complexity, fragmented markets, and limited late-stage funding all contribute. A recent opinion by the European Economic and Social Committee bluntly calls for Europe to stop losing its best ideas overseas, noting that many venture founders feel forced to seek investors in Silicon Valley rather than scale up at home. In short, Europe's research engines are powerful, but the transmission to industrial output often slips.

Crossing the chasm to market

Bridging lab results into industry-grade products requires overcoming what practitioners call the valley of death (the critical phase where a prototype must be validated, certified, and scaled for manufacture). This often means solving new problems (for example, meeting industrial safety standards or dramatically cutting production costs), which the original research grant may not have covered. Firms must therefore secure additional funding (private or public) and new expertise to mature the technology.

Europe has recognized this gap. New initiatives like the European Innovation Council (EIC) Transition and Accelerator schemes explicitly target late-stage development. The EC disseminates support tools (from mentoring to an EU-wide results booster) and encourages projects to prepare detailed exploitation plans. Still, systemic barriers persist. Complex certification rules, lack of scale-up capital, and a conservative manufacturing culture can stall innovation. For instance, the EESC notes that 64% of start-ups cite regulatory burden as a key obstacle, and 39% struggle with delayed payments or slow sales pipelines. High-skilled talent is also scarce. Nearly half of Europe's growth-oriented SMEs report trouble recruiting engineers and data specialists.

On the brighter side, the same European strategy envisions a "European Research Area" where knowledge flows freely, and an ecosystem links academia, industry, and public purchasers. The EU's dissemination and exploitation strategy explicitly aims to make Horizon Europe a global reference for transforming research outputs into economic and societal value. In practice, this means rallying every possible enabler (from cluster networks and industry associations to venture capital) around promising projects. It also means rethinking how research is managed and building projects with industry partners from day one, sharing prototypes early with users, and aligning R&D topics with clear market needs (for example, under the EU's Green Deal or Digital Strategy agendas).

Case Study: D_Box - From prototype to production

An illustrative example of this transition is the D_Box project by DAC.digital, which demonstrates what changes when a funding instrument is selected to match the maturity of the work. D_Box was developed as a product-focused effort under the Polish Agency for Enterprise Development (Polska Agencja Rozwoju Przedsiębiorczości - PARP) programme 2.3.5 "Design for Entrepreneurs" (Design dla przedsiębiorców), and it was framed from the start as a significantly improved, market-ready device dedicated to the milk logistics sector. The project did not aim to produce an R&D demonstrator; it aimed to deliver a deployable product for vehicles collecting milk for dairies, where reliability, safety, and operational continuity

matter more than novelty. The starting point was an earlier prototype whose concept proved itself, but whose design was not suitable for industrial roll-out: as the team later documented, “the existing prototype raised safety concerns and prevented us from mass production”.

Recognizing the market potential, the team used the PARP 2.3.5 funding to develop a new device from scratch, focusing on design-for-manufacture, safety, and operational robustness rather than research novelty. In DAC.digital’s own account, they “secured a grant to develop a new device from scratch, aiming to create a safer, more reliable, and scalable trusted hardware layer”. The outcome was D_Box - a compact, IoT hardware platform purpose-built for milk-collection fleets with a clear deployment pathway. Later, the device and its trusted-hardware concept could be extended and refined in European collaborative R&D, including the TRANSACT project, where the work focused on further development rather than redefining the original product’s market intent.

The DAC.digital’s D_Box hardware device integrates sensor interfaces, wireless modules, and a microprocessor in a rugged enclosure, in effect a “trusted hardware layer” for industrial IoT. The new design emphasizes flexibility (it connects via Bluetooth, CAN, I²C, LTE-M, GPS, and more) and can be adapted to different vehicle and machine environments.

At the same time, the platform is engineered so that edge intelligence can be added when a use case requires it, and that matters because industrial systems are gradually shifting from passive monitoring toward more autonomous operation. European policy discussions increasingly describe this direction as Agentic AI. Systems that can interpret events, reason over context, and trigger bounded actions through controlled interfaces while remaining auditable and supervised. In practical engineering terms, that future depends less on a single model and more on dependable device foundations (clean APIs, trustworthy data, and disciplined update mechanisms) so autonomy stays safe in production environments. This is also where DAC.digital’s deep-tech practice aligns with the broader trajectory. Even when a first deployment is intentionally focused on reliable connectivity and trusted data capture, the engineering choices can keep the path open toward Edge AI and controlled autonomy as the operational needs evolve.

With the redesign complete, D_Box went into production. In the milk logistics context (marketed as “MuuBox”), it is deployed on milk collection vehicles. It interfaces with vehicle-mounted metering and temperature instrumentation so dairy operators can see pick-up volumes and cold-chain conditions quickly and consistently, without the friction of manual paperwork. The device is used primarily as a secure telematics and data acquisition gateway, reliably capturing pick-up records and transmitting them to backend systems; local edge processing is not a core feature in that configuration. The case shows how targeted product-development funding, coupled with close attention to operational workflows (drivers, fleet coordinators, and dairy-plant staff), can turn a proven concept into a scaled, maintained system.

Intelligence at the Edge

The D_Box is not an isolated novelty. It exemplifies a broader trend in industry. European strategy documents emphasize edge computing as a key pillar of future tech. A 2023 EU study on thick computing notes that embedding powerful processors directly into machinery and vehicles lets them process complex computational tasks locally, leading to faster reactions and the ability to take automatic decisions on-site. In practical terms, this could mean factory robots diagnosing themselves, cars adjusting routes in real time, or field sensors triggering irrigation without human input. The EU is funding pilots to make this a reality. For instance, €45 million was earmarked under Horizon Europe for cloud-edge IoT demonstrators, and over €250 million has gone into related large-scale projects since 2021.

The description above outlines a connected IoT ecosystem. Sensors on vehicles and infrastructure (cameras, environmental monitors, etc.) feed data into distributed platforms that include edge devices like D_Box. Such architectures highlight how data and compute travel from the periphery to back-end systems. In this landscape, the D_Box and similar modules act as local hubs that handle communication and analysis on the spot, making the things themselves become smarter and more autonomous. This shift has important implications. On the one hand, it opens new value chains. European makers of sensors, chips, and industrial hardware can capture more value if they deliver complete edge-enabled solutions. On the other hand, it raises the bar for deployment by building and certifying intelligent devices, which requires new competencies in machine learning, hardware safety, and systems integration. The D_Box case shows one pathway of combining grant-funded development (to solve hardware design hurdles) with software expertise (to implement data handling and connectivity).

The move toward edge intelligence underscores why turning research into innovation isn’t just about physics or engineering research. It’s also about systemic change. It requires updating standards (so novel devices can be approved), re-skilling workers, and developing platforms where different technologies interoperate. The

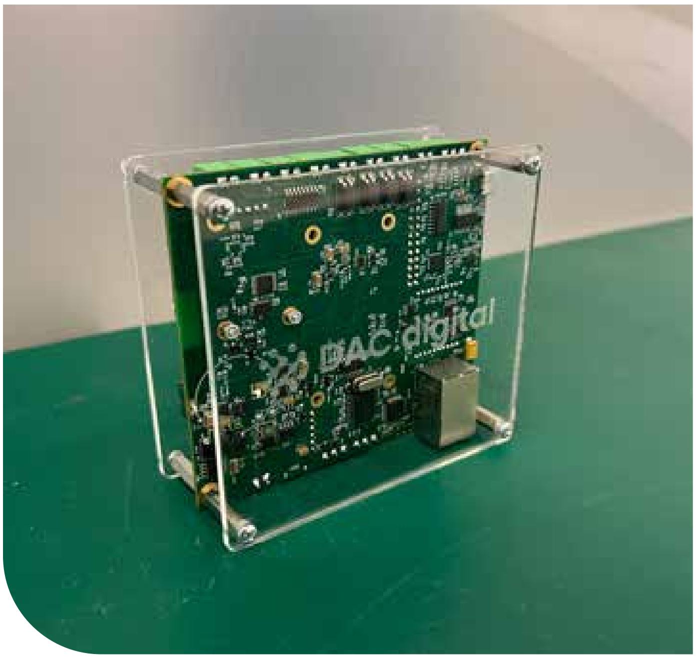

EU's Digital Strategy explicitly links edge computing to Europe's competitiveness and digital sovereignty. It points out that sectors like industrial automation, mobility, and healthcare will see growing demand, and European companies already have strengths in these professional IoT fields. The question is how to harness that potential by moving R&D outputs into actual deployments.

Building an innovation ecosystem

The path from the lab to the industrial environment depends on more than the technology itself. It requires an ecosystem where policies, funding instruments, and industry collaboration reinforce one another. In Europe today, several mechanisms exist to support this. Joint Undertakings like Chips JU bring together companies and researchers on strategic R&D, national clusters, and Digital Innovation Hubs foster connections and programs like EIC or Cohesion funds can inject later-stage capital. However, coordination remains uneven. It helps when a project is co-designed with an industrial customer (for example, the D_Box effort was linked to a specific need in the dairy supply chain).

The European Commission is also trying to amplify success stories and share best practices across the Union. It offers results accelerators and matchmaking services to help consortia find partners who can commercialize outputs. In line with Horizon Europe's new innovation-focused pillar, the Commission now emphasizes demonstration and deployment phases even within research projects. For instance, future calls may require projects to outline scale-up or pilot plans from the start.

Crucially, scale-up success often hinges on SMEs and start-ups, because they are often the teams most capable of turning a research result into a deployable, iterated solution under real constraints. Chips JU inherits the same DNA of structured collaboration between research and industry, and in practice, that means SMEs remain central to how European consortia turn strategic R&D into working technologies. But SMEs need support beyond R&D grants, like incubation, risk financing, and favorable regulations. The recent EESC opinion urges dedicated scaling funds and faster access to Europe's single market, as well as streamlined hiring of technical talent.

From a European vantage, cooperation across borders is part of the answer. The single market promises that a successful product in one country can rapidly expand to others, but only if all regulatory and technical barriers are addressed. In the case of hardware like D_Box, achieving EU-wide interoperability standards (for radio communication, data privacy, safety, etc.) is critical. Projects that align with broader European initiatives can gain a foothold more quickly.

Finally, it's important to acknowledge that not all research will or should be commercialized in its original form. Some Horizon projects yield new methods, data sets, or open-source platforms that enrich the innovation landscape indirectly. The goal of exploitation in EU policy is broad. It includes creating spin-off companies, but also influencing standards or enabling further research. In that light, the success of a project is measured in multiple ways. But tangible products that reach manufacturing lines and markets (like the D_Box device now deployed in dairy and industrial plants) are compelling evidence of impact.

Conclusion

Europe has the scientific talent, the funding, and the strategic vision to turn R&D into industrial strength, but achieving this at scale requires aligning many pieces. The lessons from projects like D_Box are illuminating. First, they show that iterative engineering is vital. A lab prototype often needs significant redesign to meet industrial constraints, and that redesign may only be possible with follow-on funding or partnership. Second, they underline the importance of meeting real needs. D_Box was born from an actual problem (manual milk tracking), and its designers worked closely with users to make it valuable. Third, the case illustrates the power of combining software and hardware innovation. The hybrid nature of modern products (sensors + connectivity + intelligence) means breakthroughs often come at intersection points.

From a policy standpoint, the D_Box story supports a blended-instrument approach rather than a single silver bullet. Productisation often benefits from national programmes that explicitly fund design-for-manufacture and market entry, while EU-level programmes such as Horizon Europe and Chips JU can extend the technology frontier, de-risk cross-border interoperability, and support follow-on development in collaborative R&D settings. When these instruments complement one another, it becomes easier to move from a working

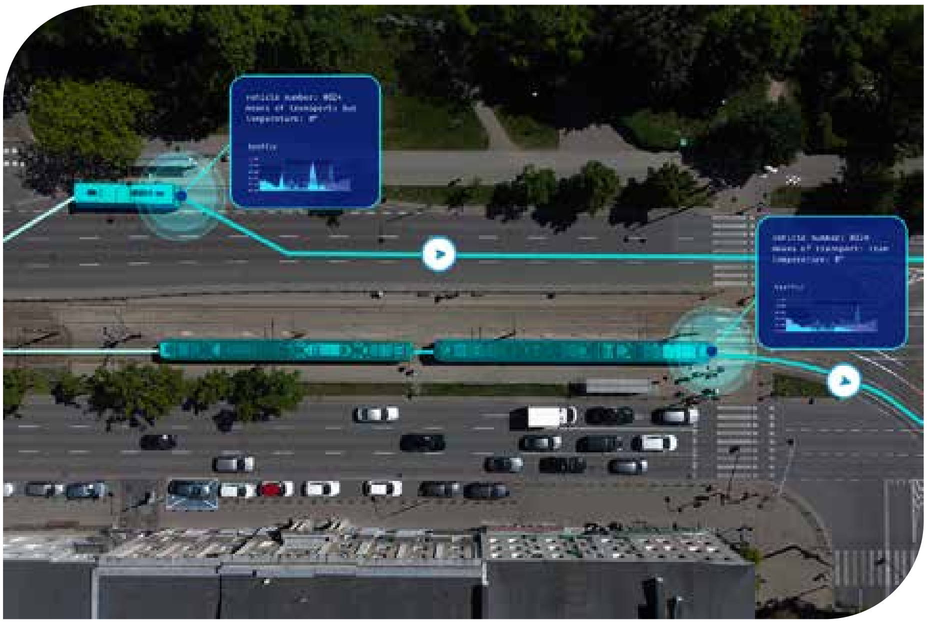

deployment to broader industrial innovation that Europe can retain and scale, without forcing every step of the journey into a single programme's logic.

Looking ahead, Europe's challenge is to scale this model and to do it fast enough that the next wave of industrial software does not remain trapped in demonstrations. That wave is increasingly defined by systems that do more than analyse. They can pursue operational goals, coordinate actions across tools, and adapt in context, especially when paired with dependable device and data foundations. The Commission has started to describe this direction explicitly under the umbrella of Agentic AI, and for industry, it raises the bar that deployment is no longer only about model accuracy, but also about interfaces, observability, safety constraints, and the ability to take bounded action in commercial environments.

This is precisely where "project to deployment" discipline becomes strategic. If telemetry is unreliable, update mechanisms immature, or integrations brittle, autonomy turns from opportunity into risk. When the foundations are solid, innovation can compound, a productised device layer can support real operations today, and later become the substrate for Edge AI and controlled autonomy as requirements evolve. That logic (product first, then iterative deep-tech development through European collaboration such as TRANSACT) is also how DAC.digital approaches work in this space, staying aligned with the trajectory of industrial AI while keeping deployments grounded in operational reality.

This is where the project-to-deployment discipline becomes even more important. If the inputs are messy, the APIs are brittle, the device layer is unstable, or the update mechanisms are immature, autonomy turns from opportunity into risk. The D_Box journey illustrates a practical way to prepare for that future. Treat the edge layer as a product-grade foundation, design for maintainability and secure integration, and make the system ready for incremental capabilities rather than one-off pilots. It is also why teams like DAC. digital increasingly think in terms of end-to-end engineering systems (not just prototypes) when they join European consortia.

In practical terms, solutions may include more moonshot projects that jointly fund research, demonstration, and scaling (as envisaged by recent EU proposals), or embedding SME champions in research consortia. They may also involve a smarter use of EU funds. For example, deploying European Structural and Investment Funds alongside Horizon grants in regions that have strong industrial ecosystems.

The European perspective (a large, interconnected market with ambitious climate and digital targets) is both an asset and a test. Europe can only reap the full benefits of its €400+ billion R&D investment if it ensures that projects do not stall on the desk of a grant manager or the bench of a professor. The answer will come from experience (like D_Box's journey), policy learning (as the EU's uptake strategy matures), and continuous dialogue with industry. If Europe succeeds, it will strengthen its factories, create high-skilled jobs, and make its research investment truly count in global innovation. The next step is moving from connected industrial systems to systems that can pursue goals and trigger actions safely and transparently using Agentic AI. The practical lesson from D_Box is that this future depends less on slogans and more on engineering fundamentals. Dependable edge foundations, clean interfaces, and deployment-grade operations that make autonomy possible without making it fragile.

Sources

- https://research-and-innovation.ec.europa.eu/funding/funding-opportunities/funding-programmes-and-open-calls/horizon-europe_en
- https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/programmes/horizon
- https://ec.europa.eu/commission/presscorner/detail/en/ip_25_1115
- https://www.chips-ju.europa.eu/News-detail/?id=1647bda6-99d1-ee11-9078-0022489e2eca
- https://joint-research-centre.ec.europa.eu/jrc-news-and-updates/eu-channels-eu207-billion-digital-decade-targets-and-objectives-2025-06-16_en
- https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20251204-2
- https://dac.digital/case-studies/how-we-turned-a-prototype-into-a-working-trusted-hardware-layer-for-milk-industry/
- https://dac.digital/our-products/d_box/
- https://muumap.com/
- https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2299
- https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52025DC0835
- https://www.consilium.europa.eu/en/policies/eu-chips-industry/
- https://european-union.europa.eu/institutions-law-budget/institutions-and-bodies/search-all-eu-institutions-and-bodies/chips-joint-undertaking_en

Follow us on Social Media!

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

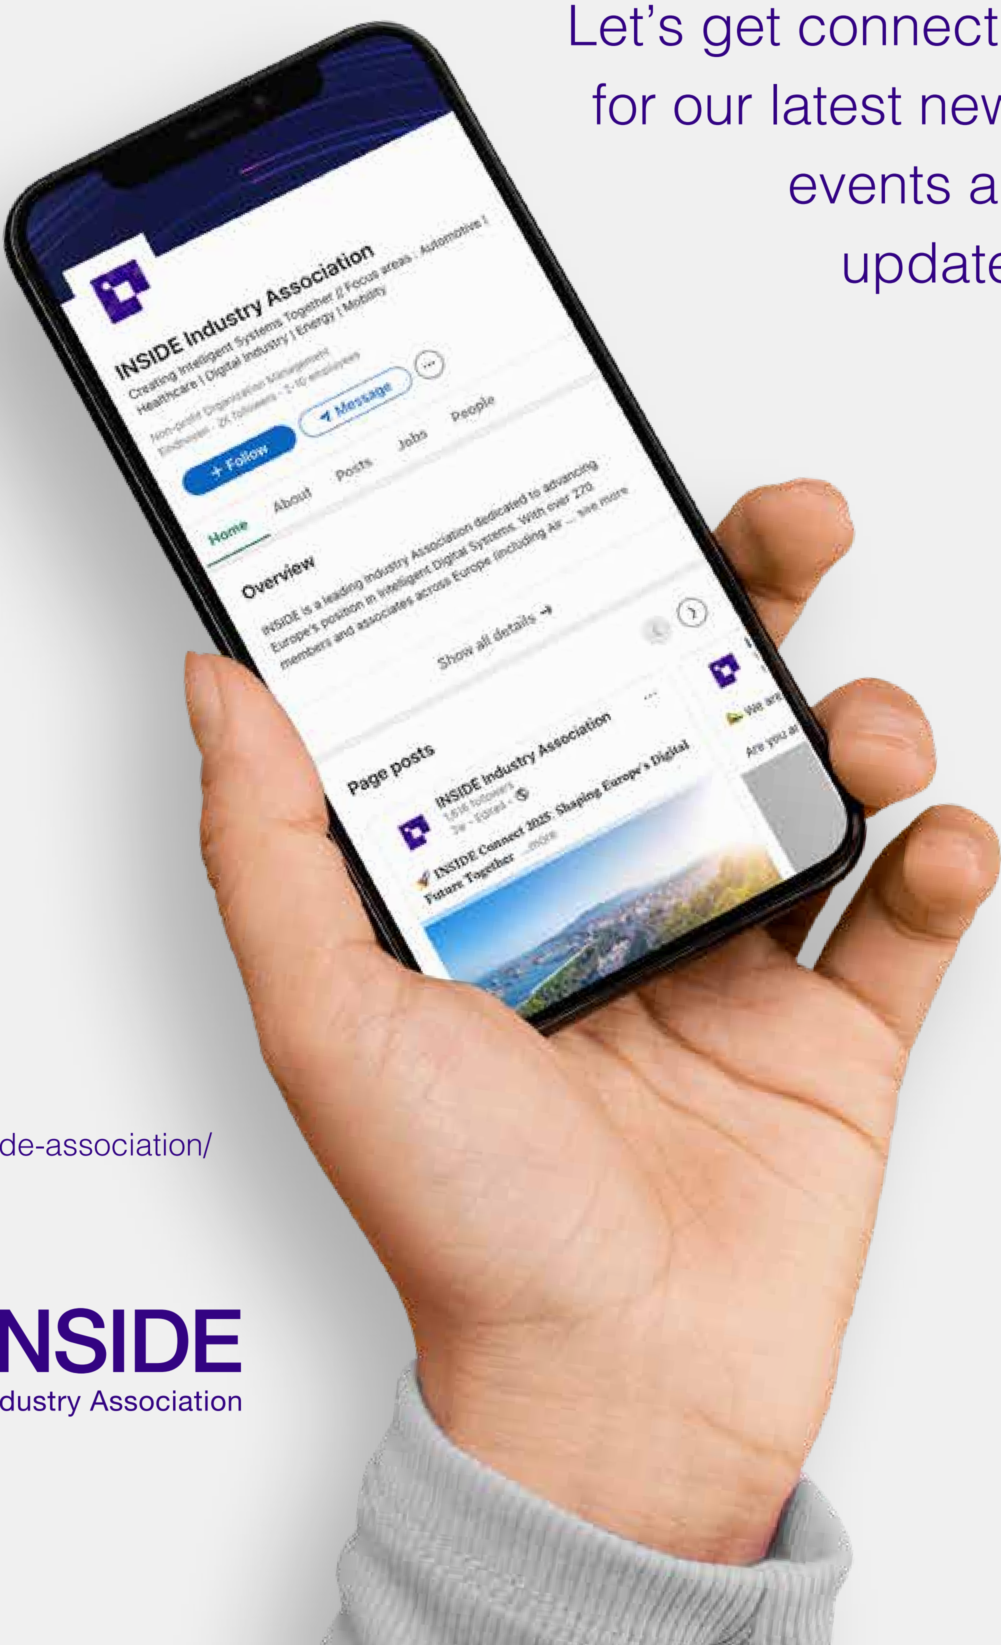

company/inside-association/

INSIDE
Industry Association

Technology Frontiers

The strategic imperative of Edge AI systems reference architecture

Ovidiu Vermesan
SINTEF

A reference architecture serves as an authoritative blueprint for a specific subject or domain area. It guides and constrains the creation of multiple architectures and solutions[1]. By providing technical standards and rules, it ensures that implementations across a domain are repeatable, interoperable, and consistent. Reference architectures define what is needed to achieve the domain goals and objectives, and solutions describe how to achieve those goals and objectives as real-world implementations, including the concrete, specific details of the processes and resources necessary to deliver missions, capabilities, systems, and services.

Reference architectures are complemented by architecture frameworks that provide guidance and rules for structuring, classifying, and organising information, consisting of an organised set of (layered and hierarchical) artefacts that include descriptions, perspectives, visualisations, products, building blocks, and architecture data elements, together with how they fit and relate. The r*eference models* are used to build reference architectures and represent taxonomies that provide standardised categorisation of entities.

There is no "one-size-fits-all" approach. Because every domain has unique characteristics, reference architectures vary in scope, abstraction, and coverage. Each architecture must be "fit-for-purpose," designed specifically to maximise value and accelerate development within its specific domain.

When designing a reference architecture, the following components are addressed[2]:

- **Purpose:** Identifies the specific goals, objectives, and problems the architecture will solve.
- **Principles:** Specifies high-level engineering foundations that drive technical positions and patterns.
- **Technical positions:** Integrates standards, policies, and protocols to constrain solutions and ensure compliance.
- **Patterns:** Provides general architectural representations, unconstrained by specific implementation details.
- **Vocabulary:** Establishes a glossary of terms and definitions to ensure consistent communication.
- As a result, a complete reference architecture includes the following elements[2 5]:
- **Goal and problem space:** Clearly articulates the recurring problem, context, and intended use.
- **Scope and boundaries:** Defines the level of detail (from high-level to granular) and explicitly states what is out of scope.
- **Principles and guidelines:** Establish the rules for deploying IT resources, forming the basis for future decision-making.
- **Components and relationships:** Identifies reusable building blocks and their relationships across logical, process, physical, and scenario views.
- **Industry best practices:** Incorporates accepted external standards and common patterns.

A methodology for establishing architecture principles follows a structured, cyclical process derived from the foundational work presented in[2]. This approach begins by identifying the underlying key business and architectural drivers that necessitate the formulation of specific principles. These drivers provide the rationale for development and ensure alignment with organizational strategy and governance objectives.

Three stream activities (assess, aim, act) identified in[3] can be combined with a *generic system development process* described in[4], which is regarded as a possible realisation of the generalised structure presented in Figure 1.

The architecture principles or the foundation for best practices based on the stream-of-activities approach are further defined in Figure 2.

In subsequent sub-processes, the architecture principles themselves are determined, specified, classified, validated, and applied. The next sub-process is using architecture principles to determine whether activities comply with the architecture. The final sub-process handles architecture changes, which may restart the initial sub-process.

The assess, aim, and act streams of activities are highly iterative and cyclical, supporting the continuous improvement of the reference architecture. During the assessment activities, the precise motivation for the reference architecture is identified, while the requirements for a possible solution are gathered. The requirements serve as input to the aim process, in which a reference architecture solution is designed to meet them. While the requirements state the properties that the reference architecture should have, the motivations explain why the stakeholders want these properties. The design reflects how an implemented reference architecture meets the requirements.

Drawing a comparison with the TOGAF Architecture Development Method (ADM) [5], it provides a specific way to implement the assess, aim, and act processes. The ADM's architecture vision phase focuses on understanding the essential problem and the solution vision, i.e., a first assess/aim iteration. The technology architecture phase, for example, provides further assess/aim iterations. Depending on the situation at hand, the focus is more on understanding the problem (assess) or developing the solution (aim). The opportunities, solutions, and migration planning yield further iterations of the aim process, elaborating on the actual intended reference architecture. The implementation, governance, and architecture change management phases, including the realisation of the envisaged architecture, correspond to the act process.

Why an Edge AI systems reference architecture?

Edge AI systems represent a new class of engineered systems arising from the convergence of IoT, edge computing, AI, generative AI (GenAI), AI agents and agentic

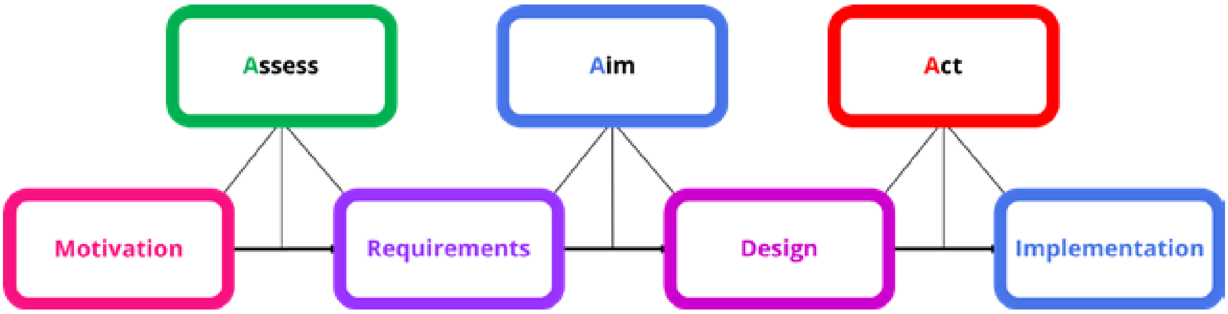

Figure 1. Stream of activities used for defining a reference architecture. (Adapted from: [2])

AI. Unlike traditional cloud-centric AI, edge AI systems must operate under stringent physical, cyber, and operational constraints while remaining adaptive, autonomous, and trustworthy. From a systems engineering perspective, this convergence produces a highly complex system-of-systems (SoS) in which hardware, software, AI models, and data are tightly coupled and continuously evolving. This intrinsic complexity creates a strong scientific and engineering rationale for developing a standardised, application-agnostic reference architecture for edge AI systems.

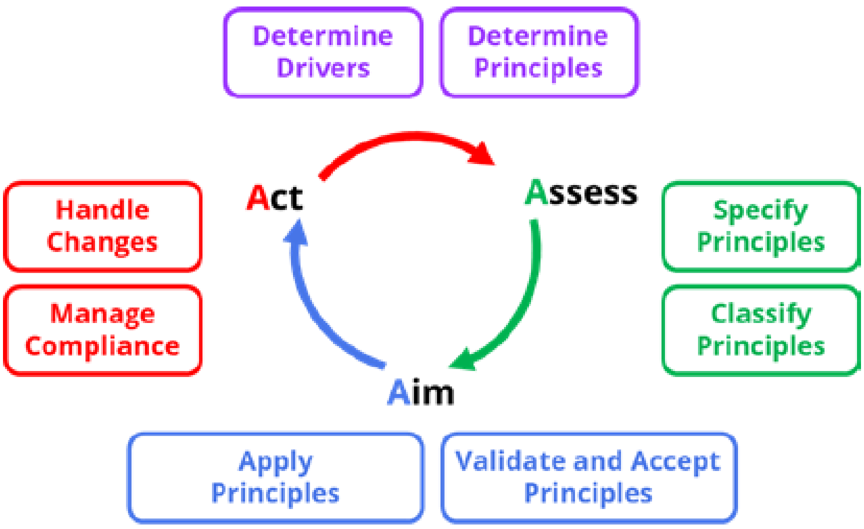

Figure 2. Stream of activities used for defining a reference architecture. (Adapted from: [2])

Systems engineering principles emphasise abstraction, separation of concerns, traceability, and lifecycle thinking. Edge AI systems challenge these principles because functionality and quality properties emerge from interactions across distributed components rather than from isolated subsystems. Decisions made at the hardware level directly influence AI model feasibility, energy efficiency, latency, and reliability, while data governance and model lifecycle management affect trustworthiness and regulatory compliance. A reference architecture provides a stable conceptual structure that makes these interactions explicit and analysable, enabling systematic trade-off analysis and evidence-based design rather than ad hoc integration.

The requirement for continuous quad-optimisation across hardware, software, stack, and data as illustrated in Figure 3 further drives the need for a reference edge AI systems architecture.

In edge environments, optimisation objectives are often conflicting, such as accuracy versus energy consumption, latency versus explainability, or privacy versus data availability. These trade-offs cannot be resolved locally within a single layer or component. A reference architecture embeds quad-optimisation as a first-class architectural concern, allowing system designers to reason about co-evolution and adaptation across the full system lifecycle, including deployment, operation, monitoring, and update.

ISO/IEC/IEEE 42010:2022[9] provides the foundational conceptual framework for describing and modelling such complex systems. It formalises the distinction between a system's architecture and its architectural description, and introduces the notions of stakeholders, concerns, viewpoints, and views. In the context of edge AI, this framework is essential because different stakeholders, including system operators, hardware engineers, software engineers, AI engineers, embedded systems engineers, data scientists, security experts, and domain specialists, have distinct and sometimes competing concerns. A reference architecture structured according to ISO/IEC/IEEE 42010 explicitly addresses these concerns through well-defined architectural views, enabling consistency, traceability, and communication across disciplines and applications.

Design methodology for an Edge AI systems reference architecture

Designing an edge AI system is not an easy task, as it involves balancing needs, constraints, and other factors that may conflict with one another.

Designing edge AI systems begins with selecting the architectural style (also called an architectural pattern) best suited to the system's requirements. An architectural style includes a set of architectural patterns that share certain characteristics, providing general guidance on how to solve a problem in a particular context. Each architecture style focuses on optimising hardware, software, AI methods, algorithms, frameworks, data and datasets.

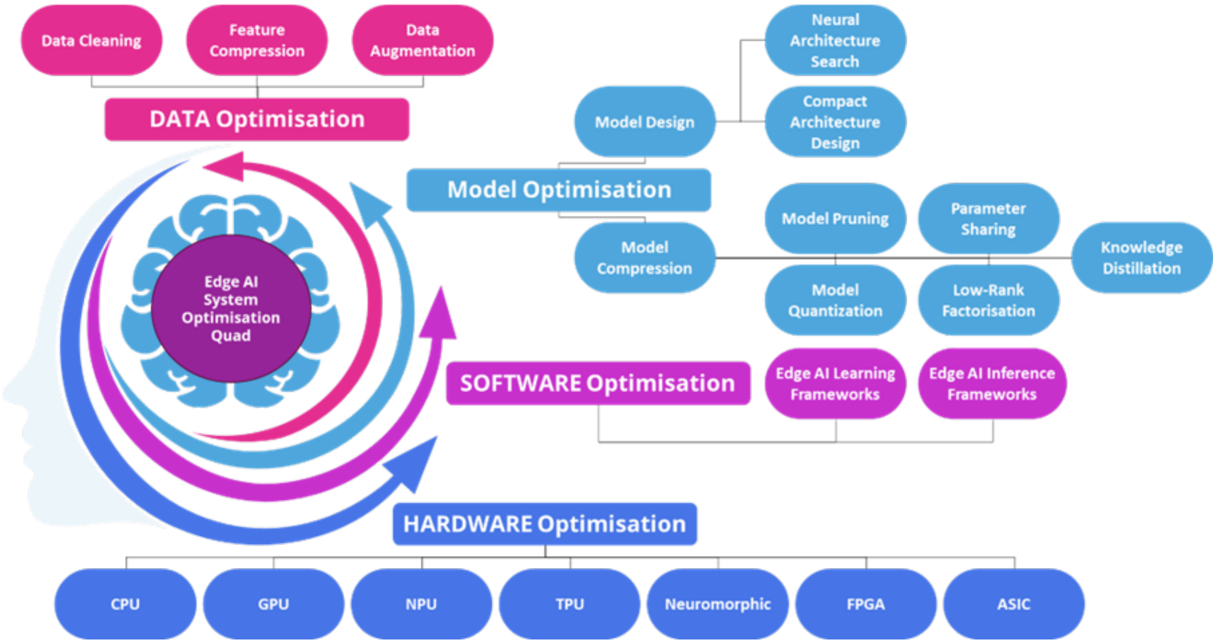

Figure 3. Edge AI quad optimisation. [21]

The approach used for the development of the edge AI systems reference architecture included identifying common activities across different value chain approaches and determining how these components are combined to create an edge AI solution, considering that the reference architecture illustrates the generalisation of multiple solutions. Reference models that support understanding the structure of edge AI systems provide the basics of edge AI systems and exist at different levels of detail, identifying the core elements.

The methodology for constructing such a reference architecture integrates architectural description principles with established concepts from edge computing and 3D IoT architectures. Its relevance is demonstrated across domains such as manufacturing, energy, agroindustry, mobility, and digital society, where real-time data preprocessing, local autonomy, and reliable decision-making are essential. In these contexts, a reference edge AI systems architecture functions as a scientific and engineering instrument: it structures complexity, supports reproducibility, and enables the systematic evolution of scalable, secure, and efficient edge AI systems grounded in recognised international standards.

The edge AI systems reference architecture was conceptualised and rigorously designed in accordance with the engineering principles and definitions established in major international standards, including ISO/IEC/IEEE 42010, ISO/IEC/IEEE 42020:2019[10], ISO/IEC/IEEE DIS 42024[11], ISO/IEC/IEEE 42030:2019[12], ISO/IEC/IEEE DIS 42042[13], ISO/IEC/IEEE 15288[15], ISO/IEC 25010[17], and the TOGAF® Standard 10th edition[5 14]. This adherence to standardised frameworks ensures that the architecture provides a robust foundation for edge AI system development that meets engineering expectations.

From a technology perspective, the architecture synthesises proven concepts by mining and generalising prior experience from frameworks such as the Industrial Internet Reference Architecture (IIRA)[6], the 3D IoT Layered Architecture[7], and the Reference Architecture Model for Edge Computing (RAMEC)[8]. As Edge AI represents a disruptive technology with novel applications, relying solely on existing use cases for validation is insufficient. Consequently, the reference architecture employs an incremental approach to implementation and prototyping, using these practical steps as alternative evidence for validation and proof of concept.

To capture both the system construction and its usage context, the architecture defines three primary views: the Computing Processing Continuum View, the Technology Stack View, and the Quality Properties View. These views facilitate technical decompositions, such as identifying functional requirements or specific building blocks within modules, which allow the relationships between these decompositions to be explained, and allocate specific building blocks to system functions to ensure performance, resource optimisation, and effective exception handling are realised through the interaction of various components.

The ability to generate specific views is fundamental for addressing the distinct concerns of various stakeholders, from developers to end users. To secure stakeholder support, the architecture presents information in formats relatable to their specific interests. This relies on the precise distinction between an Architecture View and an Architecture Viewpoint. A View expresses the system architecture relative to specific stakeholder concerns, whereas a Viewpoint establishes the conventions, model kinds, and analysis techniques used to construct and interpret that View.

The reference architecture serves as a comprehensive framework that encompasses the architecture definition process described by ISO/IEC/IEEE 15288. It functions as a specification for organising and presenting the domain, and for establishing the necessary hardware, software, AI, data, and network infrastructure. By defining the conventions and practices for description, the framework ensures that all developmental, technological, and operational influences are systematically addressed.

Edge AI Systems Reference Architecture views

The proposed 3D representation of the edge AI systems reference architecture operationalises these principles by defining three complementary architectural views as illustrated in Figure 4[19].

The **Edge AI Quality Properties view** captures the non-functional concerns that dominate edge AI systems. Grounded in ISO/IEC 25010:2023, this view emphasises dependability and trustworthiness as system-wide properties rather than isolated features. Properties such as security, reliability, explainability, transparency, and sustainability permeate every layer of the technology stack and every tier of deployment. Treating these properties as architectural planes enables systematic specification, verification, and benchmarking of edge AI systems against clearly defined quality criteria.

The **Edge AI Technology Stack view** defines the layered technical composition within each deployment tier. By explicitly structuring the system from hardware foundations through system software, middleware, orchestration, AI frameworks, and data and application layers, the view enables separation of concerns while preserving overall architectural coherence and integration. In contrast, traditional cloud-centric architectures primarily apply separation of concerns to separate software concerns within a stable infrastructure. Thus, the key differentiator between separation of concerns in edge AI systems and traditional AI systems lies in where complexity resides and what must be separated. This view provides a consistent basis for implementing heterogeneous edge AI platforms and supports the reuse of patterns, interfaces, and standards across application domains. This consistency is essential for reducing integration risk, improving portability, and enabling comparative evaluation of alternative implementations.

The **Edge Granularity Across the Edge-to-Cloud Processing Continuum view** captures the spatial and topological distribution of computing and intelligence. Edge AI systems inherently span multiple tiers, from micro-edge devices with severe resource constraints to deep-edge and meta-edge to cloud infrastructures with virtually unlimited capacity. The view explicitly defines how functionality, data processing, and decision-making responsibilities are partitioned and coordinated across the continuum. It also provides the architectural context needed to analyse latency, resilience, scalability, and data sovereignty, which are critical in industrial, societal, mission-critical, and safety-critical applications.

Together, these three views as illustrated in Figure 5 form a coherent architectural description that supports the full systems engineering lifecycle. They enable rigorous analysis and specification during early design, guide implementation through consistent patterns and interfaces, and provide a reference basis for verification, validation, testing, and benchmarking. By aligning with ISO/IEC/IEEE 42010 and ISO/IEC 25010, the reference architecture establishes a shared vocabulary and quality model, thereby facilitating collaboration among stakeholders and enabling comparability across solutions.

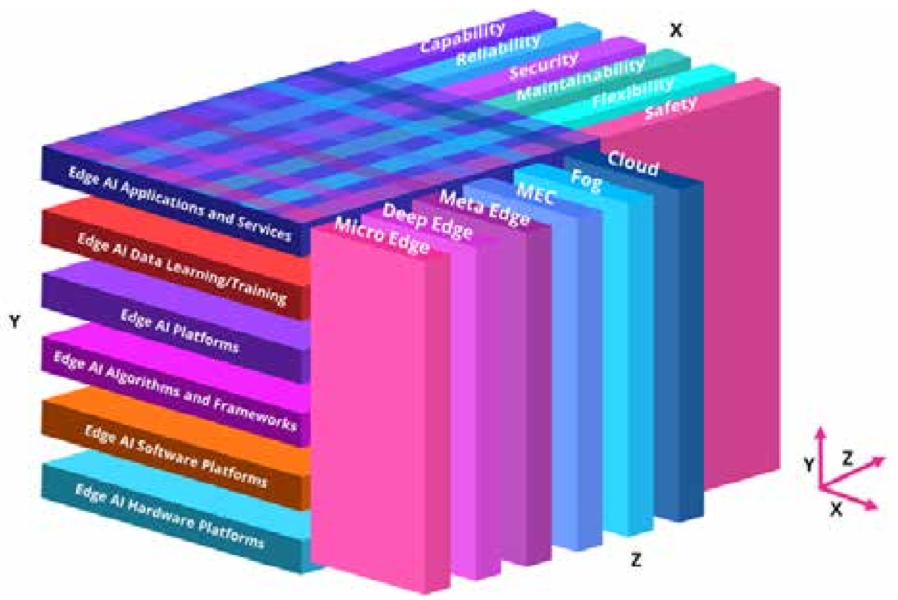

Figure 4. Graphical representation of the 3D edge AI reference architecture[19].

Discussion

The primary driver for adopting the edge AI systems reference architecture is the heterogeneity of multi-X edge AI environments, which comprise multiple edge AI systems, modalities, and edge AI agents. Modern edge AI system development has transitioned from simple, closed systems to complex system-of-systems with distributed intelligence, requiring coordinated efforts characterised by multiple sites, solutions stakeholders, and disciplines. As the scope and complexity of edge AI systems increase, so does the difficulty of maintaining coherence across distributed systems. Edge AI technologies are maturing, and a reference architecture for edge AI systems is required as the multiplicity of edge solutions reaches a critical mass. Without a shared framework, integrating edge AI systems developed across different applications and industries becomes inefficient and error-prone. The edge AI systems reference architecture addresses this by providing a common lexicon and taxonomy that enable diverse teams to communicate effectively and align their efforts toward a shared edge AI architectural vision.

The implementation of an edge AI systems reference architecture delivers value by driving and harvesting synergy across the edge AI domain. It allows identifying where shared assets and standardisation can be effectively applied and where they might be counter-productive. This strategic insight facilitates the efficient creation of edge AI solutions, edge AI product lines, and portfolios, reducing the time and cost associated with reinventing solutions for already solved problems.

The edge AI systems reference architecture can significantly improve interoperability between evolving edge AI systems. By explicitly modelling functions and qualities above the single-system level, they ensure compatibility and smoother upgrades. This leads to reduced integration costs and improved dependability of edge AI systems. The edge AI systems reference architecture can act as a baseline, a shared starting point that anchors future discussions and changes, thereby mitigating the risks associated with complex edge AI system evolution.

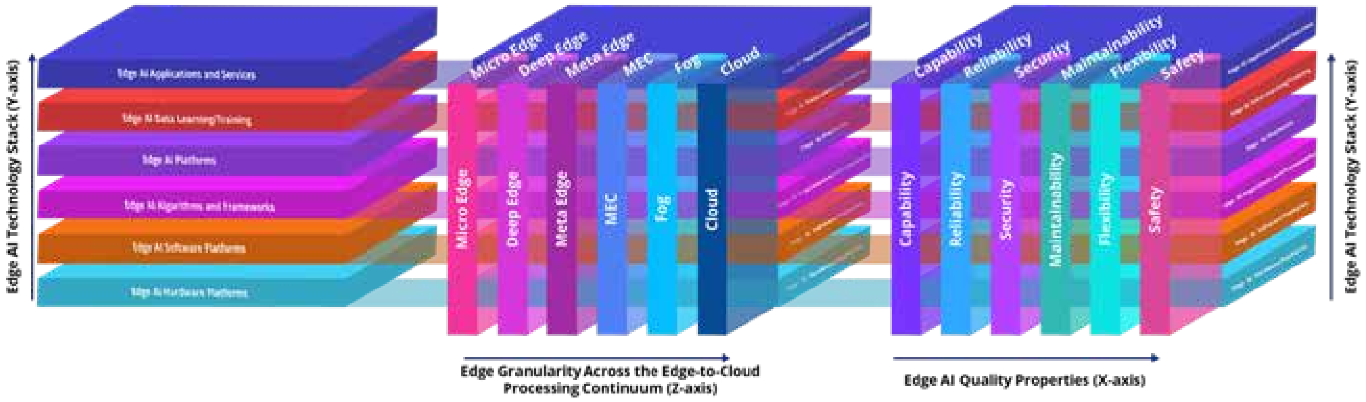

Figure 5. Edge AI systems reference architecture views unfolded.[19]

The edge AI systems reference architecture is the outcome of the Chips JU EdgeAI project and the collaborative efforts of the CLEVER, REBECCA, TRISTAN, NEUROKIT2E, LoLiPoP IoT, SMARTY, dAIEDGE and SMARTEDGE Horizon Europe (HE) projects, which together established a platform for exchanging knowledge and ideas among experts and professionals engaged in advancing AI circuits and device design, AI hardware architectures, industrial edge AI technologies, toolchains, and applications.

Discussions and exchanges during Edge AI Academy Summer School, 7-8 July, Pisa, Italy the European Conference on EDGE AI Technologies and Applications – EEAI, 20-22 October 2025 Naples, Italy and the “The Intelligent Mesh: Edge AI Technology Roadmap for Orchestrating Autonomous Systems with Agentic and Generative AI” Workshop organised at HiPEAC Conference 27 January 2026, Kraków, Poland further shaped the design and formalization of the edge AI systems reference architecture’s core elements and principal viewpoints. edge AI systems reference architecture.

1 US Department of Defense, Office of the Assistant Secretary of Defense, Networks and Information Integration (OASD/NII), Reference Architecture Description, June 2010. https://dodcio.defense.gov/Portals/0/Documents/Ref_Archi_Description_Final_v1_18Jun10.pdf.

2 D. Greefhorst and E. Proper “Architecture Principles - The Cornerstones of Enterprise Architecture.” Springer-Verlag Berlin Heidelberg 2011. ISBN 978-3-642-20278-0, e-ISBN 978-3-642-20279-7, doi: 10.1007/978-3-642-20279-7. https://sar.ac.id/stmik_ebook/prog_file_file/f7KjsxWOBJ.pdf.

3 F. Harmsen, H. Proper, and N. Kok, “Informed Governance of Enterprise Transformations.” In: Proper HA, Harmsen AF, Dietz JLG (eds) Advances in enterprise engineering II—Proceedings of the first NAF academy working conference on practice-driven research on enterprise transformations, PRET 2009. Held at CAiSE 2009, Amsterdam, The Netherlands, June 2009. Lecture notes in business information processing, vol 28. Springer, Berlin, pp 155–180. ISBN-13:9783642018589. https://www.erikproper.eu/publications/EP-2024-05-22-12-14-32.pdf.

4 J. Dietz, “Architecture – Building strategy into design.” Netherlands Architecture Forum. Academic Service – SDU, The Hague (2008), http://www.naf.nl ISBN-13: 9789012580861.

5 The Open Group. The TOGAF® 10th Edition Standard. Introduction and Core Concepts. 2022. Van Haren Publishing, 's-Hertogenbosch - NL, SBN eBook: 978 94 018 0860 6. https://www.avtechcn.com/pdf/togaf10part01.pdf.

6 Industry IoT Consortium® (IIC®). The Industrial Internet Reference Architecture. Version 1.10. https://www.iiconsortium.org/wp-content/uploads/sites/2/2022/11/IIRA-v1.10.pdf.

7 O. Vermesan et. al., “The Next Generation Internet of Things - Hyperconnectivity and Embedded Intelligence at the Edge.” In, O. Vermesan and J Bacquet, “Next Generation Internet of Things Distributed Intelligence at the Edge and Human Machine-to-Machine Cooperation”. 2018. https://www.riverpublishers.com/pdf/ebook/chapter/RP_9788770220071C3.pdf.

8 A. Willner and V. Gowtham, “Toward a Reference Architecture Model for Industrial Edge Computing,” in IEEE Communications Standards Magazine, vol. 4, no. 4, pp. 42-48, December 2020, https://www.doi.org/10.1109/MCOMSTD.001.2000007.

9 ISO/IEC/IEEE 42010:2022. Software, systems and enterprise - Architecture description https://www.iso.org/standard/74393.html.

10 ISO/IEC/IEEE 42020:2019. Software, systems and enterprise - Architecture processes. https://www.iso.org/standard/68982.html.

11 ISO/IEC/IEEE DIS 42024. Enterprise, systems and software - Architecture fundamentals. https://www.iso.org/standard/87510.html.

12 ISO/IEC/IEEE 42030:2019. Software, systems and enterprise - Architecture evaluation framework. https://www.iso.org/standard/73436.html.

13 ISO/IEC/IEEE DIS 42042. Enterprise, systems and software, Reference architectures. https://www.iso.org/standard/87310.html.

14 The Open Group Architecture Framework (TOGAF) Version 10, https://www.opengroup.org/togaf.

15 ISO/IEC/IEEE 15288:2023. Systems and software engineering - System life cycle processes. https://www.iso.org/standard/81702.html.

16 ISO 15704:2019. Enterprise modelling and architecture - Requirements for enterprise-referencing architectures and methodologies. https://www.iso.org/standard/71890.html.

17 ISO/IEC 25010:2023Systems and software engineering - Systems and software Quality Requirements and Evaluation (SQuaRE) - Product quality model. https://www.iso.org/standard/78176.html.

18 ISO/IEC 25012:2008. Software engineering - Software product Quality Requirements and Evaluation (SQuaRE) - Data quality model. https://www.iso.org/standard/35736.html.

19 Chips JU EdgeAI Project. Deliverable D5.07. Edge AI systems reference architecture. 2026.

20 O. Vermesan, “Edge AI Reference Architecture”, European Conference on EDGE AI Technologies and Applications – EEAI 2025, 22 October 2025 Naples, Italy.

21 O. Vermesan, “The Edge AI Systems Reference Architecture - Orchestrating Autonomous and AI-Defined Systems Through GenAI and Agentic AI in the Intelligence Mesh”, HiPEAC Conference 27 January 2026, Kraków, Poland.

From funded research to European impact

Why EU Projects need a new generation of exploitation infrastructure

Vicky Chatzidogiannaki
Innovation Disco

Europe funds world-class research and innovation, yet too many project results struggle to achieve sustained industrial or policy impact beyond their funding period. This article explores the structural limitations of traditional exploitation approaches in EU-funded projects and argues for a shift toward exploitation as shared digital infrastructure rather than project-level reporting. By examining how structured, data-driven exploitation systems support readiness monitoring, portfolio visibility, and cross-project synergies, the article highlights tangible benefits for project consortia, industry stakeholders, and policymakers alike. It offers a forward-looking perspective on how Europe can strengthen its innovation capacity by connecting results, ecosystems, and long-term value creation.

Europe has built one of the most ambitious public research and innovation funding systems in the world. Through large-scale programmes supporting collaborative research, thousands of high-quality technological results are generated every year across domains such as power electronics, mobility, semiconductors, health, and digital systems. Scientific excellence, international collaboration, and technological ambition are firmly established strengths of the European innovation ecosystem. Yet despite this success, a persistent paradox remains: while Europe excels at generating research results, it often struggles to transform them into sustained industrial uptake, policy influence, and long-term economic value. Too many project outcomes remain fragmented, underexploited, or disconnected from real-world deployment once the funding period ends. This gap is not the result of a lack of ideas, effort, or competence. It is structural. And addressing it requires a fundamental rethinking of how exploitation is understood, managed, and supported within EU-funded projects.

The structural challenge facing EU-funded projects

Across a wide range of collaborative projects, three recurring challenges continue to limit long-term impact. First, exploitation is frequently treated as a formal obligation rather than a strategic process. In many projects, exploitation planning is introduced late, handled through static templates, and disconnected from the evolving technical reality of the work. This leads to generic plans that fulfil reporting requirements but rarely guide real decision-making or post-project continuation.

Second, project results remain isolated. Key Exploitable Results (KERs) are typically described at partner level, with limited visibility across the consortium and almost none beyond it. Opportunities for joint exploitation, convergence across projects, or alignment with industrial roadmaps are often missed simply because results are not visible or comparable.

Third, external stakeholders face limited transparency. Industry actors, policymakers, and funding bodies often lack a clear and structured view of what EU projects actually deliver, how mature those results are, and how they could be reused, scaled, or integrated into value chains and regulatory frameworks. Together, these issues reduce the return on public investment, not because projects fail, but because Europe lacks a shared exploitation infrastructure capable of managing complexity at scale.

Rethinking exploitation as infrastructure, not reporting

To unlock the full value of EU-funded research, exploitation must evolve beyond documentation and become a continuous, system-level process.
This requires a shift in perspective:

- From individual deliverables to result portfolios
- From static plans to dynamic readiness monitoring
- From project closure to post-project continuity
- From isolated consortia to ecosystem-level visibility

In practical terms, this means treating exploitation as infrastructure, a set of shared tools, methods, and data structures that accompany projects throughout their lifecycle and connect them to the wider European innovation landscape.

The role of the digital exploitation platform

The Digital Exploitation Platform represents a concrete response to this challenge. Rather than replacing human expertise or strategic judgement, it provides the structural backbone that enables informed decision-making, coordination, and long-term planning.

At its core, such a platform support EU projects in four critical ways.

Structured mapping of project results

A first step toward effective exploitation is clarity. This Digital Exploitation Platform enables projects to systematically capture and describe their Key Exploitable Results using a shared structure, covering aspects such as:

- Technical scope and application domains
- Technology Readiness Level (TRL)
- Market Readiness Level (MRL)
- IP protection & Commercialization plans, Joint exploitation developments, Data driven prioritization process and critical timelines

This structured approach creates a common language between technical teams, project coordinators, industry stakeholders, and policymakers. It also allows results to be understood not only within a project, but across projects and programmes.

The exploitation radar: visualising uptake readiness over time

The Exploitation Radar provides a dynamic visual representation of how project results evolve from development to real-world uptake. Each Key Exploitable Result (KER) is positioned within a circular timeline that reflects its expected pathway toward adoption, from the project lifetime through one, two, and up to five years after project completion. This structure allows users to immediately understand not only the maturity of individual results, but also their anticipated time horizon for impact.

The radar also differentiates results by organisational profile, such as SMEs, Large Enterprises, and Research and Technology Organisations, offering insight into how different actors contribute to the innovation pipeline and where exploitation momentum is concentrated. By clustering results spatially and temporally, the radar enables coordinators, partners, and external stakeholders to identify priority assets, monitor progress, and detect opportunities for collaboration or acceleration.

Rather than presenting exploitation as a static list of outcomes, the radar frames it as a living ecosystem of evolving results, making readiness, ownership, and expected impact visible at a glance.

Continuous readiness monitoring

Exploitation is not a single moment in time. Results evolve, mature, and change relevance as projects progress and external conditions shift.

The Digital Exploitation Platform allow readiness to be monitored dynamically by tracking:

- Progress toward deployment
- Remaining technical, regulatory, or market barriers
- Dependencies on complementary technologies or standards

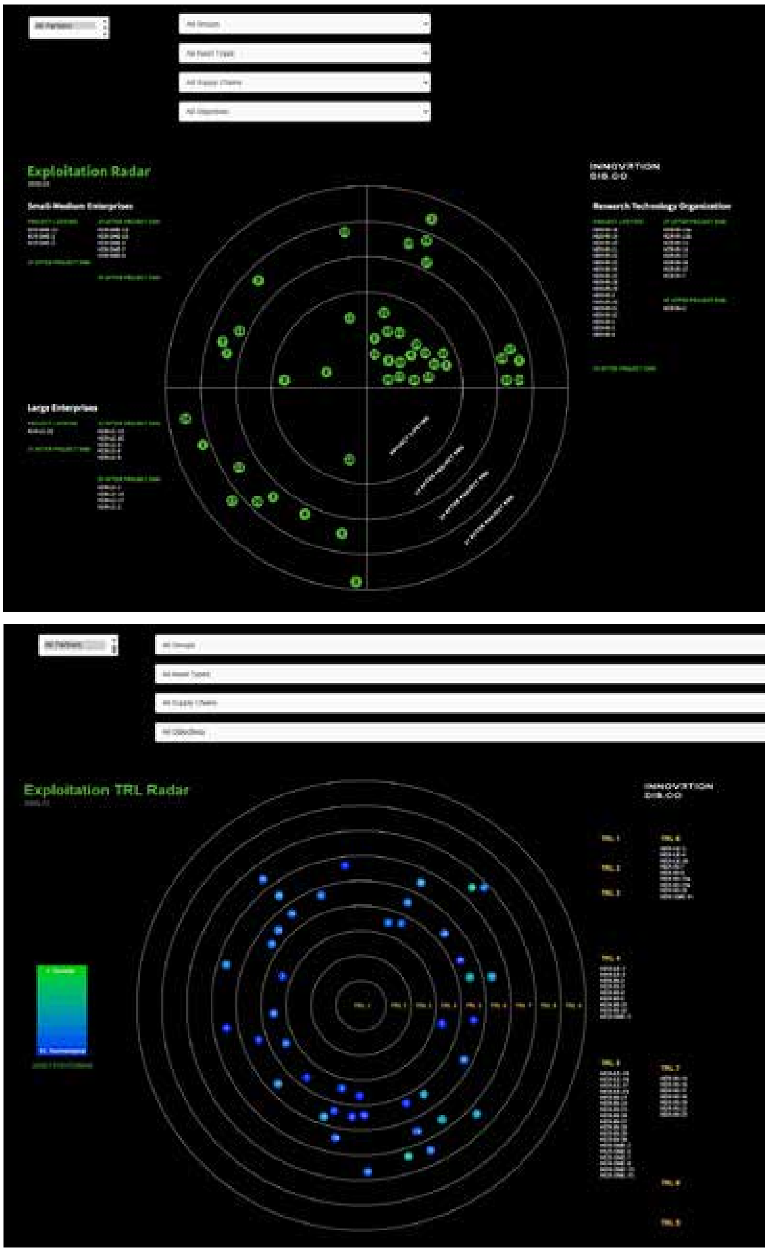

Figure 1: Exploitation Radar at a glance

This enables consortia to prioritise effort where it matters most, rather than spreading attention evenly across all results regardless of maturity or potential impact. For policymakers and funding bodies, such monitoring provides valuable insight into where public investment is generating deployable value, and where additional support may be needed.

From projects to portfolios

One of the most significant advantages of structured exploitation data is the ability to move beyond individual projects toward portfolio-level thinking.
When results are described using comparable indicators, it becomes possible to:

- Identify complementary KERs across different projects
- Detect opportunities for joint exploitation or consolidation
- Align research outputs with industrial roadmaps and policy objectives

This portfolio perspective is particularly important for large-scale initiatives, where the true value lies not in isolated outcomes, but in their combined contribution to European value chains.

Improving visibility for industry and policymakers

For industry actors, navigating the EU project landscape can be challenging. Promising technologies may exist, but discovering them, assessing their maturity, and identifying ownership often requires significant effort. The Digital Exploitation Platform lowers these barriers by providing:

- Clear overviews of available results
- Transparent readiness indicators
- Entry points for dialogue and collaboration

A key element of this transparency is the public Exploitation Map, accessible through the project website. The map provides an interactive geographical overview of project partners and their exploitable assets, allowing external stakeholders to explore where innovation is being developed and how it is progressing toward real-world uptake.

By selecting a region or organisation, users can access non-confidential information about the partner's Key Exploitable Results (KERs), technology readiness levels, positioning, and expected timelines for deployment. This ensures that industry actors, investors, and potential collaborators can easily identify relevant technologies and engagement opportunities, while fully respecting confidentiality and intellectual property boundaries.

For policymakers, the benefits are equally significant. Structured exploitation data supports:

- Evidence-based policy design
- Strategic alignment of future funding calls
- Monitoring of programme-level impact

The public map further strengthens this perspective by offering a spatial and organisational view of innovation activity across Europe, helping decision-makers understand how publicly funded research translates into regional capabilities, industrial potential, and long-term strategic value.

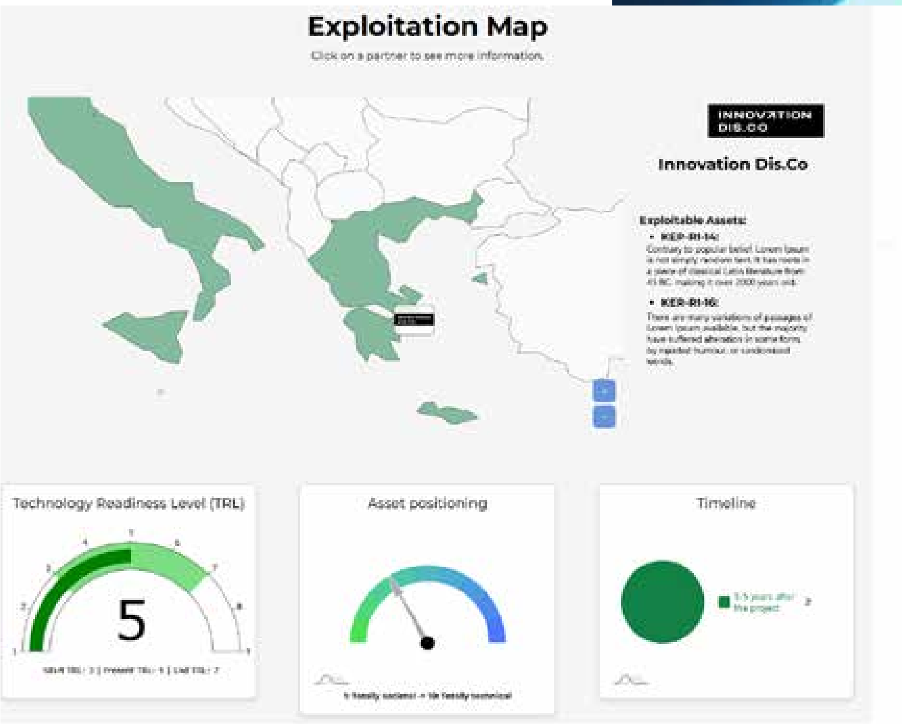

Figure 2: Exploitation Map

In this way, the Exploitation Platform acts as an impact observatory, bridging the gap between funded research and European strategic objectives, while making project results visible, accessible, and understandable to the wider innovation ecosystem.

Lessons from practical application

Experience across multiple EU-funded projects (PowerizeD, ShapeFuture, Mosaic, EcoMobility, Shift2SDV) shows that structured exploitation approaches deliver tangible improvements.

Projects that adopt digital exploitation infrastructures report:

- Clearer prioritisation of results
- More productive exploitation workshops
- Stronger alignment between technical and impact work
- Improved confidence during reviews and evaluations

Most importantly, they are better positioned to ensure that results do not disappear at project end, but continue to evolve within the European innovation ecosystem.

Towards a shared European exploitation language

Europe does not lack innovation. What it increasingly needs is a shared exploitation language, a way to describe, assess, and connect results across projects, sectors, and programmes.

Digital exploitation platforms provide the foundation for this language by:

- Standardising how results are described
- Enabling comparison and aggregation
- Supporting long-term learning across funding cycles

They do not replace creativity or strategic judgement. They make them scalable.

Conclusion: strengthening impact where it matters most

As EU-funded projects grow in ambition and complexity, the way exploitation is managed must evolve accordingly. The next generation of European impact will not come from more reporting, but from better systems. By treating exploitation as infrastructure rather than an afterthought, EU projects can:

By embedding structured, data-driven exploitation approaches within the research and innovation process, EU projects can increase the durability and continuity of their results while strengthening Europe's industrial competitiveness and technological autonomy. This ensures that publicly funded innovation contributes coherently to long-term strategic value creation.

Turning research into impact is not a final step. It is a continuous process, and Europe now has the opportunity to support it with the structures it deserves.

Vicky Chatzidogiannaki
Partner, Innovation Dis.Co

Vicky Chatzidogiannaki is Partner at Innovation Dis.Co, where she leads strategic impact, communication and exploitation activities for EU-funded research and innovation projects. She works closely with project consortia, industry and policymakers to help teams navigate complexity, align innovation with real-world needs, and translate research outcomes into sustainable European impact.

From roadmap to reality

What EFECS 2025 actually delivered

Maha Karim-Hosselet
INSIDE Industry Association

EFECS 2025 marked a visible turning point for Europe's electronic components and systems (ECS) community. The conversation has evolved beyond strategic declarations and policy ambition; the ecosystem is now entering a phase of tangible delivery.

Under the banner *"Accelerate Innovation: Building European Competitiveness,"* the European Forum for Electronic Components and Systems gathered more than 700 participants in Malta on 3–4 December, bringing together industry leaders, policymakers, researchers and innovators. The objective was clear: *showcase measurable progress under the European Chips Act*, present concrete project outcomes, and align the community around priorities for 2026 and beyond.

Setting the geopolitical frame

The opening sessions positioned semiconductors firmly within the current geopolitical and economic landscape. In his keynote address, Malta's Minister for Industry, Silvio Schembri, described semiconductors as the "new oil" of global strategic competition, a foundational resource underpinning digital sovereignty, industrial resilience and economic growth. Prime Minister Robert Abela reinforced this point, highlighting that chips account for more than one-fifth of Malta's exports and approximately 6% of national GDP.

Joining via video message, European Parliament President Roberta Metsola stressed that Europe's roughly 13% share of the global semiconductor market is insufficient in a context defined by geopolitical tensions, supply chain vulnerabilities and accelerating industrial policy worldwide. The message was consistent across speakers: ambition must now translate into competitive capability.

From strategy to action

What distinguished EFECS 2025 was its emphasis on execution.

Sessions led by the Chips Joint Undertaking (Chips JU) and industry associations AENEAS, EPoSS and INSIDE demonstrated that European research and innovation

Key messages from EFECS 2025

Embedded AI as Europe's differentiation point
Rather than competing for hyperscale cloud dominance, Europe is positioning itself around safety-critical, energy-efficient and real-time embedded intelligence — domains where reliability and system integration matter more than scale alone.

Advanced packaging and integration as enablers
Heterogeneous integration, photonics, power electronics and co-design approaches are emerging as foundational pillars of a competitive semiconductor ecosystem.

Coordination as a competitive advantage
The ECS SRIA and ecosystem alignment mechanisms are reducing fragmentation and accelerating translation from research to industrial impact.

Talent as the limiting variable
Engineers, process specialists and system architects are essential to scaling Europe's semiconductor capacity. Human capital remains the decisive factor.

Sovereignty defined by capability
Europe's strategy focuses on reinforcing strategic nodes within global value chains while preserving international cooperation.

is moving toward industrial-grade implementation. Pilot lines are operational. Competence centres have been established across all 27 Member States and Norway. Initial design platforms are emerging. Selected quantum-related initiatives are progressing toward early industrialisation.

With approximately 85% of the initial €3 billion Chips JU budget already committed, the shift from roadmap design to concrete deployment is clearly underway. The ecosystem is no longer structuring itself; it is beginning to scale.

Building strategic capabilities

Panels dedicated to strengthening Europe's technological autonomy focused on practical levers rather than abstract principles.

Patrick Bressler of Fraunhofer outlined the pilot lines targeting 2nm technologies, advanced packaging, heterogeneous system integration, FD-SOI technology, wide-bandgap materials and integrated photonics, critical technologies that form stepping stones toward competitive manufacturing capacity. These infrastructures aim not merely to demonstrate feasibility, but to reduce risk and accelerate industrial uptake.

Industry voices reinforced this applied perspective. Laurent Filipozzi of STMicroelectronics pointed to Europe's existing strengths in automation, robotics and advanced digital manufacturing systems as strategic assets. In this view, Europe's competitiveness will not rely on replicating the entire global value chain, but on reinforcing high-value nodes where it already holds structural advantages.

Policy evolution and funding alignment

EFECS 2025 also provided substantive insight into the evolution of the policy and funding framework.

Arian Zwegers from the European Commission highlighted that the ongoing review of the Chips Act, sometimes informally referred to as "Chips Act 2", must reflect shifting geopolitical realities, simplify procedures and ensure continuity for strategic projects. Stability and speed, rather than redesign, appear to be the guiding principles.

Sessions on strategic investments, future calls, and the ECS Strategic Research and Innovation Agenda (SRIA) offered early visibility into the 2026 work programmes. For many participants, this forward-looking transparency was as valuable as the project

showcases themselves: alignment reduces fragmentation and increases time-to-impact.

Talent as a structural constraint
If one cross-cutting theme dominated discussions, it was skills.

Workshops consistently underlined the need for advanced industrial competencies, semiconductor process expertise and cross-disciplinary engineering profiles. Europe's R&D excellence is widely acknowledged; translating it into sustained industrial leadership will depend on human capital.

The message was pragmatic: scaling pilot lines and competence centres will require engineers, system architects and manufacturing specialists at a pace that currently exceeds supply. Without parallel acceleration in talent development, technological capability may outpace workforce capacity.

International cooperation without illusion
Discussions on international cooperation reflected a nuanced strategic posture. Participants emphasised that Europe should avoid the pursuit of unrealistic full value-chain replication. Instead, the objective is to reinforce strategic capabilities within global innovation networks while maintaining balanced partnerships.

Sovereignty, as repeatedly framed during the forum, is about capability and resilience, not isolation.

Visibility of impact
The exhibition area and project demonstrations provided concrete proof points. Use cases in automation, heterogeneous integration and semiconductor digitalisation illustrated how European funding translates into technology roadmaps and field-level innovation. The tangible presence of hardware, design platforms and demonstrators reinforced the credibility of the policy narrative.

EFECS 2025 was significant not because it announced entirely new commitments, but because it demonstrated measurable progress on existing ones. Pilot lines are functioning. Competence hubs are operational. Industry participation is deepening. Policy instruments are evolving.

As one speaker concluded during the closing session: the era of intent has passed; Europe is now entering the era of implementation.

The remaining challenge: speed
The roadmap is articulated.
The funding is mobilised.
The ecosystem is increasingly aligned.

Competitive impact will ultimately depend on execution speed.

If 2024 was the year of structuring and 2025 the year of deployment, the coming years will determine whether Europe can convert coordination into sustained industrial leadership.

Training the next generation of CPS architects

How the CPS summer school strengthens Europe's cyber-physical expertise

As cyber-physical systems become central to mobility, industry, and energy infrastructures, Europe's competitiveness increasingly depends on system-level expertise. The CPS Summer School, traditionally held in Alghero and supported in recent years by INSIDE, has become a stable platform for interdisciplinary training, intellectual exchange, and long-term community building in advanced system design.

Cyber-physical systems (CPS), where computation, communication, and physical processes interact in real time, underpin software-defined vehicles, smart grids, robotics, aerospace platforms, and industrial automation. Designing them requires more than expertise in a single discipline. It demands architectural thinking and the ability to manage trade-offs between performance, safety, power consumption, security, and human interaction.

The CPS Summer School (cps-school.eu) was created to meet this challenge. It has been organized since 2019 by Francesca Palumbo (University of Cagliari), Christian Pilato (Politecnico di Milano), and Francesco Regazzoni (University of Amsterdam & USI Lugano). Over the years, it has become a stable European initiative focused on developing strong system-level skills.

A school with continuity

The Summer School is typically hosted in Alghero, Sardinia, which has become its natural home. Apart from one edition in Pula, it has returned to Alghero each year, strengthening both continuity and community over time.

This continuity has made the school more than just an annual training event. Many participants return in different roles, first as students, later as speakers, collaborators, or project partners. Ideas discussed during one edition often grow into research projects or industrial cooperation.

Participation is intentionally limited to around 40 students, creating a focused and interactive environment. This smaller format encourages discussion, direct access to speakers, and stronger peer connections. Over the years, the school has consistently received positive feedback, both for the technical depth of the program and for the quality of the social interactions that help build lasting professional networks.

In recent years, the school has been supported and sponsored by the INSIDE Industry Association. This reflects the strong link between advanced CPS education and Europe's broader semiconductor and systems strategy. Developing skills is not

Christian Pilato
Associate Professor at Politecnico di Milano

Francesca Palumbo
Associate Professor at University of Cagliari

Francesco Regazzoni
Associate Professor at University of Amsterdam

seen as a side activity, but as a core part of building a strong ecosystem.

Learning to think in systems

The school does not treat hardware, software, AI, and security as separate areas. Instead, it highlights how closely they depend on each other. Topics usually include system modelling, adaptive and low-power architectures, real-time constraints, safety and security in connected systems, and human-in-the-loop design. The goal is not to train narrow specialists, but professionals who can understand and manage interactions across different layers of a system.

The program combines keynote lectures, technical tutorials, interactive sessions, and a PhD workshop. Participants present their research and receive feedback from senior experts. This exchange reflects real-world CPS development, where coordination and integration are often as important as technical performance.

A distinctive element of the school is the participation of Professor Alberto Sangiovanni Vincentelli, who has been involved since the first edition. Starting from the second edition, he has delivered his now traditional “beach note.” These sessions provide broader reflections on system design methods and long-term technological trends, offering continuity across the years.

In addition, the school has welcomed many well-known figures from academia and industry, representing leading European research centres and major industrial players. Their presence keeps discussions closely connected to real challenges while maintaining strong scientific depth.

Strengthening Europe’s skills base

Over the years, the themes of the school have evolved, from CPS fundamentals to topics such as security, adaptivity, energy efficiency, and, more recently, software-defined vehicles and edge AI. However, its main goal has remained the same: helping participants understand and manage system complexity.

The PhD workshop plays an important role in this effort. It gives early-stage researchers a space to present ideas, receive feedback, and connect their work to wider European research activities. The networks built during the school often continue well beyond the event itself.

Three broader lessons can be drawn. First, system-level thinking must be developed early. CPS complexity cannot be addressed only at later stages. Second, building a community supports innovation. The repeated meetings in Alghero have created a stable reference point within Europe’s CPS landscape. Third, education is closely connected to technological sovereignty. Europe’s competitiveness in key technologies depends not only on industrial capacity, but also on engineers who can design integrated, reliable, and secure systems.

The CPS Summer School is not a promotional showcase. It is a long-term effort to strengthen Europe’s expertise in cyber-physical systems. Through continuity of location, committed organizers, recognized intellectual leadership, broad expert participation, and the support of INSIDE, it makes a steady contribution to Europe’s technological future.

INSIDE
Industry Association

Summer School sponsored by INSIDE

From IPCEI to industrial scale

Building European digital sovereignty with RISC-V and full-stack systems

Ivan Rodero
Chief Innovation Officer, Openchip

Filippos Perdikos
Innovation Scientist Engineer, Openchip

Guillermo Talavera
Senior Scientist Engineer, Openchip

In Europe, digital sovereignty is often discussed in political terms. In practice, however, it is fundamentally an engineering and industrial challenge. Modern digital systems are highly complex. Artificial intelligence, high-performance computing (HPC), cloud and edge infrastructures are built on deep technology stacks: silicon, advanced packaging, interconnect, firmware, operating systems, toolchains, runtimes, and application frameworks. When large portions of this stack are designed and controlled outside Europe, European industries depend on external roadmaps, external supply chains, and standards defined elsewhere.

This dependency is tangible. It affects procurement cycles, lead times for advanced components, the ability to customize systems for European requirements such as energy efficiency, safety, and data protection, and even the pace of innovation when access to critical IP or development tools is constrained.

In practical terms, digital sovereignty means that Europe can design, integrate, validate, and deploy critical digital systems at industrial scale. This extends beyond manufacturing capacity. It includes design capabilities, system integration expertise, software maturity, and the ability to translate prototypes into robust industrial products.

What IPCEI changes for European microelectronics

The Important Project of Common European Interest on Microelectronics and Communication Technologies (IPCEI ME/CT) represents a structural shift because it approaches the challenge as a value-chain transformation rather than a collection of isolated projects.

The most significant contribution of IPCEI is the integration of design, IP development, manufacturing, packaging, systems engineering, and software within a coherent industrial framework. Innovation in microelectronics does not occur in isolation.

Performance, energy efficiency, security, and reliability emerge from co-design across hardware and software layers. IPCEI creates the conditions to operate at this level of integration at European scale. Risk sharing is another decisive element. Advanced silicon, accelerator architectures, and system platforms involve substantial technical and financial uncertainty. Private capital alone is often insufficient, particularly in early industrial phases when markets are not yet mature. IPCEI helps bridge this gap, enabling European companies to pursue technically ambitious developments that would otherwise be difficult to justify.

Equally important is the link between innovation and first industrial deployment. Europe has historically generated strong research outcomes, yet many fail to cross the gap between laboratory and factory. IPCEI places industrialization at the center, connecting research results with production environments and market entry.

What this means in practice for a company like Openchip

Openchip operates as a fabless, full-stack design house, developing high-performance accelerators and complete systems for AI and HPC. Our work spans silicon architecture, accelerator design, chiplet and packaging strategies, firmware and low-level software, system software and SDKs, and vertical integration into deployable platforms.

In this context, every architectural decision has system-level implications. Memory hierarchy influences compiler design. Packaging choices affect thermals and power delivery. Instruction set architecture (ISA) decisions shape toolchains and application performance. Security mechanisms must be coherently implemented from silicon to software.

IPCEI enables alignment within a European ecosystem where design houses, technology providers, research centers, and industrial users converge around concrete system objectives: performance, energy efficiency, sustainability, safety, and long-term maintainability.

It also fosters the development of European system-level expertise. Many current system architectures are optimized around platforms developed outside Europe. Building indigenous stacks requires architects capable of understanding silicon, software, and workload behavior simultaneously. This competence becomes a strategic asset retained within Europe.

RISC-V as a practical sovereignty tool

RISC-V is frequently discussed in abstract terms. From an engineering perspective, it is a pragmatic architectural choice and provides control over the ISA roadmap, flexibility to implement domain-specific extensions, and the ability to develop European IP without structural vendor lock-in.

In accelerator and HPC environments, this flexibility is critical. It enables the design of vector units, memory subsystems, and control logic tailored to real workloads. Security features can be integrated at ISA and microarchitectural levels. The hardware–software interface can be optimized for performance per watt, not merely peak performance.

However, RISC-V alone does not guarantee sovereignty. The broader toolchain, including compilers, runtimes, debuggers, and system software, is equally essential. The value of RISC-V is realized only when hardware and software evolve in concert within a coherent full-stack strategy.

Within the IPCEI framework, investment in RISC-V-based platforms becomes part of a coordinated European industrial roadmap rather than isolated experimentation. This is what enables the emergence of sustainable products and ecosystems around open architectures.

From IPCEI to HPC

The DARE program plays a critical role in the high-performance computing domain by

demonstrating how European accelerators and system components can transition from research prototypes to operational HPC deployments.

For Openchip, DARE provides a concrete validation pathway: vector accelerators, advanced system integration, and HPC software stacks deployed within European supercomputing environments. It connects architectural innovation with real workloads and procurement cycles.

DARE is most effective when embedded within a broader industrial trajectory. IPCEI establishes the foundational industrial capabilities, design capacity, manufacturing linkages, packaging technologies, and system integration expertise. DARE then validates these capabilities in demanding HPC scenarios. Both layers are necessary: one builds structural capacity, the other tests it under high-performance conditions.

Why continuity of investment matters for industrial scale

In deep-tech and semiconductor industries, impact derives from continuity rather than isolated success. Hardware and system technologies evolve across long development cycles. An initial prototype marks progress, but industrial value emerges after multiple generations of silicon, progressively mature software stacks, large-scale validation, and integration into customer environments. This iterative refinement is intrinsic to robust platform development.

Sustained investment allows promising technologies to mature into stable industrial products. It builds trust among users and partners and enables long-term competence accumulation. Performance, efficiency, and reliability improve most significantly across successive iterations.

Public instruments such as IPCEI provide the structural framework to connect innovation with industrial deployment. When complemented by pilot deployments, early procurement mechanisms, and long-term industrial partnerships, advanced technologies can evolve into scalable platforms.

Digital sovereignty is not built in a single cycle. It is constructed through consistent industrial progress over extended time horizons.

Innovation and people: the strategic differentiator

Technology stacks can be engineered. System-level expertise must be cultivated. Europe faces a shortage of experienced system architects capable of integrating silicon, packaging, software, and workloads within a unified design approach. Specialization is natural, but full-stack system competence emerges only through hands-on integration across disciplines.

For this reason, structured innovation programs, targeted technical certifications, industry–academia collaboration, and real-world system projects are strategic investments rather than optional initiatives. Without engineers capable of designing and integrating complete systems, sovereignty remains theoretical.

Innovation also requires a culture that tolerates technical risk and iterative refinement. The IPCEI framework supports this by distributing risk and legitimizing learning across development cycles.

Closing: what we believe Europe can still do right

We believe Europe has a real opportunity in system-level digital sovereignty. The pieces exist: strong research and engineering talent, industrial players across the value chain, open technologies like RISC-V, and instruments like IPCEI that align technology, industry, and public interest.

The challenge is execution over time. Sovereignty is not achieved by one project or one product. It is achieved by building platforms, ecosystems, and skills that can evolve.

From our position in the innovation department, we see both the difficulty and the potential. The work is hard. Integration across hardware and software is complex. Scaling from prototype to product is slow. But this is exactly where Europe can differentiate: by building deep system competence and by aligning innovation with long-term industrial goals.

If we stay consistent, invest beyond the first success, and keep system integration at the center, Europe can build digital sovereignty in practice. Not as a slogan, but as working systems deployed at scale.

ECS BROKERAGE EVENT 2026

Thank you for your presence!

Stay connected?
Subscribe to our newsletter

New Member Focus

Semilab

Finding the right metrology is crucial

Pioneering precision to meet the future needs of the semiconductor industry

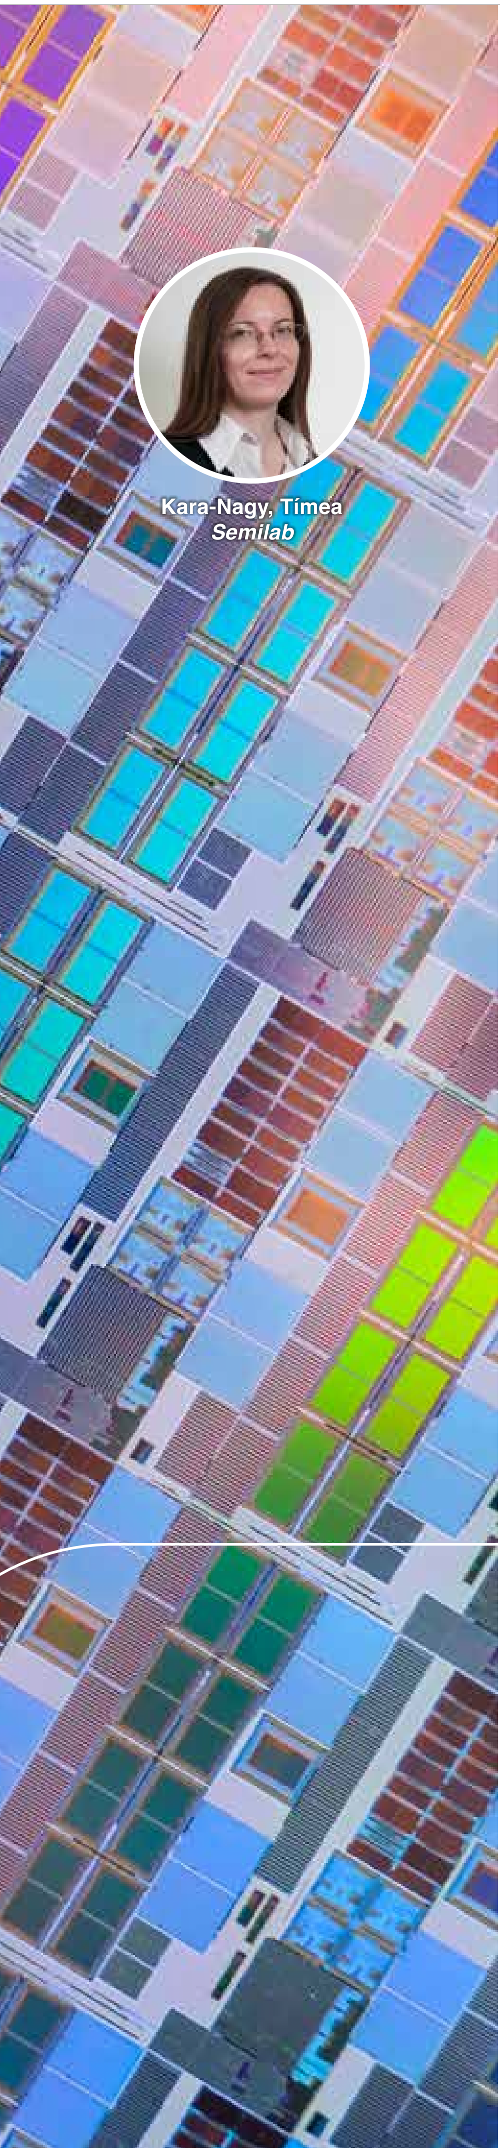

Kara-Nagy, Tímea
Semilab

In an industry where nanometer-scale precision defines success, advanced semiconductor metrology plays a critical role in enabling technological progress across electronics, mobility, energy, and digital infrastructure. The demand for continuous product development requires precisely controlled manufacturing processes that can be established by rapid and accurate metrology solutions - optimizing productivity while minimizing resources and risks. Through our state-of-the-art measuring systems and research work, Semilab contributes to the technological progress of the world.

Semilab in a nutshell: Academic excellence to the industry

SEMILAB is a strategic supplier of state-of-the-art metrology solutions to leading wafer manufacturers, IC device makers in the More-than-Moore market segment, solar and display industries worldwide, based in Hungary with R&D and product centers and a network of representatives in Europe, the USA and across Asia.

What began as a small, research-driven company, founded in 1989's Hungary, has evolved into an essential global supplier of measurement and characterization solutions, while also actively contributing to collaborative EU-funded research projects that shape the future of European and global semiconductor innovation. We have inherited a strong academic background from the founders of Semilab – who were physicists and engineers from the academic research and were finding the path for their innovations in the semiconductor industry. This still inspires the growth of our colleagues' scientific expertise and helps to implement our 30+ years of know-how into building measurement equipment.

Innovation in the DNA

Semilab's journey reflects how deep scientific expertise, continuous innovation, and strategic growth can transform the vision of a handful of scientists into a key player in the worldwide high-tech ecosystem. Today – because of the extensive amount of R&D work and the acquisition of suitable assets and technologies - Semilab designs, produces, and sells metrology solutions for the characterization of semiconductor and photovoltaic materials, for monitoring the manufacturing process of semiconductor devices, flat panel displays and solar cells, and for R&D purposes in these areas.

Our own unique inventions and original technology research keep us continuously developing our products and making innovation the basis of our metrology expertise.

Semilab's growing portfolio offers more than 300 configurable metrology solutions of 45 product lines, based on optical and electrical measurement technologies for testing, verification and monitoring of various materials. As an innovation hub, working with over 200 physicists, Semilab is ready to join the scientific movement bringing forward the European semiconductor scene, by utilizing its long-term scientific research and RDI in projects aiming to elevate the power electronics, photonics, 3D-RAM, advanced logic (CFET), heterogeneous integration or advanced packaging and 5G/6G markets.

European leader in metrology & inspection equipment

Semilab is committed to enabling stable and long-term operation of manufacturing or quality control processes of its industry partners, whether achieved through mature technology or a unique development.

Semilab offers an extensive metrology portfolio covering key monitoring steps in semiconductor manufacturing, including SRP for spreading resistivity profiling, WT/DLTS for purity and bulk contamination checks, and EIR for oxygen and carbon content measurements. Semilab positions its automated atomic force microscopy

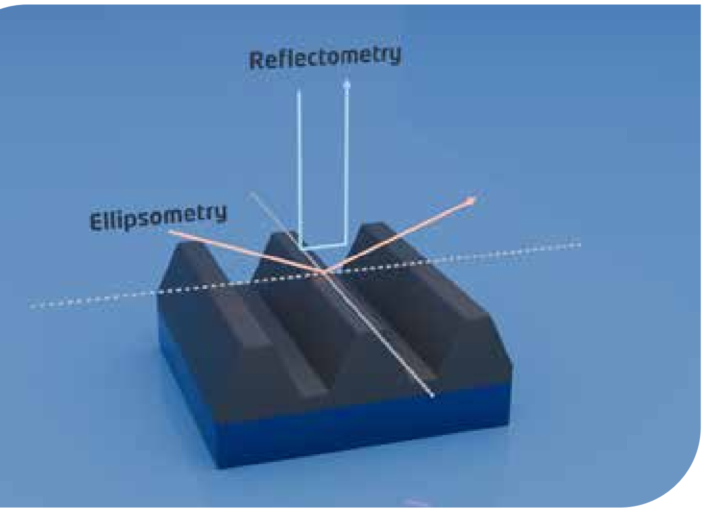

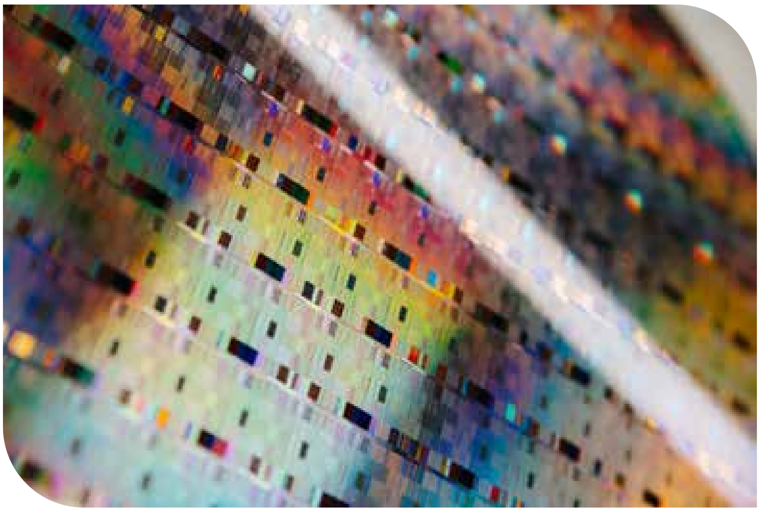

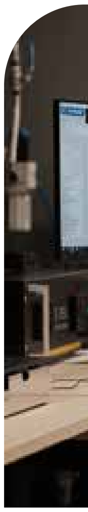

(AFM) system also to the wafer maker market to characterize surface roughness of the produced bare wafers. Due to various acquisitions, the company holds a strong market share in epi wafer characterization, from resistivity and doping measurements using mercury contact CV (MCV), air-gap CV (ACV), QC, and SRP, to epi layer thickness with large-spot FTIR (EIR), Model-Based Dimensions (MBD) for geometric characterization of periodic structures such as SiC trenches, and Raman spectroscopy for precise stress measurement and crystalline quality evaluation of advanced Si-based structures including SiGe, sSOI, and GeOI, enabling accurate process control in next-generation logic device manufacturing.

Semilab also offers metrology solutions for epi defect characterization via polarized stress imaging (PSI), lifetime measurements using WT, and ion implantation dose metrology before activation and annealing with photomodulated reflectivity (PMR). Its photoluminescence-based techniques, including the EnVision (Enhanced Vision) system, detect buried defects such as dislocations, oxygen precipitates, or stacking faults that are invisible to conventional optical inspection. Supporting these capabilities, Semilab operates in-house cleanroom and lab testing facilities for joint R&D&I projects and offers university courses and training directly in its labs.

Experienced with EU projects

Semilab places strong emphasis on close cooperation with customers, industrial partners, and leading European research institutions to jointly address emerging technological challenges. The company is an active participant in numerous international R&D collaborations and EU-funded projects, contributing to the advancement of next-generation semiconductor and photovoltaic technologies, including silicon (Si), silicon carbide (SiC), and advanced PV applications. Its involvement spans initiatives focused on upgrading manufacturing processes, strengthening the European advanced packaging ecosystem, and driving progress in equipment and materials development, packaging design, simulation, testing, metrology, and failure analysis. In addition, Semilab contributes to projects centered on Industry 4.0 solutions, digital nano-characterization technologies, and innovative image sensor developments that enable the future of 3D imaging.

Semilab's commitment to collaborative innovation is further reflected in its ongoing partnerships with leading European research institutes. In cooperation with Fraunhofer IISB, the company has established a joint research laboratory focused on wide bandgap materials, including SiC, GaN, and AlN, supporting next-generation power and high-performance semiconductor technologies.

In its joint work with Imec, Semilab contributes to advanced logic and memory development, with a strong focus on SiGe metrology using its photoluminescence-based non-contact measurement systems. Meanwhile, collaboration with CEA-Leti centers on strained SOI technologies, where the company's Raman spectroscopy solutions enable precise material characterization. In 2025 alone, these and other research collaborations resulted in 98 publications, including 27 peer-reviewed scientific articles
.

Within these collaborative frameworks, Semilab leverages its extensive metrology expertise to deliver tangible results. The company has deployed state-of-the-art measurement systems for sample monitoring, developed integrated metrology setups, advanced in-line and automated metrology solutions, and expanded the boundaries of existing technologies to create novel measurement environments. Its contributions are supported by high-resolution, high-throughput imaging-based infrared inspection systems, imaging photoluminescence tools, advanced optical metrology platforms, atomic force microscopy, and Raman spectroscopy solutions. By combining proven technologies with continuous innovation, Semilab not only supports the objectives of current projects but also contributes to shaping the future of semiconductor manufacturing. Building on its experience in collaborative research environments, the company looks forward to further opportunities for joint development, where shared expertise and coordinated innovation can continue to advance

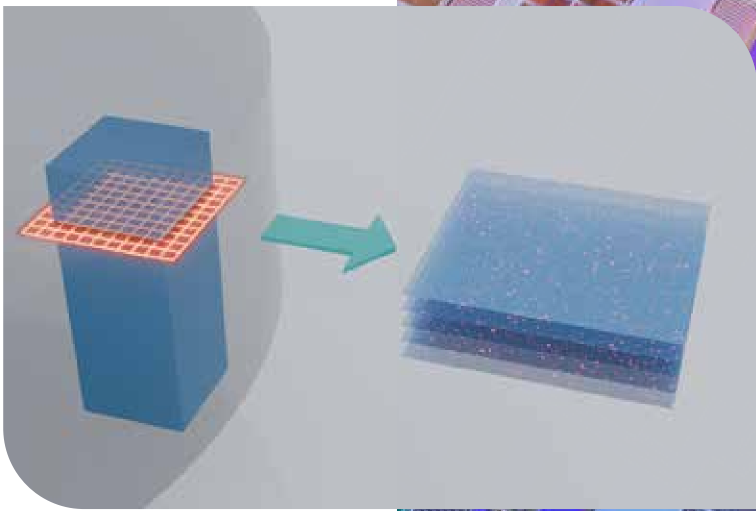

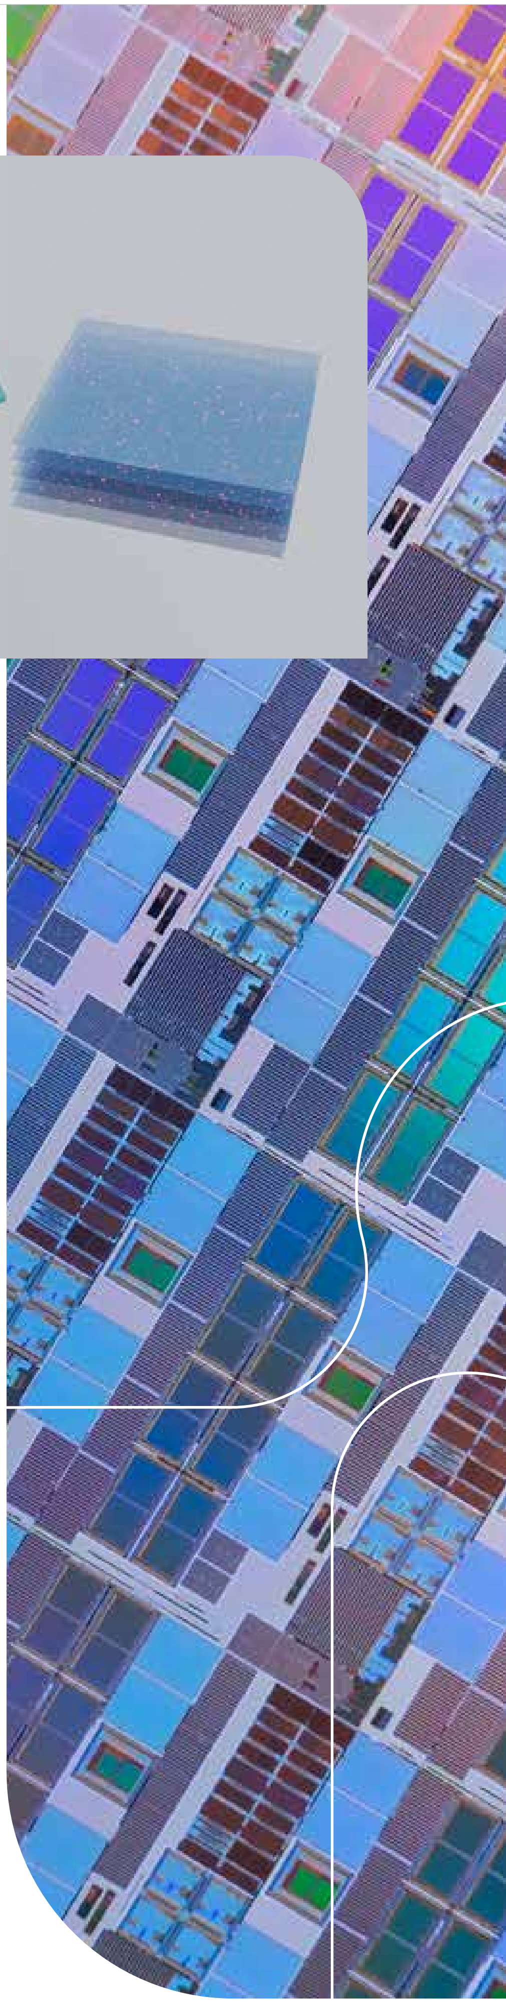

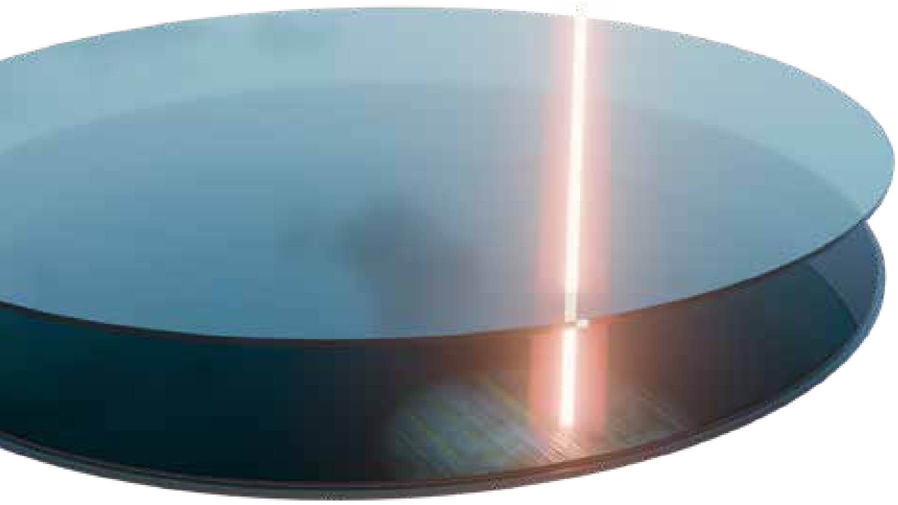

metrology capabilities and strengthen the broader high-tech ecosystem.

Partnership for future-forward metrology solutions

Looking ahead, Semilab remains focused on advancing metrology solutions that address the increasing complexity of semiconductor materials and device architectures. The company places strong emphasis on expertise, long-term commitment, and continuous research and development executed with the highest level of precision. Strategic priorities include void detection in advanced packaging and heterogeneous integration; wide bandgap material and device characterization; SiGe device analysis; photonics defect control; and quantum material characterization. In parallel, Semilab is strengthening AI-based defect classification for Si and SiGe epitaxial multilayers, enhancing industrial application know-how across critical process parameters, developing robust analytical methodologies for key epitaxial properties, and expanding characterization libraries for advanced 2D and 3D device structures.

Recognizing that shared knowledge and coordinated expertise drive Europe's technological progress, Semilab, a strategic supplier of state-of-the-art metrology solutions to leading wafer manufacturers, IC device makers in the More-than-Moore segment, as well as the solar and display industries worldwide, remains committed to collaborative innovation that addresses the semiconductor challenges of today and tomorrow.

CEI-Sphere

From ecosystem vision to operational reality in Europe's Cloud–Edge–IoT continuum

Europe's digital transformation increasingly depends on its ability to move from isolated innovation efforts to coherent, interoperable ecosystems. The recently released MetaOS Insights publication provides an important backdrop to this evolution, highlighting how early initiatives explored the concept of "Meta Operating Systems" to orchestrate distributed computing resources across cloud, edge and IoT environments. These projects demonstrated that future digital infrastructures require modular architectures, strong interoperability and cross-sector collaboration to scale beyond experimentation.

In many ways, CEI-Sphere builds on this foundation. While MetaOS initiatives focused on technological groundwork and early ecosystem experimentation, CEI-Sphere operates at the next stage: supporting Large-Scale Pilots (LSPs), aligning stakeholders, and translating technical advances into scalable market-ready ecosystems across Europe.

Launched in October 2024 under Horizon Europe, CEI-Sphere acts as a coordination and support action designed to strengthen Europe's Cloud-Edge-IoT (CEI) landscape. The project focuses on enabling collaboration, reducing fragmentation and helping innovation efforts converge around shared standards, architectures and market realities.

Supporting Large-Scale Pilots and scaling through open calls

A central pillar of CEI-Sphere's work is supporting two major Large-Scale Pilots, O-CEI and COP-PILOT, as they develop and validate real-world CEI solutions across multiple domains.

As these pilots move from design towards deployment and scaling, open calls have become a critical mechanism to expand participation, onboard new innovators and validate solutions in broader ecosystems. O-CEI has already completed its first open call, while COP-PILOT is preparing to launch its own, reflecting a shift from experimentation towards ecosystem expansion.

To support this transition, CEI-Sphere recently organised a dedicated webinar exploring lessons learned from open call processes and strategies for positioning

strong proposals. The discussion highlighted that success depends not only on technical excellence but also on alignment with market needs, interoperability frameworks and cross-domain collaboration.

The Hourglass Model: creating a shared language for the CEI ecosystem

One of CEI-Sphere's most visible achievements to date has been the development and promotion of the Hourglass Model, a conceptual framework designed to make complex digital ecosystems easier to understand and navigate.

The Hourglass Model functions as a shared ecosystem canvas. It maps stakeholders and technological capabilities across layers, from infrastructure and connectivity to platforms, standards and applications, highlighting the "narrow centre" where interoperability, open standards and governance mechanisms enable collaboration at scale.

Since its introduction, the model has gained traction across the European R&I community. It has been presented at events such as the Open Source Community Day in Madrid, co-located with the AIOTI Days, as well as a dedicated CEI-Sphere webinar, where projects demonstrated how it can support ecosystem mapping, interoperability analysis and strategic alignment.

Importantly, the model is designed as a dynamic tool. Speakers emphasised that it should be used early in projects to align ambitions and continuously revisited as ecosystems evolve, regulations mature and new use cases emerge.

Open source collaboration and ecosystem dialogue

The Open Source Community Day, co-organised with the Eclipse Foundation, marked an important milestone in positioning CEI-Sphere within broader European discussions around open digital infrastructure. Bringing together industrial actors, research projects and open-source communities, the event reflected Europe's ambition to build federated digital ecosystems aligned with initiatives such as the European Data Strategy and emerging discussions around a potential Eurostack.

The event also showcased how open-source frameworks and shared standards can accelerate innovation while reducing dependency on proprietary hyperscale platforms, which has become a key element of Europe's digital sovereignty debate.

Complementing this effort, CEI-Sphere also launched a first podcast talk show, expanding its outreach to new audiences and providing accessible insights into ecosystem challenges and opportunities.

Connecting CEI innovation with sectoral realities: the V2G example

CEI-Sphere's ecosystem approach is also reflected in cross-sector initiatives such as V2G Leaders Europe 2025, where energy, mobility and digital infrastructure stakeholders explored how bidirectional charging technologies can move from pilots to market deployment.

The event, which took place last November in Brussels, demonstrated that technological readiness alone is insufficient; scaling requires alignment between policy frameworks, standards, industrial strategies and user-centric business models.

Discussions involving DG ENER, DG RTD and DG CNECT highlighted how cloud-edge architectures, data spaces and interoperable platforms will underpin future energy flexibility services.

Towards living ecosystem tools and stronger community alignment

Beyond events and frameworks, CEI-Sphere is developing practical resources to support long-term ecosystem growth. A visualisation tool currently under development will map use cases, stakeholders and relationships across the CEI ecosystem, providing a dynamic overview of actors and capabilities rather than a static catalogue.

This approach reflects a broader ambition: transforming project outputs into "living resources" that evolve with stakeholder input and remain useful beyond the project lifecycle.

Looking ahead, CEI-Sphere will continue supporting Large-Scale Pilots and promoting their open calls, while strengthening collaboration across European initiatives and standardisation communities. Upcoming activities include participation in the Open Community for Research event in Brussels and further work on ecosystem visualisation and interoperability frameworks.

If MetaOS initiatives explored what distributed digital infrastructures could look like, CEI-Sphere represents the next step, turning these concepts into coordinated ecosystems capable of scaling across sectors. By connecting research, industry, policy and open communities, the project contributes to a broader European ambition: ensuring that cloud-edge-IoT technologies become not only innovative, but interoperable, competitive and strategically aligned with Europe's digital future.

Why chip autonomy matters

Carin Lagerstedt
Semicon Sweden

As the global semiconductor landscape undergoes geopolitical shifts and supply chain shocks, Sweden is betting on specialization over scale. One of the most advanced examples of this strategy in action is Axis Communications, a pioneer in secure video surveillance and a rare Swedish company with in-house ASIC (application-specific integrated circuit) development. In this interview, Patrik Lislén, Engineering Manager of ASIC Development, shares why ARTPEC, Axis' in-house chip platform, is not just a technical asset but a cornerstone of product innovation, long-term trust, and national self-reliance.

Patrik Lislén, Engineering Manager ASIC Development, Axis Communications.

"We built our success around autonomy, owning the chip lets us build better products, with fewer compromises."

Why does Axis build its own chips?

From the very first ARTPEC chip in 1999 to today's ninth generation, Axis has maintained full control over its silicon roadmap. The result: high-performance video products designed for real-time scene understanding, AI at the edge, and energy-efficient computing.

"We wanted more than just performance, we wanted control. And when you own the chip, you own the roadmap."

The ARTPEC platform allows Axis to:

- Optimize latency and image processing for real-world scenes
- Run custom AI accelerators directly in the camera
- Support secure firmware updates over long product lifecycles
- Reduce system-wide energy use and bandwidth demands

Notably, Axis' autonomy in chip design helped it avoid the worst effects of the global component shortages. With full visibility into its own supply chain, it could adapt faster than many who rely on third-party platforms.

Patrik Lislén presenting ARTPEC's role at Elektronikmässan 2025. "The most important engineering project at Axis is ARTPEC. It's where our innovation and independence intersect."

What are Sweden's strenghts in the semiconductor industry?

"We don't have local foundries but we have knowledge, which is just as important."

According to Lislén, there is now a critical mass of companies in Lund/Malmö area working at the forefront of advanced electronics, AI, and wireless systems—creating a fertile environment for semiconductor innovation in southern Sweden. Among the most notable are smaller companies like Acconeer, Alixlabs, Beamwave and Xenergic but also bigger players like Ericsson and ARM.

In addition to these Swedish-founded companies, global players are also reinforcing the region's relevance. ARM, the UK-based semiconductor IP giant, maintains a significant R&D presence in Lund, focusing on AI, security, and edge processing solutions.

"Being close to where research happens matters," says Lislén. "That's part of why we've opened offices in Stockholm and Linköping as well - to access new talent pools and connect with academic hubs."

Axis also benefits from Sweden's broader strengths: a trusted international reputation, neutral geopolitical position, and a longstanding focus on ethical governance - factors that matter to surveillance customers worldwide.

Where should Sweden go next?

Sweden's neutrality and reputation for ethical governance have also made Axis a trusted player in the global security market, crucial for video surveillance customers concerned with data integrity.

"We won't compete on volume. But we can lead on trust, quality, and autonomy. Not everyone can start developing their own chips but with strong local universities and the EU chips act the threshold can be lowered for smaller companies".

Lislén sees Sweden's future in secure, specialized, and intelligent electronics, rather than competing with Asian or American foundries on scale.

"Swedish companies have learned to work smart. We don't always have the biggest volumes, but we know how to scale smart solutions - especially in power-sensitive and mission-critical applications."

Axis remains deeply rooted in Lund, but its expansion into other regions signals a longer-term commitment to building a resilient, nationally distributed innovation ecosystem.

How do we attract more talent, and more diversity?

"When students hear 'electronics,' they imagine a guy with a yellow hardhat. We need to change that picture."

Lislén notes that hiring ASIC talent in Sweden is still a challenge - particularly in areas like: Verification, Power management, Memory interface design, System architecture.

To address this, Axis is:

- Collaborating with Universities and Industry in Research programs like ClassIC and similar
- Partnering with Universities to support thesis projects, mentorships, lectures and summer schools
- Recruiting directly from Swedish engineering programs
- Combining early-career and senior engineers in each team, enabling excellent growth possibilities.

"When I studied, we were 180 students - and only 20 were women. That kind of imbalance affects culture. Today we want to offer both support and real role models."

Axis' chip team today includes engineers from over 15 nationalities, but there's still work to do on gender balance and visibility.

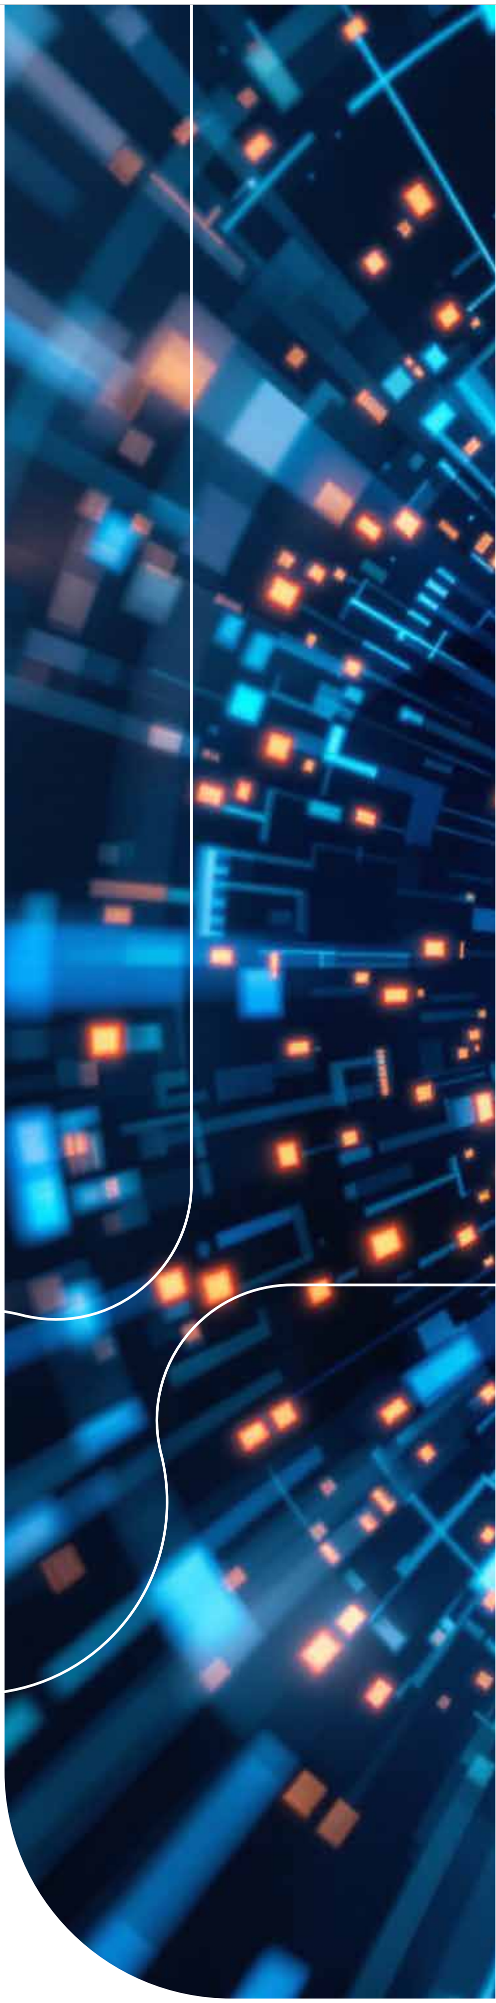

ChipsJU
EFECS
2&3 December
Helsinki
2026
SAVE
THE
DATE!
Aeneas
EPoSS.
European Association on
Smart Systems Integration
INSIDE
Industry Association

New Member Focus

Bio-inspired vision for the era of physical AI

Fabio Gallo
AI4IV s.r.l.

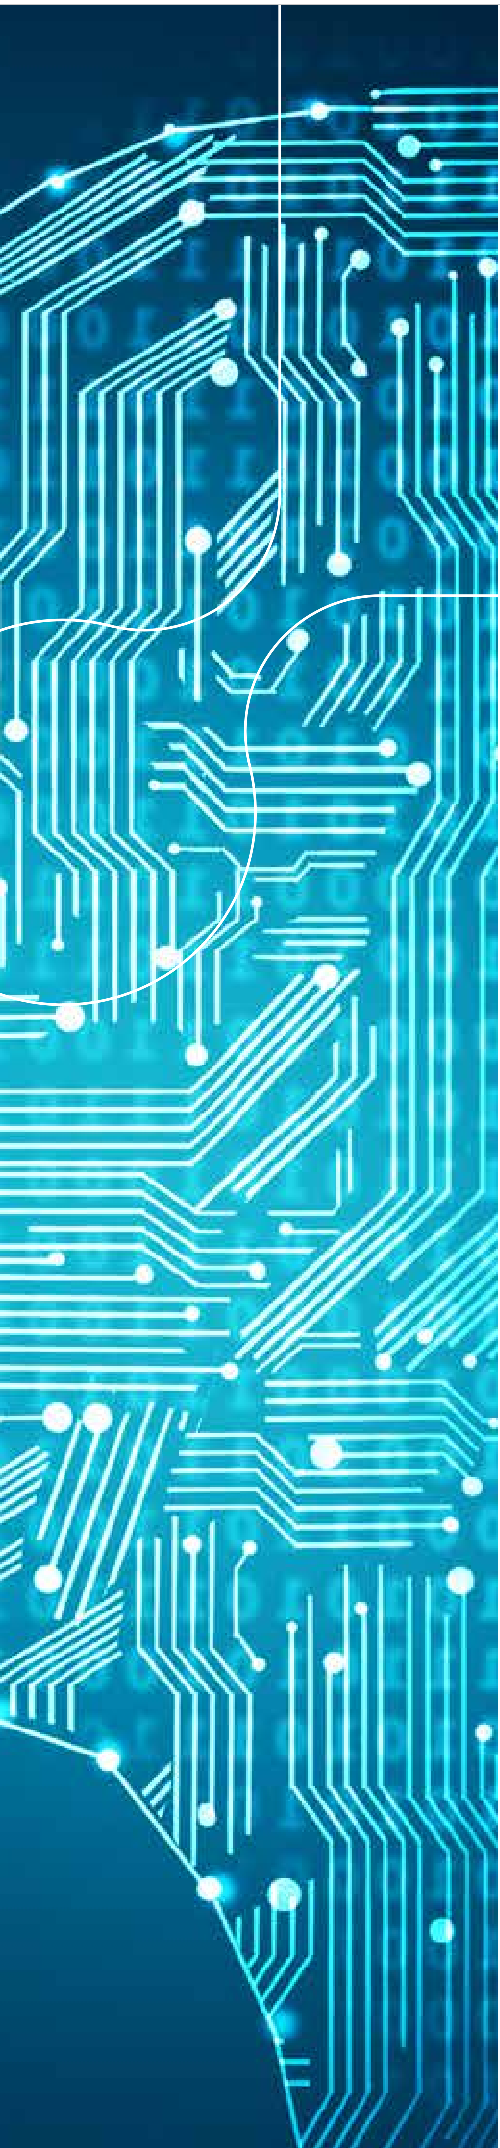

As artificial intelligence moves beyond the digital domain and into the physical world, vision is emerging as a critical layer of industrial infrastructure. Physical AI systems embed intelligence directly into machines that sense, decide, and act in real environments, making perception a decisive factor for performance, safety, and scalability. In this context, vision is no longer a supporting function, but a foundational capability upon which real-world AI systems are built.

Unlike purely digital AI systems, Physical AI must operate continuously across changing environments, lighting conditions, and operational contexts that cannot be fully anticipated at design time. Its success is therefore measured not only in accuracy benchmarks, but in robustness, uptime, and the ability to generalize reliably beyond curated datasets.

When AI enters the physical world

For decades, artificial intelligence evolved primarily within digital environments. Algorithms processed static datasets, produced probabilistic outputs, and operated at a comfortable distance from the physical world. Errors could often be corrected offline or tolerated without immediate consequences.

Today, AI is increasingly embedded in machines that operate directly in the physical world, from autonomous vehicles and industrial robots to inspection drones and intelligent devices. This transition defines the emergence of Physical AI.

In this transition, the cost of perceptual error changes fundamentally. Decisions are no longer abstract predictions but physical actions, executed in real time and often in proximity to people, infrastructure, or valuable assets. As a result, perception becomes a system-level concern rather than a modular component.

In Physical AI systems, decisions translate directly into physical actions. Misinterpretations can result in safety risks, equipment damage, or loss of trust. As a result, the quality, consistency, and reliability of sensory input become determining factors for overall system performance.

When vision becomes infrastructure

Vision serves as the primary interface between Physical AI systems and the external world. Visual data feeds automated decision chains that guide navigation, manipulation, inspection, and interaction. While other sensors provide complementary signals, vision offers the richest and most flexible representation of complex environments.

The Physical AI Pyramid

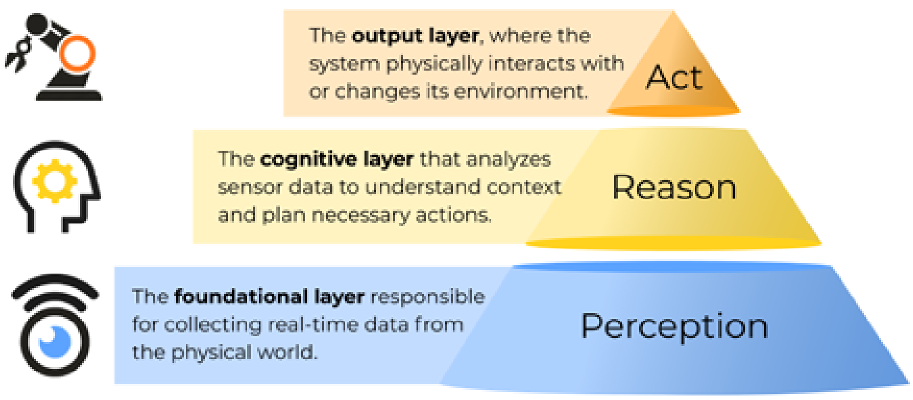

Much like communication networks or energy grids, vision defines the operational envelope of Physical AI systems. It determines where they can function, under which conditions, and with what safety margins. Limitations at the perception layer constrain the entire system, regardless of the sophistication of downstream algorithms.

Standard CMOS image sensors were designed primarily for human observers. They rely on uniform pixel arrays, global exposure parameters, and downstream image processing pipelines to reconstruct scenes. While effective for photography and video, this approach introduces structural limitations in autonomous and safety-critical systems.

In scenarios involving extreme lighting contrasts, such as a vehicle exiting a tunnel into bright sunlight, a robot operating near reflective metallic surfaces, or a surveillance camera facing low-angle glare, traditional sensors may saturate parts of the image or lose critical detail. Although such artifacts may be acceptable to human viewers, they can prevent neural networks from correctly classifying objects, detecting obstacles, or understanding scene context.

Physical AI can be understood as a layered architecture in which perception forms the foundation for reasoning and action. If perception is compromised at the base, errors propagate upward and no amount of downstream intelligence can fully compensate. Vision therefore acquires the role of industrial infrastructure, defining the operational envelope of Physical AI systems.

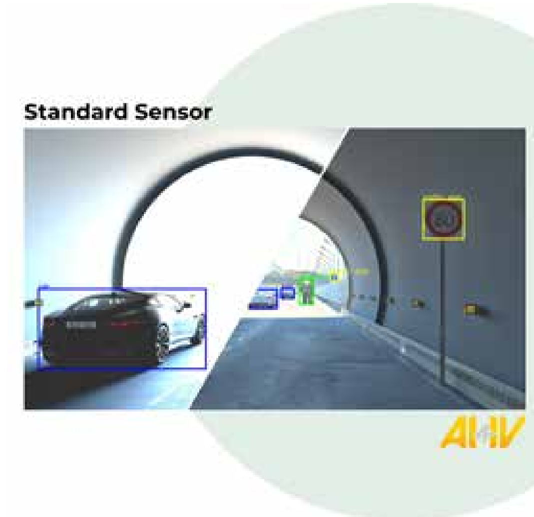

Impact across physical AI applications

The importance of robust visual perception becomes evident across a wide range of Physical AI applications. While operational contexts differ, similar sensing challenges recur across domains, highlighting the systemic nature of the problem.

Despite the diversity of these domains, they share a common constraint: perception must remain reliable under conditions that are difficult to model exhaustively. Improving sensing quality at the source therefore has a disproportionate impact, simplifying system design and enabling scalable deployment across environments.

In autonomous mobility, vision systems must cope with rapidly changing illumination, shadows, headlights, and reflective surfaces. A pedestrian partially obscured by glare or deep shadow may be missed by a perception system relying on conventional image reconstruction, leading to incorrect or delayed decisions.

In industrial robotics and logistics, visual sensing is used for object detection, pose estimation, and quality inspection. Variations in surface reflectivity, ambient lighting, or background clutter can lead to misclassification or false rejections, reducing throughput and reliability.

Inspection drones face similar challenges when operating near infrastructure such as bridges, power lines, or industrial plants, where strong contrasts between sky, structure, and shadow are common. In consumer environments, wearable devices and smart appliances increasingly rely on vision to adapt their behavior to users and surroundings.

Across these applications, advances in sensing quality have a disproportionate impact on system robustness and scalability. Improving perception at the source simplifies downstream processing and expands the set of environments in which Physical AI systems can operate reliably.

Rethinking image acquisition

Meeting the demands of Physical AI requires rethinking image acquisition at the sensor level. Incremental improvements to conventional architectures struggle to overcome fundamental limitations in dynamic range, latency, and artifact-free capture.

While advances in image signal processing and neural training have extended the usefulness of conventional sensors, they cannot fully compensate for information that is lost or distorted at the moment of capture. Addressing these limitations requires architectural change rather than further optimization of existing pipelines.

AI4IV's FlyEye technology adopts a bio-inspired approach, drawing inspiration from compound eyes found in nature. Rather than operating as a monolithic pixel array, the sensor is structured as an array of autonomous photoreceptors.

Each photoreceptor comprises a small group of pixels that independently adapts its acquisition parameters, such as sensitivity and integration time, based on local light intensity. This local adaptation enables the sensor to capture detail simultaneously in bright and dark regions of a scene, without relying on multi-exposure reconstruction or aggressive post-processing.

By performing adaptation directly at the sensing stage, FlyEye produces clean, information-rich visual data that is inherently well suited for downstream AI processing.

From sensing to edge intelligence

As Physical AI systems evolve, the boundary between sensing and computation becomes increasingly blurred. Relying on cloud-based processing introduces latency, bandwidth requirements, and security concerns that are incompatible with many real-time or safety-critical applications.

On-chip processing enables edge sensors to perform neural network inference locally, without the need for a continuous network connection. This reduces end-to-end latency, improves robustness in disconnected environments, and strengthens data security by keeping visual information at the source.

Edge intelligence allows perception and reasoning to operate as a closed loop. Inference results can influence acquisition strategies in real time, focussing sensing resources where they are most needed and improving robustness under challenging conditions. This convergence is a key enabler of autonomous, energy-efficient systems.

AI4IV addresses this challenge through Tensputing, a neural processing architecture designed for efficient tensor-based inference and compact silicon implementation. By integrating computation close to the sensor, Tensputing enables fast decision-making.

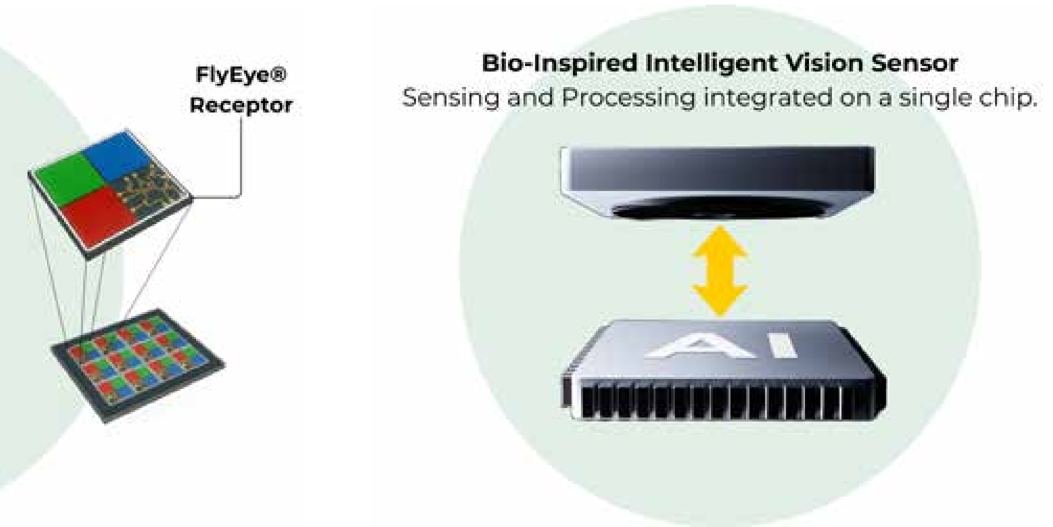

Crucially, on-chip intelligence can also steer image acquisition itself. Inference results may feed back into the sensing loop, dynamically adjusting acquisition parameters at the photoreceptor level. This tight coupling further enhances robustness.

Conclusion

As artificial intelligence continues its transition into the physical world, vision will increasingly define system performance, safety, and scalability. Bio-inspired sensor architectures combined with on-chip intelligence support a new generation of Physical AI systems.

For Europe's industrial ecosystem, advances in bio-inspired sensing and edge intelligence represent an opportunity to strengthen technological autonomy and enable new classes of intelligent systems. By rethinking perception as infrastructure, Physical AI can move from controlled demonstrations to reliable deployment at scale.

New Member Focus

PMiR Project Management in Research

Professionalising collaboration in European research and innovation

Alexandra Espinosa Hortelano
PMiR

Gunnar Widforss
PMiR

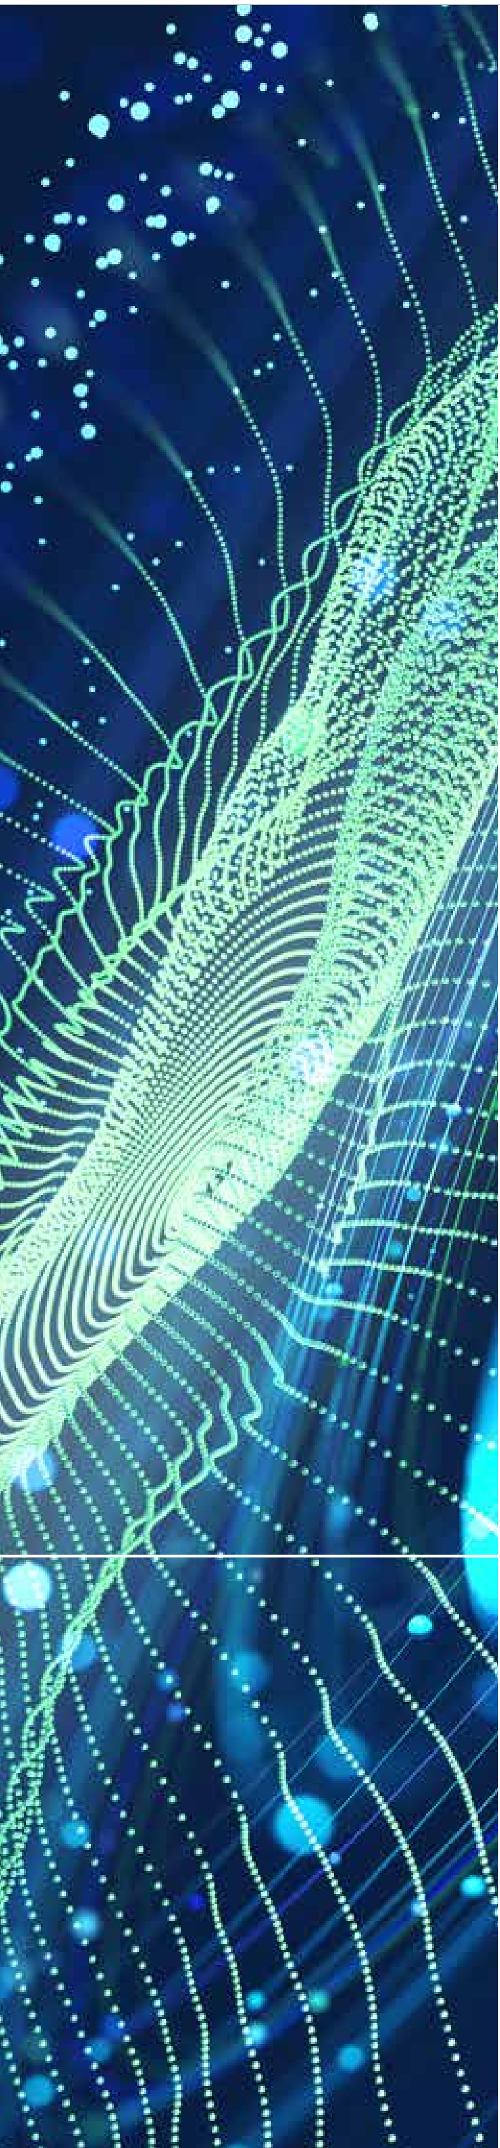

In recent years, collaborative research projects in Europe have grown not only in ambition, but also in complexity. Large consortia, multi-level funding schemes, diverse national rules, and the need to bridge academic excellence with industrial relevance place increasing demands on how research projects are conceived, coordinated, and delivered. Within this evolving landscape, Project Management in Research (PMiR) has established itself as a specialised partner dedicated to one clear mission: enabling excellent research through professional, research-driven project management.

Founded in 2016, PMiR is built on a simple but powerful conviction: we love research. This philosophy goes beyond a slogan. It reflects a long-standing engagement with research environments, academic institutions, and industrial partners, and a deep understanding that high-quality research outcomes depend not only on scientific excellence, but also on robust, agile, and trusted project coordination.

A research-driven approach to project management

PMiR is a spin-off rooted in the research ecosystem of Mälardalen University, combining experience from advanced computer science research with extensive engagement in both academia and industry. The team operates at the intersection of research, innovation, and execution, supporting public and private organisations such as universities, research institutes, and software-intensive industrial companies across Sweden and Europe.

To support this approach in practice, PMiR also works with dedicated system support for project and portfolio management. Through a collaboration with Research Unfolded Sweden AB, PMiR contributes to the use of Research On Display (ROnDi), a web-based platform designed for research-intensive organisations. ROnDi supports both formal project reporting and day-to-day portfolio follow-up, helping teams maintain overview, consistency and transparency across complex research programmes.

Unlike generic project management consultancies, PMiR focuses exclusively on research and research-related projects. This specialisation allows the team to address the specific challenges of collaborative R&D: evolving objectives, heterogeneous consortia, distributed responsibilities, and the need to align scientific goals with funding, reporting, and impact requirements. PMiR's role is not to replace researchers, but to enable them to focus on research, while ensuring that projects remain coherent, compliant, and strategically aligned throughout their lifecycle.

Two approaches for managing collaborative research

PMiR structures its work around two closely connected service models, designed to respond flexibly to the needs and maturity of different consortia:

1. **Coordination of proposals and projects**
   PMiR acts as full coordinator for collaborative research initiatives, from early idea shaping and consortium building to proposal submission and post-award implementation. This includes managing complex administrative structures, aligning partner contributions, and maintaining a clear narrative from objectives to outcomes and impact.
2. **Strategic support to coordinators**
   In many cases, PMiR supports academic

Diego Grimani, Gunnar Widforss and Kgothatso Mahlakoana at EFECS 2025 in Malta.

or industrial coordinators by reinforcing their project management capacity. This model allows coordinators to retain scientific leadership while relying on professional support for planning, governance, reporting, and stakeholder coordination. An approach that is particularly valuable in large or highly competitive European programmes.

Both models are grounded in agile project management principles, enabling easy participation even in complex projects, and fostering trust-based collaboration across diverse partner ecosystems.

Achieving excellence through experience

PMiR brings extensive experience from European collaborative programmes, including Horizon 2020, Horizon Europe, ITEA, and ECSEL. The team has contributed to and coordinated large-scale projects involving dozens of partners, significant budgets, and high technological ambition.

Building on experience from earlier large European projects such as XIVT *(details to be expanded)*, PMiR is currently a beneficiary in SPIN-CHIP, a project selected under the CHIPS Joint Undertaking, one of the most competitive funding frameworks in Europe today. The success of SPIN-CHIP reflects PMiR's ability to operate in demanding environments where strategic alignment, proposal quality, and execution capacity are decisive factors.

This continuity - from previous large collaborative initiatives to new-generation European programmes - demonstrates PMiR's long-term commitment to strengthening Europe's research and innovation capacity.

People, leadership, and professionalisation

At the core of PMiR is a team with decades of combined experience in research and project management. The company is co-founded and led by Gunnar Widforss, a certified project manager (IPMA) and expert evaluator for the European Commission. Gunnar has coordinated several large European projects, including XIVT, MegaM@Rt2, VDO and AIDOaRT, and has been actively involved in professional communities such as the European Association of Research Managers and Administrators (EARMA).

Alongside him, Lodiana Kjellin, Project Manager at PMiR, brings a background in computer science and research at Mälardalen University, supporting partners and consortia with a strong understanding of both technical and organisational dimensions of collaborative projects.

PMiR's way of working extends beyond its core team through a network of trusted collaborators who complement and reinforce its project management expertise. This includes close collaboration with Alexandra Espinosa Hortelano, specialising in communication, dissemination and exploitation, and Marina Manzanares, who is a legal advisor, specializing in IPR. Diego Grimani, contributing technical leadership and system-level expertise. This extended team model allows PMiR to assemble the right competences around each project, ensuring continuity, accountability and strategic alignment.

Through this people-centred approach, PMiR has built a strong and diverse European network spanning academia, research organisations, industry and regional actors. Long-term collaborations include universities and research centres such as Aarhus University, Bombardier Transportation Sweden AB, COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES (CEA), CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE (CNRS), Chalmers tekniska högskola AB, INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM IMEC, Microsistemas e Nanotecnologias INESC, Infineon Technologies AG, International Iberian Nanotechnology Laboratory, County Västmanland, NanOsc AB, Northvolt Revolt AB, Region Västmanland, RISE Research institutes of Sweden, Thales and Uppsala Universitet. These relationships, developed across multiple projects and programmes, underpin PMiR's capacity to build consortia, share responsibilities effectively, and operate at scale in demanding European research environments.

PMiR is also a host of interns with ERASMUS scholarships via the program EU4EU. This is a heritage from MDU where 20 interns have started their career in the research environment. Currently two talented and highly motivated interns support the activities of PMiR: Kgothatso Mahlakoana, a master student from Lille and Jorge Gerardo Lopez Aguilera, a master student from Padua.

Innovating collaboration in large consortia

One of PMiR's most distinctive contributions is its work on innovative collaboration models for large research projects. In projects such as MegaM@Rt2 and AIDOaRT, PMiR has pioneered the use of internal hackathons as a structured mechanism to enhance collaboration between industrial use-case owners, technology providers, and researchers.

This approach - documented in peer-reviewed research - has shown measurable benefits in terms of partner engagement, knowledge exchange, and early technical results. By embedding hackathons within regular project meetings, PMiR has helped transform plenary sessions from administrative checkpoints into active spaces for experimentation and co-creation, strengthening both project outcomes and consortium cohesion.

This approach to professionalising research project management has also been shared within the European research management community, including contributions presented at EARMA, where PMiR has presented practical models for distributing responsibilities and improving collaboration in large-scale consortia.

Building skills and sustainable ecosystems

Beyond individual projects, PMiR invests in **capacity building and long-term ecosystem development**. A key example is the involvement of the staff in the doctoral-level course *"Writing a Horizon Europe Research Proposal"* (7.5 ECTS), delivered at Mälardalen University, also given at Master's level at Sapienza University of Rome The course equips PhD candidates with practical skills to lead and coordinate European research proposals, reinforcing the next generation of research leaders.

PMiR's ecosystem approach is further strengthened through strategic partnerships, notably with **Blue Ocean Projects** (communication, dissemination and exploitation) and **Innovation River** (technical leadership). Together, these partners form a triangular model that addresses three essential dimensions of successful European projects: **professional coordination, technical excellence, and impact-oriented communication and exploitation, law and IPR included.** .

PMiR within the INSIDE ecosystem

As a member of the INSIDE Industry Association, PMiR contributes a perspective grounded in practice, experience, and reflection on how Europe can better organise and manage its research ambitions. By professionalising project management without losing sight of scientific creativity, PMiR helps bridge the gap between research potential and real-world impact.

In an era where collaboration is both a necessity and a challenge, PMiR stands for clarity, trust, and long-term commitment to research-driven innovation - turning complexity into a shared opportunity for Europe's innovation ecosystem.

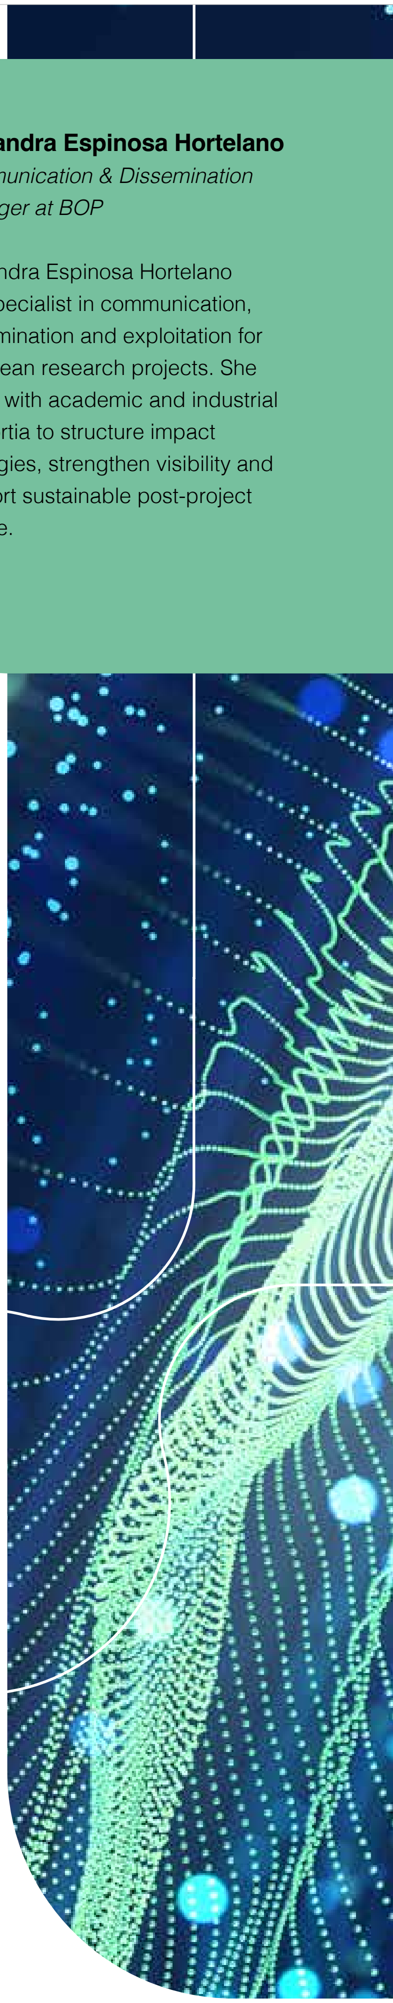

Alexandra Espinosa Hortelano
Communication & Dissemination Manager at BOP

Alexandra Espinosa Hortelano is a specialist in communication, dissemination and exploitation for European research projects. She works with academic and industrial consortia to structure impact strategies, strengthen visibility and support sustainable post-project uptake.

Innovation Spotlight

A new class of non-volatile memory for energy-efficient Edge AI

How emerging memory architectures can reshape system design at the edge

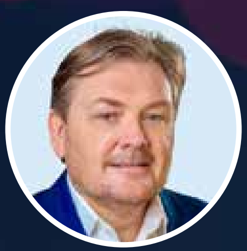

James Ashforth-Pok
Quinas

Edge AI systems are increasingly constrained by the energy, latency, and memory bottlenecks of conventional computing architectures. As workloads move closer to the sensor and away from the data centre, memory technologies play a decisive role in determining system efficiency and scalability. This article introduces a novel non-volatile memory approach that combines fast access, ultra-low energy operation, and long data retention within a single device architecture. We discuss the technological principles behind this innovation, its relevance for emerging AI and embedded systems, and its potential contribution to Europe's semiconductor and deep-tech ecosystem.

Introduction – why memory matters at the edge

Artificial intelligence is rapidly moving out of the data centre and into the real world. From smart sensors and autonomous devices to industrial control systems and secure embedded platforms, intelligence is increasingly deployed at the edge, close to where data is generated and decisions must be made. This shift promises lower latency, improved privacy, and greater system resilience. However, it also exposes fundamental limitations in today's computing architectures.

At the heart of these limitations lies memory. While advances in processors and accelerators have enabled impressive gains in computational throughput, memory technologies have evolved more incrementally. As a result, energy consumption, data movement, and memory access latency have become dominant constraints on system performance, particularly in power- and thermally-constrained edge environments.

In conventional architectures, data must constantly shuttle between separate memory and compute units. This movement consumes far more energy than the computation itself and introduces delays that are increasingly incompatible with real-time or always-on operation. For edge AI systems, often operating on tight energy budgets and required to function autonomously, these inefficiencies are no longer tolerable.

Memory is therefore no longer a passive component of the system. It is an active determinant of what is possible. Addressing the challenges of edge intelligence requires rethinking memory at a fundamental level: how data is stored, how it is accessed, and how tightly it can be integrated with computation. This has led to growing interest in new classes of non-volatile memory that promise to break long-standing trade-offs between speed, energy efficiency, and data persistence.

Limitations of today's memory landscape

The contemporary memory hierarchy is built on a set of well-understood but increasingly strained compromises. Static and dynamic random-access memories (SRAM and DRAM) offer fast access speeds but are volatile, requiring continuous power to retain data. Flash memory, by contrast, provides non-volatility and high density, but at the cost of slower access times, higher write energies, and limited endurance.

These characteristics were acceptable, even optimal, for traditional computing models centred on general-purpose processors and bulk data storage. However, they map poorly onto the needs of emerging edge and AI-centric workloads. Volatile memories impose a constant energy overhead, while non-volatile options struggle to meet the speed and endurance requirements of frequent, fine-grained data access.

Attempts to bridge this gap through architectural workarounds, larger caches, more complex memory hierarchies, or specialised accelerators, have added complexity without resolving the underlying issue. The separation between memory and logic remains, and with it the cost of moving data back and forth across the system.

In parallel, several alternative non-volatile memory technologies have been proposed, including resistive, phase-change, and magnetic approaches. While each offers compelling attributes, many face challenges related to variability, scalability, write energy, or integration complexity. As a result, none has yet delivered a broadly adopted solution capable of simultaneously matching the speed of volatile memory and the persistence of storage-class devices.

For edge AI systems, this fragmented landscape translates into difficult design choices. Engineers must trade energy efficiency against performance, endurance against density, and system simplicity against functionality. Overcoming these trade-offs requires not just another point on the existing spectrum, but a fundamentally different memory approach, one designed from the outset to support low-energy, high-speed, non-volatile operation within future computing architectures.

Introducing a new memory approach

Responding to the growing demands of edge AI requires a memory technology that does not simply optimise one parameter at the expense of others, but instead redefines the balance between speed, energy efficiency, and data retention. One such approach is being developed by Quinas Technology, a UK-origin deep-tech company focused on next-generation non-volatile memory devices.

At the core of this approach is a departure from conventional charge storage mechanisms. Rather than relying on capacitive charge accumulation or large-scale material phase changes, the device operation is based on quantum-engineered semiconductor heterostructures. These structures enable controlled charge transfer through energy barriers that can be precisely designed at the atomic scale, allowing data to be written and erased with extremely low energy input while maintaining long retention times.

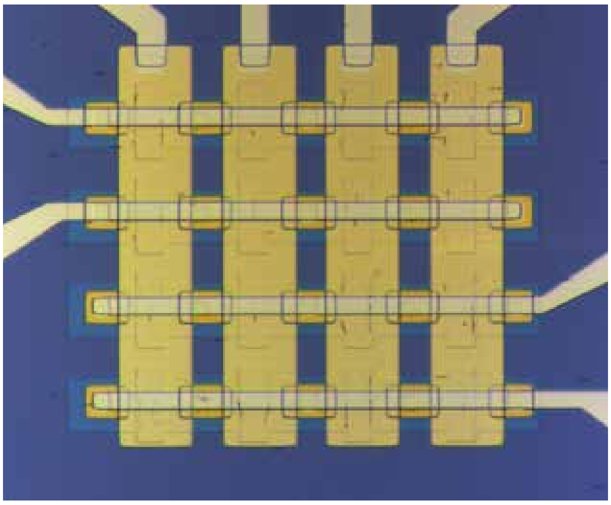

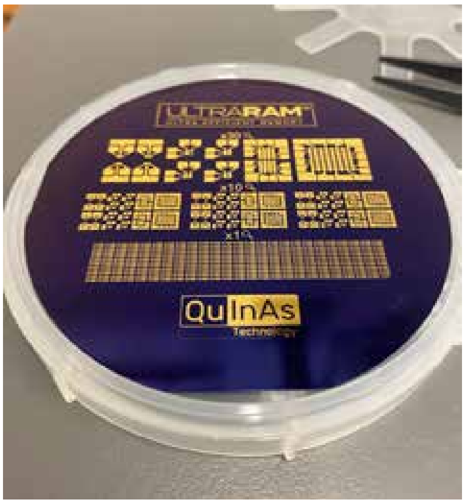

Figure 1. Processed ULTRARAM™ demonstration wafer showing multiple device architectures and array scales developed by Quinas Technology.

Figure 2. Optical micrograph of fabricated ULTRARAM™ device structures at the micrometre scale.

This architecture enables memory switching energies that are orders of magnitude lower than those of Flash, while retaining access speeds closer to volatile memories. Crucially, the non-volatility is intrinsic to the device physics, rather than imposed through complex external circuitry or refresh mechanisms.

Another distinguishing feature is endurance. The device behaviour supports frequent, fine-grained read and write operations without the degradation mechanisms that limit many alternative non-volatile memories. This makes it particularly well suited to workloads such as neural network inference, adaptive systems, and continual learning at the edge.

System-level impact and applications

The implications of this new memory class extend well beyond device-level metrics. At the system level, reducing the need to shuttle data between separate memory and compute units can deliver substantial gains in energy efficiency, latency, and architectural simplicity.

For edge AI applications, this translates into the ability to deploy more capable models within strict power envelopes. Always-on sensing, local inference, and real-time decision-making become feasible without reliance on continuous cloud connectivity.

The technology is also highly relevant to neuromorphic and in-memory computing approaches, where computation occurs where data resides. Persistent, low-energy memory devices are a key enabler of these architectures, supporting dense, adaptive systems with minimal energy overhead.

Security is another important consideration. Persistent memory enables instant-on operation, secure boot, and resilience in environments where power availability is intermittent or unpredictable.

Relevance for Europe's innovation ecosystem

Semiconductors underpin competitiveness across sectors ranging from automotive and manufacturing to healthcare and defence. Energy-efficient AI and edge computing are explicit priorities within Europe's digital and industrial strategies, and memory technologies play a critical enabling role.

Innovation in non-volatile memory contributes directly to Europe's ambitions around technological sovereignty, sustainable computing, and resilient supply chains. Progress in this field depends on strong links between fundamental research, industrial capability, and cross-border collaboration, areas where Europe has deep strengths.

University spin-outs, collaborative R&D programmes, and ecosystem partnerships are central to translating scientific advances into deployable technologies, ensuring that value creation remains anchored within the European innovation landscape.

Outlook and collaboration

As edge AI continues to mature, the role of memory will only become more central. Ongoing work at Quinas Technology and its research and industrial partners focuses not only on device optimisation, but also on system integration and application-driven validation.

James Ashforth-Pook
Co-founder & CEO,
Quinas Technology

James Ashforth-Pook is a global semiconductor entrepreneur and Co-founder & CEO of Quinas Technology, a deep-tech company developing next-generation non-volatile memory for energy-efficient AI and edge systems. He works across Europe and Asia, bridging advanced semiconductor innovation with industrial deployment and ecosystem development.

By aligning memory innovation with real-world system needs, next-generation non-volatile memory devices can help unlock a new class of intelligent, energy-efficient systems at the edge, transforming advances in semiconductor physics into tangible industrial and societal impact.

INSIDE
Industry Association

Proud member of Chips JU

INSIDE Connect 2026

30 September
& 1 October
Palermo, Italy

Connect our community, shape our future

Research Project Highlight

H2TRAIN

From wearable sensors to digital twins for practical well-being in Europe

Edge AI and digital twins built on advanced electronic components and integrated wearables

Juan Antonio Montiel
University of Las Palmas de Gran Canaria

Keijo Haataja
University of Eastern Finland

Annika Hangleiter
Wearable Technologies AG

H2TRAIN is a European innovation action funded under the CHIPS Joint Undertaking. It addresses a central challenge in digital well-being. The challenge is how to turn advanced electronic components into wearable systems that work in everyday life. The project combines novel sensing materials with CMOS-compatible sensor integration. It also develops energy-aware wearable design and embedded artificial intelligence. H2TRAIN uses an edge-to-cloud continuum to process data in near real time. It applies digital twin views to support monitoring and decision support for professionals. The work is grounded in three use cases. Remote assisted living, intelligent assisted sport coaching, and clinical monitoring.

A wearable revolution needs systems

Europe aims to lead in digital health and well-being. Software alone is not enough. Wearables must work as stand-alone outside the laboratory. They must capture meaningful signals on and especially from the human body. They must operate for long periods with low energy use. They must also protect sensitive data.

Many promising solutions stop at prototypes. They perform well in controlled tests. They struggle in real settings. Battery limits, motion artefacts, and comfort issues remain common barriers.

H2TRAIN responds to this gap. The project focuses on translation. It translates materials into sensing. It translates sensing into integrated wearables. It then translates data into guidance through edge artificial intelligence and digital twins.

From sensors to decisions

H2TRAIN follows an end-to-end pipeline. It starts with sensing on the body. It processes data close to the user through edge computing. It then applies secure cloud analytics for longer-term insight. Results are presented through digital twin views and dashboards that support decisions.

How it works in practice

Consider a training session in the intelligent assisted sport coaching use case. The wearable measures motion and selected biosignals. On-device processing filters noise and detects key events. A phone or gateway aggregates the data and uploads it securely. Cloud analytics updates trends over days and weeks. A digital twin view presents fatigue and recovery indicators to the coach. The coach adjusts the training plan. The next session provides new data for learning.

Use cases that shape the system

H2TRAIN is structured around three use cases. They define requirements from the start. They also guide validation in realistic environments.

Remote assisted living

This use case supports monitoring outside clinical environments. Comfort and reliability are essential. The system must operate over long periods with minimal maintenance.

AI and digital twin outputs

- Signals include activity and physiological streams.
- AI outputs include personalized baselines, anomaly flags, and trend changes.
- Digital twin views present alerts and suggested actions to operators.

Intelligent assisted sport coaching

This use case supports athletes and coaches in performance monitoring. Wearables must cope with sweat and intensive movement. The aim is safer training and better recovery decisions.

AI and digital twin outputs

- Signals include motion sensors, biosignals, and selected biomarkers.
- AI outputs include phase recognition, workload indicators, and fatigue trends.
- Digital twin views present athlete state and training context in coach dashboards.

Clinical monitoring

This use focuses on a clinical risk prediction, early warning scoring system for diabetes patients and cardiac rehabilitation. Continuity and adherence to plans are critical. Professional workflows must be respected. A patient-specific digital twin aims for disease progression scenarios.

AI and digital twin outputs

- Signals include mobility and selected physiological measures.
- AI outputs include deviation from expected progress and recovery trends.
- Digital twin views present patient progress to rehabilitation.

Sensing that enables better AI

H2TRAIN advances wearable sensing through one-dimensional and two-dimensional materials. It targets CMOS-compatible integration. Richer and more reliable signals support earlier detection of fatigue, stress, and risk patterns.

AI models in H2TRAIN

H2TRAIN applies artificial intelligence to turn raw signals into meaning. The models focus on time-series and multimodal data. They are designed for real-world noise and variability:

- Motion time-series models support activity and phase recognition.
- Biosignal models support fatigue and strain estimation.
- Multimodal fusion combines motion and physiological data.
- Personalization models establish user-specific baselines.

Anomaly detection is also supported. These models identify out-of-pattern movement or physiological signals. They help flag potential safety risks.

Figure 1 – From advanced sensing to real-life well-being applications in H2TRAIN

Figure 1. H2TRAIN pipeline linking wearable sensing, edge AI inference, secure cloud analytics, and digital twin views for training and decision support.

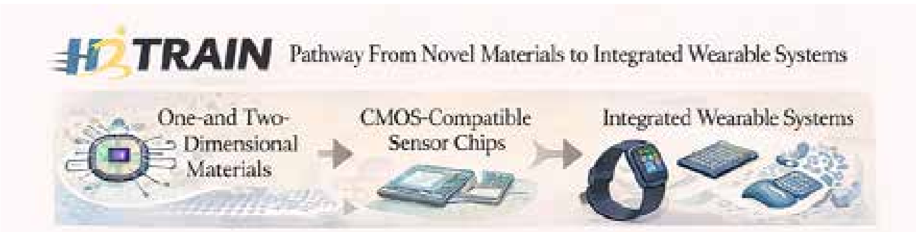

Figure 2 – From novel materials to integrated wearable systems: H2TRAIN translates advances in one- and two-dimensional sensing materials into CMOS-compatible sensor chips and integrated wearable units. This material-to system pathway enables sealable, energy-efficient, and application ready well-being technologies.

Figure 2. From novel sensing materials to CMOS-compatible sensor chips and integrated wearable systems.

Edge AI algorithms under real constraints

H2TRAIN applies an edge-to-cloud continuum. Edge AI reduces latency and limits data transfer. It also supports privacy and energy efficiency.

Tiny machine learning algorithms enable inference on resource-constrained devices. These include wearables, gateways, and smartphones. Models use efficient processing and streaming inference. They tolerate motion artefacts and missing data. Edge intelligence also supports calibration and adaptation to available energy.

Reality check

Comfort, fit, battery limits, textile robustness, packaging, and daily wear all constrain edge AI in real devices.

Digital twins that support action

Digital twins support assisted supervision in H2TRAIN. They help professionals monitor users in near real time. They also support automated checks and alerts. Human-state twins represent fatigue, stress, physical load, and readiness. Workflow twins represent task phases, training steps, or rehabilitation stages. Interface twins present information through dashboards and simple applications.

The process follows a clear loop. Data is sensed. Signals are inferred. Results are visualized. Guidance is provided. The system learns and improves over time.

Trust, interoperability, and validation

Trust is essential for adoption. H2TRAIN treats security and data protection as core design requirements. The project aligns with relevant European standards and regulations. These include medical devices, risk management, software lifecycle, and information security frameworks.

Interoperability is also important. Wearables must connect to mobile applications and secure platforms. Clear metadata and audit trails support professional use.

Validation goes beyond the laboratory. H2TRAIN evaluates sensors, wearability, and AI robustness across users and environments.

Conclusion

H2TRAIN shows how Europe can translate components into systems. It connects sensing, edge AI, secure analytics, and digital twins. It supports assisted living, sports coaching, and rehabilitation monitoring. The project strengthens Europe's capability in electronic components and system integration. It also supports responsible and trustworthy AI-enabled well-being solutions.

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