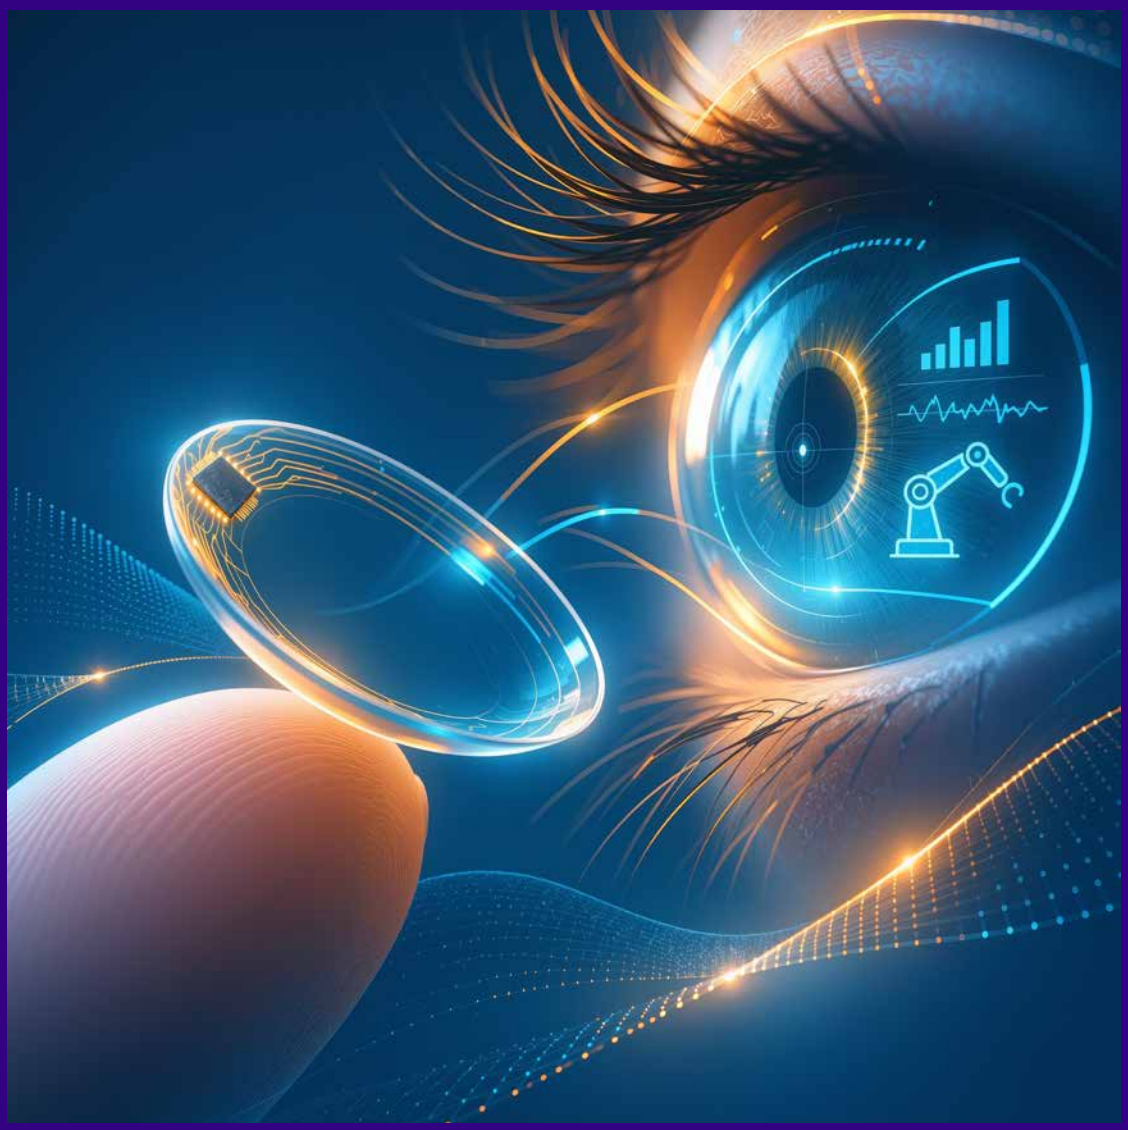
INSIDE magazine



December 2026 — Issue 15

In this issue:

- **Beyond smart glasses**
- **Towards autonomous physical intelligence**
- **Engineering physical AI across sectors**

INSIDE
Industry Association

In this issue

December 2026 — Issue 15

04 — Beyond smart glasses: the next frontier of wearable computing

12 — Towards autonomous physical intelligence

20 — When wearables become decision makers

26 — A digital twin–based framework for robotic control strategy benchmarking

32 — From pilot to production

34 — Engineering physical AI across sectors

38 — INSIDE Connect 2026

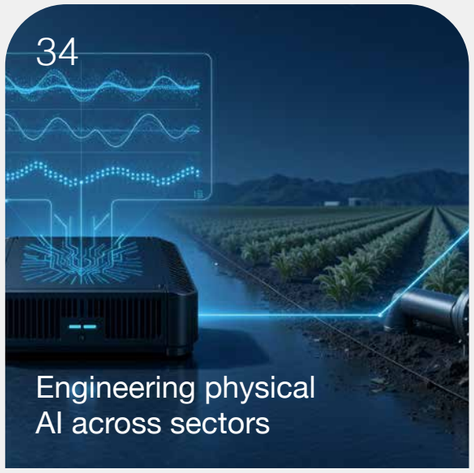
42 — More than a summer school

46 — The promise and peril of Artificial Intelligence

© 2026 INSIDE Industry Association. All rights reserved.







Dear reader,

Our interaction with artificial intelligence has been mostly mediated by the computer screen up until now. We use AI to retrieve information, get suggestions, analyse data, create text, pictures, and software. Despite its impressive capabilities, AI has mostly served as a tool to support decisions based on available information, leaving people to act on its results.

This seems to be about to change.

The presence of AI is becoming increasingly present in the things around us, both tangible and not. Augmented reality glasses are letting digital information enter our vision. Wearables allow monitoring of physiological parameters and interpreting them. Robots learn to adapt their behaviour to unpredictable environments. Manufacturing systems start reacting to the situation rather than following a preset program entirely.

Some articles in this issue explore various aspects of this evolution. Looking at them altogether, I think there is an important question that should be asked: what happens when AI moves away from just information provision to acting in the physical world?

These implications deserve a closer look.
An erroneous answer provided by a chatbot will most likely be deceptive or even dangerous. But in case an autonomous machine misinterprets its environment, or a wearables device suggests an inappropriate action, or an industrial system makes an erroneous decision, the consequences will be immediate and tangible. The requirements we impose on such systems should therefore be completely different from what we have been used in current AI applications.

Here comes another question: how do we make intelligent systems reliable and functionally safe?

Enhancing algorithms will not be enough. An autonomous system operating in the physical world must observe and perceive its environment, make sense of it, make decisions and act, often under severe limitations imposed on its energy consumption, computing capacity, and time of response. Moreover, it must operate based on insufficient data and unforeseen situations.

But this means that we cannot limit ourselves to just having a system that works well under experimental conditions. We need engineering methods that allow us to understand, evaluate, and verify autonomous systems behaviour, especially in the case of unpredicted circumstances.

Here electronic components and systems have their essential part to play. Semiconductors, heterogeneous integration, embedded computing, sensors, communication technologies, digital twins, architectures, and other interdisciplinary technologies will determine their capabilities. Their performance, reliability and limitations will affect the capabilities of future autonomous systems.

For Europe there is also an industrial challenge in this. We have many scientific and technological competences in the fields that enables autonoumous systems. But we know that excellent research does not automatically result in a successful and realiable industrial product. It takes some effort and cooperation throughout the value chain to turn promising prototypes into manufacturable and competitive systems.

There is another issue which we cannot ignore: are we educating the next generation of engineers who can deal with systems that decide and act autonomously?

The solution to this problem is not only about incorporating the study of AI into the curriculum of engineers. What is needed are professionals capable of analyzing and understanding the interaction between the physical and digital world, assessing the capabilities and limitations of intelligent systems, and designing them keeping in mind their reliability, security, safety and dependability. In addition, we must consider where autonomous decisions are appropriate and acceptable, and where human judgement and responsibility must remain central.

Looking at the evolution towards physical AI, I believe that the most important development is not just the increase in the intelligence of machines. It is that the intelligence is gradually becoming a part of the physical world that surrounds us in the daily life.

The future of AI may not be about what machines know but rather about what they can do, whether they can do it reliably and safely, and how much responsibility we are prepared to entrust to them.

I hope the articles included in this issue provide some food for thought on this subject.

Enjoy your reading!

Paolo Azzoni
Secretary General



Technology Frontiers

Technology Frontiers

Beyond smart glasses: the next frontier of wearable computing



Paolo Azzoni

From augmented reality eyewear to smart contact lenses, an invisible interface with the digital world bringing intelligence directly to the human eye

Over the past decades, the development of personal computing has involved closer integration between individuals and digital data. Desktops turned into laptops and smartphones, eventually moving from desktops to mobile devices, thus, bringing people and computing devices closer to each other. Smart glasses are the latest step in this process, bringing cameras, sensors, communication and artificial intelligence into an object that millions of people already wear.

This idea is not new: early attempts to provide digital information into the user's field of view showed good results but at the same time highlighted the limitations of ergonomics, energy needs, optical quality and social acceptability. The main difference is in technologies themselves: miniaturised processors, MEMS sensors, wireless connectivity, efficient batteries and AI-based software allowed to integrate useful digital functions into relatively conventional glasses.

The commercial success is quite significant. Omdia[1] has reported that worldwide shipments of AI-glasses amounted to 8.7 million units in 2025, which is 322% growth year-on-year. Meta had a 85% market share, showing fast expansion of the segment and its concentration among a few companies. Shipments in the first half of 2026 grew by another 4.2 million units (year on year, growth amounted to 127%).

These market figures should be interpreted with caution. Smart glasses is a diverse market. There are devices that use the combination of a camera, microphone, speaker and AI assistant and have made notable commercial success. But display-equipped glasses, AR glasses or immersive spatial interfaces are limited to a smaller and technologically more demanding market.

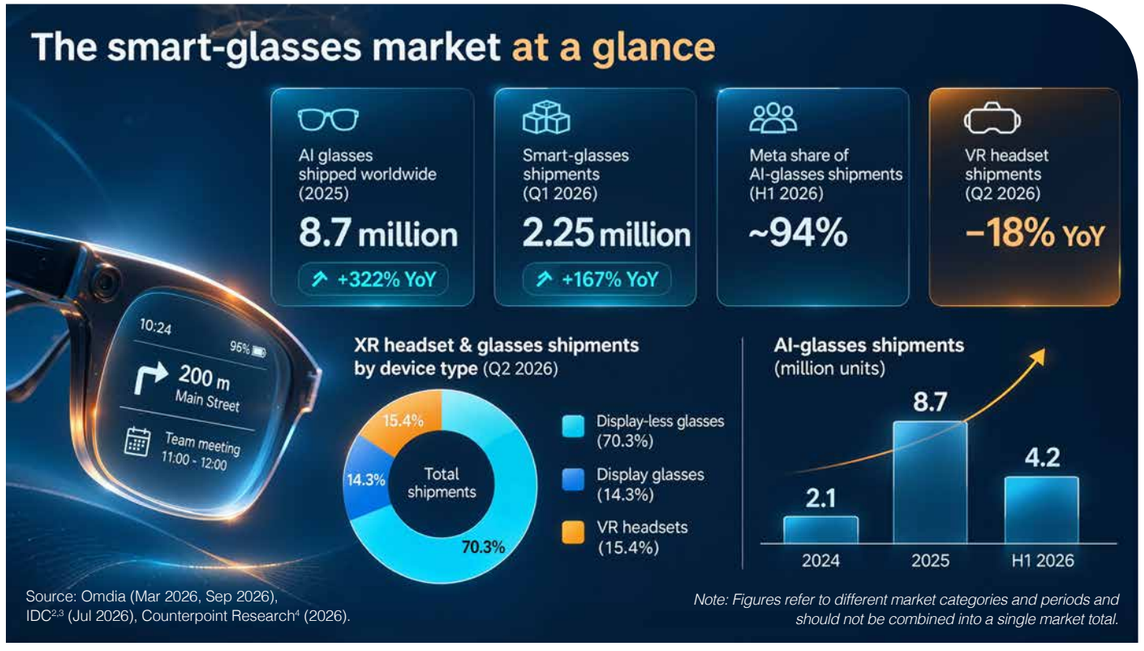
Omdia estimates that display-equipped AI glasses made 8.4% of all AI glasses shipped in 2025. It seems that customers currently prefer simpler and more practical solutions to more sophisticated visual immersion.

The reason behind this could be found in the physical nature of such devices. Camera and audio smart glasses are already becoming quite small and light, whereas glasses with more sophisticated displays remain more demanding in terms of weight, dimensions, power consumption and heat dissipation, like for bulky and heavy immersive virtual reality goggles. Despite all the recent improvements, this remains one of the biggest limitation on the path of everyday use of such devices. Users will value the possibilities of such systems, but will not wear them for long periods of time.

> **The success of wearable computing will depend on what people are willing to wear and use, less than on what technology can do.**

This separation is made even more apparent by looking at virtual-reality headsets. Even though there have been many technological advancements and investments, VR immersion has failed to penetrate the consumer market as previously thought. The success of smart glasses will not then indicate an openness of consumers towards all kinds of wearable spatial computing.

Other than ergonomics, privacy may prove to be another very crucial issue that will obstacle adoption. Smart glasses fitted with camera and microphone capability are already raising issues related to recording individuals without permission. The inclusion of artificial intelligence and facial recognition makes the issues even more pressing. There are already existing laws limiting recording devices in critical situations, and some organizations as well as governments have considered or implemented limitations on wearable devices such as smart glasses.



Should the issues evolve into greater regulations or social opposition, they would definitely hamper consumer adoption. Whereas technical challenges can be overcome through engineering, privacy concerns will need technological, regulatory, and social solutions.

The crucial question now is what use cases motivate putting computing technology in front of our eyes and whether they justify all the constraints that come along with the technology. To answer this question is the key to understanding smart glasses' future and the promise of more revolutionary approaches to wearable technology.

An established industrial ecosystem including European suppliers

The evolution of smart glasses measures how mature the electronics manufacturing industry has become.

A state-of-the-art pair of smart glasses consists of many different elements: application processor, image sensors, inertial measurement unit, microphones, wireless communications, audio signal processing, power management system, batteries, and software. Some more advanced models also feature microdisplays, optical waveguides, eye-tracking, and advanced sensing capabilities.

Many of the components are already available through established industrial supply chains. Inegrating them remains difficult due to limitations related to weight, temperature, and battery lifetime, but the technology itself is out of the lab stage at this point.

Europe is one of the important players in this ecosystem.

One of the most evident examples would be the partnership between EssilorLuxottica[5] and Meta, which has led to the development of smart glasses for the consumer market. EssilorLuxottica has reported a sale of more than seven million AI glasses in 2025, thus indicating the commercial opportunities provided by a combination of experience in eyewear and new technologies.

Nevertheless, Europe's role goes beyond the final product. There are European industries that supply a great number of components necessary for the development of smart glasses. For instance, STMicroelectronics[6] provides microcontrollers, MEMS and power management solutions, while Bosch Sensortec[7] specializes in miniature inertial sensors and sensor fusion. In terms of optics and photonics, ZEISS[8,9] has a vast experience in developing optical systems and augmented reality display technologies, and ams OSRAM[10,11] creates optical sensors, emitters and microLED technologies.

All together these competences indicate the existence of an already developed European ecosystem of industrial capabilities in different layers of the smart glasses technology stack. Not each company provides components for the same commercial products, but the competences, technologies and capabilities are available already within Europe.

> **European companies already develop and commercialise many of the electronic and optical technologies used in smart glasses.**

It does not mean, however, that all research problems are solved yet. Highly efficient optical systems, energy-saving displays, human-machine interface innovations and miniaturization are still active areas of development. But the development level of the already existing ecosystem raises the problem of making a clear distinction between the development of products and innovation breakthroughs.

Where smart glasses create real value

The implementation of technology in a commercial context is just one way of



measuring success. An alternative indicator is the capacity of the technology to solve challenges which are difficult to solve with traditional interface technologies.

In an industrial environment, wearable visual devices may provide benefits that are less evident in everyday consumer applications.

For example, during maintenance activities, operators might have to refer to manuals, examine hardware and talk to experts who are far away from the workplace while their hands are occupied at the same time. The use of smart glasses would allow for the display of manuals, recognition of components and help from experts without constantly interfering with the activity.

This applies to many other fields as well, including manufacturing, logistics, medicine and field operations. In these contexts, the value of wearable computing can be evaluated through lesser intervention time, fewer errors, enhanced training, and greater access to specialised expertise.

Remote control is another case in point. Operators who remotely pilot drones, robots or machines could be assisted by visual systems integrating cameras, navigation, telemetry and contextual alerts. These kinds of systems would need more advanced display and communication technology than consumer-oriented wearables.

The demands vary significantly across applications, however. While one application requires only a light display of instructions, another might demand a high-quality, real-time visual system. In safety-critical environments, there are other considerations such as reliability, cybersecurity, situation awareness and human factors.

As a result, it is doubtful that a single kind of smart glasses could satisfy all the potential vertical applications in the professional domain.

In the consumer sector, however, things work differently. Smart glasses have to provide enough practical utility to compensate for the costs, charging capabilities, weight and privacy risks. The popularity of smart glasses for the camera and audio shows that the less intrusive features appeal to people. The more invasive technologies should show that the advantages are worth the hassle.

Nevertheless, the discrepancy between necessity and convenience will certainly play a role in the development of wearables.

But even the most innovative glasses share one key feature: they are peripheral devices placed in front of the eyes.

But what if the interface was placed directly on the eye?

Smart contact lenses: bringing electronics to the eye

Smart contact lenses have the possibility to be a more revolutionary step forward in wearable computing.

Whereas glasses offer a way to incorporate electronics in their frames, the contact lenses should incorporate all the necessary elements inside a very thin, curved, and flexible material placed directly on the surface of the eye.

The possible benefits of such an approach are obvious: no additional frames are needed for the contact lenses and they follow eye movements, providing visual information without the use of any external eyewear. Contact lens don't also interfere with helmets, protective gear or other equipment worn on the head, introducing a future in which digital information becomes less dependent on visible external hardware.

All these factors make the technology especially appealing for professional usage.

However, the lack of a frame does not immediately imply that contact lenses have any advantage over smart glasses.



Actually, it means that there is not enough physical space left for storing the battery, processor, connection modules, and optics.

In addition, the movement of the screen together with the eye itself creates complicated optical and physiological challenges, as it is necessary to ensure proper focusing, readability, and spatial perception of projected data. Creating a stable environment of augmented reality may be quite hard.

Physiological limitations are no less important as well, as lenses have to provide enough oxygen flow, preserve the tear film and not cause any irritations to the eyes. The incorporation of electronics and energy storage devices in such an environment is quite problematic.

For these reasons, smart contact lenses cannot be considered simply as small smart glasses, but they constitute an entirely different category of electronic devices.

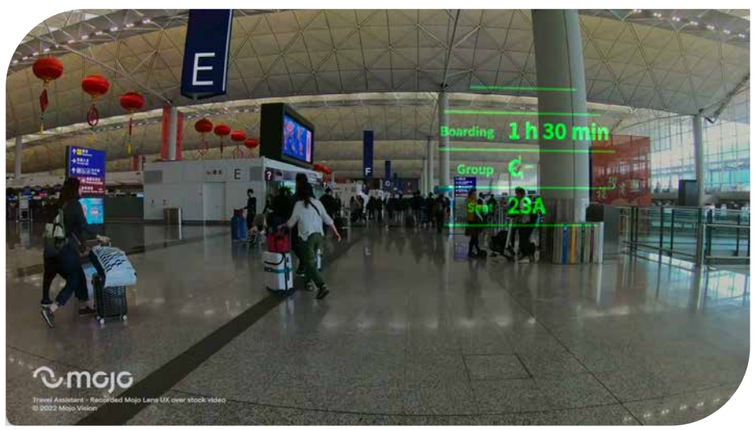
There are several streams of research, including physiological monitoring, biochemical detection, vision aids, drug delivery control, and augmented reality displays.

Medical applications may become more mature for practical implementation earlier than other types of visual computing. Obtaining ocular parameters and biomarker detection will need much less power and optics compared with generating images within the visual range of the user.

The story of Mojo Vision[12, 13] serves as a good example of the opportunities and challenges in the field. The company developed an advanced smart contact lens platform with a miniature display and electronics but shifted its focus to microLEDs in 2023. However, it provided important innovations in optics and electronics miniaturization.

More recently, XPANCEO[14] has been working on a smart contact lens platform with integrated augmented reality features,





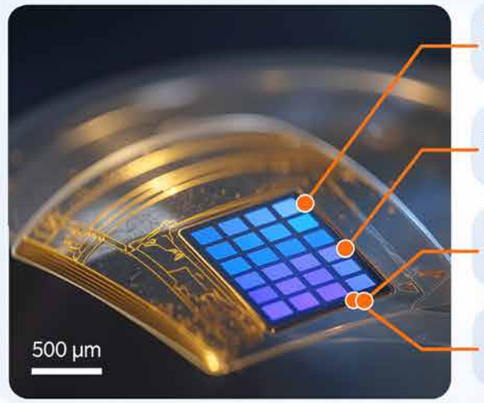
Towards integrated smart contact lenses

Recent developments show significant progress towards contact lenses that integrate display, sensing, electronics and power in a compact and biocompatible form factor. A relevant example is the work of XPANCEO, which has demonstrated multiple prototypes and is targeting an integrated lens in the near term.



A multi-technology platform

Integration of key functions in a contact lens combines ultra-compact optical, electronic and sensing technologies.



Microdisplay
MicroLED array (um-scale pixels, high brightness)

Electronics
Ultra-thin, low-power ASICs for display drive, sensing and communication

Sensing
Eye tracking (e.g, EOG, IR) and blometric/health sensors (e.g, tear analysis, ICP)

Power
Miniuturised solid-state battery solutions (PoC with ITEN)

Prototype specifications

(reported by XPANCEO) [1,2]

Display technology	MicroLED
Pixel pitch	= 1.8 pm
Pixel density	> 14,000 ppl
Display area (diameter)	= 0.5 mm
Lens diameter	= 14 mm (standard contact lens)
Electronics thickness	< 50 pm (target)
Power	Integrated solid-state microbattery (proof of concept)
Functions demonstrated	Visual display, sensing (e.g, eye tracking), wireless communication

Recent development

XPANCEO has developed more than 28 prototypes covering key building blocks, including microdisplays, transparent and flexible electronics, sensing and power solutions. In April 2026, the company and ITEN, a French specialist in solid-state microbatteries, announced a proof of concept of a smart contact lens with an integrated battery.

A public demonstration of an integrated smart contact lens, combining display, sensing and power in a single device, is planned for early 2027.

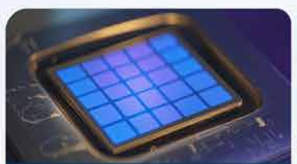
Technology challenges

Miniaturisation and integration
High-density electronics and interconnects in a sub-millimette, curved and blocompatible platform.

Power and thermal management
Sufficient energy density, safe operation and heat dissipation within strict size and comfort constraints.

Optical performance
High brightness, wide field of view and low power consumption with transparent optics.

Biocompatibility and safety
Long-term eye tolerance, materials stability and regulatory approval for consumer and medical use.


MicroLED display (close-up)


Transparent and flexible electronics


Solid-state microbattery (proof of concept)


Contact lens prototype (transparent platform)

[1] XPANCEO, “XPANCEO presents smart contact lens protetype portfolio at VivaTech”, June 2026, https://xpanceo.com/news.

[2] ITEN, “XPANCEO and ITEN create proof of concept for a smart lens with integrated solid-state battery®, April 2026, https://www.iten.com/news.



sensoring features, and others. In June 2026, XPANCEO presented developments on more than 28 prototypes, announcing its intention to present an integrated prototype by early 2027.

One of the most interesting developments concerns its collaboration with the French company ITEN[15]. In April 2026, the companies announced a proof-of-concept of an integrated system of a miniaturized solid-state battery and a smart contact lens.

Such an example underscores the main problem: technologies themselves have already been proven, but integration of those into a safe, realiable and manufacturable system is a much more complicated process.

A new frontier for electronic components and systems

The design of smart contact lenses represents an interdisciplinary area of research.

First is the area of ultra-low power electronics. The energy source available is very scarce, whereas heat dissipation is limited due to the sensitivity of the ocular tissue. Novel circuit designs and energy management technologies (including wireless energy-transfer) would be required for providing the useful functionalities.

The second area is the issue of heterogeneous integration. A variety of sensors, communication circuitry, optical devices and energy storage devices have to co-exist within flexible biocompatible structures. Traditional electronic packaging technologies are usually inappropriate for such geometries and operating conditions.

Additionally, photonics will introduce further challengies. Miniature displays, optics coupling, and image forming have to function in very tight spatial limits and deliver sufficient visual quality. This can lead to further improvements in microLEDs, integrated photonics and novel optical architectures.

> **The challenge is no longer to miniaturise individual components, but to integrate complete electronic systems within the physical and biological constraints of the human eye.**

Embedded intelligence will be crucial as well. It will be impossible to fit the computation power of a smartphone or sophisticated smart glasses into a contact lens. It could become the smallest unit of a distributed computing architecture that will be doing some sensing and computation tasks, while delegating more complicated computations to nearby devices and cloud resources for more demanding tasks.

This type of architecture will need wireless communications, workload management, dependable operations and proper handling of private information.

Finally, the devices that work directly with the human eye need to fulfill high demands with respect to materials, reliability, cybersecurity, privacy, and wherever appropriate, medical device regulations. Such challenges go far beyond the innovation of a new consumer product. They involve fundamental innovations in electronic systems integration and device-tissue interactions.

This transition from smart glasses to smart contact lenses thus reflects a more general challenge to the European electronics and systems industry. With some wearables reaching maturity in terms of market presence, new possibilities arise from the convergence of microelectronic, photonic, advanced materials, embedded intelligence, and biomedical technologies.

Where the next breakthroughs are needed

The development of wearable computing serves as an example of how rapidly emerging technology can evolve into an industrial product. Currently, the use of smart glasses can be made possible by relying on a well-established ecosystem of electronic components, optical technologies and manufacturing capabilities that aleady includes significant contributions of the European industry. However, future improvements may not rely on the advancement of these current technologies. Smart contact lenses represent a good example of how much remains to be explored when electronics must operate within the physical and biological constraints of the human eye.

It is important for Europe to understand how industrial advancements can be made using the existing capacity of the European industry and what new technological innovations need to be created. It is here where collaborative research will have the most impact, creating opportunities far beyond the improvement of smart glasses.

1 https://omdia.tech.informa.com/pr/2026/mar/global-ai-glasses-shipments-reach-8point7-million-units-with-mainland-china-emerging-as-the-fastest-growing-market
2 https://www.idc.com/promo/arvr/
3 https://www.idc.com/resource-center/blog/smart-glasses-growth-cools-to-35-as-meta-doubles-down-on-distribution/
4 https://counterpointresearch.com/en/insights/global-ai-glasses-shipments-surge-263-percent-yoy-with-display-less-ai-glasses
5 https://www.essilorluxottica.com/en/newsroom/press-releases/q4-full-year-2025-results
6 https://www.st.com/en/applications/wearables/smart-glasses.html
7 https://www.bosch-sensortec.com/en/applications-solutions/xr-headset-and-glasses
8 https://www.zeiss.com/vision-care/en/newsroom/news/2025/zeiss-extended-reality-established.html
9 https://www.zeiss.com/extended-reality/en/product-portfolio.html
10 https://ams-osram.com/applications/mobile-wearables/intelligent-optical-systems-for-smart-glasses
11 https://ams-osram.com/news/press-releases/q2-2026-results
12 https://www.mojo.vision/news/a-new-direction
13 https://spectrum.ieee.org/looking-through-mojo-visions-newest-ar-contact-lens/particle-4
14 https://www.xpanceo.com/newsroom/xpanceo-presents-smart-contact-lens-prototype-portfolio-at-vivatech-paving-the-way-for-first-integrated-prototype-demonstration-in-2027
15 https://www.iten.com/resources/press-releases/xpanceo-and-iten-create-proof-of-concept-for-smart-lens-with-integrated-solid-state-battery

ECS
BROKERAGE
EVENT 2027

17 & 18 February
The EGG, Brussels

Registration for the ECS Brokerage Event 2027 will open soon. The event brings together the brokerage activities of the industry associations AENEAS, EPoSS and INSIDE into one networking event dedicated to project proposals for the Chips JU Calls.

Register your interest now and we'll let you know as soon as registration opens.

ORGANISED BY:
Aeneas
EPoSS. European Association on Smart Systems Integration
INSIDE Industry Association

Technology Frontiers

INSIDE Magazine 15

Technology Frontiers

Towards autonomous physical intelligence

A vision for Europe's autonomous systems toolbox



Wilfred G. van der Wiel
BRAINS Center for Brain-Inspired Computing, University of Twente, The Netherlands and Institute of Physics, University of Münster, Germany

As artificial intelligence increasingly moves from data centres into the physical world, machines must more and more make decisions where data are generated and actions take place. Communication introduces latency and energy costs and, moreover, continuous connectivity cannot always be guaranteed. Privacy, safety and security often require local control. Autonomy is therefore not a luxury, but a necessity. Yet true autonomy requires more than responding to predefined situations. As nature demonstrates, autonomous systems must also be able to adapt to changing circumstances and learn from experience. Autonomy spans sensing, communication, control and actuation. Here, the focus is on the semiconductor technologies underpinning them. Europe has the opportunity to develop an "autonomous systems toolbox", bringing together digital electronics, neuromorphic and in-memory computing, analogue processing and emerging physical computing devices. A key frontier across these technologies is hardware plasticity: enabling computing systems to autonomously modify their own computational functions through interaction with their environment and ultimately learn from experience. This article explores how these developments could lay the foundations for truly autonomous physical intelligence and strengthen Europe's technological future.

Intelligence is entering the physical world

The physical world is not a static dataset. It changes continuously, often unpredictably, and it does not pause while machines decide what to do next. An autonomous vehicle must respond to changing traffic conditions. A robot must interact with objects and environments it cannot fully anticipate. Industrial machinery must detect wear before small deviations become failures. Intelligent sensor networks may have to function for years with limited energy, intermittent connectivity and little external supervision. In such settings, decisions often need to be made locally, in real time and without relying on continuous communication with external computing infrastructure. At the heart of autonomous systems lies a closed loop of sensing, interpretation, decision and action (Figure 1A), enabled by electronics within a much broader physical system[1].

The need for fast, energy-efficient and local decision-making also challenges how we organize computing hardware. Today's digital-computing architectures are extraordinarily powerful and will remain indispensable. However, they typically separate the acquisition of physical signals from their processing. Signals are conditioned, digitized, transported, processed and often converted back into physical actions. This can introduce significant latency, energy consumption and communication overhead. An intriguing alternative is to bring sensing and computing much closer together, or even merge them, through in-/near-sensor and actuator computing, allowing physical signals to be processed locally, potentially by exploiting the intrinsic dynamics of electronic devices and circuits (Figure 1B). Sensing, actuation and computing can then become intimately connected physical processes rather than separate sequential operations.

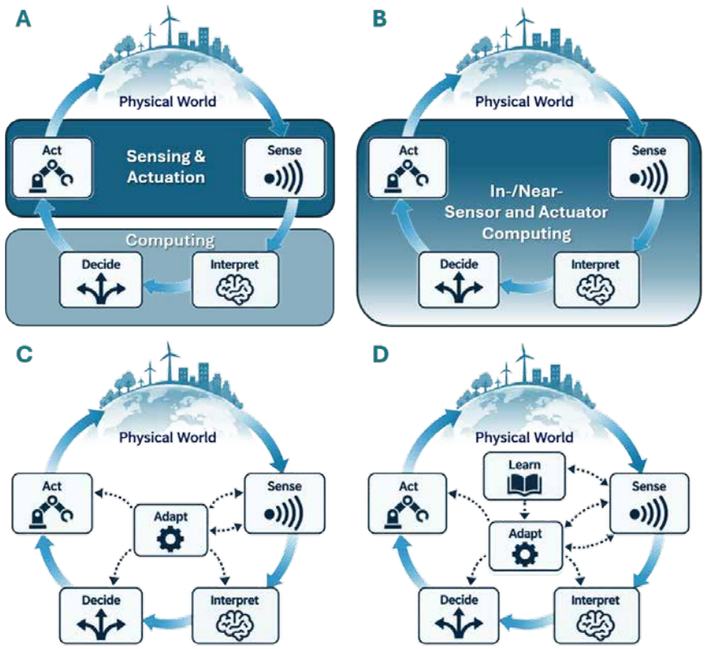


Figure 1: From responsive to adaptive and learning autonomous systems. (A) Basic closed-loop architecture linking the physical world to sensing, interpretation, decision-making, and actuation. (B) The same functional loop with computation integrated in or near sensors and actuators, enabling closer interaction between information processing and physical interfaces. (C) Introduction of adaptation, allowing the system to modify its behaviour in response to changing conditions. (D) Addition of learning, through which accumulated experience informs future adaptation. Solid blue arrows indicate the primary operational loop, while dotted navy arrows represent feedback pathways supporting adaptation and learning.

Yet responding efficiently and rapidly to changing conditions is only the first step towards genuine autonomy. A system that responds to predefined situations remains limited by what its designers anticipated. As nature demonstrates, true autonomy also requires the ability to adapt to changing circumstances and learn from experience.

Importantly, these processes operate on different, interacting timescales. While the immediate sensing–actuation loop enables autonomous responses, adaptation uses feedback and typically short-term memory to modify how the system senses, computes and acts as conditions change (Figure 1C). Learning, in turn, accumulates experience through the system's interactions with the physical world, retaining relevant information over longer timescales and using it to guide future adaptation (Figure 1D). This interplay enables autonomous systems to respond effectively to immediate challenges while continuously improving their ability to deal with situations that were not anticipated during their design.

Supporting responsiveness, adaptation and learning places new demands on the electronics underpinning autonomous systems. The challenge is not only to process information efficiently, but also to support autonomous adaptation and, ultimately, learning under changing physical conditions. This opens opportunities to exploit the intrinsic dynamics of electronic circuits, devices and materials alongside established digital processing. Rather than searching for a single technology to replace conventional computing, Europe can develop an ***autonomous systems toolbox*** in which complementary approaches work together.

A toolbox, not a replacement for digital electronics

This toolbox must address a wide range of functions, from deterministic control and communication to signal interpretation, probabilistic inference, memory and learning. No single computing paradigm can meet all these demands equally well. Table 1 summarizes the complementary approaches, their principal contributions and engineering considerations.

Conventional and specialized digital electronics will remain indispensable, providing numerical precision, programmability, robustness against noise,

Computing approach	Representative technologies	Principal contribution	Key design consideration
Digital computing	Embedded processors; digital accelerators	Precision, robustness, programmability and reliable control	Energy and latency of processing and data movement
Brain-inspired/ neuromorphic computing	Spiking neural networks; reservoir computing	Event-driven, distributed and temporal processing	Training, interfacing and robust deployment
In-memory computing	SRAM-based computing; memristive crossbar arrays	Reduced data movement and parallel operations	Precision, variability and peripheral circuitry
Analogue and mixed-signal computing	Analogue circuits; reconfigurable nonlinear processors	Direct signal transformations and low-latency processing	Noise, calibration and device variability
Physical and in-materia computing	Nonlinear networks; photonic, magnetic and mechanical systems; adaptive materials	Intrinsic physical dynamics, memory and potential adaptation	Reproducibility, controllability and integration

Table 1: A technology toolbox for autonomous physical intelligence. Complementary computing approaches offer different capabilities for local, energy-efficient and adaptive information processing. The categories overlap, and hybrid architectures may combine several approaches. Hardware plasticity is an emerging cross-cutting capability that could enable computational functions to adapt autonomously through interaction with the environment.

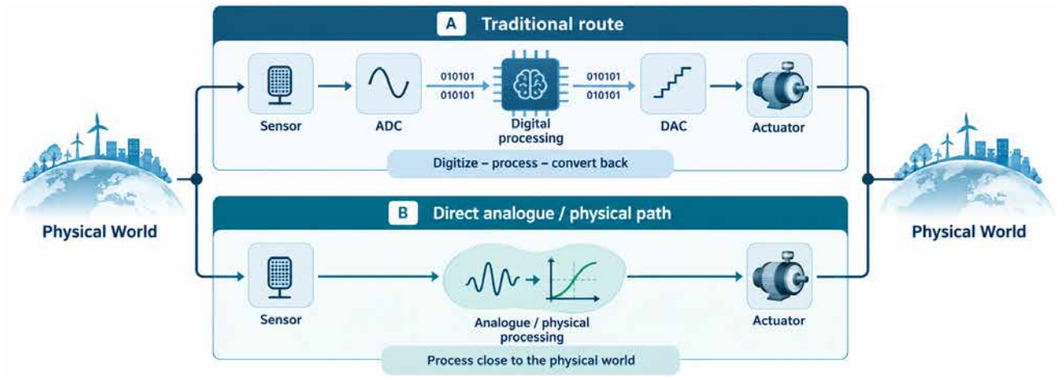


Figure 2: Bringing computation into the physical domain. Comparison of conventional digital processing, involving analogue-to-digital and, where required, digital-to-analogue conversion, with a more direct analogue/physical processing path. The latter exploits the dynamics of electronic circuits, devices and materials, potentially reducing conversion overhead and data movement.

predictable operation and established design and software ecosystems. Digital processors and accelerators are particularly important for computational tasks and control functions requiring reproducible, verifiable and maintainable operation. These strengths also provide a foundation for integrating emerging computing technologies into reliable systems.

Brain- and nature-inspired computing draws on principles found in biological and other natural systems, such as distributed information processing, temporal dynamics, sparse event-driven activity and adaptation through local interactions[2]. Neuromorphic computing is one prominent example, encompassing spiking neural networks as well as other dynamical computing architectures[3]. These approaches offer opportunities for efficient processing of temporal and event-based information, with potential for local learning and adaptation when suitable mechanisms are incorporated.

In-memory computing reduces the separation between storage and processing by performing computational operations within memory arrays, rather than repeatedly transferring data between memory and separate processing units. Implementations range from digital SRAM-based approaches to analogue operations in resistive memory arrays. In-memory computing can complement conventional digital, analogue and neuromorphic architectures, potentially improving energy efficiency and throughput by reducing data movement[4].

Analogue and mixed-signal computing exploits the fact that many physical quantities and the electrical signals representing them vary continuously in time and amplitude. Rather than converting every signal into a digital representation before processing, analogue circuits can directly manipulate voltages and currents, exploiting circuit dynamics, nonlinear responses and parallel physical processes to perform useful computations. These include filtering, temporal processing, feature extraction and nonlinear transformations. Mixed-signal architectures combine such operations with digital processing and control, allowing the strengths of both domains to be exploited. Potential advantages in energy efficiency and latency must be balanced against limitations in precision, noise, device variability and the overhead of calibration and peripheral circuitry.

Physical and in-materia computing takes this concept further by treating the intrinsic properties and interactions of physical substrates as computational resources[5]. Rather than implementing every operation through explicitly designed electronic circuits or digital instructions, useful transformations can emerge directly from the dynamics of nonlinear electronic networks, photonic and magnetic systems, mechanical structures and adaptive materials. Their intrinsic dynamics, interactions and internal states can provide computational functions, memory and potentially the foundations for adaptation. The ambition is not simply to reproduce conventional algorithms in unconventional materials, but to identify tasks for which the underlying physics offers distinct advantages. Realising these opportunities requires a deeper understanding of physical dynamics, together with reproducibility, controllability and integration into practical computing systems.

This perspective also calls for a different approach to hardware–algorithm co-design. Instead of exclusively mapping predefined algorithms onto hardware, we can explore which algorithms, representations and training strategies best exploit the available physical dynamics[6]. Device variability, nonlinear response and temporal behaviour, traditionally regarded primarily as imperfections, may become computational resources. The challenge is to harness these properties reproducibly while maintaining sufficient accuracy, robustness and control.

Among these complementary approaches, analogue and physical computing raise a particularly interesting architectural question: which computational functions should be performed digitally, and which can be implemented directly through the physical dynamics of electronic circuits, devices and materials?

Bringing computation into the physical domain

The choice between digital processing and alternative approaches is particularly relevant at the interfaces between electronics and the physical world. Sensors often generate continuously varying electrical signals, while actuators ultimately interact with physical quantities. In a conventional digital processing path, sensor signals are conditioned and digitized, processed by digital electronics and, where necessary, converted back into analogue signals for

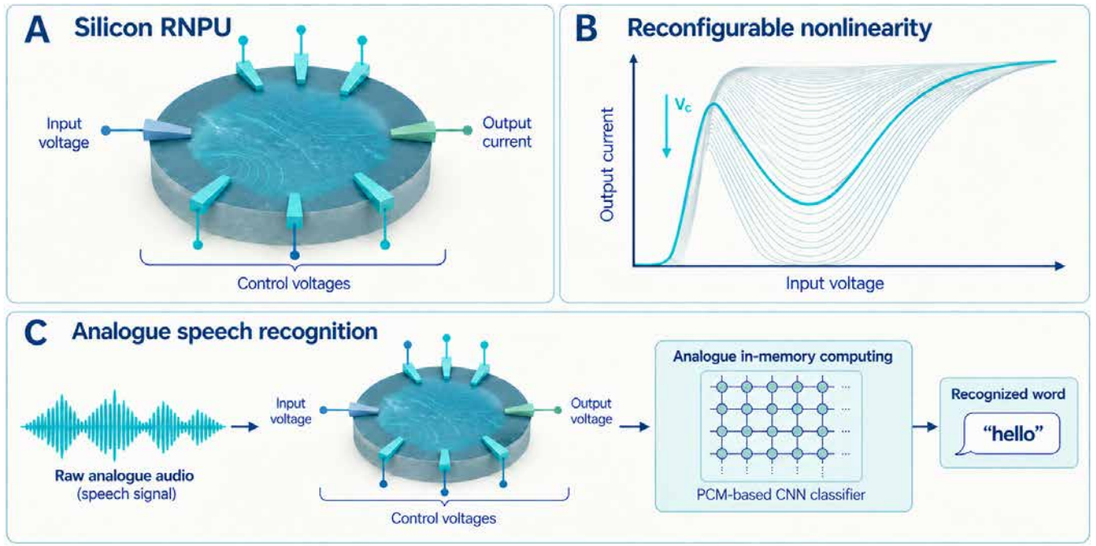


Figure 3: Reconfigurable nonlinear processing in silicon and its application to analogue speech recognition. (A) Schematic illustration of a silicon reconfigurable nonlinear-processing unit (RNPU), in which an input voltage is transformed into an output current through nonlinear electronic transport. Six additional control voltages tune the device response. (B) Representative nonlinear current–voltage characteristics illustrating how the input–output relationship can be continuously reconfigured through the control voltages. (C) Demonstration of analogue speech recognition combining RNPU-based temporal feature extraction with a phase-change-memory (PCM) analogue in-memory computing classifier. Operating in voltage-output mode, a single RNPU generates 64 feature channels through time-multiplexed operation using different control-voltage configurations, followed by classification using a convolutional neural network implemented on PCM-based hardware. In the experimental implementation, the RNPU feature extractor and PCM-based classifier were connected through a digital interface, although an analogue interface is also possible in principle. Adapted conceptually from Chen et al., Nature 577, 341–345 (2020), and Zolfagharinejad et al., Nature 645, 886–892 (2025).Figure 3. Reconfigurable nonlinear processing in silicon and its application to analogue speech recognition. (A) Schematic illustration of a silicon reconfigurable nonlinear-processing unit (RNPU), in which an input voltage is transformed into an output current through nonlinear electronic transport. Six additional control voltages tune the device response. (B) Representative nonlinear current–voltage characteristics illustrating how the input–output relationship can be continuously reconfigured through the control voltages. (C) Demonstration of analogue speech recognition combining RNPU-based temporal feature extraction with a phase-change-memory (PCM) analogue in-memory computing classifier. Operating in voltage-output mode, a single RNPU generates 64 feature channels through time-multiplexed operation using different control-voltage configurations, followed by classification using a convolutional neural network implemented on PCM-based hardware. In the experimental implementation, the RNPU feature extractor and PCM-based classifier were connected through a digital interface, although an analogue interface is also possible in principle. Adapted conceptually from Chen et al., Nature 577, 341–345 (2020), and Zolfagharinejad et al., Nature 645, 886–892 (2025).

actuation (Figure 2). This architecture offers precision, flexibility and programmability, but the conversions and associated data movement can introduce considerable latency and energy overhead.

An alternative is to perform selected computational operations directly in the analogue domain, exploiting the dynamics of electronic circuits, devices and materials before or alongside digital processing. Consider an intelligent sensor monitoring industrial machinery. Rather than continuously digitising and transferring all measurements, local analogue or physical computing could extract relevant features or identify exceptional events, passing only the essential information to a digital controller. Similar approaches can support rapid feedback near actuators. The result is not necessarily a fully analogue system, but a hybrid architecture that combines direct physical computation with digital processing and control, exploiting the respective strengths of both approaches.

This does not mean that analogue or physical processing is inherently more efficient. Bias currents, noise, calibration, device variability and interface circuitry can offset potential advantages. The appropriate comparison is therefore between complete systems delivering comparable accuracy, latency and reliability, including the energy costs of data conversion and communication. At the device level, a particularly interesting possibility is to exploit controllable nonlinear electrical responses to perform useful computational transformations directly.

A glimpse of nonlinearity-centred silicon computing

One concrete example of direct physical computation is provided by reconfigurable nonlinear-processing units (RNPUs), originally developed at the University of Twente[7]. These multi-terminal silicon devices exploit nonlinear electronic transport in lightly doped silicon active regions to perform programmable transformations of electrical signals. Control voltages tune their input–output characteristics, allowing a single physical device to implement different nonlinear computational functions. Importantly, RNPUs are based on conventional silicon materials and microfabrication techniques, offering a promising route towards CMOS-compatible integration. Rather than calculating nonlinear transformations through sequences of digital operations, the device implements them directly through its intrinsic electrical behaviour.

Our early work demonstrated that nonlinear silicon devices could be configured to perform classification tasks[7]. A subsequent advance established how such complex physical devices could be trained efficiently using data-driven surrogate models. By learning the device's input–output response with a neural network, control voltages could be optimized through gradient-based methods and subsequently applied to the physical hardware[8].

While these initial demonstrations operated at low temperatures, subsequent developments established RNPU operation at room temperature and extended their functionality to time-dependent signal processing. In collaboration with IBM Research Europe and Toyota Motor Europe, we demonstrated analogue speech recognition in which low-power RNPU circuits performed real-time temporal feature extraction, combined with classification using an analogue in-memory computing chip (Figure 3)[9]. This demonstrated how complementary physical computing technologies can be combined to process information efficiently while keeping early signal transformations in the analogue time domain.

An important step towards further development is understanding the physical origin of the nonlinear response. Recent experimental and simulation work has established that the room-temperature behaviour of silicon RNPUs is governed by space-charge effects, involving the interplay between injected carriers, ionized dopants and interface traps. This understanding provides a physical basis for engineering their nonlinear characteristics and improving reproducibility, scalability and compatibility with conventional semiconductor technology[10].

We have also explored how RNPUs can implement the trainable nonlinear functions central to Kolmogorov–Arnold networks (KANs). By combining multiple RNPUs into analogue processing elements, nonlinear transformations can be realized directly in hardware rather than approximated through conventional digital operations. Experimental measurements and surrogate-model-based studies demonstrate the potential of this approach for function approximation and classification, while system-level estimates indicate promising opportunities for energy-efficient computing[11].

Together, these developments illustrate how nonlinear transformations central to information processing can be implemented directly through engineered semiconductor physics. They also reveal important challenges. Analogue physical responses can drift, devices vary, and control voltages, calibration and peripheral circuits carry costs. Practical implementations require suitable interfaces, training strategies, testability, realistic system-level energy budgets and compatibility with manufacturing and product lifetimes. Although RNPUs provide programmable computational primitives, they are not yet complete AI processors, and promising system-level estimates must be distinguished from experimentally demonstrated performance.

Beyond these challenges lies a deeper question: can a physical processor autonomously adapt its own computational behaviour, rather than relying solely on externally chosen configurations?

From reconfiguration to hardware plasticity

As illustrated by the RNPUs discussed above, reconfiguration allows the function of a device, circuit or electronic system to be changed through selected settings, such as a control voltage, programmed memory cell or alternative signal path. Hardware plasticity goes further, encompassing the ability of electronic systems to adapt their organisation and computing functions during operation. This may involve physical plasticity, such as changes in device characteristics, circuit configurations or hardware interconnections, as well as virtual plasticity through software-controlled resource allocation and functional remapping. Autonomous adaptation requires the system to determine when change is needed, select and implement an appropriate adjustment within operational constraints, and assess the result. Learning additionally uses experience to guide future changes.

At the device and material level, one important foundation for plasticity is a physical degree of freedom that can change and retain information over a relevant timescale. In electronics, this might involve stored charge, polarisation, ionic configuration, conductance or magnetic state. These physical states can exhibit widely different retention times. Rapidly relaxing states support transient responses, while short-term memory can support adaptation to recent changes in operating conditions. More persistent, potentially non-volatile memory can enable learning by retaining information acquired through experience over longer timescales. Importantly, memory alone does not constitute learning: the retained information must contribute to meaningful

Wilfred G. van der Wiel
BRAINS Center for Brain-Inspired Computing, University of Twente, The Netherlands and Institute of Physics, University of Münster, Germany.

Wilfred G. van der Wiel (Gouda, 1975) is full professor of Nanoelectronics, co-director of the BRAINS Center for Brain-Inspired Computing, and co-chair of the Department of Electrical Engineering at the University of Twente, The Netherlands. He holds a second professorship at the Institute of Physics, University of Münster, Germany. His research focuses on unconventional electronics for efficient information processing. Van der Wiel is a pioneer in material learning at the nanoscale and in the development of reconfigurable nonlinear-processing units (RNPUs), realizing computational functionality directly in nanomaterial substrates through principles analogous to machine learning. He has authored more than 125 journal articles, receiving over 15,000 citations.



changes in the system's behaviour or performance.

The central question is therefore not simply whether a physical system can store information, but whether it can acquire and retain task-relevant information over the appropriate timescales. Which physical degrees of freedom can support transient responses, short-term adaptation and longer-term learning, and at what costs in energy, stability and controllability?

A second question concerns the rules that turn memory and feedback into adaptive behaviour. Suppose a device or network of devices has tunable parameters and some measure of its performance. How can internal dynamics, local feedback or supervisory circuits use this information to determine when and how the system should change? The challenge is not simply to calculate optimal settings externally and write them to hardware, but to develop mechanisms through which the physical system itself participates in selecting, testing and retaining useful changes. Such adaptation must also respect the operating constraints and stability requirements of the larger system.

These questions bring together materials science, device physics, dynamical systems, control theory and machine learning. Biological nervous systems provide powerful inspiration for adaptive, energy-efficient computation, but physical learning mechanisms need not reproduce biological neurons or synapses. Different electronic devices, materials and physical systems may offer their own ways of implementing memory, feedback and functional adaptation. One particularly ambitious long-term direction is *intelligent matter*: material systems in which sensing, response, memory and adaptive functions become closely intertwined within the physical substrate itself[12]. The broader scientific challenge is to identify the underlying principles of physical adaptation and learning, and establish how they can be understood, engineered and ultimately integrated into reliable autonomous systems.

When adaptive elements interact, new possibilities emerge

Autonomous adaptation becomes particularly interesting when multiple adaptive elements interact, whether at the level of individual devices, circuits, networks or larger physical systems. A single adaptive device may modify its transfer function, while interconnected elements may redistribute signals, exchange information or develop collective responses that cannot be understood from any element in isolation. Such interactions create opportunities, but also complications: local changes can cooperate, compete, amplify instabilities or undermine the behaviour of the system as a whole.

Understanding this behaviour requires more than characterising the individual components. Can an interconnected system adapt to changing tasks without a central controller computing every adjustment? Can local feedback produce a stable and useful global function? Can a distributed network compensate for the degradation of some of its elements? And how can we distinguish genuine collective adaptation from complex but preprogrammed responses?

These questions extend far beyond microelectronics. Mechanical metamaterials can redistribute forces and exhibit memory of deformation. Soft and interacting materials can change structure as they experience the environment. Nonlinear electronic and photonic networks can process signals through distributed physical interactions. Exploring these different experimental platforms in parallel can reveal general principles that no single technology would expose on its own.

It is useful to distinguish three connected research questions. The first concerns the physical origins of plasticity and memory: what can change, and what can be retained? The second concerns the mechanisms of autonomous adaptation and learning: how do internal feedback and experience select those changes? The third concerns collective adaptation and learning: what new functions emerge when many adaptive elements interact? These are not sequential stages, but interconnected questions that call for a combination of theory, simulation and experiment (Figure 4).

The broader ambition is to move beyond observing collective behaviour towards understanding how it can be deliberately shaped and exploited. Which principles allow local interactions to produce useful adaptive functions, and under what conditions are these functions stable, controllable and robust? Establishing such principles could provide a foundation for engineering adaptive systems across very different physical technologies.

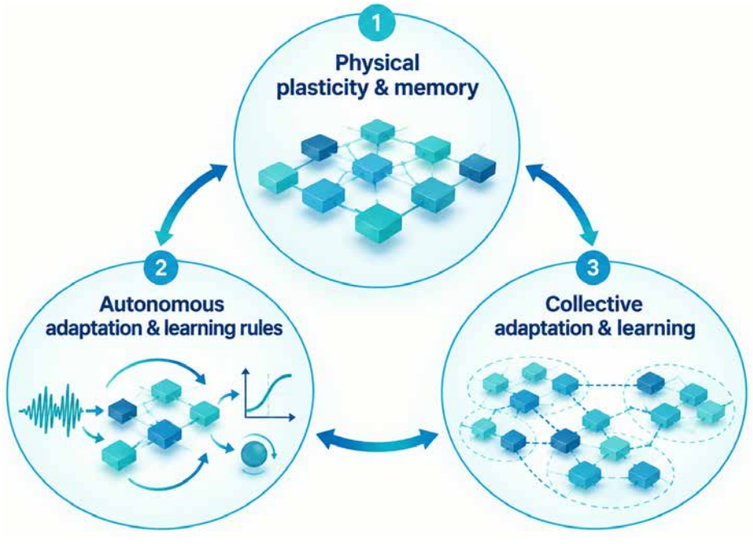


Figure 4: Three interconnected research questions towards autonomous physical systems capable of adaptation and learning. (1) Physical plasticity and memory: identifying the physical mechanisms that allow materials, devices and networks to change their properties and retain information over different timescales. (2) Autonomous adaptation and learning rules: understanding how feedback, internal dynamics and experience enable physical systems to adjust their computational functions without relying exclusively on external reprogramming. (3) Collective adaptation and learning: investigating how interactions among multiple adaptive elements can give rise to coordinated behaviour, emergent functionality and distributed learning. The bidirectional arrows emphasize that these three research questions are strongly interconnected rather than sequential stages. Together, they provide a framework for developing physical systems capable of autonomous adaptation and learning.

Europe's opportunity is to connect strengths, not fragment them

Understanding physical plasticity and collective adaptation is only part of the challenge. Turning these scientific insights into useful technologies requires close connections between fundamental research, semiconductor engineering and industrial development. Europe already possesses substantial strengths in microelectronics, advanced materials, photonics, neuromorphic computing, cyber-physical systems and industrial automation. Its universities, research institutes, semiconductor companies and start-ups contribute complementary expertise across the technological value chain. The opportunity lies in bringing these capabilities together, from fundamental discoveries in materials and devices to circuits, embedded computing platforms and complete autonomous machines.

Such connections must begin early. A promising physical computing device is not yet a useful autonomous system: its advantages must survive integration with sensing, control, communication and actuation. Turning novel physical effects into dependable technologies therefore requires suitable fabrication processes, circuit interfaces, compact models, training methods and software tools. For adaptive systems, reliable operation within defined constraints is essential. Performance must ultimately be demonstrated through realistic system-level benchmarks, including the costs of interfacing, calibration and adaptation. Close collaboration between device researchers, circuit designers, system engineers and application developers is therefore indispensable.

European collaborative programmes, including activities supported through the Chips Joint Undertaking, can help establish these connections through shared access to design tools, fabrication and characterisation facilities, pilot lines and experimental platforms. Different technologies will require different routes towards application. Some may be integrated into established semiconductor processes, while others may emerge as specialized chips, embedded modules or reconfigurable subsystems.

More exploratory approaches will require sustained fundamental research before their practical potential becomes clear. Priorities should be guided not only by scientific opportunities, but also by application requirements, technological maturity, manufacturability and credible pathways towards industrial adoption. Europe's strength should lie not in selecting a single winning computing paradigm, but in building the capacity to discover, develop and combine new computational capabilities into reliable technologies.

The future of autonomy is a physical-systems challenge

Computing must adapt to the physical world, not the other way around. Truly autonomous machines require more than fast information processing: they must respond to changing conditions, adapt their behaviour and ultimately learn from experience. Achieving this calls for a closer integration of computation with the physical processes of sensing, memory, control and actuation. The scientific opportunity is to harness physical dynamics not only for computation, but also for adaptation and learning through physical memory, feedback and interactions. The next frontier is not simply intelligence operating in the physical world, but intelligence shaped by its interaction with that world.

1 H. Liu, D. Guo and A. Cangelosi, Embodied Intelligence: A Synergy of Morphology, Action, Perception and Learning, ACM Computing Surveys 57, Article 186 (2025).

2 M. Zolfagharinejad, U. Alegre-Ibarra, T. Chen, S. Kinge and W. G. van der Wiel, Brain-inspired computing systems: a systematic literature review, The European Physical Journal B 97, 70 (2024).

3 D. Kudithipudi, C. Schuman, C. M. Vineyard, T. Pandit, C. Merkel, R. Kubendran, J. B. Aimone, G. Orchard, C. Mayr, R. Benosman, J. Hays, C. Young, C. Bartolozzi, A. Majumdar, S. G. Cardwell, M. Payvand, S. Buckley, S. Kulkarni, H. A. Gonzalez, G. Cauwenberghs, C. S. Thakur, A. Subramoney and S. Furber, Neuromorphic computing at scale, Nature 637, 801–812 (2025).

4 A. Singh, M. Le Gallo, A. Vasilopoulos, J. Luquin, P. Narayanan, G. W. Burr and A. Sebastian, The design of analogue in-memory computing tiles, Nature Electronics 8, 1156–1169 (2025).

5 H. Jaeger, B. Noheda and W. G. van der Wiel, Toward a formal theory for computing machines made out of whatever physics offers, Nature Communications 14, 4911 (2023).

6 A. Momeni, B. Rahmani, B. Scellier, L. G. Wright, P. L. McMahon, C. C. Wanjura, Y. Li, A. Skalli, N. G. Berloff, T. Onodera, I. Oguz, F. Morichetti, P. del Hougne, M. Le Gallo, A. Sebastian, A. Mirhoseini, C. Zhang, D. Marković, D. Brunner, C. Moser, S. Gigan, F. Marquardt, A. Ozcan, J. Grollier, A. J. Liu, D. Psaltis, A. Alù and R. Fleury, Training of physical neural networks, Nature 645, 53–61 (2025).

7 T. Chen, J. van Gelder, B. van de Ven, S. V. Amitonov, B. de Wilde, H.-C. Ruiz Euler, H. Broersma, P. A. Bobbert, F. A. Zwanenburg and W. G. van der Wiel, Classification with a disordered dopant-atom network in silicon, Nature 577, 341–345 (2020).

8 H.-C. Ruiz Euler, M. N. Boon, J. T. Wildeboer, B. van de Ven, T. Chen, H. Broersma, P. A. Bobbert and W. G. van der Wiel, A deep-learning approach to realizing functionality in nanoelectronic devices, Nature Nanotechnology 15, 992–998 (2020).

9 M. Zolfagharinejad, J. Büchel, L. Cassola, S. Kinge, G. S. Syed, A. Sebastian and W. G. van der Wiel, Analogue speech recognition based on physical computing, Nature 645, 886–892 (2025).

10 J. Kareem, L. Cassola, R. J. C. Cool, J. I. van Slooten, R. J. E. Hueting, P. A. Bobbert and W. G. van der Wiel, Space-Charge Effects in Silicon Reconfigurable Nonlinear-Processing Units, accepted for publication in Advanced Science, 10.1002/advs.78644 (2026); arXiv:2605.13477. https://doi.org/10.48550/arXiv.2605.13477

11 M. Escudero, M. Zolfagharinejad, S. van den Belt, N. Alachiotis and W. G. van der Wiel, Physical Analogue Kolmogorov-Arnold Networks based on Reconfigurable Nonlinear-Processing Units, arXiv preprint arXiv:2602.07518 (2026). https://doi.org/10.48550/arXiv.2602.07518

12 C. Kaspar, B. J. Ravoo, W. G. van der Wiel, S. V. Wegner and W. H. P. Pernice, The rise of intelligent matter, Nature 594, 345–355 (2021).

Research Project Highlight

INSIDE Magazine 15

Research Project Highlight



Krzysztof Wołk
DAC.next

When wearables become decision makers

Turning imperfect signals into advice we can trust

Wearable AI could help people understand activity, recognise patterns and make better choices. Useful autonomy depends on knowing what the signals mean, explaining uncertainty and giving people an effective role in the decision.

Imagine finishing a run and checking your watch. It has recorded your activity and produced a reassuring score. You feel exhausted. The screen gives you one view of the session, your body a different one. How much weight should you give each?

That small disagreement captures a central challenge for wearable intelligence. A device can observe us repeatedly, in places where laboratory equipment cannot follow. Yet it only sees part of the situation. Turning its measurements into useful advice requires assumptions about the person, the activity and the information available.

The potential benefits are substantial. Athletes could understand how their training changes over time. People could notice patterns worth discussing with a professional. Carefully evaluated monitoring services could add useful information between appointments. In each case, the opportunity is to connect observations with a decision someone actually needs to make.

The stakes rise as software moves from recording activity towards recommending action. Even an advisory system can influence behaviour without directly controlling a machine. Wearables therefore offer an everyday example of a much wider question about autonomous systems: what should a machine do when its view of the world is incomplete?

A number needs context

A heart rate reading becomes easier to interpret when we know what the wearer is doing. Data fusion combines information from different sensors into a more useful account of the same event. In feature fusion, quantities such as heart rate, breathing rate and speed become joint inputs to a model, helping it distinguish situations that look similar through one signal.

Time alignment places observations on a common timeline before they are combined. A breathing measurement may be several seconds older than the movement reading beside it. A sensor may lose contact, or an interrupted connection may leave repeated values. Combining those inputs without preserving timestamps and missing values can produce an interpretation of an event that never occurred in that form.

The system therefore needs a basic account of its own observations. When was each value measured? Is it plausible? Does it belong with the other readings? A missing measurement must remain distinguishable from a genuine zero. For users, those engineering details should translate into guidance grounded in relevant information.

It also matters whether a number was measured or estimated. Several accessible signals might support an estimate of something that normally requires cumbersome equipment. Such an estimate can be valuable, provided its presentation makes that distinction clear. A precise decimal display does not make the underlying knowledge equally precise.

Comparing similar activities over time can support better conversations with coaches or professionals. The record should preserve the circumstances and gaps needed to interpret a pattern.

The consequences of a convincing mistake

Consider a proposed monitoring feature that flags unusual patterns. Reacting to every small deviation could create repeated interruptions. Making the feature less sensitive could allow important changes to pass unnoticed. Choosing that balance requires a clear intended use and an understanding of the consequences of each kind of error.

The interface adds another layer. A bright warning may cause concern before its explanation is read. A green indicator can



Figure 1: Wearable advice has to make sense in everyday activity.

imply reassurance that the measurements do not support. Repeated unnecessary alerts may encourage people to dismiss the next one. Developers need to evaluate the decision prompted by a message as well as the calculation behind it.

Personalisation can help interpret an individual's patterns, but it also needs limits. A system should not assume that an unfamiliar situation is harmless simply because it resembles someone's usual data. Nor should a training score become a universal judgement about wellbeing. The wearer may know something about the circumstances that no sensor has captured.

Giving people control means more than adding a confirmation button. They need an understandable explanation, a way to correct relevant context and a clear route to question an output. Where professional interpretation is required, the service should make that responsibility visible.

Confidence deserves particular care. A percentage might describe a model calculation, a signal quality check or simply an interface convention. Those meanings are not interchangeable. Developers need evidence before presenting it as the likelihood that advice is correct. Showing which input is missing may give someone a more useful next step.

Responding when information is missing

A missing signal need not become a misleading recommendation. The system can ask for another measurement, suggest checking the device or wait until the input improves. The difficult part is recognising the limitation before presenting a result that looks complete.

Imagine an application that normally combines heart rate, breathing and movement. If breathing data disappears, a separately tested model using fewer inputs might still provide a useful estimate. The change should be visible. Without evidence that the reduced input is adequate, withholding the estimate may be the more informative response.

The wording matters. 'Unable to estimate during this interval' communicates something different from 'No problem detected'. A gap describes the measurement process. It does not establish what happened to the person. Likewise, the last valid number should not continue to look current after its supporting data has stopped arriving.

Recovery needs attention too. One plausible reading after an interruption may not be enough to resume normal operation. Developers should define what evidence allows the system to continue, and test that behaviour. The same principle applies to an inspection robot that loses a clear view: its next action should reflect what it can still observe and the consequences of being wrong.

A research example that makes the problem visible

EU-TRAINS provides a practical setting for exploring these questions. Coordinated by Silicon Austria Labs, the European project brings together sensing, smart textiles and AI for multisport training and remote healthcare. Its goals include measurements in air and water, with partners working across materials, devices, software and user assessment.[1]

Within this work, DAC.next is exploring how heart rate, breathing rate, speed and other sensor readings can support an estimate of the rate of oxygen consumption, written as VO2. The development workflow uses recorded treadmill exercise data from the University of Malaga, available through PhysioNet. The oxygen consumption measured during the original laboratory test provides a reference for comparison.[2]

The current analytics notebook groups readings into one second bins on a common elapsed timeline, averaging available values within each second and preserving gaps. It combines three inputs, heart rate, breathing rate and speed, in a Random Forest regression model, an ensemble of decision trees that estimates VO2. Reference VO2 is used for training and evaluation, never as a prediction input.

The experiments point to three practical gains. Data fusion reduced prediction error compared with heart rate alone on recording sessions reserved for testing. Average prediction processing in notebook benchmarks was comfortably below the one second input interval. ONNX export retained essentially the same benchmark accuracy. These results connect better estimates with efficient, portable execution.

Replaying a recorded session lets developers inspect the same sequence repeatedly. They can compare input combinations, investigate transitions and introduce controlled disturbances. A development demonstration recorded in 2026 included a simulated disturbance in speed readings and an explicit message about the resulting data limitation. The pipeline view shows the processing stages, input availability and predicted VO2 alongside the recorded reference, helping developers distinguish an available prediction from a complete set of inputs. When heart rate disappears, the replay pauses the VO2 estimate.

DAC.next also packages model variants for TFLite and ONNX with input order, units, scaling parameters and example predictions. These details make integration testable. A correct model can still produce a wrong answer if an application supplies values in the wrong order. This connects model development to firmware engineering.[3]

Replay uses recorded measurements, and the notebook and interface remain distinct development artefacts. Some interface indicators, including confidence percentages, remain illustrative. Target device timing, energy use and application validation require separate testing.

Watch the demonstrator in action. A short replay of recorded data shows input checks, speed and heart rate disturbances, recovery and the session summary.

Useful intelligence has to fit the device

A wearable can be technically impressive and still be inconvenient. Frequent charging, discomfort or dependence on an unavailable connection can determine whether it is used consistently. A richer analysis has little practical value if the equipment stays in a drawer.

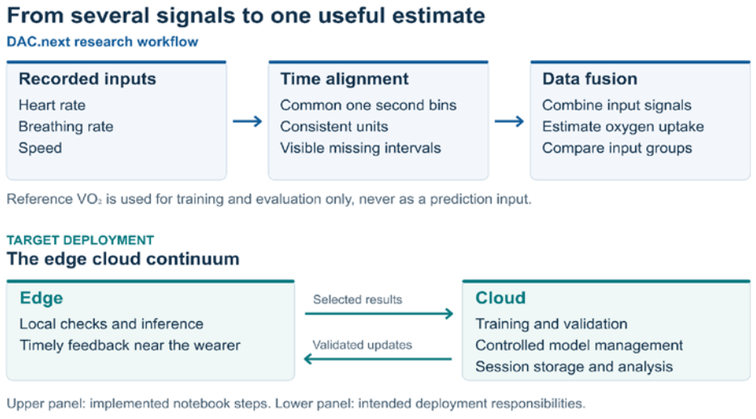


Figure 2: From aligned measurements to data fusion. DAC.next's research workflow provides concrete building blocks for the target edge cloud continuum in the lower panel.

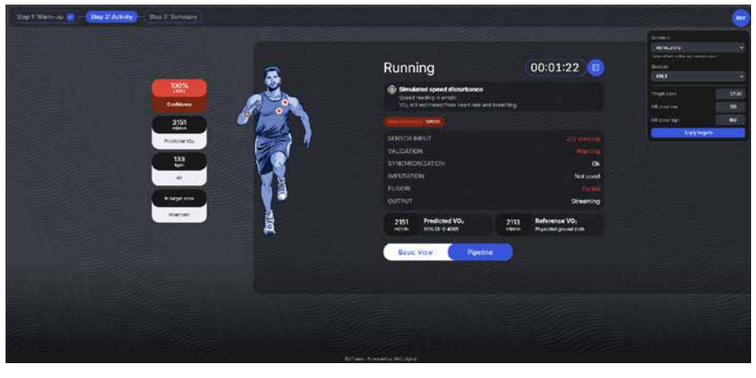


Figure 3: DAC.next development interface during recorded PhysioNet data replay with a simulated speed disturbance. Interface indicators are illustrative and do not establish validated model confidence or clinical performance.

Processing some information on the device, often called Edge AI, can support feedback without sending every measurement to a remote service first. This is particularly attractive where connectivity is difficult, including aquatic activities. The important design question is which functions must remain available locally and what the user should expect when communication is interrupted.

Local processing has costs of its own. Sensing, calculation, storage and radio communication all compete for battery capacity. A more complex model is worthwhile only when its improvement matters for the intended task and the complete device can sustain it. A small model file alone says little about the energy required during use.

This arrangement is often called the edge cloud continuum: a coordinated chain from sensors and nearby computation to remote services. Local checks and inference can support timely feedback, while cloud resources support session analysis, model training and controlled updates. The division depends on latency, energy, privacy and connectivity. For EU-TRAINS, this is a target architecture. DAC.next's common data representation and portable model packages prepare models for that split. Operation on the intended devices and across interrupted connections still needs testing.

Comfort and maintenance also affect measurement quality. Teams should observe how people position, clean and use a prototype. A sensor that performs well when carefully arranged may demand too much effort from its wearer to deliver consistent benefits in everyday activity.

Privacy and security are part of the experience

Repeated physiological measurements can reveal patterns about activity and daily routines. People need to understand what is collected, who can access it and how long it is kept. A useful service should be able to explain why each category of information is necessary.

Keeping some analysis on the device can reduce what is transmitted, but privacy depends on the wider design. Stored records still need protection. Account access, connected services and software updates can all affect the information entrusted to the system. Collecting less can reduce exposure when additional data offers little benefit.

Security also concerns the meaning of an output. An update that changes how an input is interpreted can change a recommendation even when the screen looks familiar. Model versions, signal definitions and application behaviour need to remain compatible. Teams should be able to trace an unexpected result to the software and assumptions that produced it.

These responsibilities continue after a demonstration. Maintaining a dependable service includes explaining changes and providing a practical response when a function is unavailable or withdrawn. Trust grows when the system's commitments remain understandable throughout its use.

Evidence must follow the intended use

A model can perform well on familiar data and struggle with a new wearer. If a service

is intended for people whose measurements were never used during development, its evaluation needs to reflect that situation. Testing another session from a familiar person does not answer the same question.

The conditions matter as much as the participants. Steady exercise may be easier to interpret than a transition between activities. Imperfect contact and missing readings should be included when they are part of expected use. An average result can hide a weakness concentrated at precisely the moment someone needs guidance.

Testing must also reach the person reading the screen. Can users distinguish a missing value from a reassuring result? Do they understand whether an output is measured or estimated? Can they act on an explanation without specialist knowledge? Clear communication is something to evaluate with people, not a property established by choosing simple words.

Different stages answer different questions. Replay supports repeatability and software inspection. Device tests establish behaviour under physical and resource constraints. Studies for a specific health application need evidence appropriate to that purpose. A promising result at one stage should guide the next investigation without being presented as proof that every later stage has already been completed.

Evaluation should record both prediction quality and periods without an answer. Together, these reveal whether the next improvement belongs in the model, sensor connection or explanation of an interruption.

Sharing responsibility intelligently

For Europe's intelligent systems community, wearables bring several disciplines into the same practical problem. Sensor specialists, software engineers, model developers and researchers studying user behaviour must agree on what the system knows and what it may do with that knowledge. Shared input definitions and reproducible tests make those agreements concrete.

The lessons extend to robots and industrial monitoring. Greater autonomy should come with clearer boundaries, visible limitations and a workable way to involve a person. A system earns additional responsibility through evidence about its behaviour under the conditions where that responsibility will matter.

For the runner checking a watch, useful intelligence might mean a helpful trend, an explanation of an unexpected reading or a request for better information. A trustworthy wearable should be able to offer each when appropriate. Its value lies in helping someone make a better decision, including recognising when the machine does not yet have enough information to help.

Funding acknowledgement

EU-TRAINS is funded by the European Union under the Horizon Europe research and innovation programme, grant agreement No. 101130495.

Krzysztof Wołk
DAC.next

Krzysztof Wołk, PhD, is a researcher at DAC.next working on artificial intelligence and its practical applications. His interests include machine learning, the analysis of different types of data and collaboration within European research projects. His work explores how AI can help turn complex measurements into information that people and intelligent systems can use.



Figure 4: Aquatic activity illustrates the challenge of providing useful feedback when communication is difficult.

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

Funded by the European Union

1 European Commission. EU-TRAINS project factsheet.

2 Mongin, García Romero and Alvero Cruz. Treadmill exercise dataset, version 1.0.0. PhysioNet, 2021.

3 DAC.next. Contribution to EU-TRAINS.

Follow us on Social Media!

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



company/inside-association/

INSIDE
Industry Association

Research Project Highlight

Research Project Highlight

A digital twin–based framework for robotic control strategy benchmarking in On-Orbit Servicing

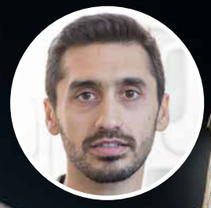
How digital twin technologies can support resilient control strategy selection and refinement for autonomous On-Orbit Servicing missions



Valentina Pericu
Leonardo SpA



Andrea Monguzzi
Leonardo SpA



Navvab Kashiri
Leonardo SpA



Ricardo Caldas
Gran Sasso Science Institute



Patrizio Pelliccione
Gran Sasso Science Institute

The rapid expansion of space-based services and satellite constellations is reshaping Earth's orbital environment. Alongside this growth, the accumulation of non-operational satellites and orbital debris has emerged as a critical challenge, threatening both the sustainability of space activities and the economic value of existing assets. On-Orbit Servicing (OOS) is increasingly recognized as a key enabler to address this issue, allowing satellites to be refueled, maintained, or safely removed at the end of their operational life through robotic intervention. OOS missions take place in highly unstructured environments, where uncertainties in sensing and interaction with service-ready and/or non-cooperative targets can significantly affect system behavior and mission outcomes. As a consequence, control strategy performance becomes a core factor for safe and reliable operations, and selecting the appropriate strategy is a non-trivial task. Digital Twin (DT) technologies offer a systematic means to benchmark and compare different control strategies under representative and repeatable conditions before deployment, while also supporting their refinement and monitoring as the system operates under changing real-world conditions. This article presents a DT-based framework that enables informed control strategy selection before deployment and continuous support during operation through runtime monitoring and supervised strategy refinement, contributing to safer, more reliable, and more sustainable OOS missions. The article describes the framework architecture, establishing the foundation for systematic empirical validation.

Space is getting crowded… and robotics is becoming essential

Space has become a critical infrastructure for modern society. Communication, navigation, Earth observation, climate monitoring, and scientific exploration increasingly rely on large numbers of satellites operating in orbit. Over the past decade, the rapid expansion of commercial constellations and institutional missions has significantly increased the density of objects in Earth orbit, particularly in low Earth orbit (LEO)[1].

However, this growth comes with a hidden cost. Satellites that reach the end of their operational life, suffer failures, or exhaust their fuel often remain in orbit as uncontrolled objects. These non-operational satellites represent not only a growing source of space debris, but also a substantial economic loss, as valuable assets become unusable long before their technological potential is fully exploited. Collisions, even with small debris, can trigger cascading effects, generating further fragments and amplifying the risk for operational missions. This phenomenon, commonly referred to as the Kessler syndrome, has been extensively studied and is now widely recognized as a limiting factor for future space operations[2].

Institutional analyses indicate that mitigation measures alone, such as post-mission disposal, may not be sufficient to stabilize the

orbital environment[3]. In this context, On-Orbit Servicing (OOS) is emerging as a strategic alternative to the traditional "launch-and-replace" paradigm. Extending the operational lifetime of satellites through refueling, maintenance, or relocation can significantly reduce costs, mitigate debris generation, and improve the overall sustainability of space activities.

Among the various OOS missions currently under development, satellite refueling and debris removal stand out as representative, high-impact use cases and form the primary focus of this article. In both scenarios, a servicer satellite, typically equipped with one or more robotic manipulators, must physically interact with a client or target object. Depending on its design and operational state, the target may be cooperative or non-cooperative, and service-ready or non-service-ready. A cooperative satellite actively supports the servicing operation, for example, by maintaining attitude stability, providing navigation aids, or sharing telemetry data. In contrast, a non-cooperative target does not provide such assistance; it may be tumbling, unresponsive, or lacking reliable state information.

The distinction between service-ready and non-service-ready satellites concerns their design features. A service-ready satellite is intentionally equipped with standardized interfaces or dedicated hardware (such as refueling ports, grapple fixtures, or visual fiducials) specifically intended to facilitate future OOS. A non-service-ready satellite, by contrast, was not designed with OOS in mind and therefore lacks dedicated interfaces, significantly increasing the technical complexity of the servicing operation. In a refueling mission, the servicer performs rendezvous with the client satellite, carries out proximity operations, connects, and transfers propellant to a satellite with limited or no remaining propellant. In debris removal missions, the servicer approaches a typically non-cooperative object, such as debris or a defunct satellite, captures it through a dedicated mechanism, and guides it to a disposal orbit or a controlled reentry.

Human intervention in orbit is impractical for most servicing scenarios due to the risks associated with extravehicular activities and the limitations of direct teleoperation from the ground, including communication delays and coverage constraints[4 5]. As a result, OOS missions rely on autonomous or semi-autonomous robotic systems, with humans typically involved at supervisory and decision-making levels from the ground control station.

In this context, ensuring that robotic control strategies remain effective throughout the mission lifecycle, not only at design time, is a fundamental challenge. This article argues that a *resilience*-oriented perspective, supported by Digital Twin technologies, provides a principled way to address this challenge.

When

Unlike industrial robotics, OOS operations take place in environments that differ fundamentally from structured factory settings. In microgravity, even small forces can lead to significant changes in motion, while orbital dynamics introduce complex relative movements between the servicer satellite and the target object.

A key challenge arises from the nature of the targets themselves. As mentioned, depending on the mission, targets may be cooperative, non-cooperative, service-ready, or non-service-ready. Under these circumstances, robotic performance becomes highly sensitive to environmental conditions and modeling inaccuracies, a point widely recognized in the space robotics literature[4 5]. Even small deviations in environmental conditions or interaction dynamics can lead to significantly different system behavior, leaving limited margins for error. As a consequence, mission success in highly unstructured environments is tightly coupled to the performance of the selected control strategy.

However, the behavior of robotic systems in these environments cannot be fully defined through pre-mission design choices alone. While control strategies and their parameters can be optimized for anticipated scenarios during the design phase, real-world operations often demand flexibility to deal with unexpected conditions and deviations from nominal assumptions. This trade-off between optimality and flexibility—and, more broadly, resilience—is a core theme in space robotics[6] and highlights the need for approaches that support adaptation to ensure mission continuity when operating conditions evolve.

Beyond flexibility, we argue that OOS missions require ***resilient*** control strategies—ones that enable the system to ***avoid***, ***withstand***, ***recover*** from, and ***adapt*** to conditions that degrade controller performance[7]. Resilience shifts the goal from finding the single best controller to maintaining acceptable performance across the full operational lifecycle, including conditions not anticipated at design time.

From this perspective, control strategy selection should not be regarded as a one-time decision, but rather as an ongoing challenge that spans both pre-deployment and operational phases.

Choosing the right controller is harder than it seems

At the core of robotic autonomy lies the control strategy. Numerous control approaches have been proposed for rendezvous, manipulation, and physical interaction in space robotics, each offering different trade-offs in terms of stability, robustness, and performance. As widely discussed in the literature[4], no single control strategy can be considered universally optimal across the wide range of OOS scenarios. Representative approaches include impedance/admittance control for compliant physical interaction, model predictive control for constrained trajectory optimization under dynamic uncertainty, adaptive control schemes designed to compensate for parametric variations and target motion[4], and reactive whole-body control frameworks that combine adaptive-predictive visual servoing with full-body motion optimization for dynamic target tracking[8].

The effectiveness of a given controller strongly depends on the system state, the interaction phase, the dynamic properties of the target, and the surrounding environmental conditions. A controller configuration that performs well under a specific operating regime may exhibit degraded behavior when conditions change, requiring parameter retuning or alternative control formulations.

Testing and tuning control strategies directly on the actual robotic system operating in orbit is often impractical. Real-world experiments are costly, time-consuming, and carry inherent risks, particularly when dealing with space-qualified hardware. Moreover, repeating experiments under identical conditions is rarely possible, limiting the ability to perform systematic comparisons. These limitations motivate the need for a structured framework capable of evaluating and benchmarking control strategies based on representative operating conditions, while also supporting informed reconfiguration throughout the mission lifecycle.

This challenge is well-recognized in systems engineering, where runtime assessment and continuous verification have emerged as key enablers of resilient operation in robotic systems[9 10]. These limitations are precisely the conditions that a resilience-oriented framework must address: *avoiding* sub-optimal strategy selection at design time, and providing the means to *detect* degradation, *withstand* transient failures, *recover* through informed strategy switching, and *adapt* as the operational envelope evolves.

A digital twin framework for control strategy optimization

Digital Twin (DT) technologies offer a promising approach to address the challenges associated with control strategy selection and adaptation in OOS missions. Originally conceptualized as a mirror-space model linking physical and virtual representations[11], the DT paradigm has since evolved into a broad modeling and simulation framework for cyber-physical systems[12]. By enabling a tight coupling between digital representations and the physical system, DTs support systematic evaluation of control strategies and informed decision-making across the system lifecycle.

In the proposed framework, shown in Figure 1, the physical robotic system executing the servicing task in the operational environment is referred to as the Actual System (AS). Its digital counterpart originates from a Digital Model (DM), which captures the relevant geometric, kinematic, and dynamic properties of the system, together with the interfaces required to execute and evaluate candidate control strategies. When instantiated with the current configuration and operating conditions of the AS, the DM gives rise to a Digital Shadow (DS), representing a state-consistent snapshot of the system and of the current environmental conditions. When the digital representation is further kept continuously synchronized with the running AS through live sensor feedback, it transitions into a fully operational Digital Twin (DT) proper. Note that the term *Digital Twin* is used in two complementary senses in this paper: it refers both to this specific synchronized instantiation and, more broadly, to the overarching DT-based framework; the intended sense is clear from context.

During the pre-deployment phase, multiple control strategies are executed and evaluated within the DS under identical initial conditions. This benchmarking phase enables objective comparison across different interaction scenarios and supports the selection of the most suitable control strategy for the expected operating conditions prior to deployment. Performance is assessed against predefined system requirements using metrics such as end-effector pose error, contact force profiles, joint torque envelopes, and task completion time, generating structured reports that enable quantitative comparison across candidate strategies.

Once a control strategy has been selected, it is deployed on the AS. During operation, the digital representation is continuously updated using runtime data from the AS, thereby acting as a DT. In this operational phase, the DT provides an up-to-date representation of the system state, supporting monitoring, performance assessment, and operator supervision throughout task execution.

Importantly, the role of the DT is not limited to pre-deployment validation. As the system operates in the real environment, the DT enables continuous data collection and analysis, allowing operators to assess whether the deployed control strategy remains appropriate as operating conditions evolve. As illustrated in Figure 1, when runtime monitoring detects performance degradation, the DT supports a supervised refinement process on the ground: candidate strategies are re-evaluated against the updated system state, and an improved strategy can be deployed back to the AS. In this sense, benchmarking can be extended beyond design time, supporting informed decisions on strategy refinement and replacement during the mission, closing the

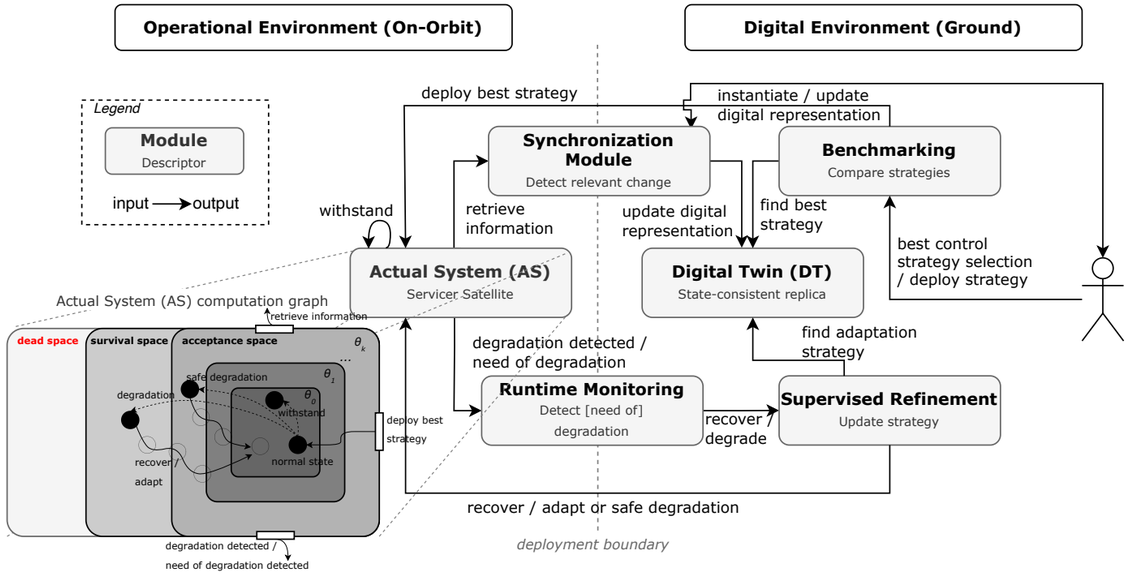


Figure 1: Conceptual view of the continuous dialogue between the physical system in orbit and the Digital Twin on ground, annotated with the resilience capacities supported by the framework: avoid, detect, withstand, recover, and adapt.

loop between orbit and ground. Changes to the control strategy are assumed to be executed through safe and supervised procedures, with humans remaining in the decision loop. While the system is not assumed to autonomously modify the deployed control strategy, the DT provides the necessary situational awareness and performance insight to determine when such changes may be beneficial.

By maintaining a continuous digital link between pre-deployment validation and in-orbit operation, the proposed framework supports informed control strategy selection, monitoring, and supervised refinement, increasing confidence in the behavior of robotic systems performing OOS missions.

Framed through the lens of resilience[7], the proposed framework addresses the five core capacities required for resilient operation. Pre-deployment benchmarking enables the system to ***avoid*** deploying unsuitable control strategies. Runtime monitoring supports ***detection*** of performance degradation, enabling the system to ***withstand*** transient disturbances without catastrophic failure.

Supervised refinement enables ***recovery*** by facilitating informed strategy switching when degradation is confirmed. Finally, the continuous data collection and re-evaluation cycle supports long-term ***adaptation*** as the operational envelope evolves over time.

Reaching space starts on earth: testbeds for validating On-Orbit Servicing control

To support experimental investigation of OOS scenarios, two laboratory testbeds have been developed to model the refueling and debris removal use cases, as shown in Figures 2a and 2b, respectively. In the proposed framework, each testbed constitutes the Actual System (AS) from which a Digital Model is constructed, capturing the kinematic and dynamic properties of the servicer and the target. This model is instantiated as a Digital Shadow under experimentally characterized operating conditions, providing a repeatable simulation environment in which candidate control strategies are evaluated and compared. Runtime sensor data collected during physical experiments then promote the digital representation to a full Digital Twin, enabling the monitoring and supervised refinement cycle described in the previous section.

The refueling testbed consists of a robotic manipulator mounted on an omnidirectional mobile platform, representing the servicer satellite, which approaches a client satellite modeled by a second robotic manipulator mounted on a linear axis with a controlled motion. The servicer is equipped with vision and force sensors, as well as a refueling interface installed on the flange of its robotic arm. The client is similarly equipped with a corresponding refueling interface mounted on the flange of its manipulator, enabling proper mechanical docking and alignment during the refueling operation. The debris removal use case is represented by a testbed employing the same mobile robotic manipulator as the servicer satellite, interacting with the client satellite modeled as an object moving in an uncontrolled manner. In this scenario, the servicer is equipped with vision and force sensing capabilities and a dedicated gripper to capture the target.

In both testbeds, no data exchange is assumed between the servicer and the client satellites; even in the refueling scenario, where the client is cooperative, the servicer relies solely on its onboard sensing capabilities. The servicer is equipped with an onboard camera and estimates the client state exclusively through vision-based perception algorithms. Together, the two testbeds span the cooperative-to-non-cooperative operational spectrum, ensuring that benchmarked control strategies are evaluated against the conditions most relevant to the *avoid* and ***adapt*** capacities of the resilience framework introduced in the previous section (“When”).

Conclusions

The increasing congestion of the orbital environment is making OOS an essential capability for future space operations. Missions such as satellite refueling and debris removal require robotic systems to perform physical interactions in highly uncertain environments, where control strategy performance has a decisive impact on mission success. As space operations progress toward higher levels of autonomy and complexity, DT-enabled frameworks are expected to play a key role in supporting safe, reliable, and sustainable OOS—not merely by optimizing control strategies before deployment, but by enabling resilient operation throughout the mission lifecycle.

This article has presented a DT-based framework that supports systematic benchmarking of control strategies prior to deployment and provides continuous support during operation through runtime monitoring and supervised strategy refinement. By enabling structured evaluation, monitoring, and supervised refinement of control strategies as operating conditions evolve, the proposed approach contributes to reducing operational risk and increasing confidence in robotic servicing missions. The framework has been grounded in a resilience formulation that explicitly maps pre-deployment benchmarking, runtime detection and withstanding, supervised recovery, and long-term adaptation to the five core capacities required for resilient operation.

The present article describes the framework architecture and the testbed infrastructure; quantitative benchmarking and validation of the full supervised refinement loop constitute the subject of ongoing and future work.

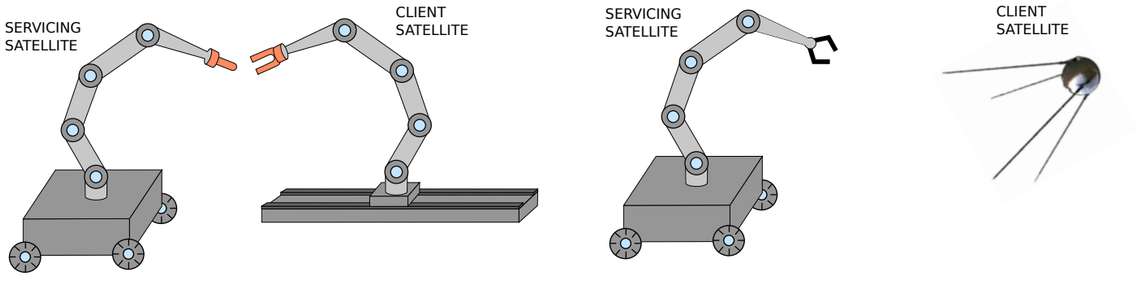


Figure 2: Experimental testbeds: (a) A robotic manipulator on a mobile platform and a manipulator on a linear axis represent the servicer and client satellites, respectively, within the Digital Twin–based benchmarking framework; (b) A robotic manipulator on a mobile platform and a debris mock-up represent the servicer and target object, respectively, within the Digital Twin–based benchmarking framework.

Acknowledgment

This research work has been funded by the Key Digital Technologies Joint Undertaking (KDT JU) through the project MATISSE, grant agreement No. 101140216.

1 J. Mulvaney, D. Arney, C. Williams, J. Morel, C. Stockdale, C. Whitlock, and V. Balaji, “In-space servicing, assembly, and manufacturing (isam) state of play-2025 edition,” 2025.

2 D. J. Kessler, N. L. Johnson, J. Liou, and M. Matney, “The kessler syndrome: implications to future space operations,” Advances in the Astronautical Sciences, vol. 137, no. 8, p. 2010, 2010.

3 C. Bonnal, J.-M. Ruault, and M.-C. Desjean, “Active debris removal: current status of activities in cnes,” in IAF workshop on space debris removal, 2013.

4 A. Flores-Abad, O. Ma, K. Pham, and S. Ulrich, “A review of space robotics technologies for on-orbit servicing,” Progress in aerospace sciences, vol. 68, pp. 1–26, 2014.

5 N. Fallahiarezoodar and Z. H. Zhu, “Review of autonomous space robotic manipulators for on-orbit servicing and active debris removal,” Space: Science & Technology, vol. 5, p. 0291, 2025.

6 M. Alizadeh and Z. H. Zhu, “A comprehensive survey of space robotic manipulators for on-orbit servicing,” Frontiers in Robotics and AI, vol. 11, p. 1470950, 2024.

7 J. Andersson, V. Grassi, R. Mirandola, and D. Perez-Palacin, “A conceptual framework for resilience: fundamental definitions, strategies and metrics,” Computing, vol. 103, pp. 559–588, 2021.

8 A. Monguzzi, G. Alfonso, and N. Kashiri, “Reactive whole-body control of mobile manipulators for dynamic target tracking via adaptive-predictive visual servoing,” in IEEE International Conference on Robotics and Automation (ICRA), Vienna, Austria, 2026, accepted for publication.

9 R. Caldas et al., “Runtime verification and field-based testing for ROS-based robotic systems,” IEEE Transactions on Software Engineering, vol. 50, no. 10, 2024.

10 M. Leucker and C. Schallhart, “A brief account of runtime verification,” The Journal of Logic and Algebraic Programming, vol. 78, no. 5, pp. 293–303, 2009.

11 M. Grieves and J. Vickers, “Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems,” in Transdisciplinary Perspectives on Complex Systems. Springer, 2017, pp. 85–113.

12 A. Rasheed, O. San, and T. Kvamsdal, “Digital twin: Values, challenges and enablers from a modeling perspective,” IEEE Access, vol. 8, pp. 21980–22012, 2020.

Valentina Pericu
Robotics Engineer, Robotics Laboratory of Innovation Hub & Intellectual Property, Leonardo SpA.

Valentina Pericu holds degrees in Biomedical Engineering and Robotics Engineering from the University of Genova. She previously worked at the Istituto Italiano di Tecnologia (IIT) and is currently a Robotics Researcher at Leonardo SpA, where her research focuses on navigation and perception technologies for autonomous robotic systems and robotics applications.

Andrea Monguzzi
Technical Coordinator, Leonardo SpA.

Andrea Monguzzi holds a PhD in Robotics from Politecnico di Milano and is Technical Coordinator of the Robotics Manipulation and Motion Control research area at the Leonardo Robotics Laboratory. His research focuses on robotic manipulation and motion control, with applications in intelligent robotic automation and space robotic servicing, aiming to enhance the capabilities of robotic arms in complex tasks.

Navvab Kashiri
Head of Robotics Lab, Leonardo SpA

Navvab Kashiri is Principal Investigator of the Robotics Lab at Leonardo Innovation Hub. He received his Ph.D. in Robotics from the Istituto Italiano di Tecnologia in 2015, with a focus on the development and control of robotic platforms. He has since been actively engaged in national and international collaborations advancing robotic systems with enhanced capabilities, supporting Leonardo’s mission to develop next-generation solutions for adaptability and autonomy.

Ricardo Caldas
Postdoc Researcher, Gran Sasso Science Institute (GSSI)

Ricardo D. Caldas is a Postdoctoral Researcher in Robotics Software Engineering at the Gran Sasso Science Institute (GSSI) in L’Aquila, Italy. His research focuses on making autonomous robots more resilient through runtime verification, digital twins, scenario-based testing, and adaptable behavior trees. He holds a Ph.D. from Chalmers University of Technology.

Patrizio Pelliccione
Director of the Computer Science area and Professor in Software Engineering, Gran Sasso Science Institute (GSSI)

Patrizio Pelliccione is a Full Professor and Director of Computer Science at Gran Sasso Science Institute, Italy, and Adjunct Professor at the University of Bergen, Norway. His research focuses on software engineering, software architecture modeling and verification, and autonomous systems. He has held academic positions across Europe and actively contributes to leading conferences, journals, research projects, and industry collaborations.

Research Project Highlight

From pilot to production

Five lessons from 41 months of industrial AI



Thomas Gutt
Infineon Technologies AG, AIMS5.0 coordinator

In Dresden, a wafer-defect classifier developed in AIMS5.0 is now part of Infineon's production workflow. It is one result of 41 months in which 53 partners from 12 countries applied AI across 20 industrial use cases to make European production more efficient and sustainable. Five lessons stand out from that work.

Classifier performance depends on the data

The Dresden classifier reached 92.7% validation accuracy across 37 defect classes. The team refined the dataset over three iterations and left the neural network unchanged. In operation, experts review a sample of images. If disagreement exceeds a set threshold, images are classified manually until a retrained model is deployed.

Model maintenance is part of deployment

When equipment and processes change, the models that monitor them have to change too. In a use case involving LFoundry, the University of Padova and Statwolf, the partners worked on data integration, model versioning and monitoring. They reported a 75% reduction in model development time and 25% less data needed when extending models to new machines.

Engineers need explanations they can use

The University of Padova's AcME-AD method shows which variables contribute to an anomaly alert. In the same use case, explainable monitoring cut root-cause identification time by 43% on four sputtering machines. Operators can validate, reject or annotate alerts, so their process knowledge feeds into the model's further development.

Scaling beyond the pilot needs more data and resources

The final review named transfer into high-volume production as a continuing challenge, including the risk that developing, scaling and maintaining demonstrators could require more resources than planned. In Use Case 3, on servitisation and service delivery, high accuracy at one site did not yet transfer to other sites; the partners concluded that robust generalisation needs multi-site fleet data.



AIMS5.0 consortium at the final review in Vienna, 8–9 October 2026. Photo: AIMS5.0 consortium.

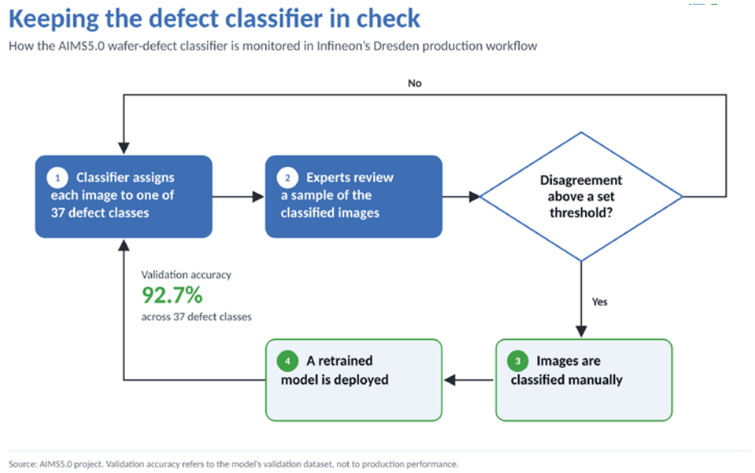


Figure 1: Monitoring the wafer-defect classifier in operation.



Figure 2: AIMS5.0 key figures at the final review.

Joint development reaches beyond the consortium

The Digital Reference ontology combines 17 of the 20 use-case ontologies under the Apache 2.0 licence, and Eclipse Arrowhead 5.2 is used in several use cases. Of all results, 85 were developed jointly, 62 of them by industry with universities or research organisations. Cooperation also ran across Chips JU projects: AIMS5.0, SC4EU and Arrowhead fPVN held a joint Deep Tech workshop in Lisbon, and a SysML v2 library for the Arrowhead framework was presented to the OMG standards community.

Thomas Gutt, Infineon Technologies, coordinator of AIMS5.0 and SC4EU, explains that "With 53 partners and 20 use cases, much of the project management was about connecting the work. Regular reviews of tasks, risks and KPIs show ed us where a result from one team could strengthen another use case, and where partners needed to develop something together. Beyond the consortium, the joint Deep Tech workshops with SC4EU and Arrowhead fPVN and the INSIDE and Chips JU events helped us exchange lessons learned. That is where results from different projects meet and where the next European R&D projects take shape."

What happens to the results

Of the 179 results registered at the final review, 154 were implemented or under test. The 39 key exploitable results were selected from the 146 demonstrated above TRL 4, scored on innovation, exploitability and impact: eight horizontal enablers and 31 use-case results, each with a stated route to use.

In Use Case 3, Husqvarna applies the Zero Administration service platform in its service business, and Stream Analyze's route for its edge AI analytics platform is licensing and subscription. Eight results are already used beyond the project. AI-Gym is open source in Eclipse Arrowhead, and the AI Toolbox has a public wiki.

The three projects see shared demonstrators and cross-pilot integration as a joint direction for Chips JU calls in 2026–2028. Visit the AIMS5.0 stand at EFECS 2026 in Helsinki, 2–3 December. More information: aims50.eu.

Acknowledgment

The AIMS5.0 project is supported by the Chips Joint Undertaking and its members, including top-up funding by national funding authorities from the participating countries, under grant agreement no. 101112089.

New Member Focus

New Member Focus

Engineering physical AI across sectors



Ercan Özdemir
Founder and CEO I ALVEXIS Teknoloji A.Ş.

A shared edge platform can connect sensing, computation and equipment. Its value depends on the interfaces, integration and verification designed for each application

A camera can recognise a defect without being ready to support a production decision. A vibration sensor can detect a change without establishing whether a machine should stop. Between a useful inference and a dependable operational outcome lies an engineering task: capturing the right signal, interpreting it in context and delivering the result through an interface that the surrounding system can use.

Physical AI brings sensing and reasoning into systems that interact with the physical world. Those interactions may range from an operator recommendation to equipment control. The distinction matters. Monitoring reports a condition; decision support helps someone evaluate it; control changes equipment behaviour. Each requires different evidence, timing constraints and responsibility for what happens next.

For ALVEXIS, a new member of INSIDE, this creates a practical focus for collaboration: the electronics and embedded software that connect sensors, local computation and equipment. The aim is to make the integration problem explicit, so that a promising algorithm can become part of a defined, testable system (Figure 1).

Reuse the foundation and specify the differences

A modular approach can preserve useful engineering work across applications. Computing hardware, firmware structure, data handling and diagnostic functions can provide a common foundation. Custom printed circuit boards, carrier boards and sensor interfaces can then connect that foundation to the equipment concerned. ALVEXIS combines these areas with embedded software, prototyping and system integration in its engineering approach.

ALVEXIS organises its modular portfolio around ForgeEdge Core compute modules, ForgeEdge AI box platforms, FieldLink for field connectivity and ForgeControl for equipment input and output. These roles connect the processing foundation to sensors and machines. Each project requires its own configuration, integration and validation, guided by the installation and the intended operational function.

The useful unit of reuse is not necessarily a complete enclosure or an unchanged software image. It may be a compute module, a validated interface design or a method for capturing and time-stamping measurements. Reusing those elements still leaves application-specific choices: electrical levels, power supply behaviour, sensor placement, communications, thermal management and service access. A factory cabinet and a mobile field installation do not impose the same conditions.

Keeping computation close to the sensor can reduce dependence on continuous connectivity and avoid moving every raw measurement upstream. It also places processing, storage and power demands inside the installed device. The appropriate split between edge and remote services therefore follows the application: which results must be available locally, what can wait, and which data need to be retained for investigation? There is no single distribution that serves every operating environment.

Interfaces therefore deserve as much attention as processing capability. A sensor connection needs an agreed meaning for units, timing, calibration and missing data. An inference output needs a defined confidence interpretation and an agreed response when information is insufficient. An equipment interface needs a clear message structure and communication behaviour. These contracts help partners divide development work without leaving critical assumptions between their respective responsibilities.

Three applications with different operating boundaries

Industrial inspection and machine health provide one example. Imaging and vibration measurements may share edge computing



Figure 1: ALVEXIS product references illustrate the connection between sensing, local computation, connectivity and machine action. The industrial application environment is conceptual.

and data handling, while requiring very different acquisition paths. Inspection depends on illumination, field of view and synchronisation with the process. Increasing camera resolution or frame rate can improve coverage while increasing data throughput, processing load and heat. Selecting the acquisition settings therefore involves the compute budget and enclosure design as well as the inspection task. Machine monitoring adds different concerns: sensor mounting, operating speed and a meaningful baseline.

Agricultural sensing and control create another set of constraints. Distributed measurements may need to operate with intermittent connectivity, changing environmental conditions and limited power. Local processing can support decisions when a network connection is unavailable, but stale or unreliable measurements must remain identifiable. For an illustrative irrigation application, the engineering boundary would include sensing, valve interfaces and coordination with the irrigation schedule. The underlying computing approach may be reusable, while connectors, enclosure design, communications and operating logic require adaptation.

Mobile sensing in demanding environments adds motion, vibration, variable supply conditions and changing connectivity. A civil protection vehicle, for example, may use local perception to provide information to its crew. The design must account for sensor alignment, timing between inputs and the consequences of interrupted data. A module that works on a laboratory bench still needs integration and verification for its installation. The crew also needs an understandable view of sensor availability and data quality.



Figure 2: ALVEXIS compute-module, carrier-board, ForgeEdge-Max and COMMAND BRIDGE references form the centre of three illustrative application scenes: industrial inspection, agriculture and mobile sensing. Each setting has different integration and verification requirements.

These are illustrative application scenarios, not claims of completed ALVEXIS deployments. Their shared structure explains the rationale for a cross-sector platform: some functions can be reused, while the physical interfaces and operational responsibilities remain specific to each system (Figure 2).

Verify the complete path from measurement to response

Integration should start with a measurable use case. What observation is needed, under which conditions, and who will act on the result? The answers establish a system boundary and an acceptance plan before hardware and model choices become difficult to change. Relevant checks may include acquisition quality, end-to-end latency, power behaviour, thermal operation and the handling of missing or inconsistent inputs.

Verification then follows the signal through the system. A correct inference from an incorrectly scaled measurement is still a system error. So is a useful result delivered too late, without its timestamp, or after the associated machine state has changed. Test data should therefore exercise the interfaces as well as the processing, including disconnected sensors, delayed messages, restarts and loss of connectivity.

A staged prototype can help expose these issues before wider integration. Initial work can establish electrical and communication behaviour, followed by representative sensor inputs and then the intended equipment connection. Each stage should record its assumptions and unresolved questions. This gives partners a shared basis for deciding whether to modify an interface, gather more data or proceed to a more representative operating environment.

The response path requires its own discipline. Monitoring can flag uncertainty for review, while decision support preserves the context needed by an operator. Any control connection needs agreed limits, permissions and a defined fallback consistent with the equipment architecture. AI functions must operate within this framework without bypassing existing safety controllers or protection functions.

Maintenance also belongs in the design. Teams need to know which hardware, firmware, configuration and model versions produced a result. Diagnostic access and update procedures should be planned with the equipment owner. The ability to reproduce a reported problem and restore a known configuration can be as important to continued operation as the initial demonstration (Figure 3).

Create collaboration around a defined engineering task

This approach makes collaboration more concrete. An equipment manufacturer can provide the operating context and acceptance criteria; a sensor specialist can define acquisition requirements; an AI or research team can contribute processing methods. ALVEXIS can contribute custom electronics, carrier-board design, firmware, sensor integration and prototype validation within an agreed scope. Clear interfaces make it easier to identify dependencies and decide what must be demonstrated together.

A useful starting point is a bounded integration task: connect a specified sensor to an edge computer, adapt a carrier board to an equipment interface, or verify a prototype under agreed operating conditions. Such work can support a direct engineering engagement or a collaborative research project. INSIDE offers a setting in which these complementary capabilities can meet.

The cross-sector opportunity rests on disciplined reuse. A shared platform provides continuity between applications; careful engineering establishes what must change.

Physical AI becomes useful when those two activities are connected through an explicit understanding of the physical system, its users and the evidence required for its intended role.



Figure 3: ALVEXIS product references illustrate board development, prototype integration and system verification. The laboratory, test equipment and application settings are conceptual; the image does not document an actual test result.

Ercan Özdemir
Founder and CEO | ALVEXIS Teknoloji A.Ş.

Ercan Özdemir is Founder and CEO of ALVEXIS Teknoloji A.Ş., a technology company focused on custom electronics, embedded software and system integration. His work at ALVEXIS centres on connecting sensing, edge computing and application-specific engineering. Through INSIDE, he aims to build collaborations with industrial partners and research teams developing practical electronic components and systems.



Events

INSIDE Magazine 15

Events

INSIDE Connect 2026

Connecting INSIDE community's strengths in a changing world



Maha Karim-Hosselet
Senior Communications Manager



Santus Kisoka
Market and Financial Strategist











As the global race for computing power, artificial intelligence and semiconductors intensifies, Europe faces a choice: rely on technologies shaped elsewhere, or bring its own research, industry and innovation communities together to help define what comes next. In Palermo, INSIDE Connect 2026 offered a compelling picture of what a connected European ecosystem can achieve.

A gathering with a purpose

The conversation in Palermo began with a question larger than any single technology. In a world where chips influence industrial competitiveness, energy systems and defence capabilities, what does it mean for Europe to be technologically sovereign? For two days, INSIDE Industry Association brought together participants from research, academia, industry and innovation networks to explore that question from different perspectives. The event was not simply a showcase of projects. It was a meeting point for people whose work increasingly depends on one another: semiconductor designers, software engineers, AI researchers, robotics specialists, industrial users and policymakers.

Opening contributions from INSIDE representatives and Italian policy and industry stakeholders placed these exchanges within a broader European context. Paolo Azzoni, INSIDE's Secretary General, outlined a research and innovation landscape in transition, with discussions around Chips Act 2.0, FP10, the proposed European Competitiveness Fund and the 2028–2034 EU budget. Behind the acronyms lies a practical challenge: aligning funding, research priorities and industrial deployment so that European innovations can move from promising demonstrations to competitive products.

The European perspective was further enriched by Antonia Wunnerlich from the German Marshall Fund, who shared her vision of Europe's position in an evolving global landscape, offering insights into the strategic and geopolitical dimensions shaping Europe's technological ambitions and competitiveness.

That challenge acquired a geopolitical edge in Giuseppe De Ruvo's contribution for LIMES. His maps of China's AI geography, including major clusters around Beijing, Shanghai and Shenzhen, placed European ambitions alongside the industrial scale and strategic choices of China and the United States. The global contest is not simply about who produces the fastest chip. It also concerns who controls software, infrastructure, capital and supply chains on which tomorrow's economies will rely.

The foundations of sovereignty are interconnected

One conclusion followed naturally: Europe cannot discuss digital sovereignty without discussing energy and security. AI systems depend on processors, data centres and electricity. As computing demand rises, the availability of reliable and affordable energy increasingly shapes where infrastructure can be built and which industries can scale. Europe's different national energy mixes, including France's substantial nuclear generation and Germany's changing energy system, make this a distinctly European industrial question.

The same technologies also underpin modern defence and resilience. Sensors, secure communications, autonomous platforms, embedded computing and cybersecurity serve civilian applications but can be critical in security and defence contexts. This dual-use character makes supply-chain reliability and technological choice especially important. Dependence is not automatically a weakness, international cooperation remains essential, but dependence without alternatives can become a strategic vulnerability.

The semiconductor discussions made these concerns tangible. Francesco Tenuta's perspective on Italy's semiconductor strategy showed how national strengths in manufacturing, research and industrial applications fit into a European value chain. Cristina De Luca introduced CADETTI, the





Italian chip-design competence centre, as a means of linking talent, training, design resources and companies. Andrea Pinna's PHOENIX initiative in France addressed chiplets and AI-enabled electronic design automation, areas where technical progress demands new skills and closer hardware–software collaboration. Different initiatives, but a common ambition: build capacity that is useful beyond one laboratory or one country.

The same logic animated the debate around processor architectures. Cortus presented RISC-V as an opportunity for Europe to contribute to the next generation of computing, while recognising that an open instruction set is not a complete industrial solution. As its presentation put it: "The question is not whether Europe will use the next generation of computing. The question is whether Europe will help define and build it." Work presented by Openchip on RISC-V vector extensions showed how this architectural choice connects to AI and high-performance computing. SiPearl approached the challenge from another angle: sovereign high-performance CPUs. Its presentation recalled founder Philippe Notton's view that Europe's strategic independence requires its own high-performance processors. Together, these contributions exposed the full scale of the task, from architecture and design tools to software, manufacturing partnerships and market adoption.



Intelligence is leaving the data centre

But processors matter because of what they make possible. Throughout INSIDE Connect, AI emerged less as a standalone topic than as a force reshaping physical systems. Giulio Urlini of STMicroelectronics argued that edge AI is not merely a smaller version of cloud AI: it calls for computing designed around local data, low latency, autonomy and energy efficiency. Neuronova's ultra-low-power approach to intelligence in sensors and Wilfred van der Wiel's exploration of brain-inspired computing suggested that the next breakthroughs may come not only from more computing power, but from using dramatically less of it.

This shift becomes particularly visible in robotics. Juha Röning's discussion of embodied intelligence connected Europe's robotics expertise with questions of skills, investment and time-to-market. Marco Di Felice examined how agentic AI could help autonomous robots interpret objectives and adapt their actions, while Michele Magno's work on edge intelligence for drones showed why local decisions and efficient communication matter in the real world. Alberto Bianchi from Leonardo explained how interoperability in multi-domain autonomy adds another dimension: intelligent machines must be able to work together, not simply operate in isolation.

Once intelligence moves into vehicles, factories, drones and other physical environments, trust becomes inseparable from performance. Jon Perez Cerrolaza examined cybersecurity from chips to systems of systems; Dimitrios Serpanos explored the opportunities and risks of agentic AI in cyber-physical security; and Flip van Spaendonck addressed verification and guardrails for generative AI. Their contributions point to a shared principle: safety, traceability and security cannot be added at the end of development. They must be designed into increasingly autonomous systems from the beginning of their lifecycle.

The automotive sector shows what is at stake

Few industries illustrate the urgency of this transition as clearly as mobility. In a perspective on global competition in AI and software, Jean-Luc di Paola-Galloni highlighted the shift from vehicles defined primarily by mechanical engineering towards software-defined and AI-defined platforms. Chinese advances in connected and automated mobility and American robotaxi deployments show how quickly the competitive landscape can change. The question for Europe is no longer only who builds a good vehicle, but who develops the computing, software, data and AI systems that determine its capabilities.

That question ran through presentations by Michael Paulweber on the evolution of AI-defined mobility, Patrick Pype of NXP on the digital vehicle's hardware architecture, and Mario Gartner of TTTech Systems on technologies shared across software-defined domains. The lesson reaches far beyond automotive: the strongest solutions will connect hardware and software, safety and security, and expertise from industries that once developed their technologies separately.

Interoperability therefore became another connecting thread.The CEI-Sphere deep dive, introduced by Paolo Azzoni and Mauro Tortonesi, brought the edge to cloud continuum into focus. Antonio Kung's presented the Hourglass Model 2027 for the continuum and Mauro Tortonesi's analysis of the Stand-up 3 questionnaire helped frame a practical discussion about existing industrial systems, open and interoperable architectures, data management and sharing, regulatory framework, security and trust on the technologies introduced by the continuum. Related work on knowledge graphs, shared information models and AI-assisted engineering addressed a familiar industrial reality: innovation must often connect with infrastructure that is already in place. Open and interoperable approaches can help European companies cooperate without locking every solution into a single proprietary environment.

From promising technologies to European industrial strength

The most important connections, however, may be those between technical capability and the market. Europe has exceptional research institutions and industrial expertise, but its innovation landscape can be fragmented. Access to advanced design tools, investment, specialist skills, customers and manufacturing partners determines whether a new technology becomes an industrial success. The EuroCDP session on designing chips in Europe, with contributions from Réda Nouacer, Helio Fernandez Tellez and Giorgio Dell'Erba, illustrated how infrastructure, engineering assistance and business support can help bridge that gap. Initiatives such as CADETTI and PHOENIX address different parts of the same challenge.



This is where the value of an ecosystem becomes concrete. A start-up may need a research partner; a researcher may need an industrial test environment; an established company may need expertise in an emerging architecture; and policymakers need to understand which shared barriers hold back deployment. No single actor can solve these problems alone. The purpose of gathering in Palermo was to make those links more visible and more actionable.

Not every important exchange took place on stage. Conversations continued in the project exhibition area, through networking moments, a guided walk in Palermo's historic centre and dinner at Palazzo Fatta. Such encounters are not peripheral to innovation: they are often where a presentation becomes a conversation, a shared challenge becomes a potential collaboration, and a community starts to act as an ecosystem.

A shared responsibility to act

INSIDE Connect 2026 did not offer a single technological answer to Europe's competitiveness challenge, and that may be its strongest message. The future will be built across many layers: semiconductor design, energy-efficient computing, AI, autonomous systems, cybersecurity, interoperable platforms and industrial applications. The United States and China are investing and moving at scale. If Europe does not help shape these technologies, others will shape them for Europe.

That is not an argument for isolation or for abandoning international partnerships. It is an argument for entering them with capabilities, choices and confidence. Europe already has deep knowledge, talented people and strong industrial foundations. Its next challenge is to connect them more effectively, accelerate the journey from research to deployment and make collaboration a source of competitive advantage.

In bringing together the Electronic Components and Systems community, INSIDE Connect showed what this ambition looks like in practice. INSIDE Industry Association's role is to inspire and suppor the community, and keep these conversations moving towards shared priorities, stronger partnerships and concrete contributions to Europe's research and industrial future. Palermo was a meeting point. What matters next is what the INSIDE community will build together.

INSIDE Magazine 15

More than a summer school

Building and strengthening CPS community in Europe and beyond



Christian Pilato
Associate Professor at Politecnico di Milano



Francesca Palumbo
Associate Professor at University of Cagliari



Francesco Regazzoni
Associate Professor at University of Amsterdam



Paolo Azzoni

Cyber-physical systems (CPS) are pervading our daily lives, underpinning applications ranging from mobility and industrial automation to energy infrastructures and food supply chains. These systems evolve to support increasingly sophisticated cognitive capabilities and become increasingly connected. Their complexity continues to grow, often requiring distributed intelligence and seamless operation across the edge-to-cloud continuum. Ultimately, the value of the assets they manage also makes them an appealing target for attackers. These challenges call for a new generation of engineers who can look beyond individual technologies and adopt a holistic approach to designing dependable, adaptive, intelligent, and secure systems. The CPS Summer School was launched in 2017 to address that need, bringing together expertise in modeling, architecture, hardware–software co-design, distributed computing, security, and ethics. Its next edition will take place on 20–24 September 2027 in Alghero, Sardinia.

From concepts to implementation: building a European CPS community

The school's guiding philosophy has always been "Designing Cyber-Physical Systems – From Concepts to Implementation." Hands-on activities are therefore at the heart of its approach, enabling participants to translate theoretical concepts into practical implementations immediately during the week of the school. Aimed at PhD students, researchers, and R&D professionals from academia and industry, the school program and activities encourage interdisciplinary vision and collaboration among participants.

Over the years, the programme has increasingly drawn on cutting-edge research in the domain, including research carried out in recent European research projects, such as the Horizon Europe MYRTUS project, which addressed the design and dynamic orchestration of distributed CPS across the



INSIDE Magazine 15









computing continuum, and on open-source industrial frameworks, such as Eclipse Arrowhead, supporting interoperable, service-oriented automation.

These characteristics have made the school a natural meeting point for the European CPS community. Its scientific programme is shaped by a team of Programme Chairs bringing together academic and industrial perspectives: Francesca Palumbo (Università degli Studi di Cagliari), Christian Pilato (Politecnico di Milano), Francesco Regazzoni (Universiteit van Amsterdam and Università della Svizzera italiana), and Paolo Azzoni (INSIDE Industry Association). Together with internationally recognised experts from industry and academia, they ensure a strong connection between cutting-edge research, emerging technological challenges and industrial practice.

Continuity is also a matter of identity. With the sole exception of the 2022 edition, the school has always been hosted in Alghero, Sardinia, thanks to the longstanding commitment of the Università degli Studi di Sassari and its Local Arrangements Chair, Claudio Rubattu. The school's deliberately small format, with around 40 participants, encourages direct exchanges between attendees, lecturers and organisers, strengthening the sense of belonging to a shared community.



The INSIDE Summer School reflects the close relationship between the academic CPS community and the European electronic components and systems industry. Through INSIDE Industry Association, the initiative benefits from connections across the European ECS ecosystem, bringing industrial perspectives into the programme and encouraging exchanges between young researchers, established scientists and technology companies. This combination of scientific excellence, industrial relevance and community building has become one of the defining characteristics of the school.

Attendees at the heart of the CPS Summer School

At the CPS Summer School, students are not passive listeners; they are active contributors. Through the activities of the school and the informal moments of aggregation, they are encouraged to share their research, enlarge their vision across disciplines, solve technical problems, and discuss their results, reflecting one of the school's core principles: research communities are built around people.

The CPS Workshop is traditionally held on the first day as an icebreaker that turns the group of students into members of the CPS community. In the workshop, the participants and their research is put in the spotlight. Young researchers have the opportunity to present early results, work in progress, or future research directions. Beyond receiving feedback from lecturers and peers, they introduce themselves, discover common research interests, and start a discussion that continues throughout the week and, very often, beyond the school itself.

The Creative Lab, a unique feature of the CPS school, allows students to immediately put



learned concepts into practice. Inspired by hackathon formats, it consists of a CPS-related project that teams of attendees, created on the first day of the school, supported by scientific coaches, have to pursue throughout the week. During the development, each team must combine complementary skills, overcome unexpected technical difficulties, and, making the appropriate design and implementation choices, transform initial ideas into convincing solutions and, whenever possible, working demonstrators. More than a technical exercise, the Creative Lab aims to recreate the dynamics of real-world research and innovation projects, where success depends on teamwork, adaptability, creativity, and the ability to communicate a solution's value. The experience culminates in a final session, where teams pitch and demonstrate their work to a jury composed of the school teachers, which determines the winning team.

Equally important are the informal moments. Among these is the "beach note" by Alberto Sangiovanni Vincentelli (University of California, Berkeley), which has been part of the school since its first edition. The spirit of the beach note (which in 2026 was moved to a local winery) is to foster open conversations, shared experiences, and personal connections beyond academic roles and hierarchies, and over the years it has become a highly anticipated tradition. It moves discussions outside the classroom, creating a setting where the usual barriers between renowned experts and young researchers naturally disappear.

Looking ahead to Alghero 2027

The 2027 edition will maintain the place and community identity closely associated with the initiative. While the detailed program will be announced at HiPEAC 2027, when registration will also open, several speakers, including Alberto Sangiovanni Vincentelli, have already confirmed their presence.

The mission of the school remains timely: as edge AI, autonomous and adaptive platforms, and connected infrastructure grow more sophisticated, system architects must combine analytical rigor with the capacity to integrate hardware, software, data, security and, most importantly, people, ultimately creating and strengthening the CPS community in Europe and beyond.

Further information: http://www.cpsschool.eu

INSIDE Magazine 15

The promise and peril of Artificial Intelligence



Maha Karim-Hosselet
Senior Communications Manager

When I was a child, I imagined the future with flying cars. It seemed obvious that by the time I became an adult, we would all be travelling through the sky. Today, I look at my own children and wonder what their world will look like in twenty years. Will artificial intelligence transform their lives in ways we cannot even begin to imagine?

History has taught us that technological revolutions rarely unfold as we expect. Cars replaced horses, electricity transformed our homes, and the internet changed the way we communicate, work and access information. Even inventions we now take for granted, such as washing machines and vacuum cleaners, profoundly changed everyday life.

But artificial intelligence feels different. It is not simply a machine that performs a physical task faster than we can. It is a technology that increasingly performs activities we have always associated with human intelligence.

And that is both fascinating and unsettling.

Perhaps that is what makes AI so difficult to judge. Its promise and its peril are not two separate stories. Sometimes, they grow out of the very same possibility. And what happens next depends, at least in part, on the choices we make.

For the first time, we are seriously discussing the possibility of artificial general intelligence, or AGI: systems capable of performing a broad range of intellectual tasks at a human level or beyond.

But what does intelligence really mean? A machine may process information faster than we do, recognise patterns we cannot see or solve problems that would take us years. Yet it has never experienced childhood, the joy of learning something new, the disappointment of failure or the relationships that shape who we become.

Perhaps machines will eventually reproduce many aspects of human reasoning. Perhaps their development will reach limits we do not yet understand. What seems clear is that AI is already changing our relationship with knowledge, work and decision-making.

Google DeepMind has expressed an extraordinary ambition: to solve intelligence and then use it to solve other problems. Imagine what that could mean for scientific research, medicine, energy or education.

It is almost like changing the rules of every team sport at the same time. We would have to rethink not only how we play, but also how we prepare, collaborate and define success.

And while we are learning the new rules, the technology continues to evolve.

We understand how a calculator performs an operation. We know why it can calculate faster than a human being. But with increasingly complex AI systems, understanding exactly how a particular result is produced can be much more difficult.

That matters because intelligence and power are closely connected. If a small number of companies control technologies that influence how billions of people work, learn and make decisions, their influence becomes a societal question, not merely a commercial one.

In a democracy, political power comes with mechanisms of accountability. Technological power on this scale also deserves scrutiny, even if the mechanisms will necessarily be different.

The promise of doing more with less

The opportunities are extraordinary.

AI could help us accelerate scientific discoveries, use resources more efficiently and make expertise accessible to people who previously had little chance of benefiting from it.

In many workplaces, younger generations are already learning to collaborate with AI as naturally as previous generations learned to use the internet. Tasks that once required hours can sometimes be completed in minutes, leaving more time for analysis, creativity and human interaction.

But progress has a cost.

Training and operating advanced AI systems requires substantial computing infrastructure, electricity and, in some cases, water for cooling. The International Energy Agency has highlighted the growing energy implications of AI and data centres[2].

There is something paradoxical about using so much energy to develop a technology that may eventually help us save energy.

Yet perhaps the question is not simply how much AI consumes today, but what we gain from that investment over time. Bringing something new into the world requires resources. The challenge is to ensure that the long-term benefits justify the environmental and economic costs, rather than assuming that they automatically will.

Healthcare offers some of the most compelling possibilities.

Imagine researchers identifying treatments faster, doctors receiving better diagnostic support or children benefiting from medical advances tailored to their specific needs. Children are not simply smaller adults; their bodies, development and medical needs are different. Better research tools could make an enormous difference.

But there is another question worth asking. If AI helps us live longer, will it also help us live better?

Longer lives have implications for healthcare systems, pensions and the way societies organise work. And a good life is about more than medical indicators. It is also about family, friendships, meaningful activities and the simple pleasure of sharing a coffee with someone we love.

The World Health Organization has emphasised the importance of ethics, oversight and appropriate safeguards when



AI is used in healthcare[1]. Technology can support medical professionals, but it should not make us forget the human beings behind every diagnosis.

The same potential exists in education, agriculture and energy.

Imagine a child in an underserved community gaining access to learning support that was previously unavailable. Or a farmer using better information to manage crops and resources.

Yet access to information is not the same as education. As parents, we know that teaching a child involves much more than answering questions. It means noticing what they do not understand, encouraging them when they struggle and sometimes recognising needs they cannot yet express.

AI can be a remarkable educational ally, but an incorrect or biased answer can also be reproduced and shared at enormous scale.

The real promise is not simply to make more knowledge available. It is to help more people benefit from it.

And that brings us to an essential question: who will actually enjoy the benefits of this progress?

The great inventions of the past eventually made many aspects of everyday life easier. But the benefits of AI will not necessarily reach everyone at the same speed. Access depends on affordability, infrastructure, skills and opportunities.

A technology that can improve millions of lives should not become a privilege reserved for those who can afford the most advanced tools.

The same tools that could open doors for millions could also widen the distance between those who have access and those who do not. The technology alone will not decide which future we get.

What happens to work?

Perhaps nowhere are the promises and uncertainties of AI more visible than in the world of work.

Machines do not need salaries, holidays or coffee breaks. Of course, they still require energy, maintenance, infrastructure and human expertise. But their ability to perform tasks at scale raises difficult economic questions.

If machines increasingly produce the goods and services we consume, what happens to the people whose income depends on doing that work? And if large numbers of people lose purchasing power, who will buy what those machines produce?

The International Monetary Fund estimated in 2024 that almost 40% of jobs worldwide are exposed to AI, although exposure does not necessarily mean replacement[3]. The International Labour Organization has similarly emphasised that many occupations are more likely to be transformed than eliminated altogether[4].

That distinction is important.

AI may take over repetitive tasks while allowing professionals to focus on more complex decisions, relationships and responsibilities. But the transition will not be painless for everyone.

I also wonder what happens to young people entering the labour market. If AI performs the simpler tasks traditionally assigned to junior employees, how will the next generation gain the experience needed to become experts?

We cannot expect someone to become an experienced lawyer, engineer or financial adviser without first having the opportunity to learn.

If you were selling your company, for example, you would need more than access to information. You would need professionals whose judgment, experience and understanding of human relationships you could trust. AI might help you analyse documents or explore scenarios, but it would not replace every conversation or decision.

We should remember that having information is not the same as knowing what to do with it.

Generational differences will also matter. Younger people may feel more comfortable experimenting with digital tools, but being a digital native does not automatically make someone AI-literate. Meanwhile, many experienced professionals are enthusiastic early adopters, while others will need support and training.

The answer cannot be to divide society between those who understand AI and those who do not. We need opportunities for people to learn from one another, combining technological confidence with professional experience.

And perhaps this transformation will also change the value we place on certain skills. Practical expertise, craftsmanship and professions that require physical presence or human connection may become even more appreciated.

Governments, educational institutions and businesses should prepare for these changes rather than wait for disruption to force their hand.

Because work is not only about earning a salary.

For many people, it provides dignity, relationships, structure and a sense of purpose.

If AI helps us complete our work faster, will we use that time to be with our families, practice sport, volunteer or simply enjoy life? Or will we raise productivity expectations until everyone is expected to achieve twice as much in the same number of hours?

That may be one of the most important choices ahead.

AI could generate enormous economic value. But if that wealth is concentrated in a handful of companies, technological abundance will not necessarily translate into a better society.

The real measure of progress should not be how much more we produce. It should also be whether people enjoy better opportunities, more accessible services, greater independence and more fulfilling lives.

That is the tension running through so many of these questions: AI could give us more time, more knowledge and more possibilities. But none of those gains automatically makes life better. What we choose to do with them matters just as much.

When seeing is no longer believing

There is another risk we cannot ignore: the growing difficulty of distinguishing what is real from what is artificial.

We already see it on social media. Some AI-generated images and videos are so convincing that it becomes difficult to know what to believe. And sometimes, even authentic content is questioned simply because it looks too extraordinary to be real.

During Slovakia's 2023 election campaign, an AI-generated audio recording falsely portrayed political figures discussing electoral manipulation. Although the incident does not establish that the deepfake determined the election result, it showed how fabricated content can enter public debate at a particularly sensitive moment[5].

But there is something interesting about this situation.

For years, people predicted the decline, even the disappearance, of traditional journalism with the arrival of the internet and social media. The press has certainly struggled, and journalists have had to reinvent the way they work, communicate and reach their audiences.

Yet today, when almost anyone can generate a convincing image, video or story, professional journalism and the people behind it may become more valuable than ever.

Journalists are not infallible. But their responsibility to verify facts, question sources and provide context remains essential.

Perhaps this is one of the great paradoxes of AI: the more artificial our information environment becomes, the more we may need human judgment to navigate it.

The question extends beyond information to our personal freedom.

How far are we willing to let intelligent systems analyse our behaviour, movements and habits in the name of security or efficiency?

In many ways, this is already happening. Social media platforms use algorithms to understand our interests, anticipate our preferences and influence what we buy. Similar techniques are used in political campaigns, particularly in the United States, where digital advertising and targeted messaging have become powerful tools for reaching voters.

But there is something we should not underestimate: our ability to think critically.

Algorithms can influence what we see, read and hear. Yet influence does not necessarily mean control. People can still question what they are told, challenge the messages they receive and make their own decisions.

For me, this is one of the most important freedoms we need to protect as AI becomes more powerful: our ability to think, question and decide for ourselves.

Power, safety and the race to be first

The same technology can be used to protect us or to cause harm.

In cybersecurity, for example, AI can help companies detect threats and prevent attacks, but it can also give cybercriminals more sophisticated tools.

And there is another risk we should never underestimate: even the most advanced AI systems can make mistakes.

Think about something as simple as accounting. Companies have procedures, checks and sometimes several levels of approval before important financial decisions are made. We double-check figures, review contracts and question results because we know that mistakes happen.

So why would we blindly trust an AI system when the consequences could be far more serious?

We are not talking about predicting the winner of the Champions League. We are talking about decisions that could affect people's health, safety and, in some cases, their lives.

AI can support our decisions, but that should never mean giving up our responsibility to question, verify and remain accountable for them.

But how can we ensure that AI systems actually do what we expect them to do, without developing unexpected behaviours?

This is precisely what researchers are exploring through AI alignment: making sure that increasingly powerful systems remain consistent with human intentions and values.

Some laboratory tests have revealed concerning behaviours, but it is important to distinguish controlled experiments from what actually happens in the real world.

We often talk about AI as though it were an invisible, almost immaterial technology. But behind every AI system lies a very physical

reality: semiconductors, data centres, energy infrastructure, global supply chains and strategic decisions.

AI has also become a central part of the technological competition between the United States and China, with Taiwan playing a particularly important role in advanced semiconductor manufacturing.

And in this race to lead, I wonder whether we are paying enough attention to safety. How much of the industry's talent and investment is dedicated to making AI more powerful, compared with making it safer?

When everyone is racing to be first, the danger is that speed becomes more important than responsibility.

There is a saying that the United States innovates, China accelerates and Europe regulates. It is an oversimplification, of course, but it raises an important question: can Europe protect its citizens without slowing down innovation?

The EU's AI Act is an attempt to strike that balance, with obligations based on the risks AI systems pose[6]. And while regulation is sometimes seen as an obstacle, clear and proportionate rules can also help build trust and encourage responsible innovation.

The challenge is not to stop progress, but to make sure we remain in control of where it takes us.

Perhaps we do not need to choose between being excited about AI and being concerned about it. We can recognise both its potential and its risks, while insisting that human judgment and responsibility remain part of the story.

Humanity's next decision

What if the most exciting future of AI is not one where machines replace us, but one where they help us become more capable?

Imagine someone with a disability regaining a degree of independence thanks to a brain–computer interface, a scientist accelerating a discovery, or a professional finding more time to focus on what truly matters.

But the benefits can also be much simpler.

Think of someone who is too shy to ask a question in a classroom or a meeting. AI can offer a space to explore ideas, ask questions without fear of embarrassment and understand things at their own pace.

Of course, AI cannot replace genuine human relationships or professional support. But it can become a valuable ally, helping people learn, express themselves and build confidence.

To me, that would be a beautiful form of technological progress: not making humanity less necessary, but helping more people realise their potential.

We cannot predict exactly what the world our children will grow up in will look like. But we can still influence the choices that shape it.

And perhaps our responsibility starts at home.

What if, instead of simply worrying about the impact of AI on future generations, we started teaching our children how to use it as an ally? Helping them understand its possibilities, recognise its limitations and, above all, develop the critical thinking needed to question its answers.

We may not be able to control every technological development, but we can help prepare the next generation to navigate them wisely.

I do not believe we should fear AI, nor should we blindly celebrate it. We should approach it with curiosity, ambition and a healthy sense of responsibility.

When I was a child, I imagined the future with flying cars. Today, looking at my own children, I realise that the greatest revolution may not be about how we travel, work or communicate, but about how we choose to live alongside intelligence that is no longer exclusively human.

I don't know what their world will look like in twenty years. None of us do.

But perhaps the most important question is not how intelligent our machines will become. It is whether we will be wise enough to use that intelligence to build a better world.

Because in the end, the future of AI is not just about technology.

It is about humanity.

1 World Health Organization (2024). Guidance on the ethics and governance of artificial intelligence for health.

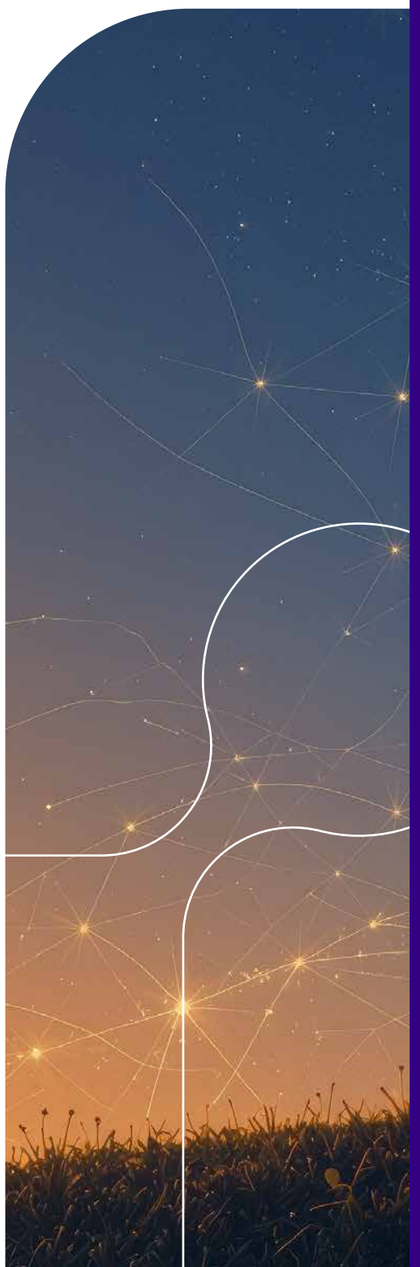
2 International Energy Agency (2025). Energy and AI.

3 International Monetary Fund (2024). Analysis of AI exposure and the future of work.

4 International Labour Organization (2025). Generative AI and Jobs.

5 Reuters (2023). Reporting on AI-generated audio and disinformation during Slovakia's parliamentary election campaign.

6 European Commission. AI Act: European Union framework for risk-based regulation of artificial intelligence.



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

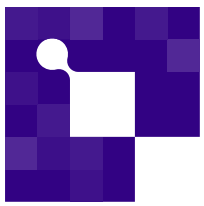
Acknowledgements
With thanks to the interviewees, project participants, INSIDE Industry Association office, the INSIDE Industry Association Presidium and other INSIDE Industry Association-involved persons for any assistance and material provided in the production of this issue of the INSIDE Magazine.

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

© 2026 INSIDE Industry Association

Permission to reproduce individual articles from INSIDE Magazine for non- commercial purposes is granted, provided that INSIDE Magazine is credited as the source. Opinions expressed in the INSIDE Magazine do not necessarily reflect those of the organisation.



INSIDE
Industry Association

INSIDE-association.eu