

HiPEAC Vision 2026

HIGH PERFORMANCE, EDGE AND CLOUD COMPUTING



Editorial board:

Marc Duranton (editor in chief), Gaël Blondelle, Koen De Bosschere, Paul Carpenter, Madeleine Gray, Thierry Goubier, Axel Nackaerts, Charles Robinson, Tullio Vardanega, Olivier Zendra.



Funded by
the European Union

HiPEAC

This document was produced as a deliverable of the Horizon Europe HiPEAC (“High Performance, Edge And Cloud computing”) CSA under grant agreement 101069836.

The editorial board is indebted to Dr Max Lemke and Dr Rolf Riemenschneider of the Directorate-General for Communication Networks, Content and Technology of the European Commission for their active support of this work.


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

Design: www.magelaan.be / Cover design: Roger Castro, Monzón Studio / Cartoons: Arnout Fierens, Arnulf.be

© 2026 HiPEAC

ISBN 9789078427087

Contents



Introduction	4
Recommendations	12
The Next Computing Paradigm	20
Artificial Intelligence	25
New Hardware	45
Tools	51
Cybersecurity	59
Open Source	71
Sustainability	79
State of the European Union	89
Process	100
Acknowledgements	100

Introduction



Welcome to the HiPEAC Vision 2026

It is clear that artificial intelligence (AI) is the most disruptive technology in the domains covered by HiPEAC during the last 20 years. AI, particularly through neural network-based approaches revived by Geoffrey Hinton's work in 2012, first transformed data analysis (images, text) and inspired companies such as NVIDIA early on, whose valuation now exceeds \$4.6 trillion thanks to a strategy focused on accelerated computing for AI. The turning point for (generative) AI came in 2017 with Google's paper "Attention Is All You Need," [2] which introduced the "transformers" approach (which was implemented by OpenAI's GPT – "Generative Pre-trained Transformer" – models as early as 2018), propelling applications such as ChatGPT, which in 2023 became the fastest-growing consumer tool in history, and extending AI to the generation of images, videos, and even scientific content.

Since 2024, the trend has been towards models that are not necessarily larger but which employ "reasoning" (i.e. optimization in the inference phase). Since late 2025, we have also seen a shift towards agentic AI, where systems such as Google DeepMind's Aletheia orchestrate specialized tasks to solve complex problems (e.g. advanced mathematics) without modifying the underlying models, illustrating a new era of autonomy and collaboration between intelligent agents that can use computer tools directly.

The explosion of these agentic AI approaches (with their strengths and weaknesses, especially in terms of cybersecurity) can be illustrated by the rapid success

It is hard to communicate how much programming has changed due to AI in the last 2 months: not gradually and over time in the "progress as usual" way, but specifically this last December. There are a number of asterisks but imo coding agents basically didn't work before December and basically work since – the models have significantly higher quality, long-term coherence and tenacity and they can power through large and long tasks, well past enough that it is extremely disruptive to the default programming workflow. <...>

As a result, programming is becoming unrecognizable. You're not typing computer code into an editor like the way things were since computers were invented, that era is over. You're spinning up AI agents, giving them tasks "in English" and managing and reviewing their work in parallel. The biggest prize is in figuring out how you can keep ascending the layers of abstraction to set up long-running orchestrator Claws with all of the right tools, memory and instructions that productively manage multiple parallel Code instances for you. The leverage achievable via top tier "agentic engineering" feels very high right now.

It's not perfect, it needs high-level direction, judgement, taste, oversight, iteration and hints and ideas. It works a lot better in some scenarios than others (e.g. especially for tasks that are well-specified and where you can verify/test functionality). The key is to build intuition to decompose the task just right to hand off the parts that work and help out around the edges. But imo, this is nowhere near "business as usual" time in software.


Andrei Karpathy, February 25, 2026 [1]

of OpenClaw [3], initially an open source agentic artificial intelligence solution developed by a single person. Just as ChatGPT opened the public's eyes to the potential of large language models (LLMs), OpenClaw made agentic AI accessible to the masses [4]. Is this the realization of the "Guardian Angel" (or "Digel") discussed in editions of the HiPEAC Vision since 2020?

Looking back at the last few months (at the time of writing in March 2026) really illustrates the speed of the progress in AI. The performance of AI in code generation, which is also at the core of AI technology itself, has improved enormously, as illustrated by the opening quote of Andrei Karpathy [1]. In this period, the explosion of agentic AI (embodied by OpenClaw) shows that AI is escaping from a window

in a browser to really interacting with a computer environment to carry out real tasks.

Over the previous year (2025), it became clear that the concept of the "next computing paradigm" (NCP), proposed in previous editions of the HiPEAC Vision, was both increasingly relevant and increas-



ingly realizable. The HiPEAC Vision 2025 proposed the definition of **communication protocols** between agents. This was partially achieved in 2025 with protocols such as the Model Context Protocol (MCP) and Agent2Agent (A2A) protocol, etc. developed by large AI companies based in the United States (US). However, these protocols are still limited to functional requirements, rather than including non-functional requirements as proposed in our vision.

At the core of the NCP is the **orchestrator**, which is now also key for **orchestrating agentic AI**. **Distribution of AI across the computing continuum** is also already starting with local AI running on smartphones and laptops, and we are even seeing the emergence of dedicated accelerators as add-ons to consumer devices [5]; all this was made possible thanks to the development of smaller AI models that can run locally.

The next step can be summarized by the phrase “**AI is becoming physical**”, allowing AI systems to be integrated with the physical world and controlling factories, self-driving cars and robots. Adding non-functional requirements such as latency, time, and safety as proposed in the NCP is a necessary enabler for this transition to the physical world.

The current AI race is focused on ever-larger data centres. This creates growing problems in term of energy consumption, grid access, cooling, consumption of natural resources (such as water), availability of components (such as graphics processing units (GPUs) and memory, leading to the current “RAMapocalypse” or shortage of random access memory (RAM) for other application domains), and reliability (the mean time to failure for a data centre with half a million GPUs is in the order of few minutes).

But there is an alternative approach proposed by the NCP: **distributing computing on demand** instead of systematically requesting cloud resources. This alternative approach is indeed a risk because it is not yet available, but it is a



bet on the future: will Europe take its own path adapted to its structure, culture and ethics; or will it blindly copy what others are doing, inevitably lagging behind due to its late start?

This is the core of the HiPEAC Vision 2026: “*Europe should not imitate others; it should create its own future*”.

The ripple effects of AI across the computing ecosystem


The AI revolution is having evident impact on various domains within the HiPEAC ecosystem, which spans the computing continuum from edge to cloud to high-performance computing (HPC), while also having a major impact on society.

The industry’s centre of gravity is moving from training to inference. As AI systems transition from demonstrations to services used by millions, inference becomes the dominant consumer of compute – and the competitive battleground for hardware architectures, data-centre design, and software stacks.

As producing inference accelerators is less complex than making relatively general-purpose GPUs, smaller (European) companies can propose attractive solutions based on innovative concepts, and the associated software ecosystem is also

relatively simpler. Architectures that are more specialized, hence less flexible and less programmable, are inherently more energy efficient. However, a good compromise between flexibility and efficiency can be found in architectures that are reconfigurable or based on coarse-grain reconfigurable architectures (CGRA); near- or in-memory computing solutions can be very attractive as well; all of these are in reach of Europe and in several cases are being commercialized by European companies. Chiplet platforms allow a flexible combination of such tailored processing elements, and can facilitate specialization over a basis of common functionality. (See the hardware chapter for further discussion of hardware directions for Europe).

Thinking about solutions that can be inserted across the continuum from edge to cloud is also a shift from the data centre-centric way of thinking. It can enforce European policy objectives, such as



privacy, and be implemented by a diversity of providers, etc.

The current direction in AI is also producing a bifurcation that policymakers, companies, and researchers must address explicitly. On one trajectory sits “AI for everybody”: capable, affordable models embedded into products, workplaces, and public services, increasingly executed close to data sources and users. On the other trajectory sits the pursuit of “superintelligence”: a belief-driven, capital-intensive effort that concentrates compute and talent in a small number of actors, and that is already reshaping geopolitics and industrial policy. Europe will have to explicitly address these trajectories.

This is precisely where HiPEAC’s long-running focus on the computing continuum and the “next computing paradigm” becomes more than a conceptual framework. **The future we must design for is not a monolithic intelligence running in a single place; it is a federated fabric of specialized models, tools, and services that cooperate under orchestration.**

In the foreword to the HiPEAC Vision 2025, we already argued for open interoperability and distributed agentic AI as a pragmatic subset of this paradigm. This logic strengthens in 2026: AI models can be fine-tuned, validated, and deployed on heterogeneous devices; their composition can be adapted to latency and privacy requirements; and their lifecycle can be managed with clearer safety and compliance boundaries than a single opaque system. Europe’s industrial structure – rich in small and


medium enterprises (SMEs), with significant embedded systems and sector-specific expertise – maps naturally to this “many specialists, one orchestrator” direction, provided we define the protocols, tooling, and evaluation culture to make it real.

Moreover, the NCP is not incompatible with the pursuit of “superintelligence”. Europe still ranks high in terms of global computing resources, but these resources are distributed over multiple HPC centres and medium-sized data centres. The NCP can federate European HPC centres, data centres and even smaller compute resources into a virtual distributed supercomputer that could compete with monolithic multi-gigawatt data centres proposed by other countries.

One other impact of AI can be illustrated by a less well-known quote from NVIDIA Chief Executive Jensen Huang:

You know that in the future, the vast majority of content will not be retrieved, and the reason for that is because it was pre-recorded by somebody who doesn’t understand the context, which is the reason why we had to retrieve so much content ... If you can be working with an AI that understands the context – who you are, for what reason you’re requesting this information – and produces the information for you, just the way you like it, the amount of energy you save, the amount of network and bandwidth you save, the waste of time you save, will be tremendous. <...> The future is generative.

Jensen Huang, CEO of NVIDIA,
March 2024 [6]



Two years ago, this was science fiction; now it is becoming more and more realistic. Generative AI is becoming increasingly capable in multiple domains: of course code generation (one of the main emphasis of the major AI players), text generation, video generation (Kling, Veo 3, ...), audio (music) generation (Suno, Udio or if you want to have it locally SongGeneration Studio for example), image generation (becoming mainstream now with all chatbots, e.g. “Nano Banana 2” for Google Gemini, locally possible with many tools, one precursor was Stable Diffusion) and possibly, in the near future, interactive games (Genie 3 from Google Deepmind is an interesting example of the current potential).

Despite the societal impact it may have, **it will be possible to generate most of the digital content that was previously stored in large servers using AI.** It could therefore be possible to move from the current era of (stored) data (that drove the cloud era) to an era of generated, on-demand data, therefore trading storage capacity by (potentially local) computing capacity. An AI model can “condense” a lot of data in its parameters and compute, and could be executed locally (at least for code, text, image, audio, perhaps not yet for full video).

Today it is perfectly possible to generate code for a bespoke application for personal use. For medium-complexity applications, it is now more efficient to ask a chatbot to generate a dedicated application than to try to fetch it from the web (moving therefore from stored code to on-demand generation of code). It is possible that in the future, AI will be at the core (we don’t mean generating everything, but used more and more) of entertainment and its various media, with all potential societal, social, cultural and ethical impacts that this would entail. (Side

A small thought experiment (not realistic, but just as an eye opener):

- 500 M ($5 \cdot 10^8$) smartphones in Europe in 2025
- Typical mid-range 2025 phones deliver roughly (≈ 10 – 25 TOPS) (10^{13}) (these will be most of the market in three to five years from now)
- 40% of them are recharged at night (so they might be remotely available for computing)
- Therefore, if we could combine all of these smartphones (through the NCP) into a super distributed compute fabric in 2030 during their recharging time, the potential compute power would theoretically be 2 ZOPS ($2 \cdot 10^{21}$). Obviously not everybody would contribute, but even if 10% were to participate, this would still provide impressive computing power without direct investment in large data centres.

note: The cartoons and illustrations in this Vision were generated by a talented human (Arnulf). Most of the text was written by humans, but chatbots were nonetheless used as helpers in various parts of this document).

We have witnessed this first hand in the HiPEAC Vision editorial board. One member wanted a particular application running on a laptop. After spending more than one hour in looking for available (stored) software on the web, then on open source repositories, he gave up due to incompatibility of libraries, compilers, etc. Then he gave to a chatbot the specification of what he wanted and in less than 30 minutes the chatbot generated code that ran in a web browser container (therefore safely) and that perfectly satisfied the requirements. The chatbot even proposed adding additional useful features.

Another major potential impact is that **AI could be the entry point of users to cyberspace**, as proposed in the concept of the “Guardian Angel” (or “Digel”) presented in previous editions of the HiPEAC Vision. AI could replace all the applications in a smartphone by one single portal: the AI chatbot interface. AI could interface with services or do what all other applications are doing, i.e. look for information on the web – something which

is already well underway, with 37% of consumers saying they begin their searches with AI tools rather than traditional search engines [7]. Having provided the interface to various services – something which is already possible e.g. by ChatGPT through MCP for example – AI could then generate audio, video, games and even new (virtual) applications, as explained above.

The HiPEAC Vision website is already ready to be accessed by AI: the Vision’s retrieval-augmented generation (RAG) tool is now a part of a MCP server [8]. Browsers are evolving to integrate AI like OpenAI’s Atlas or Chrome+Gemini; soon we will be able to show these agents that an MCP exists and they will use the correct tools depending on where on the website they are. This is another step towards the global NCP concept.

We are witnessing an emerging battle on this “Super-App” that will replace all others. This could also have impact on the device itself: perhaps a smartphone with a touchscreen would not be the best device for this single chatbot interface, and the most suitable form factor and functionality has not yet been defined (glasses, pins, stylus, watch, pendant,...) but this could certainly be a topic for innovative solutions (from Europe?).

Of course, AI is not the only topic in the computing domain, but it has drastic impacts. Some of these impacts are a direct result of AI, such as the way we will develop software and hardware, and the direction in which hardware is going – for example GPUs will support less double-precision floating-point computation – necessary for classical simulation but not necessary for AI-related computation. Others are indirect: for example, AI is currently cannibalizing a lot of resources, in terms of people, money, energy and components. The scarcity of RAM (and flash memories) will have direct impact on non-AI devices, including consumer devices such as smartphones, personal computers (PCs), and game consoles, where the price will rise (Figure 1).

Implications for how computing systems are developed

As AI systems improve at code generation and reasoning over large codebases, the traditional model of development where human engineers manually write most of the implementation is rapidly becoming obsolete. As noted in the opening quote from Andrej Karpathy, developers increasingly supervise AI agents that generate and modify code rather than writing every part themselves. Engineers therefore focus more on specifying goals, constraints and evalu-






Figure 1: Increase in DRAM price due to the demand for building AI data centres



ation criteria, while AI systems generate implementations, explore design alternatives, run tests and propose improvements.

This transformation moves the scarce human contribution upward in the abstraction stack. As more code is generated automatically by AI systems, the **specification becomes the primary human-owned artefact**. Humans must remain in the loop to ensure correctness, safety, security and compliance. This shifts the **main bottlenecks to specification, verification** (checking that the generated code satisfies the specification) **and validation** (ensuring that the system's behaviour in real-world use matches the intended outcome). In this context, development toolchains themselves become strategic infrastructure. Control over compilers, development environments, hardware design tools, verification frameworks and deployment systems increasingly determines how engineering is practised and which ecosystems dominate (The tools chapter of this edition provides further discussion of these issues).

There is also a rapid evolution of digital systems interacting with the physical world (cyber-physical systems – CPS). In particular leadership in safety-critical AI is seen as essential for maintaining European strategic autonomy. While recent developments suggest that AI may increas-



ingly support or even automate parts of the cognitive work involved in system design and assurance, this transition is not straightforward. CPS are typically engineered across multiple specialized domains [9] – such as safety, security, and performance – each with distinct methodologies, timescales, and risk models.

These fragmented approaches create challenges for both human designers and AI systems, which require **consistent frameworks, shared criteria, and sufficient structured data** to operate effectively. This requires **coordinated advances on cross-cutting challenges between stakeholder knowledge domains**. Assurance of safety-critical systems depends on **traceability, explainability, and independent validation** – properties that must be preserved even as AI becomes more capable. In this context, **Europe's long-standing strengths in rigorous engineering methodologies, certification processes, and regulatory frameworks represent a key opportunity**. They have always been slow and expensive, but that is precisely the kind of problem AI is good at solving, not by replacing rigorous methodology with something opaque, but by accelerating well-defined processes so they scale.

Detailed recommendations are provided in the supporting report *Cyber-Physical Systems – Investing for our Technology Sectors and Future European Markets* [10]. Three encompassing ones, supporting NCP physical interactions and European leadership of CPS technology markets, are provided here.

First, Europe will need to establish **governance and coordination frameworks** that are specifically focused on the creation of **composite technologies for physical interaction**, including mechanisms for collaborative research, alignment with emerging regulatory environments, and shared approaches to managing risk and responsibility across complex value chains.

Second, research and innovation efforts must **advance integrative technologies for the safe and effective deployment of the NCP**, including tools for cross-domain

system integration, life-cycle management of AI components, and improved techniques for automating safety, security, and performance interactions in complex systems.

Third, Europe must invest in **scalable and resilient infrastructures supporting highly distributed and interconnected systems, and advanced methods for monitoring system performance and isolating hazards in real time**.

Cybersecurity too is currently undergoing a profound transformation driven by **the rapid integration of AI into both attack and defence**. AI enables unprecedented automation, scalability, and speed of cyberattacks, ranging from large-scale vulnerability exploitation to highly targeted disinformation and influence campaigns, and even AI-supported high-level attacks. This evolution challenges traditional human-in-the-loop-based cybersecurity approaches, which are intrinsically not designed to operate at such scale and / or speed. Consequently, **cybersecurity must evolve toward automated, scalable, and AI-empowered defence methods**, capable of matching and even outmatching AI-enabled adversaries.

The trustworthiness and security of the AI systems themselves is also a key challenge. Indeed, examples of AI such as **LLMs and autonomous agents present new attack surfaces**, such as prompt injection and adversarial manipulation. As explained just above, these vulnerabilities raise concerns about the reliability of AI in critical applications, i.e. especially when systems interact with the external physical world or with sensitive data. Addressing these risks requires a combination of **secure-by-design AI architec-**



tures, robust evaluation methodologies, and new protection mechanisms capable of separating trusted instructions from untrusted inputs. More broadly, the spread of AI must be accompanied by advances in formal verification, monitoring, and resilience techniques, ensuring that AI systems can be trusted even in adversarial environments.

The growing dependence on complex software and hardware supply chains further amplifies cybersecurity risks. Reinforcing IT supply chain security, covering software components to distribution infrastructures, is another crucial action for the EU. The complexity of software supply chains poses significant risks as vulnerabilities can be propagated quickly and widely, thus compromising integrity across applications and domains worldwide. Key strategies include robustly identifying software components and developing high-quality software bills of materials (SBOMs) to track dependencies, which shall facilitate transparency and quick response to vulnerabilities. Finally, improving source-code sanitation methods (e.g. for packages) and tools in software-supply repositories (including open source) is necessary to improve code cybersecurity, requiring automated techniques for vulnerability detection and automated software repair.

Lightweight yet robust authentication mechanisms are necessary to address the distributed nature of the NCP, ensuring strong security with minimal overhead. Secure inter-service communications within the NCP must guarantee confidentiality, integrity, and authenticity. Additionally, the NCP's distributed nature must be leveraged to increase fault tolerance and resilience, thereby fostering greater trust in the system.

Cybersecurity is also closely linked to European strategic independence and sovereignty, in addition to the technical problems listed above. The EU's current digital infrastructure relies heavily on technologies from outside Europe, such as cryptographic trust anchors, cloud services, and hardware and software supply

chains. This reliance poses both a technical and geopolitical risk. To strengthen EU sovereignty, we must avoid measures that weaken security, like requiring cryptographic backdoors. Instead, we should support trust infrastructures, verifiable supply chains, and secure open source ecosystems that are controlled by Europe. In this context, transparency, auditability, and control over critical components become essential to ensure long-term resilience and independence.

Finally, addressing these challenges requires sustained investment in human capital and organizational capabilities in cybersecurity and AI. Raising population-wide awareness of basic cybersecurity practices and expanding cybersecurity education, developing specialized training programmes, with the tools of today and tomorrow – including AI, are therefore critical priorities. For a comprehensive discussion of these issues, refer to the cybersecurity chapter in this edition.

Baking sovereignty and sustainability into European computing systems

Given its context, Europe could and should leverage open source, beyond its important role in supply-chain security, as a strategic anchor to achieve true digital autonomy. Indeed, open source is a fundamental pillar of Europe's sovereignty and sustainability in a software-defined world.

Europe must make sure that software resources, including all open source components hosted in package registries outside of the continent, remain available to European citizens and companies through sovereign infrastructures that support the software supply chain.

Furthermore, pooling resources to co-develop industrial-grade, European-led open source platforms reduces industry-wide duplication of effort and creates de facto standardization. This collaborative open strategy has demonstrated efficiency in software and should be extended into hardware, through continued strategic investments in mature,

open source RISC-V designs and electronic design automation (EDA) toolchains. This should be promoted in the context of establishing the NCP.

Open source can also be a powerful sustainability enabler: mandating open APIs and open source implementations for end-of-life consumer electronics can combat programmed obsolescence and grant functional hardware a virtually indefinite second life. The chapter on open source examines all of these themes in depth.

Solutions like this are important given that the sustainability of computing systems has become an increasingly urgent topic. The rapid growth of AI data centres and edge devices has prompted significant concerns regarding the environmental sustainability of AI technologies. The environmental impact of AI encompasses three primary areas: (i) embodied emissions associated with device manufacturing, (ii) operational emissions incurred during usage, and (iii) indirect emissions – or, conversely, reductions in emissions – resulting from the application domains where AI is deployed. Indirect emissions are largely influenced by customers, rather than vendors, while operational emissions depend on the carbon intensity of local electricity sources, which utility companies control and are anticipated to decrease in the future. Embodied emissions remain challenging both to model accurately and to allocate fairly among multiple applications sharing a single device.

Debates about sustainability can only be useful if they are based on evidence. Therefore, it is crucial for the computing industry to establish standard method-



ologies for measuring impact through life-cycle assessments and to incorporate these metrics into design tools. For many applications, the embodied emissions are dominant (either because the device is very power efficient, or because the device is not intensively used). Since semiconductor dies (i) are causing the bulk of the embodied emissions, and (ii) are notoriously difficult to recycle, their environmental impact is minimized by (i) using them as long as possible, and (ii) using them as intensively as possible before discarding them. The industry should strive to considerably extend the lifetime of devices and find ways to distribute the embodied emissions fairly over the different applications running on the device. In parallel, it will also have to adapt its business models to make them more sustainable, which means that they cannot be based on sales volume as the primary metric. For a thorough discussion of sustainability in computing systems, refer to the chapter on sustainability.

What does all this mean for Europe?

Europe lacks companies with the scale of the so-called “Magnificent Seven” in the US or the large players in China. On the other hand, Europe has in the last decade developed a competitive HPC infrastructure, it has more founders than the US, it has a net brain gain of researchers, and, in 2024, more startups were created in Europe than in the US. Europe has many strengths, which should be leveraged to counter the narrative that Europe is weak, that it cannot compete with the rest of the world or defend itself. Focusing on the areas in which Europe could excel, rather than blindly copying what others are doing, Europe should in parallel strengthen its scientific, industrial and innovation capacity.

As a first step, Europe should become more entrepreneurial, starting during secondary education and continuing during higher education. For example,



Europe should set a target and facilitate the creation of 100,000 startup companies per year, spread over the European Union, which should be incubated in science and technology clusters where they find all the resources (talent, coaching, capital, ...) to grow. At the same time, Europe should make it easier to found a company and to scale it up and it should create challenge-based funding instruments similar to DARPA or SPRIN-D, fast and agile, and with clear potential for commercialization.

Finally, “invest in Europe” and “buy European” policies should be promoted, while goods and services that still need

The Japanese Rapidus approach*

- RUMS (Rapid and Unified Manufacturing Service)
- “Raads” (Rapidus AI-Assisted Design Solution) supports design optimization using AI, and “DMCO” aims to mutually optimize design and manufacturing
- “GAA for the front-end process and “chiplets” for the back-end process were developed to deliver the world's fastest cycle time.
 - The **single-wafer process** can flexibly respond to the production of a wide variety of specialized products.

On July 18, 2025, Rapidus unveiled its first wafer featuring a gate-all-around (GAA) transistor fabricated using a 2 nm process. This milestone was reached just over three months after transporting an extreme ultraviolet (EUV) lithography system in December 2024 by air from the Netherlands. By June 2025, the first production lot had been processed, and the resulting GAA transistor wafer was showcased at a customer event in July. Rapidus attributes this exceptionally rapid establishment of production capability to its defining manufacturing strength.

* From <https://www.rapidus.inc/en/business/#technology>

** From <https://www.rapidus.inc/en/interview/it0003/>




Figure 2: Example of rethinking semiconductor manufacturing

to be developed should be ordered from within Europe.

Without an entrepreneurial mindset, the resources to found a company and to grow locally, and scale up in Europe and beyond, Europe will become less competitive in the future. Since growing an entrepreneurial mindset and building a science and technology cluster takes time, Europe should start working on it today in order to reap the benefits in 10-15 years.

If Europe is tempted to “copy what the leaders are doing”, 2026 offers a caution and a blueprint. The caution is that copying, by definition, makes you late – especially in domains with multi-year capex cycles and tight supply constraints. The blueprint is to compete on the dimensions where different

choices can be pivotal: agility, integration, vertical focus, and time-to-market.

Japan’s Rapidus (Figure 2) is an instructive example of potential differentiation. Rapidus is a government-backed attempt to re-enter leading-edge manufacturing with an aggressive roadmap, and industry reporting highlights its emphasis on process agility using AI and faster turnaround cycles compared to traditional batch paradigms. The lesson for Europe is not to replicate Rapidus, but to internalize the strategic pattern: choose a segment of the value chain where “**thinking differently**” yields a structural advantage, then align research, engineering, and industrial policy behind it.

In conclusion, the HiPEAC Vision 2026 promotes a **systems-oriented direction**.

HiPEAC assumes that impact in the next wave of computing (closely linked to AI) will be delivered by a heterogeneous continuum: **specialized services (models) collaborating under orchestration, running across edge, cloud, and everything in between; AI becoming increasingly intertwined with the physical world** (hence needing to cope with its constraints); **specialized accelerators coexisting with flexible platforms; and memory-centric designs shaping what is feasible outside hyperscale data centres**. Europe can lead here, but only if it **stops optimizing for imitation and starts optimizing for coherence**: with an innovative mindset, through the creation of innovation and startups, shared roadmaps, shared tooling, shared benchmarks, shared protocols, and shared infrastructure – aligned with Europe’s strategic autonomy objectives and industrial realities.

References

- [1] <https://x.com/karpathy/status/2026731645169185220>
- [2] <https://arxiv.org/abs/1706.03762>
- [3] <https://github.com/openclaw/openclaw>
- [4] <https://www.technologyreview.com/2026/03/11/1134179/china-openclaw-gold-rush/>
- [5] <https://www.kickstarter.com/projects/tiinyai/tiiny-ai-pocket-lab>
- [6] <https://www.hpcwire.com/2024/03/19/the-generative-ai-future-is-now-nvidias-huang-says/>
- [7] <https://searchengineland.com/consumers-start-searches-ai-not-google-study-467159>
- [8] <https://github.com/hipec/hipec-mcp>
- [9] Charles R. Robinson et al. (2026). Representing the Knowledge Domains that Characterise Cyber-Physical Systems. HiPEAC. <https://doi.org/10.5281/zenodo.19063577>
- [10] Charles R. Robinson et al. (2026). Cyber-Physical Systems - Investing for our Technology Sectors and Future European Markets. HiPEAC. <https://doi.org/10.5281/zenodo.19091135>

Recommendations



The recommendations of the HiPEAC Vision provide a blueprint for computing research which will allow Europe to meet its ambitions for the future.

Previous editions of the HiPEAC Vision explained that Europe is de facto in various races, mainly competing with countries like the United States (US) and China. However, competing in a race suggests that competitors share the same objectives and values. For some years, it has become increasingly clear that the European context, and the values promoted by European policy and society, differ in important ways from those of its main competitors. Hence, while the urgency of the races is still evident, the objectives, and therefore directions, may well be different in Europe, to respond to the European ecosystem

with its specific structural, economic, cultural and ethical concerns. Moreover, it is clear that if Europe tries to blindly copy the main competitors in the race, it risks lagging behind due to its late start. Therefore, as stated in the introduction to this edition, the core of the HiPEAC Vision 2026 is *“Europe should not imitate others; it should create its own future”*.

As an example, the current race in artificial intelligence (AI) is focused on bigger and bigger data centres, which cause increasing problems in terms of energy consumption, access to the grid, cooling, consumption of natural resources (water...), availability of components (graphics processing units (GPUs), memory), centralizing knowledge and power in the hands of few large technology companies, etc. But

there is an alternative approach proposed by the “next computing paradigm” (NCP) endorsed by the HiPEAC Vision: **distributing computing on demand** instead of systematically requesting cloud resources.

To enable the NCP, the technologies on which it is based, and the governance required for its operation, the following consolidated recommendations are proposed. They are clustered into **technological recommendations**, **methodological and standardization recommendations**, and **policy recommendations** and include both technological and operational advancements necessary to realize the NCP. Detailed recommendations and their rationale can be found in the associated chapters.



Technological recommendations

Create the digital envelopes paradigm

“**Digital envelopes**” are the “anything-as-a-service”-based solution, the basic concept underlying the NCP proposal. Digital envelopes should be created to integrate “anything” (a person, company, physical entity, digital resources such as computing, storage, etc.) in the digital networked space, allowing interoperable access to the array of horizontal (collaborative) and vertical (enabling) services that can be delivered there and that can also be “twinned” to actions occurring in the physical world. Digital envelopes should

be able to live-migrate across the compute continuum, with live migration allowing seamless continuity of service in accordance with changes occurring in the physical world.

Several technologies need to be further developed to enable functional digital envelopes. A starting point is the first use case of the NCP: supporting **agentic artificial intelligence**, which will both contribute to the NCP and act as a blueprint of a more generic NCP.

Create a subset of the NCP infrastructure and ecosystem to support agentic AI and specialized action models (SAMs)

The infrastructure for agentic AI, (called here A^2I : agentic artificial intelligence infrastructure, also simplified as $(AI)^2$) should be a subset of the NCP and should act as the foundation from which the full-fledged NCP can emerge. It should include non-functional properties, allowing the dynamic selection of agents according to the user’s criteria. In view of the envisioned NCP, where digital envelopes will

“float” while seeking to achieve user-set goals, a **seamless, continuum-friendly runtime infrastructure for orchestrated (and “orchestratable”) services** to run on should be developed.

Novel and emerging software technologies that can help implement this runtime infrastructure should be scouted. Eligible technology solutions should be **open, amenable to sandboxed execution**, and capable of **near-native performance** as well as of **permissioned access to host-based services (authentication)**.

The key elements of this interoperable infrastructure are:

Orchestrators for AI agents

Systems for managing agents, helping to select and orchestrating the growing suite of AI agents they’re using from different providers. **Europe should develop orchestrators** that allow the **composition of services** provided by agents and SAMs in response to a request (“prompt”) and the **non-functional criteria** associated with the request. The orchestrators should be **trustable** (for example, open source and therefore auditable; they should work for the **benefit of their user**, and not for their creator or related third-party interests), and they should be **interoperable**, therefore their specifications should be open and should lead to standardization.

AI is a technology that will help realize more efficient orchestrators: AI-powered orchestrators will be an essential capability of the intelligent agent at the centre of the digital envelope. Orchestrators should hierarchically decompose user-given goals (where “user” may be a human, a business or administrative entity, another permissioned orchestrator) into a set of sub-orchestrations that cooperate to achieve the set goals. These orchestrations might be generated by (a federation of) generative AI engines or retrieved from trustworthy repositories, both reachable from edge nodes and capable of collaboration with other orchestrators within dynamically federated zones. Core orchestrators should especially be developed for deploy-

ment at the edge, nearest to the final user, in a manner that can dynamically combine collaborative compute components into executable applications tailored around specific user needs.

Agent communication protocol

A **communication protocol between agents**, building on existing de facto standards of 2025 (such as MCP, A2A, etc.) but also allowing support for non-functional properties, such as latency, time, etc. **Adding non-functional properties will be crucial to extend agentic AI to physically entangled systems** such as robots and self-driving cars, and further to extend it to full NCP support. The (AI)² should be **flexible in terms of its implementation** and the location of the agents, from a lightweight protocol when everything is executed within a single computing resource to totally distributed execution on the computing continuum (from edge devices to cloud and high-performance computing (HPC)), where multiple resources are interconnected by “long” distance communication protocol with enforced security.

Agent containerization

“**Containerization**” of agents and SAMs **allowing migration** (and, eventually, portability) in a secure way. The agents and SAMs should work in well-protected “silos” on various kinds of hardware, preventing them from contaminating the system they are running on. A reference software infrastructure for using agents, SAMs and orchestrators should be publicly available and open. These developments should evolve to be more generic and to support any service, allowing smooth migration to the NCP.

Dedicated hardware

Support the agents’ performance and containerization with **dedicated hardware**. **Hardware should enforce the separation of the silos** where agents are running from the underlying hardware. **Dedicated accelerators** should also be developed to increase the efficiency of the agents and SAMs (such as neural processing units

(NPU), for example). This could again be the blueprint for the close containerization of any service.

Develop neurosymbolic systems and solutions for continuous evolution of agents

Develop **agents that can evolve according to their environment** and **adapt** in a controlled way. This is still an open research topic, using either external (short-term) memory in an agentic framework, new ideas for the core structure of AI itself, or new training schemes.

These systems could marry classical systems and generative AI solutions, and add classical supervision systems as agents in the agentic AI infrastructure. Classical systems, with explicit knowledge about the limits of the requested evolution, can offer services that can be combined with ones provided by SAMs.

Use “software-defined everything (SDX)” applications as a driver to develop NCP-compliant technologies and evolve towards AI-defined everything (AIDX)

In recent years, we have seen the transition from systems based on specific hardware to a “software-defined” approach, such as that used for software-defined vehicles, where software is decoupled from hardware and can therefore be updated and optimized, offering greater flexibility and adaptability, and representing a sort of forerunner of the NCP. The “AI-defined” approach builds on this by creating systems that are dynamic and self-optimizing. SDX applications can be used as a driver to develop technologies that work within the NCP framework and pave the way towards AI-centric systems.

Prepare for the reconfigurable design era

The speed of algorithmic evolution, the urgent need for energy efficiency and the rising costs and time for the design of a custom processor, will likely lead to the adoption of reconfigurable elements in chiplet platforms. **Softcores with their**

associated middleware support will gain importance. Long term, **automated code-generation** using generative artificial intelligence should be combined with **automated architecture generation on reconfigurable platforms** to enable the ultimate hardware / software co-design. This is a direct application of the NCP concepts where services (here softcores) can be (dynamically) mapped on a hardware substrate (here reconfigurable).

This requires an understanding of what such a reconfigurable platform could look like (an evolution of the field-programmable gate array (FPGA) and coarse-grained reconfigurable array (CGRA)), how it is governed (dynamic resource allocation, just-in-time synthesis of architecture description), and how it is used and perceived by applications. It will lead to ecosystem or supply-chain changes and is an opportunity for the application of NCP.

Develop integrative technology capacity for AI interacting with the physical world

As AI becomes “physical”, i.e. interacts directly with the world, it has to keep within the world’s constraints like latency, timing, etc. in a trusted manner. It will therefore be necessary to **develop SAMs that can be used for interaction with the physical world** with a particular focus on providing **instances of the (AI)² ecosystem dedicated to robotics**. It will be also neces-

sary to increase research on **tools and methodologies that facilitate life-cycle management of AI components**, and to improve **techniques for analysing safety, security and performance interactions** in complex systems-of-systems. Progress in these areas will enable the deployment of intelligent capabilities while maintaining the levels of assurance required for critical infrastructures and industrial applications.

Ingrain trustability into the fabric of the NCP

In order to be accepted, systems based on the NCP concept should be safe, secure and trustable.

- **Develop scalable, distributed and resilient infrastructures.** The increasing scale and interconnectivity require infrastructures capable of supporting highly distributed and continuously evolving systems. Europe should therefore invest in technologies that enable scalable and resilient infrastructures for system (composite) technologies operating in the real world, including **hybrid edge-cloud data analytics, trustworthy fleet-scale industrial internet of things (IoT) deployments**, and advanced methods for **monitoring system performance and containing failures in interconnected environments**.
- **Promote lightweight yet robust authentication mechanisms for NCP services.**

Develop scalable authentication mechanisms that provide strong security guarantees while remaining lightweight, efficient and practical for high-frequency usage in large-scale distributed systems like the NCP.

- **Secure inter-service communications in the NCP.** Similarly, ensure confidentiality, integrity, and authenticity of communications between the numerous services in the NCP through strong encryption and mutual authentication techniques that can scale up.
- **Exploit the NCP’s distributed nature for resilience and degraded modes.** Leverage distribution to provide fault tolerance and maintain essential functionality under failures or attacks, through the capacity to function in isolation in degraded mode.

Secure AI systems against cyberattacks and malicious influence

As AI systems become ubiquitous, tampering with their proper functioning has a huge impact. It is thus vital to ensure that AI systems are robust against cyberattacks, as well as against adversarial manipulation (e.g. poisoning, prompt injection) through **secure-by-design approaches, standardized evaluation benchmarks, and formal verification methods**. Advanced AI systems should be used to help develop robust systems and consolidate the protection of the NCP infrastructure.



Standardization recommendations

Establish governance and coordinated frameworks for European digital ecosystems

The (AI)², and eventually the full NCP ecosystem, will increasingly operate as large-scale ecosystems that integrate technologies, organizations, and infrastructures across sectors such as trans-

port, healthcare, manufacturing, energy, etc. Different communities from varying disciplines, each with their own priorities, will contribute to these ecosystems, and they need to be aligned so that they work towards a common goal.

Trustworthiness and long-term sustainability require **coordinated governance**

structures that align research communities, industry stakeholders, and regulatory frameworks. Europe should **strengthen coordination for cross-cutting capacity and knowledge exchange** between knowledge domains, including AI, dependable systems and other domains, and establish **shared approaches**

for distributed risk-sharing and certification in interconnected systems.

Set up an “NCP foundation” to ensure interoperability and trustability of the NCP ecosystem

An organization should be established to ensure the alignment of different contributions to the NCP ecosystem, similarly to the way that the RISC-V foundation, RISC-V International, aligns contributions to the RISC-V standard. The role of this organization should be, first, to help **define overall specifications for the ecosystem**; second, to **ensure that agents, SAMs and orches-**

trators conform to requirements and are trustable. This organization should have a longer-term view of supporting the full NCP requirements and technologies. It could also act as a broker between providers and users, and kickstart the business ecosystem and business cases, taking institutional (cities, administrations, etc.) drivers and procurement as a starting point.

Standardize chiplet platforms

Evolving from single-vendor to heterogeneous, multi-vendor systems-in-package (SiP) based on chiplets requires **agreements on the physical, electrical and soft-**

ware interfaces between chiplets. Europe should focus on chiplet-platform systems and application-development methodologies. As the NCP starts to shape the software side, it is important to tune the hardware architecture accordingly. As each chiplet only represents part of the system, and systems are assembled for specific applications, **standard protocols and interface intellectual property (IP) blocks** – along with a method for system self-discovery – should be adopted, with implications on the architectural partitioning of functionality (such as virtualization of shared resources) and the operating system governing the SiP.



Methodological recommendations

Develop the use of AI in software development

- Develop **agent-centric workflows** where AI systems act autonomously in human developer-like roles, planning, invoking and coordinating multiple tools (compilers, simulators, profilers, debuggers, deployment systems) with high-level guidance and learning from specific feedback.
- Use AI to reduce the cost of software development, using natural language as the primary user interface, to democratize software development workflows. This should include the whole deployment workflow for production or testing (e.g. for the NCP). Software developers should leave the coding to specialized agents, and should focus on the completeness of specifications and validation of the results.
- **Embed AI inside optimization tools**, especially for hard, large-scale hardware problems where traditional methods are slow, providing improved solution quality and throughput beyond what is economically feasible with existing methods. This

includes targeting and optimizing code for the NCP to a specific platform / device.

- Evolve **programming languages, models, and related tools and runtimes** to be more easily **targeted by AI.**

Develop the use of AI in hardware development

- Support **disruptive hardware paradigms** where **AI-native approaches can reset the competitive landscape** (e.g. CGRAs, canvas-based design), giving Europe a strategic advantage despite the lack of domestic EDA vendors.
- Use AI-driven estimation and exploration to **shift hardware analysis and optimization earlier in the hardware design flow** (e.g. 3D integration, thermal, mechanical constraints, and embodied emissions).
- Employ **AI-driven hardware / software (HW / SW) co-design agents**, enabling joint exploration of software, architecture, and hardware implementation to optimize defined metrics across the full stack, from application use case to hardware implementation.

Explore technical and operational solutions that facilitate the deployment of personalized, safe and trustable AI services at the edge

Moving AI services and embodiment to the edge of the continuum requires solutions that do not subjugate the human user to the negative and perverse phenomena that arise in the digital space propelled by “attention economy”. The solution that the HiPEAC Vision 2026 promotes is dubbed “**personal AI**” (PAI) and enacting that requires exploring technical and operational innovations such as solutions that support “natural interface” interaction with the human user. Ensure **trustworthy storage of personal data** arising from such interactions and provide **user-friendly trustworthy intermediation** in the acquisition and delivery of collaborative services.

Support the use of AI for cybersecurity and trustability, and to fight influence campaigns

- Promote **AI-powered cybersecurity** to counter AI-driven attacks. AI can enable massive automation of cyber defence techniques, allowing for enhanced threat

detection, real-time analysis and earlier identification of new attack patterns, autonomous prioritization of high-risk events, and autonomous incident response, both in the case of malware infections and of intrusions.

- Develop AI tools to detect and mitigate disinformation / influence campaigns.

Reinforce IT supply-chain security (software and hardware)

Address systemic risks arising from dependencies in software and hardware supply chains that can propagate vulnerabilities, or can be used to propagate malevolent code, at large scale.

- **Support robust identification and high-quality software bills of materials (SBOMs) for tracking and integrity verification of software components.** Enable precise, verifiable identification of software components to improve vulnerability tracking, compliance, and long-term traceability, by supporting solid and precisely defined SBOMs – including the reliance of intrinsic, tamper-proof, software identifiers such as the Software Hash Identifier (SWHID) – and their inclusion in tools.
- **Secure package managers.** The reuse of software (components) is extensive, through dependencies managed by package-management ecosystems that currently offer too many opportunities for attackers. It is thus crucial to support the design of secure yet usable software-package managers.
- **Promote research on source-code sanitation methods and tools.** Improve the security of open source code through large-scale, automated approaches that can systematically detect known vulnerabilities in code repositories, using standardized vulnerability databases and automated analysis tools. Develop reliable techniques to automatically generate, validate, and apply security patches at scale.

Address the verification challenge arising from increasing use of generative AI

- As AI increasingly performs code generation, **keep humans in the loop for acceptance, accountability and specification ownership.** In the short to medium term, human verification will become the bottleneck. Over time, a growing fraction of code may be generated and maintained by AI systems, raising challenges for IP protection and increasing the importance of **precise, auditable specifications as the primary human-owned artefact** over the system life cycle.
- **Develop specialized AI agents to guide, prioritize and scale verification tasks** (test generation, property discovery, and bug localization) for both safety-critical and non-safety-critical systems, including code generated by AI. In particular, use AI to guide and prioritize formal verification rather than replacing it. This is an opportunity for Europe, which is strong in formal verification.
- **Establish European verification testbeds and benchmarks for AI-assisted development and verification tools,** enabling systematic comparison, reproducibility and shared evaluation methodologies for competing approaches (in both academia and industry).

Develop “outside-the-box” concepts and ideas and test them rapidly in the (AI)²

This can be facilitated by **sharing knowledge on emerging new ideas and new open source solutions** in a structured and open manner (knowledge infrastructure). Easy access to tools and resources to validate proposals should be provided, with an open mindset. Once tested, the new idea can use the supporting infrastructure to enable swift transition to a real product.

Use validated life-cycle models for computing

Making informed decisions on the environmental impact of IT requires

solid data about the impact of products and services. This data is based on a life-cycle assessment (LCA). An LCA starts with defining the system boundaries (which emissions are included), makes assumptions (a device is used for x years), and is based on the actual emissions of the components of a system. Every individual product or service comes with its own environmental footprint data. The computing-systems community should be encouraged to **carry out an LCA on all the products and services it produces.** The LCA should take into account **embodied emissions, end-of-life emissions, operational emissions,** and if possible, **indirect emissions** (rebound and enablement). This information should be included in a **digital product passport (DPP)** containing information about the environmental impact. This information will help consumers to make informed sustainability choices. It should also be **used by the orchestrators in the NCP to select services that optimize the sustainability requirements** specified by the owner of the orchestrator.

Promote sustainability-focused design methodologies and business models

Accurate life-cycle models will help designers or AI-based design tools to make the most effective eco-design decisions. To be effective, **design tools should automatically provide the environmental impact of the components and technologies used in the design,** without placing the burden on the designer. Incorporating reparability, reusability, recyclability, and end-of-life processing considerations from the beginning of the product development process will also lower the environmental impact of the final design. **LCA should be part of a standard design space exploration.** Inevitably, reducing the environmental impact of a product will have an impact on companies’ business models. Designing products that last longer will reduce sales of new products and hence lower the profitability of the company. This can only be mitigated by **developing new business models,** e.g. based on extra services: maintenance, repair, disposal, ... up to completely replacing the ownership of hardware by a service contract.



Policy recommendations

Position Europe in the “artificial super intelligence” race

Some actors in the field of AI aim to develop systems that could be of higher “intelligence” than human experts to solve very complex problems. This involves large efforts (comparable to the Manhattan Project) and extreme computing power as of the current state of the art. Such systems, as of today, are self-improving: any progress done makes the next iteration faster and better, which means any delay potentially becomes exponentially costly to catch up. This is a real race between companies and countries involving enormous investments.

- **Europe should decide quickly if it wants to enter this race.** It is first a political decision. If so, Europe can leverage its position in global processing power by **developing and using a transnational (AI)² infrastructure** (in which each agent (“component”) could be a complete HPC centre), using a unified infrastructure arising from NCP concepts and intelligently harnessing existing decentralized European hardware resources and extending the AI Factories initiative into a unified computing continuum.
- This is an action that calls for **creating a physical centre like CERN** to bring together **competences within Europe in the field of AI** and coordinating actions, for example extending the RAISE initiative. This should also help to assess the state of the art and the innovations needed.
- While developing advanced models, **emphasis should be on the alignment problem:** developing approaches that control the level of model “cheating”, avoiding deceptive behaviours and the development of internal goals that are not “aligned” with the goals of the system designers. This is also an opportunity to

develop **advanced models that reflect values and culture endorsed by European society.**

Strengthen European capacity in heterogeneous chiplet design

Design flows should be adapted for compatibility with multiple physical integration flows. **Europe should strengthen its EDA capabilities for chiplet platform architectures**, and prepare for the deeper move to **optical communication and processing in combination with electrical digital processing and memories.**

Bridge the gap between open source software and open source hardware

While open source software has a reproduction cost of zero, open source hardware faces huge capital expenditures for manufacturing. To support European semiconductor sovereignty, investment should **support the creation of mature industry-grade open hardware-based designs**, and beyond that, to secure open, end-to-end toolchains from EDA design to physical silicon.

Evaluate the move to cryogenic operation in data centres

Data centres are today fundamentally power and thermally constrained. Dissipation-free circuits that operate at very high frequencies, based on Josephson Junctions, could break the trend. Superconducting logic has been known for years, and recent advances have made it manufacturable at scale. Their obvious limitation is their need to be cooled to cryogenic temperatures, which requires a more or less fixed amount of power. In operation, their variable power consumption is much lower than for CMOS-based digital circuits. The crossover point, where the cooling power and dynamic power consumption of superconducting logic is lower than for

conventional complementary metal-oxide-semiconductor (CMOS) technologies, is crossed in petascale systems.

Prioritize cybersecurity for European sovereignty

- **Do not mandate backdoors or weaken encryption and defences.** Despite recurring pressure, **preserve strong cryptographic and cyber protections** as necessary fundamental foundations for secure digital infrastructure.
- Develop and promote **EU-based public key infrastructures (PKI) and certificate authorities (CA).** Strengthen EU control over cyber trust infrastructure. PKIs and CAs are crucial to ensure the uninterrupted functioning of information technology (IT); these should be based in the EU to increase the EU’s security and sovereignty.
- Mandate **verifiability, transparency and auditability of non-EU components**, to be able to assess their trustworthiness.

Increase cybersecurity (and AI) skills and awareness

- Increase **European cybersecurity workforce capacity**, expanding cybersecurity education to train more specialists. Address the shortage of skilled professionals needed to secure modern digital infrastructures, by strengthening education and training programmes to develop a larger and more capable workforce in cybersecurity as well as in AI for cybersecurity tools.
- Raise **population-wide awareness of basic cybersecurity practices.** Improve general cybersecurity hygiene to reduce risks linked to human factors.

Establish sovereign infrastructure for the software supply chain

Real digital autonomy requires immediate action to ensure that Europe has freedom of action regarding the **software dependencies** we rely upon. The current landscape features a critical vulnerability: package repositories and registries are all hosted and legally controlled outside of Europe. The risk is that if those registries were to become inaccessible, EU software would stop working. **Establishing fully mirrored, European-hosted infrastructure for critical dependencies** is a necessary step to mitigate geopolitical risks, export controls, and unilateral technology withdrawals. This sovereign infrastructure should be treated as a **public utility**, funded and maintained to ensure uninterrupted access to the building blocks of the modern digital economy.

- **Secure critical open source components as a European common good.** Identify key open source components as critical infrastructure and support their sustained maintenance and access.
- Avoid application of foreign export control of tools and the software built using them through either (a) **open source tools** (e.g. GCC, LLVM, open source EDA tools) or (b) **fully European developer tools**.
- Support **portability and retargetability across diverse accelerators, specialized chips, domain-specific architectures**. Avoid foreign vendor control through vendor-specific programming models or application programming interfaces (APIs). It is especially important to **enable cross-vendor composability** with respect to application binary interface (ABI) compatibilities, intermediate representations (IRs), security model enforcement, verification, for the NCP or EuroStack.

Maximize impact with European-led open source platforms

Europe should **pool its talents to co-develop industrial-grade, European-led platforms in open source** as they have

the potential to accelerate standardization and drive adoption. The European Union (EU)'s research funding ecosystem presents an opportunity to maximize the impact of open source innovation. To avoid research efforts starting from scratch, and unintentionally resulting in "abandonware", software that falls out of use when research grants expire, **public funding could be leveraged to reinforce collaboration around open source platforms**. This would require organizations to be incentivized to extend existing platforms and funding for new platforms to be extended beyond the standard grant cycle to support projects to the point that they are sustainable, which usually requires five to six years.

Additionally, establishing a **European sovereign technology fund** could help support existing platforms, promote an "upstream first" collaborative culture, and empower communities to sustainably secure the critical digital infrastructure we rely on.

Leverage open source to ensure digital sovereignty in the age of AI

To ensure true digital sovereignty and prevent "open-washing", **policymakers should endorse the Open Source Initiative's (OSI) Open Source AI Definition (OSAID)**. This requires that AI systems marketed as "open source" go beyond publishing "open weights" by mandating full transparency in training code, and detailed information regarding training data sources and preparation methods. Establishing this rigorous approach ensures that developers can genuinely use, study and modify these AI systems.

Furthermore, strategic initiatives should support the **deployment of AI agents** specifically designed to **maintain, secure, and manage the growing influx of code contributions within the foundational open source ecosystems** that power the AI revolution.

Mandate open source APIs for end-of-life consumer electronics

To combat e-waste and programmed obsolescence, regulators should **mandate that consumer electronics manufacturers open source device APIs and technical documentation at product end-of-life (EOL)**. When proprietary cloud servers shut down, this requirement ensures devices can still be controlled locally rather than becoming useless. By empowering open source communities to maintain these abandoned systems, functional hardware gains a virtually indefinite second life. Ultimately, this protects consumer investments and leverages open source development to drastically reduce the environmental impact of digital products.

Aim for a target of 100,000 tech-enabled startups per year in Europe

To grow scaleups and unicorns, there must be enough startups to start from. European policy makers should **set a target and help enable the conditions for the launch of plenty of tech-enabled startups per year**. These are not only purely digital startups but also include startups that leverage digital technology to innovate in a different sector. The startups should be distributed over all the countries / regions in Europe, making each local government responsible to meet their assigned targets.

Build more science and technology clusters

The EU has only one science and technology (S&T) cluster (Paris, ranked 12) in the global ranking of the 20 largest S&T clusters of the world. Science and technology clusters are ecosystems that help startups to hatch and grow by providing affordable facilities, the proximity of a world-class higher education institution providing a talent pool, incubators and accelerators, growth capital, a favourable legislative framework, and first and foremost a vibrant community of entrepreneurs. Such an ecosystem will attract founders (including from outside Europe), it will create well-paid jobs for young graduates from the higher education institutions, and

if it is good, it will keep the scaleups close to the local ecosystem. **Europe should therefore actively promote the creation of European S&T clusters in every major urban area** and help them grow to a scale that they can support scaleup companies.

Introduce DARPA model of challenges

DARPA (the United States' Defense Advanced Research Projects Agency) funds high-risk, high-rewards projects to generate transformative technologies. DARPA focuses on radical innovation and is willing to accept failure as part of exploring new ideas. Projects are quite short (two to five years) and must show measurable progress quickly. They are led by entrepreneurial programme managers who have a vision for technology breakthroughs, scout for innovative ideas, assemble the best teams and take corrective action if milestones are not

met (including termination). This introduces a **new research and development (R&D) culture: fast, milestone-based, competitive, risk-tolerant, visionary, agile. Europe should use a similar model to tackle some of the grand challenges.** The challenges should focus on frontier technology with a clear societal benefit: clean energy, smart mobility, advanced robotics, affordable health care, ...

Invest European, buy European, and stimulate pre-competitive procurement

There is a lack of venture capital (VC) for European scaleups, but at the same time Europeans invest in US and Chinese companies, because they want to maximize their return on investment. By doing so, they are strengthening the non-European competition for the European companies.

European money is not used to strengthen European companies.

Incentives should be made to encourage investors to fund ventures which start – and stay – in Europe. Policy makers should support European technology by **amending public procurement rules to specify European providers.** Pre-competitive procurement is another tool which can help stimulate European innovation.

Keep the momentum

Of course, the recommendations of the previous HiPEAC Visions hold, and some of them are already started to be realized, so the momentum should be continued.

Some salient insights from previous editions of the HiPEAC Vision:

Agentic AI and orchestration

The HiPEAC Vision 2021 called for a “moonshot” programme to develop “Guardian Angels”: “advanced orchestrators... loyal to their users... to orchestrate the new web in a safe and secure way”. This concept also appeared in the comic *Past, present and future of the Internet and digitally-augmented humanity. A HiPEAC Vision*, published in 2020; see Figure 1.



Figure 1: Excerpt from the HiPEAC comic published in 2020

Raising the level of abstraction

The central premise of the HiPEAC Vision 2010 was “keep it simple for humans, and let the computer do the hard work”. The HiPEAC Vision 2015 recommended “[d]evelop[ing] approaches that enable domain specialists to express only what needs to be done, while tools and computers take care of how to transform this knowledge into an efficient computing representation and execution”.

Digital sovereignty

The HiPEAC Vision 2017 contained the following warning: “Many countries understand that computing is a key-enabling technology of strategic importance, and invest in their own research, products and companies. If Europe fails to do the same, it might eventually become dependent on technology which is designed, developed, produced, and controlled outside Europe.” The HiPEAC Vision 2023 noted that governments can “cause disruption or even completely cut off other nations and regions, merely because of where an essential piece of software was created or a where a critical service is operated” and recommended that Europe ensure access to digital essentials.

The Next Computing Paradigm



Recommendations for the next computing paradigm

Create digital envelopes

As agentic artificial intelligence (AI) becomes increasingly embodied and the “as-a-service” paradigm becomes the principal delivery mode of virtually everything, an operational model needs to be created to enable seamless physical-to-digital interaction in the compute continuum, in a controlled and trustworthy manner.

“**Digital envelopes**” are the “anything-as-a-service” based solution to this need. Digital envelopes should be created to integrate “anything” (a person, company, physical entity, digital resources, etc.) in the digital networked space, allowing interoperable access to the array of horizontal (collaborative) and vertical (enabling) services that can be delivered there and that can also be “twinning” to actions occurring in the physical world. Digital envelopes should be able to live-migrate across the compute continuum, with live migration allowing seamless continuity of service in accordance with changes occurring in the physical world.

Develop AI-powered orchestrators

AI-powered orchestrators will be an essential capability of the intelligent agent at the centre of the digital envelope. AI-powered orchestrators should especially be developed for deployment at the edge, nearest to the final user, in a manner that can dynamically combine collaborative compute components into executable applications tailored around specific user needs. Orchestrators should hierarchically decompose user-given



“Digital enveloping” is the technology-enabled concept by which any item of reality – human, material or immaterial – can be associated with a computable digital representation capable of delegated, autonomous and coordinated action.

Delegation implies that a human authority mandates certain actions to be carried out, at least partially, within the digital realm.

The notion of **autonomy** suggests that the execution of the required action is carried out by independently acting, executable agents that operate within the remits of delegated authority.

The notion of **coordination** suggests that the required service may involve an ad hoc (not predefined) aggregation of finer-grained services, which have first to be identified from an analysis of requirements (“understanding the request”), then located and negotiated from publish-subscribe registries (“planning”), and finally called to execution by a per-request orchestrator (“orchestrated execution”). Interestingly, as noted in the “Artificial Intelligence” chapter (with reference to the operation flow of Google’s Aletheia, shown in Figure 1 in that chapter), planning should leverage “internal feedback loops”, with multiple preparatory alternatives being explored and one being chosen on evidence-based decision.

The envisioned digital envelopes may cross the boundaries between digital and physical spheres, as in real-time twinning. This capability requires digital envelopes to be equipped with (prosthetic) sensors, which pull digitalized input from designated sources located in the physical world or in other digital envelopes, and well as (prosthetic) actuators, which push computed outputs onto designated targets (physical things or other digital envelopes).

goals (where “user” may be a human, a business or administrative entity, another permissioned orchestrator) into a set of sub-orchestrations that cooperate to achieve the set goals. These orchestrators might be themselves generated by (federated) generative AI engines or retrieved from trustworthy repositories, both reachable from edge nodes and capable of collaboration with other orchestrators within dynamically federated zones.

Facilitate the creation of a continuum-friendly runtime infrastructure

The envisioned NCP, where digital envelopes will “float” while seeking to achieve user-set goals, requires a seamless runtime infrastructure for orchestrated (and “orchestrable”) services to run on. Such infrastructure should be comprised of ephemeral opportunistic federations of computing nodes as determined by component collaboration, on-demand

resource pooling or a combination of both. Deployments required in the continuum should not be constrained by node-specific characteristics and should support state-preserving live migration of running components following changes that occur in the physical environment twinned to the application.

Novel and emerging software technologies that can help implement this runtime infrastructure should be scouted. Eligible technology solutions should be open, amenable to sandboxed execution and transportable memory, and capable of near-native performance as well as of permissioned access to host-based services.

Explore technical and operational solutions that facilitate the deployment of personalized AI services at the edge

Moving AI services and embodiment to the edge of the continuum requires

solutions that do not subjugate the human user to the negative and perverse phenomena that so naturally arise in the digital space propelled by “attention economy” and “surveillance capitalism” drivers. The solution that the HiPEAC Vision 2026 promotes to that end is dubbed “personal AI” (PAI): enacting that element of the Vision requires exploring certain technical and operational innovations that are summarized in the following recommendation. Notably, the PAI is a direct emanation of aspects of the “Guardian Angel” that was first proposed in the HiPEAC Vision 2021.

Explore the implementation of personal AI (PAI) solutions that support “natural interface” interaction with the human user. Ensure trustworthy storage of personal data arising from such interactions. Provide user-friendly trustworthy intermediation in the acquisition and delivery of collaborative services.

Background: The drive for the NCP


When a technology infrastructure becomes invisible to the user and its sole apparent manifestation is the delivery of the service that it carries, that technology stops being a tool and becomes an “environment”, the space within which the user “lives”. This unquestionably happened with utilities which gained mass adoption in the twentieth century. Take, for instance, electricity: no one today is surprised when the pressing of a switch in a wall appliance delivers electric illumination, yet society would collapse were that not to happen.

In the first 20 years of the 21st century, the above-mentioned phenomenon of “technology-turning-environment” has occurred to networked computing, i.e. the delivery of computing services via always-on internet connectivity. In fact, the mode of delivery of computing-based services envisioned at the birth of the internet, in the late 1960s, well before there was any mention of “the cloud”, explicitly evoked “computer utilities” as analogous to “elec-

tric and telephone utilities” [1]. Since then, an increasingly vast range of services have been delivered to us users virtually (as in the making of a reservation, for example) or physically (as in arranging a parcel delivery, for example), and the technology infrastructure behind those services has

become an augmentation of the environment in which we live.

As part of this transition from “tool” to “environment”, web browsers – both visible (as applications to be explicitly launched) and invisible (as “default screens” that do



not need launching) – have increasingly become analogous to “light switches” for access to services. Web browsers are both the fastest evolving runtime infrastructure (i.e. host environment) for all computing devices that has appeared in the last decade, and the primary point of access (door more than window) into the digital space for us humans. Inevitably at the outset but even more evidently now, this “point of access” is expected to possess agentic AI capabilities, not only for “elevated” interaction with human users (who increasingly use “natural interfaces” away from keyboards and specialized programming languages), but also for site-to-site autonomous collaboration typical of the “everything-as-a-service” modality intrinsic of the cloud model [2]. And this capability is beginning to be regarded as an expected default, which materializes for example, in the Model Context Protocol that currently augments Chrome, Google’s browser.

Multiple modalities of use of twentieth-century utilities have been made possible by physical “end-terminals” designed and deployed to facilitate use. The same happens, in an infinitely more versatile and flexible manner, for computing services, both at the point of delivery, which can extend to the physical world via internet-of-things (IoT) solutions, and in their composition, which is becoming increasingly collaborative, federative, and decentralized. No single commercial provider will ever possess sufficient capacity to undertake self-contained end-to-end offering of all digital services that we humans may need or want. As long as the delivery infrastructure remains public, thus separate from vertical integration, which the European Union (EU) regulatory framework terms “unbundling”, there will always be room for multiple and diversified service providers.

The “next computing paradigm” (NCP), which has been the flagship of the HiPEAC Vision since 2024, is the space where innovative and value-added manners of collaborative networked computing and digital services can be conceived, constructed, and delivered. It is easy to observe that generative AI (genAI) is going to play a massively important role in that space: from the infrastructure level, where data collection and training occur, via the space where the development and execution of the algorithms that build learned models happen, to top-tier level where the services that leverage learned models and are delivered “environmentally” in an agentic manner, i.e. capable of autonomous coordinated agency in response to user prompts.

The natural integration between IoT and genAI will equally cause the digital and the physical worlds to converge seamlessly, with digitally controlled actuators (bound to become agentic prostheses) acting on the physical environment, whose context is digitally understood by prosthetic-agentic sensors, and the control loop is closed by human decision-making augmented by agentic genAI assistant prostheses. This is what we have called “digital envelopes”. In the digital-to-physical convergence evoked in this chapter, the physical world – the one where we humans live – continues to be “centre stage”, while the digital sphere, equipped with global or specialized “world models”, provides the augmentation needed to solve problems (faster) or accomplish (difficult or tedious) tasks.

The space resulting from this convergence is exceedingly vast and malleable, essentially because it spans all places where humans and “things” can be, with both equipped with networked computing devices. This space extends vertically

across layers in stacks of heterogeneous-but-connected¹ implementation technologies that span from actual basic computing to energizing external input-and-output interfaces, and, even more so, horizontally, above the infrastructure level², in various manners of collaborative federations. The EU, with its diversity of digital-market actors, small in terms of material resources but agile in terms of adaptivity and robust in terms of intelligence and creativity, has great potential for transformative impact in that space. Leveraging these opportunities requires departing from the “imitation game” in which we follow reactively what commercial giants indicate as direction, and move towards connective federations of computing resources, solutions and services.

The horizontal space evoked above is essentially that of the *compute continuum*, which previous editions of the HiPEAC Vision evoked as the space that originated the NCP.

Containerization in the cloud has shown how beneficial it is for applications and for the resource infrastructure to be able to support flexible and dynamic deployments³. Traditional containerization eliminated the need for individual applications to be tightly bound to immutable deployment solutions. This prompted the demand for opportunistic deployment decisions to be made that help optimize execution parameters “on the fly”. This in turn allows balance to be kept between the interests of the infrastructure provider and those of the application owner or user, which can diverge at times⁴ and currently tend to default in favour of the former. The cloud model to date has been characterized by nearly static deployment of (user-side) application components and obligated transfer of application data

1 The Internet of the origin acknowledged diversity and established network protocols that would “neutrally” connect heterogeneous nodes. Those protocols were very basic and even rudimentary when looked at today. Their primary concern was the technicality of transmitting data between sources and destinations. Equivalent protocols are needed today with much more sophisticated concerns, shifted from the functional to the non-functional domain, for “do this” to “do this with certain guarantees in time, space, privacy, locality, provenance, trustworthiness, energy, etc.”

2 Infrastructure in the modern digital space is analogous to the “operating system” of the olden days. It is software acting with the sole purpose of creating and maintaining a friendly, friction-free, resourceful environment for value-added applications to run. As the digital space that we are speaking of here is inherently distributed and decentralized, so must be the associated infrastructure, which arises from networked collaboration of dynamically varying sets of node-bound components.

3 The envisioned deployments become flexible and dynamic because they happen in response to changes in the environment they are “twinned” with, be that the physical space, where humans and things move about, or digital models, which adaptively change for size, scope and requirements during execution.

4 The default situation in the software-as-a-service delivery model is that the infrastructure provider decides where deployment happens, seeking to meet infrastructure-centric optimization objectives, within the freedom afforded by the service level agreement stipulated with the service provider.

towards the “centre” of the cloud, away from user control into the possession of cloud providers.

One of the principal traits of the compute continuum is its ability to reverse that trend, allowing applications to move close(r) to the data sources of their interest, ensuring more privacy, incurring less latency, and requiring less energy. This is favoured by breaking the application down into aggregates of smaller communicating components that can be deployed opportunistically and individually, and possibly even live-migrate subsequent changes in the physical reality operated on by the application or seeking to improve certain functional (e.g. numeric accuracy) or non-functional parameters (e.g. completion time).

The NCP consists essentially of an interconnected fabric of lightweight runtime environments, augmented by trustworthy agentic AI, which can host the sandboxed execution of (small and migratable) application components and allow them to communicate efficiently and interoperably among themselves as well as with the host node.

WebAssembly (Wasm) is an exemplary technology for the NCP in that it is the bytecode target of ahead-of-time or just-in-time compilation of higher-level programming languages, capable of near-native execution performance and no cold-boot syndrome, living in a sandboxed capsule provided for by a highly portable and lightweight runtime.

To enable the NCP to function correctly on WebAssembly technology, certain enhancements to its baseline ecosystem should be furthered to production-ready maturity past the prototyping stage that current research around them has reached to date:

1. nimble orchestration that allows opportunistic ephemeral federations of compute nodes to be created where Wasm modules can be deployed as and

when needed to carry out application-level actions⁵; and

2. transparent augmentations that allow Wasm executions to live migrate, that is without needing to restart, seeking certain functional or non-functional optimizations.

Web-level protocols are used to warrant communication interoperability among networked application components. In order to support the “modernity” of the NCP vision outlined in this chapter, those protocols need to be streamlined, bypassing unnecessary, general-purpose, layers which have accumulated over time, where low-level “general-purposeness” is no longer useful. A notable example of this streamlining is the replacement of TCP with QUIC under HTTP/3, but this is only the beginning of the required transformations. Furthermore, those application components might conceivably also be the on-the-fly product of genAI agents that operate as personalized agentic prostheses, and inherently as dynamic orchestrators, for the application owner or user.

A word of caution: Personal AI

While the vision depicted above is full of innovation promises, with genuine potential to improve people’s lives, we should

not be naïve and ignore the mounting risk of behavioural manipulation implicit in allowing private actors to reach out directly, privately and manipulatively to every single connected individual. A phenomenon known as the “attention economy and surveillance capitalism” [4] has emerged from the pervasive diffusion of connected social media, which commodifies human attention and behaviour, subjugating rational agency to market objectives. The result of this phenomenon is not only epistemic corruption – the distortion, degradation, or manipulation of knowledge and the processes by which knowledge is produced, shared, and validated – but also the erosion of autonomy, deliberation, and democracy. (Further discussion of the contamination of the information environment can be found in the chapter on cybersecurity in this volume.)

One response to this risk is “personal AI” (PAI), a technical architecture that fits very well at the user-end of the NCP vision, and caters for a decentralized, self-sovereign system of digital intermediation. The concept of PAI originates from [5] and is illustrated in Figure 1.

PAI envisions a decentralized, user-centric technical architecture, which matches well the NCP. PAI functions as an autonomous mediator acting solely on



⁵ An obvious, current example of opportunistic ephemeral federation arises from the notion of multi-agent orchestration, for example, as in [3].

behalf of the individual person, to reestablish communicative symmetry between them and people-oriented public institutions. This nature makes the institutional element of PAI “culture-centric”, and not necessarily single and universal. PAI intermediates exchanges between users and service providers. As noted earlier, in several respects, the PAI can be seen as an evolution of the “Guardian Angel” / “Digel” concept which was first outlined in the HiPEAC Vision 2021 [9].

One core functional principle of PAI, in its initial incarnation, is to rely on retrieval rather than on complex procedural generation or product of dynamic orchestration⁶. The key element is authoritative verification, which may translate into retrieving verified procedures from trusted registries

or composing them in a trustworthy supervised manner. While this is an alternative approach to the embrace of the generative approach as outlined in the introduction to this HiPEAC Vision, it can be considered as an initial, cautious step towards the full implementation of the NCP with dynamic orchestration, until effective approaches to add dynamicity in a trustable and reliable way become available.

These verified procedures and the associated fragment of personal data may be injected into the PAI instance at the user end as part of a live-migratable Wasm module. In that manner, sensitive data would be processed solely at the most appropriate and secure nodes in the continuum, without having to be permanently devolved to a centralized cloud

store. That would result in a data-inclusive, state-preserving live migration, which allows PAI to “follow” the user across the compute continuum as a mobile “capsule” of both intelligence and personal context, fulfilling the role of a truly personalized agentic prosthesis while maintaining strict data sovereignty. Further, the standardization of PAI functional modules across the continuum allows for a delta-transfer migration model. When target nodes are already equipped with these standardized Wasm procedures, the migration process is streamlined to include only the transfer of personal data fragments and the program’s execution state. This minimizes overhead and facilitates near-instantaneous mobility in response to changes in the service environment.





Figure 1: An outline of personalized AI

References

- [1] UCLA. “UCLA to be first station in nationwide computer network”. Press release, 1969. (Contains Leonard Kleinrock’s prediction about the rise of “computer utilities”) <https://www.lk.cs.ucla.edu/data/files/Kleinrock/UCLA%20to%20be%20First%20Station%20in%20Nationwide%20Computer.pdf>
- [2] Pillay, Tharin. What Is Moltbook, the AI Only Social Network Going Viral? TIME, Feb. 3, 2026. (Describes thousands of AI agents conversing, coordinating, and even inventing languages.)
- [3] CAMEL: Communicative Agents for “Mind” Exploration of Large Language Model Society Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, Bernard Ghanem, arXiv:2303.17760 [cs.AI]
- [4] Zuboff, Shoshana. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. PublicAffairs, 2019.
- [5] Hasida, Kôiti. Personal AI to Maximize the Value of Personal Data while Defending Human Rights and Democracy. in Johannes Glückler and Robert Panitz (eds.) Knowledge and Digital Technology, Springer, 239-256 (2024).
- [6] Berkeley Function-Calling Leaderboard. <https://gorilla.cs.berkeley.edu/leaderboard.html>
- [7] OpenAI Developers: Function calling. <https://developers.openai.com/api/docs/guides/function-calling>
- [8] Model Context Protocol: Specification. <https://modelcontextprotocol.io/specification/2025-03-26>
- [9] Duranton, Marc., & Vardanega, Tullio. (2021). “Guardian Angels” to protect and orchestrate cyber life. HiPEAC Vision 2021 <https://doi.org/10.5281/zenodo.4719375>

⁶ As trustworthiness in the direction of genuine service to the human individual is the central goal of PAI, retrieval from trusted and pre-screened repositories is preferred to procedural generation, including dynamic orchestration, as the latter are much more difficult to be checked for trust with present technology.

Artificial Intelligence



Recommendations for artificial intelligence (AI)

Rapidly position Europe in the two “directions” of AI¹:


Direction 1: AI for everyday life, for industries, for business, for customers, for everybody

Distributed AI, mainly for inference, enabling the vision of the next computing paradigm (NCP):

Continue to develop agentic AI, and specialized action models (SAMs)

As proposed in the HiPEAC Vision 2025, these SAMs (“small and specialized models that can interact with their environment”, unlike large language models (LLMs) that focus on language) should operate in a distributed infrastructure and an ecosystem should be created to support research, development and business around them. These models need to be refined, optimized, and reduced in size to improve efficiency. They can be optimized from more general foundation models by an ecosystem of companies providing their optimized SAMs in a marketplace so that they can be dynamically discovered and used by agentic orchestrators.

- Europe should provide **basic infrastructure to develop and fine tune the SAMs**, e.g. by providing a repository of trusted “foundation models” and the compute power to fine tune them on private and / or proprietary data.



Create the infrastructure and ecosystem to support agentic AI and SAMs

The infrastructure (A²I²: agentic artificial intelligence infrastructure, also simplified as (AI)²) should include non-functional properties, allowing the dynamic selection of the agent according to the user’s criteria. The key elements of the interoperable infrastructure are:

- **Software for managing agents** (the focus of the next big AI battleground), helping businesses handle the growing suite of AI agents they’re using from different providers [1]. **Europe should develop orchestrators** that allow the activation of services provided by agents and SAMs in response to a request (“prompt”) and the non-functional criteria associated with the request. The orchestrators should be trustable (for example, open source and

therefore auditable; they should work for the benefit of their user, and not for their creator or related third-party interests) and interoperable, therefore their specifications should be open and should lead to standardization. Features should be added over time as new technology becomes available, without necessitating a major change in the infrastructure (e.g. starting with static orchestration before moving to dynamic orchestration, beginning with a predetermined set of services before moving to a process to discover services, adding trustable ledgers that can qualify the services, ...)

- **A communication protocol between agents**, building on existing de facto standards of 2025 (such as MCP, A2A, etc) but also allowing support for non-functional properties, such as latency, time, etc. **Adding non-functional properties will be crucial to extend agentic**

¹ These two “directions” of AI are now more widely recognized. India hosted the AI Impact Summit in New Delhi in February 2026; rather than compete in building artificial general intelligence (AGI), policymakers and investors say India aims to become the world’s “adoption capital,” focusing on low-cost, local-language applications and “ultra-affordable” AI tools.

AI to physically entangled systems such as robots and self-driving cars. The (AI)² should be flexible in terms of its implementation and the location of the agents, from a lightweight protocol when everything is executed within a single computing resource to totally distributed execution on multiple of edge devices interconnected by “long” distance communication protocol with enforced security.

- **“Containerization” of agents and SAMs allowing migration (and, eventually, portability) in a secure way.** The agents and SAMs should work in well-protected “silos” on various kinds of hardware, preventing them from “contaminating” the system they are running on. A reference software infrastructure for using agents, SAMs and orchestrators should be publicly available and open.

- **Support the SAM containers with dedicated hardware.** Hardware should enforce the separation of the silos where agents are running from the underlying hardware. Dedicated accelerators should also be developed to increase the efficiency of the agents and SAMs (such as neural processing units (NPU)s, for example). The development of efficient accelerators for inference of SAMs is still complex, but more accessible than the development of complex, multipurpose graphics processing units (GPU)s or NPU)s (and tensor processing units – TPU)s designed to run large LLMs and targeting servers where several of these accelerators could be combined in a larger virtual one, using a unified memory space and extra-fast low latency communication interfaces. Unified memory between the central processing unit (CPU) and the accelerators in consumer devices (smartphones, laptops, ...) facilitates the use of SAMs at the edge. The dilemma of performance / specialized hardware versus flexibility / CPU-GPU can be alleviated by hard or soft reconfigurable accelerators, for example.

- An organization should be established to ensure the alignment of different contributions to the NCP ecosystem, similarly to the way that the RISC-V foundation, RISC-V International, aligns contributions to the RISC-V standard. The role of this organization should be, first, to help **define overall specifications for the ecosystem**; second, to **ensure that agents, SAMs and orchestrators conform to requirements and are trustable**. This organization should have a longer-term view of supporting the full NCP requirements and technologies. It could also act as a broker between providers and users, and kickstart the business ecosystem and business cases, taking institutional (cities, administrations, etc.) drivers and procurement as a starting point.

Develop “outside-the-box” concepts and ideas and test them rapidly in the (AI)² infrastructure

This can be facilitated by **sharing knowledge on emerging new ideas and new open source solutions** in a structured and open manner (knowledge infrastructure). Easy access to tools and resources to validate proposals should be provided, with an open mindset. Once tested, the new idea can use the supporting infrastructure to enable swift transition to a real product. The aim is to use the (AI)² as a means to rapidly create and validate specific new agents and to improve the orchestration and communication protocols, rather as Lego bricks can be put together in different combinations to build a large variety of objects. Example of concepts that can be developed include the following:

- **Decompose complex processes into smaller process (“bricks”) that can be realized (“constructed”) with the distributed SAM infrastructure** for smooth evolution and avoiding dependence on huge data centres.
- **Develop methodology and tools allowing a large model to be “broken up” into agents that are orchestrated**

together (i.e. implementing a large model by components in a (AI)² infrastructure). For example, a first concept could be to automatically transform a mixture of experts large language model (MoE LLM) into a set of specific agents that are orchestrated together in a distributed way.

- Focus on SAMs that can be used for **interaction with the physical world** (cyber-physical systems), thus keeping within real-world constraints like latency, timing, etc., i.e. interacting directly with the world, in a trusted manner, with a particular focus in providing instances of the (AI)² dedicated for robotics.

Develop neurosymbolic systems

These systems should marry classical systems (for supervision) and generative AI solutions, and add classical supervision systems as agents in the agentic AI infrastructure. Classical systems can offer services that can be combined with ones provided by SAMs.

Develop solutions for the continuous evolution of agents

Develop agents that can evolve according to their environment and adapt in a controlled way. This is still an open research topic, using either external databases in an agentic framework, new ideas for the structure of AI, or new training schemes [2].

Direction 2: Position Europe in the “artificial superintelligence” race

Some actors in the field of AI aim to develop systems that could be of higher “intelligence” than human experts to solve very complex problems. This involves large efforts (comparable to the Manhattan Project) and extreme computing power as of the current state of the art. Such systems, as of today, are self-improving: any progress done makes the next iteration faster and better, which

means any delay potentially becomes exponentially costly to catch up. This is a real race between companies and countries involving enormous investments.

- **Europe should decide quickly if it wants to enter the race.** It is first a political decision.

If so, Europe can leverage its position in global processing power by developing and using a transnational (AI)² infrastructure (in which each agent (“component”) could be a complete high-performance computing (HPC) centre),

using a unified infrastructure arising from NCP concepts and intelligently harnessing existing decentralized European hardware resources and extending the AI Factories initiative into a unified computing continuum.

- This is an action that calls for **creating a physical centre like CERN** to bring together competences within Europe in the field of AI and coordinating actions, for example extending the RAISE initiative [3]. This should also help to assess the state of the art and the innovations needed.

- While developing advanced models, **emphasis should be on the alignment problem:** developing approaches that control the level of model “cheating”, avoiding deceptive behaviours and the development of internal goals that are not “aligned” with the goals of the system designers. This is also an opportunity to develop advanced models that reflect values and culture endorsed by European society.

Introduction

Cut-off date for this rationale: beginning of March, 2026

The rapid evolution of AI the recent months is illustrated by the Andrej Karpathy quote [4] featured in the introduction to this edition of the HiPEAC Vision. Similar constataions are shared by others, such as [5].

Progress in AI is advancing at an increasing rate. Before 2024, the race mainly concerned the size of the models: bigger was supposed to be always better. New techniques were created to improve model training, drawing on well-curated data or on synthesized data, allowing the cost of training to be reduced and therefore the size of the model to be increased. The “mixture of experts” (MoE) approach became widespread to reduce the cost of computing during inference. 2025 was the year of the **“thinking models”**: using more compute power at inference time resulted in drastic increases in model performance.

At the time of writing, in 2026, we are witnessing the **rise of the orchestrated and agentic models**, where models can explore options in parallel, discard wrong directions, refine the results, and interact through tools with the physical world or computers. An example is Google Aletheia (see Figure 1), which improves upon the

results of Google Deep Think (a “thinking” model) without changing the model itself. Figure 2 shows not only the gain of using orchestration, but also the remarkable gain in compute power for the same level of performance in six months: nearly 2^7 , or $\times 128$.

Aletheia is an illustration that the agent wrapper (the “orchestrator”) with “generate verify revise loop” is going to be more important than just throwing raw compute at the base model. In 29 out of 30 problems where Aletheia returned a solution, its conditional accuracy was 98.3%. Its success rate was about 6.5% on research-grade problems, which seems low, but is still impressive according to the complexity of the problems. This demonstrates that

the agent layer is where the real capability gains will be made, not just the base model itself. This leads to the observation that **the next big AI battleground will be centred on software for managing agents**, helping businesses handle the growing suite of AI agents they’re using from different providers.

(For a more in-depth discussion of “thinking models”, agentic and autonomous models, see “Thinking models, agentic AI and increasing autonomy”, on p.31.)

The progress is not only in the models, but also in the **hardware that allows the models to be executed**. For example, “NVIDIA GB300 NVL72 systems now deliver up to 50x higher throughput per mega-




Figure 1: The workflow of Google Aletheia from [6]

watt, resulting in 35x lower cost per token compared with the NVIDIA Hopper platform.” [7]. The H200 was launched in 2024 and the GB300 by end of 2025, showing the rapid improvement of performance of the

architecture, technology and software stack (it should be noted that this gain is also due to the move from an 8-bit floating-point representation to a mixed 4-bit floating-point representation).

The constant improvement of the (energy) efficiency of the hardware / software stack is the fuel that allow companies to deliver higher-performing models demanding more and more compute power at a similar cost, and to serve more and more users. By July 2025, ChatGPT had been adopted by around 10% of the world’s adult population [8]. We can observe a correlation between the release of new versions of ChatGPT and the availability of a new generation of GPU, with improvement in their efficiency (see Figure 4). This is an economic reason: the cost of training with “older” generation GPUs will be too high in terms of energy (and therefore in \$).

This tight connection between hardware efficiency and the availability of new LLMs is bidirectional – the hardware should be flexible enough to cope with the new innovations in the domain of AI, especially for training; often, when a new hardware is available, it was specified for algorithms that are one or two generations older, therefore a certain flexibility (programmability) is required, to the expense of less efficiency. NVIDIA’s rapid performance gains and continued dominance serving AI workloads are particularly impressive given the difference between the timescales for AI model advances (which take months) and the timescales needed for hardware development (which are currently in years). This was explained, for example, by keynote speaker Deming Chen at HiPEAC 2026 [9]. Inference seems to have reached relative maturity, so we see more specialized solutions for efficient inferences appearing, such as Groq, Cerebras, etc. All this shows the importance of co-design in hardware / model development, which NVIDIA calls “extreme code-sign” and which has been promoted in multiple editions of the HiPEAC Vision.

Only a few companies are mastering the progress of AI; most of them are in the United States (US) or China, with hardly any in Europe. Most other companies are simply “users” of AI technologies.

It can be observed that the open source community is playing an increasing role, driven by China and its “open weights”




Figure 2: As of January 2026, the latest advanced version of Deep Think had significantly outperformed the International Mathematical Olympiad (IMO) Gold version (July 2025) on Olympiad-level problems. Aletheia makes further leaps in terms of reasoning quality with lower inference-time compute. All results were graded by human experts. The figure also shows that for the same levels of performance reached in July 2025, the January 2026 version required nearly x100 less inference-time compute [6].




Figure 3: Performance improvement of NVIDIA’s latest generation of GPU over the previous one [7]

models and some US companies (with European background) like Hugging Face. However, due to the increasing costs of developing new models and stakeholder pressure, companies that previously delivered open source models have in some cases transferred to closed models (e.g. Meta, perhaps Alibaba / Qwen after the departure of key people [11]). Further discussion of open source models can be found in the chapter on open source.

The recent success of OpenClaw [12], initially developed by one man (the

Austrian programmer Peter Steinberger), shows that it is still possible for a small team to develop successful AI technology without a large company. A self-hosted autonomous AI agent platform, OpenClaw runs locally and performs actions across apps and services on a user's behalf. It can act continuously without user supervision. It is positioned as infrastructure rather than an app and its development has been unusually rapid and chaotic, with multiple renaming (it was first called Clawdbot, followed by Moltbot, before finally being named OpenClaw), viral growth, secu-

rity controversies, and institutionalization within months. It was the one of the fastest growing GitHub repositories. Finally, in mid-February 2026, Peter Steinberger was hired by OpenAI. In an announcement, OpenAI said that the open source AI tool would “live in a foundation” inside the company”, and the company's chief executive Sam Altman wrote: “We expect this will quickly become core to our product offerings,” [13]. This is also an illustration that the world of AI is controlled by few powerful companies – powerful enough to absorb any interesting innovation.




Figure 4: Graph originally from the GTC 2023 keynote by NVIDIA CEO Jensen Huang, updated with GPT-5




Figure 5: Slide promoting a global codesign, included for example in the HiPEAC Vision 2023 – see [10]

In addition to power over the development of AI models being concentrated in a few companies, access to the most advanced AI model capabilities is being restricted to an elite of paying customers for example, Google Gemini 3 Deep Think is only available to Google AI Ultra subscribers (at a cost of about €275 per month). This is an example of the segregation that we see appearing between “normal users” using AI for mundane tasks, who are offered models with limited performance which are cheaper to run, and “supermodels” that eventually could only be limited to use by an elite (and we have not yet seen “super-intelligent” models, very expensive to run, which we believe will have very restricted access).

The current access to AI for “normal” users is mainly through apps (and therefore the cloud), but smartphone makers (Apple, Google, Samsung, ...) are slowly changing the game by **allowing more and more functions to be done by local models running on the user's smartphone or laptop**. The capacity of small models is also rapidly augmenting [14] local capabilities: models that can run on smartphones can now be “thinking” models and they can use tools.

On laptops or personal computers, unified memory allowing the CPU and the GPU to share memory allows de facto bigger models to run on these machines, no longer limited by the amount of video random-access memory (VRAM) available on the GPU card. This is the architecture of Apple machines since the M series of processors,

Date	Event	Significance
November 2025	Clawdbot released	First public version
Late 2025-Jan 2026	Viral growth	Massive adoption: a “weekend hack” that unexpectedly became infrastructure.
Jan 2026	Renamed Moltbot	Trademark conflict: Anthropic raised concerns about brand confusion with “Claude.” But this opened the door for impersonation scams and malware distributions
Late Jan 2026	Renamed OpenClaw	Stable identity
Jan-Feb 2026	Moltbook ecosystem	Agent social network: Entrepreneur Matt Schlicht launched Moltbook, a social network for AI agents. OpenClaw agents could interact autonomously on it. This marked the transition from a tool to an ecosystem phenomenon.
Jan-Feb 2026	ClawHub skills	Extension platform: Skills - modular capabilities for agents. Distributed through ClawHub. Essentially executable extensions with broad system access. This dramatically increased functionality but also risk.
Jan-Feb 2026	Security incidents	Major controversies: security researchers discovered: • Hundreds of malicious skills (trojans, keyloggers, backdoors) • Lack of initial moderation • Skills could execute arbitrary code OpenClaw later partnered with security scanners to vet submissions.
Feb 2026	Founder joins OpenAI	OpenClaw began appearing in corporate workflows: • Used by companies in Silicon Valley and China • Integrated with workplace platforms • Treated as an “AI employee” in some demos Its ability to operate continuously made it appealing for automation. Institutional backing: Peter Steinberger joined OpenAI. OpenClaw moved toward governance by an open source foundation with a commitment to remain open source The project had already reached massive scale: • Over 100k GitHub stars • Millions of visitors • Major media attention
Mid-Feb. 2026	Anthropic and Google bans OAuth tokens in third-party tools like OpenClaw.	A lot of users are moving away from using Anthropic’s Claude AI and Google Gemini as engine for their (Openclaw) bot.

Figure 6: A brief history of OpenClaw

and it is why Macs are often preferred to develop local AI (on GitHub repositories, AI workloads still mainly use the NVIDIA CUDA software stack, but more and more are using the MLX stack of Apple silicon). Machines with 128 GB are now accessible by the public (until the increased cost of memory will limit this) – for instance, machines using the M5 Max from Apple, the GB10 from NVIDIA or the AMD Ryzen AI Max+ 395 processor. These machines can run quite large models locally. On smartphones, models like Qwen Qwen3.5-0.8B or 2B or 4B, Gemma-3n-E2B can run locally on modern smartphones with good levels of performance. The consumer hardware is now ready for the NCP allowing SAMs to run locally. This trend towards local execution of AI models continues a tendency we have noted and advocated for since at least the HiPEAC Vision 2023. SAMs for




Figure 7: Increase in the number of GitHub stars for the OpenClaw GitHub

use in the NCP can now be run locally on widely available consumer hardware, albeit hardware that is currently designed and produced outside Europe.

This new AI era could be an opportunity for European companies to be “back in the game” concerning consumer hardware: smartphones, laptops and computers might not be the best form factor for this new AI era. Will it be an earbud like in the movie *Her*? There are a lot of tests of new devices, such as AI pins, smart glasses, etc. Most of them are unsuccessful for now, but OpenAI, Meta, Apple (?) seems to be working on defining what could be the “post smartphone” device. Will Europe sit still and wait?

Recent directions in AI development and their relation to the HIPEAC Vision

Thinking models, agentic AI and increasing autonomy

In 2025, the field of artificial intelligence saw remarkable advancements, particularly in LLMs and their applications. Current

advances in artificial intelligence are occurring at an accelerating pace and are widely underestimated by the general public and often incompatible with the timescales currently required to develop e.g. Horizon Europe project calls, and the three-year schedule of most EU-sponsored projects, where their aim might be outdated before the end of the project.

Modern AI systems can now perform complex technical tasks end-to-end with minimal human intervention. In software development, for example, an application can be generated from a natural-language description, tested automatically, refined through iterative self-evaluation, and delivered in operational form. As an example, Figure 8 shows the increase in complexity of software tasks that LLMs can perform with a 50% chance of success (by February 2026): we can observe a drastic improvement since the end of 2025. These systems are increasingly capable of making design choices and adjustments autonomously, reducing the need for expert supervision. As an example, Spotify announced in early 2026 that “its best developers [hadn’t] written a line of code since December, thanks to AI” [15].

Even the benchmarks that are specially designed to be hard for AI are becoming rapidly obsolete, with LLMs reaching nearly 100% on some of them (Figure 9 shows the progress on ARC-AGI benchmarks, which were defined to be especially difficult for LLMs).

A strategic focus on using LLMs for programming has accelerated this transformation because software underlies AI development itself. Once systems became proficient at writing code, they could assist in building improved versions of future models. For example, OpenAI wrote “GPT-5.3-Codex is our first model that played a key role in its own creation. The Codex team used early versions to debug its own training, manage its own deployment, and diagnose test results and evaluations – we were impressed by how quickly Codex was able to accelerate its development.” [18]. This creates a self-reinforcing cycle in which each generation of AI helps produce the next, potentially leading to very rapid increases in capability.

Recent developments (for example in Google’s Gemini 3 ecosystem) illustrate a shift in artificial intelligence progress from




Figure 8: Example of the increase of complexity of software task that recent LLM can perform with 50% success [16]




Figure 9: The ARC-AGI benchmark was designed to be very difficult for LLMs. At the time of writing, the ARC-AGI1 was nearly obsolete, and ARC-AGI 2 seemed destined to be so in a few months [17]

simple model scaling to more sophisticated reasoning architectures and agent-based systems. An update to the “Deep Think” capability introduces a reasoning mode that allocates additional computation at inference time, allowing the model to elaborate longer before producing an answer. Instead of following a single linear chain of reasoning, the system explores multiple hypotheses in parallel, evaluates alternatives, backtracks from dead ends, and dynamically adjusts the number of reasoning steps according to problem difficulty. This approach enables substantial performance gains without modifying the underlying model parameters. Benchmark results indicate high performance across complex tasks, including abstract reasoning, competitive programming, and advanced scientific problem solving.

In the case of Gemini “Deep Think”, the computational cost per task has decreased significantly compared with earlier versions, showing that efficiency improvements can accompany capability gains [6]. Experimental data suggest that improvements in inference-time reasoning

can reduce the computational resources required for expert-level performance by orders of magnitude, highlighting a **new scaling paradigm based on smarter use of compute rather than larger models.**

Alongside improvement using inference-time reasoning, **research agents built on top of this reasoning framework** demonstrate how **orchestration around a model can further enhance performance.** For example, Google’s Aletheia operates through a loop composed of generation, verification, and revision stages. It proposes candidate solutions, critically evaluates them for logical flaws or unsupported assumptions, and refines or restarts the process when necessary. By grounding its reasoning in external sources such as academic literature and by acknowledging when a problem cannot be solved, the system reduces hallucinations and improves reliability. In controlled evaluations, this agent achieved very high accuracy on advanced proof tasks, outperforming approaches that rely solely on increased computational [5].

Case studies in scientific and mathematical research show that such systems can really assist in real problems rather than only in standardized benchmarks. However, success rates still remain limited, and results do not indicate consistent breakthroughs. Evaluations of hundreds of research problems show that only a small fraction is solved correctly, though this still represents a significant improvement over earlier systems.

We can identify that advances in AI in 2025 highlight three major trends:

- 1) **Allocating more reasoning time at inference can dramatically increase performance while improving efficiency.**
- 2) **The architecture surrounding a model, including tools and agent loops, can contribute more to capability gains than enlarging the base model itself.**
- 3) **AI systems are beginning to function as research assistants capable of contributing to complex problem solving, although they remain far from consistently autonomous scientific discovery.**

Orchestrating autonomous software agents

These developments suggest a transition from single-shot language models toward structured, tool-using agents that reason, verify, and iterate. Artificial intelligence development is shifting towards the **management and orchestration of autonomous software agents**, a concretization of the concept of “Guardian Angels” or “Digels” dating back to the HiPEAC Vision 2021 [19]. Businesses are increasingly deploying multiple AI agents from different providers, creating demand for platforms that can coordinate these systems, control their activities, and integrate them into existing workflows. This is exactly what the HiPEAC Vision 2025 proposed: building an ecosystem of distributed agents as a blueprint of the NCP (and therefore implementing some concepts promoted by the NCP).

Competition has intensified following the **introduction of agent-management products designed to coordinate tasks**

across various business applications through a single interface (see the success of OpenClaw, for example, but also the products released by Anthropic, Abacus, Manus (now Meta), ...). One system can direct multiple agents to perform work across different software platforms.

Broader competition is emerging over which platform will become the central dashboard for supervising AI agents. The hiring of the creator of OpenClaw by OpenAI is a clear signal. Technology companies offering cloud infrastructure, business software, and productivity tools are each promoting their own control environments, anticipating that organizations will standardize on a single system. Ownership of this control layer could determine long-term influence over enterprise software ecosystems. On the consumer side, this could transform the current multi-app ecosystems running on smartphones into a single app that will be the entry point for all functions that the user will want. As there will be only one app, the form factor of smartphone with a touch screen

allowing to access multiple apps could evolve into a new device that will perform the same functions and more but without this access to numerous diversified apps. Several companies are exploring what form this device could take (OpenAI with Jony Ive [20], Meta, Apple...).

Investors are particularly attentive to the possibility that AI agents could replace significant portions of traditional software functionality, reducing demand for existing products while accelerating disruption across the software industry [21].

Overall, the emergence of agent-management platforms represents **a new phase in computing, where the focus shifts from individual applications to coordinated networks of autonomous systems**. Success in this environment is likely to depend on standardization, integration with business processes, data governance, and operational oversight, as well as the ability to navigate an increasingly competitive and volatile market landscape.

Designing hardware for AI

Specialization versus more general-purpose accelerators: The case of NVIDIA and Groq

Due to the increasing number of users, the heavy computing demand of “thinking models” and of agentic AI, the inference part is becoming orders of magnitude more demanding in computing power than training. Therefore, the competition to have more efficient systems for inference (which will result in lower utilization costs) is heating up. Major companies are now developing their chips, with often a variant for training only (such as AWS, which is developing the “Inferentia” chip for inference only, besides its “Trainium” for training). GPUs can embed features that make them more efficient for inference, but they will never be as efficient as architectures specially designed for inference. This could explain why NVIDIA made a “non-exclusive deal” with Groq to use their infer-

ence technology (and hired most of their team).

HiPEAC had spotted Groq’s relevance in 2024. For more detail on Groq’s first chip generation, readers can refer to the Semi-Analysis report published at that time [22].

NVIDIA’s move is entirely logical. A well-known principle in compute architecture is that the more you specialize an architecture, the more efficient it becomes. As a rough rule of thumb, you can see about an order-of-magnitude improvement when moving from CPUs to GPUs, another order of magnitude from GPUs to reconfigurable dataflow-style architectures such as coarse-grained reconfigurable arrays (CGRAs) or field-programmable gate arrays (FPGAs), and another order of magnitude from FPGAs to application-specific integrated circuits (ASICs). This compounding effect is why, for a fixed task, we can end up near a thousandfold efficiency gap between a highly specialized ASIC like CEA’s Neuro-

Corgi [AI41] and a general-purpose GPU. (For further discussion of reconfigurable computing platforms and how these fit into the HiPEAC Vision, see the hardware chapter.)

The trade-off is that specialization reduces generality. AI algorithms evolve quickly, so you need a certain amount of flexibility – programmability and reconfigurability – to keep pace. GPUs are more flexible than NPUs and therefore tend to adapt faster to algorithmic innovation, but they rely on a substantial software environment to bridge the gap between new algorithms and the hardware. A more “hard-wired” solution has less flexibility and therefore does not require such a generic and complex software stack. In practice, it often needs more of a “mapper” than a full compiler. However, the downside is clear: algorithms that differ significantly from what the ASIC was designed for are either not implementable at all, or they run with a severe efficiency penalty. This limits the




Figure 10: Architecture of Groq chip [24]

specialization for rapidly evolving tasks like training, which illustrates why some level of flexibility remains necessary.

What has changed recently is that there is now more stability in algorithms, at least on the inference side. Inference still relies largely on transformer-based models (see “The rise of the transformers”, p.35), with some variants such as Mamba (see “Continuous learning and context awareness”, on p.39) and mixture of experts (MoE) (see “Sub-giga-parameter models”, p.36).

At the same time, demand for inference compute has risen sharply, often by a factor of roughly ten to fifty compared with training workloads, depending on the application and how you measure it. Under these conditions, less flexible solutions than GPUs can generate a return on investment before they become obsolete, which is precisely where Groq becomes attractive. Specialized chips for inference are also easier to realize, allowing more companies to enter the game, unlike for training, where the complexity is much higher and even large groups have to give up [25].

For a less flexible approach for inference, a dataflow architecture like Groq’s is a natural fit. Groq also made two notable first-

generation bets. First, they did not rely on the most cutting-edge process technology. Second, they leaned into “near-memory” computing by integrating memory on the chip, in order to avoid the very large energy cost of moving data off-chip, which can dominate overall power consumption. This reduces the need for external memory and can be an advantage during this current period of memory scarcity driven by AI demand (the so-called “RAM apocalypse”). The drawback is that on-chip memory capacity is limited per device, so scaling to even a “mid-sized” LLM requires very large chip counts. According to the analysis referenced above [22], running the Mixtral MoE language model requires computing power in the order of hundreds of chips – 576 are mentioned – which effectively confines this approach to data-centre deployments.

Groq also optimized heavily for low latency rather than maximizing throughput in the way GPUs typically do. Low latency is particularly valuable for real-time and interactive use cases, for agentic AI, and for systems that perform iterative inference steps, including so-called “reasoning” models. In that sense, the design aligns well with the dominant patterns in current inference demand.

Another practical point is complexity. Groq’s architecture is significantly less complex than NVIDIA’s GPU architectures, which can translate into smaller development teams on both the hardware and software sides. With NVIDIA’s cash reserves, acquiring intellectual property (IP) and integrating it can be faster than building everything organically, and this approach may also help them avoid the regulatory friction they encountered with their attempted Arm acquisition, which was ultimately abandoned due to competition concerns. There is also the talent angle: they would be bringing in Groq’s Jonathan Ross, a key architect from Google’s original TPU effort.

The next step will be to see what NVIDIA actually ships and how they position it. The expectation is that they will aim to support future inference-oriented chips through CUDA to keep a coherent ecosystem, even if the supported functionality is more constrained than on general-purpose GPUs. They may also evolve the design onto more advanced process nodes over time. Finally, because NVIDIA increasingly sells complete systems rather than only chips, they are in a strong position to build multi-chip inference systems at the scale Groq requires, potentially integrating hundreds of chips in a single data centre-grade solution.

A reminder of recent AI history: The rise of the transformers

The current “boom” of artificial intelligence, and more precisely of the generative AI can be traced back in 2017 with the paper “Attention Is All You Need” [26] by researchers at Google. Its main goal was to accelerate the training of recursive neural networks that were used for language processing in introducing more parallelism. Unlike machine vision, which was already booming due to the rebirth of neural network-based approaches (using convolutional neural networks, or CNNs) and is inherently parallel, sequential processes like text, speech, translation, etc. required a different structure than CNN, and this paper showed one way to induce more parallelism. The side effect was to unlock a complete domain.

“We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely.

Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.”

Google didn’t seem to realize the impact of its paper at that time. A startup, OpenAI, did see the potential, and started to develop more and more complex models using “transformers”. This is the series of GPT (GPT stands for generative pre-trained transformer) that led to a series of models leading to GPT 3.5, the basis for the famous ChatGPT that shocked (and changed) the world when introduced in November 2022.




Figure 11: The paper from Google that triggered the current AI boom

Model	Architecture	Parameter count	Training data	Release date	Training cost
GPT-1	12-level, 12-headed Transformer decoder (no encoder), followed by linear-softmax.	117 million	BookCorpus: 4.5 GB of text, from 7000 unpublished books of various genres.	June 11, 2018	“1 month on 8 GPUs”, or 1.7e19 FLOP.
GPT-2	GPT-1, but with modified normalization	1.5 billion	WebText: 40 GB of text, 8 million documents, from 45 million webpages upvoted on Reddit.	February 14, 2019 (initial/limited version) and November 5, 2019 (full version)	“tens of petaflop/s-day”, or 1.5e21 FLOP.
GPT-3	GPT-2, but with modification to allow larger scaling	175 billion	499 Billion tokens consisting of CommonCrawl (570 GB), WebText, English Wikipedia, and two books corpora (Books1 and Books2)	May 28, 2020	3640 petaflop/s-day, or 3.2e23 FLOP.
GPT-3.5	Undisclosed	175 billion	Undisclosed	March 15, 2022	Undisclosed
ChatGPT	Undisclosed	? (rumor 20B???)		November 20, 2022	
GPT-4	Also trained with both text prediction and RLHF; accepts both text and images as input. Further details are not public.	Undisclosed (1.8 trillion aka 1.8e12)	Undisclosed (13 trillion tokens, aka 1.3e13)	March 14, 2023	Undisclosed Estimated 2.1e25 FLOP.

From https://en.wikipedia.org/wiki/Generative_pre-trained_transformer

Figure 12: Evolution of models from OpenAI. The details of models after GPT-3 are not publicly known, some leaks exist for GPT-4, but none about the real size of GPT-5

Smaller and smaller models

While frontier models continue to increase in size (see Figure 12), progress in algorithms, in training and datasets allow smaller models to have similar levels of performance as those of larger models from a few months before, at least in more specialized domains. For example, current models of about 10 billion (10B) parameters demonstrate better performance on specific tasks (benchmarks) than the original ChatGPT of 2022 (see Figure 14) and also [14]).

This trend was predicted in the HiPEAC Vision 2023 [10] (p.88) which called for Europe to “reduce both the memory footprint and the computing power of algorithms for artificial intelligence required for inference and learning”.

What is also interesting is that smaller models are often “open weight”, meaning that they can be easily duplicated (following their respective licensing agreement) and that open source (open weight) models are

generally catching up closed models within the space of a few months.

Smaller models can also be “trained” by synthetic data generated by bigger models, as exemplified by the various DeepSeek small models, based on existing small models (often Qwen models) but fine-tuned using data generated by the largest DeepSeek model:

“We demonstrate that the reasoning patterns of larger models can be distilled into smaller models, resulting in better performance compared to the reasoning patterns discovered through RL on small models.

Using the reasoning data generated by DeepSeek-R1, we fine-tuned several dense models that are widely used in the research community. The evaluation results (see Figure 15) demonstrate that the distilled smaller dense models perform exceptionally well on benchmarks. We open-source distilled 1.5B, 7B, 8B, 14B, 32B, and 70B checkpoints based on Qwen2.5 and Llama3 series to the community.” [27].

Sub-giga-parameter models

More generally, this seems to show that **for generalization on a particular task, a small subnetwork is sufficient**: this is the confirmation of the lottery ticket hypothesis [28] of the “mixture of experts” (MoE) approach, which selects the subnetwork, and especially of the agentic approach, which can be seen as a decentralization of the MoE approach. An MoE is a model architecture where many specialized “expert” networks run in parallel, and a gating network learns to select which experts should handle each input. Instead of one big model doing everything, MoE routes different tasks or tokens to the experts best suited for them, making the system both scalable and efficient. The link to the agentic approach is the following: agentic systems also rely on specialized components (tools, skills, agents) that are dynamically invoked depending on the situation. Like an MoE gate, an agentic orchestrator decides which agent to activate for each subtask. In essence, MoE is the neural-network analogue of an agentic

Model name	Announced	MMLU Pro3*
GPT-5 (High)	August 2025	0.87
GPT-4o	May 2024	0.73
Seed-oss-36B	August 2025	0.83
Qwen3-30B	August 2025	0.81
Phi-4-14B	December 2024	0.76
GPT-oss-20B	August 2025	0.74
Gemma-3-12B	August 2025	0.61

Figure 13: Performance of various models on the same benchmarks [26]

Benchmarking LFM2-2.6B-Exp vs. GPT-4

Benchmark	LFM2-2.6B-Exp (Liquid AI)	GPT-4 (Original)
IFEval (Strict Instruction Following)	88.13	79.0
Multi-IF (Complex Instructions)	73.05	65.0
IFBench (Formatting Accuracy)	44.40	24.0
GPQA (PhD-level Reasoning)	41.31	35.7
MMLU-Pro (Hard Expert Knowledge)	53.87	63.7
AIME25 (Olympiad-level Math)	22.67	15.0

Figure 14: Comparison of performance of a 2.6B parameter model of 2025 with the “original” ChatGPT

Model	AIME 2024 pass@1	AIME 2024 cons@64	MATH-500 pass@1	GPQA Diamond pass@1	LiveCodeBench pass@1	CodeForces rating
GPT-4o-0513	9.3	13.4	74.6	49.9	32.9	759
Claude-3.5-Sonnet-1022	16.0	26.7	78.3	65.0	38.9	717
DeepSeek-RI-Distill-Qwen-7B	55.5	83.3	92.8	49.1	37.6	1189
DeepSeek-RI-Distill-Qwen-4B	69.7	80.0	93.9	59.1	53.1	1481

Figure 15: Performance of fine-tuned models [27]. Note : The American Invitational Mathematics Examination (AIME) is a selective and prestigious 15-question 3-hour test given since 1983 to those who rank in the top 2.5% on the AMC 10

system: both achieve greater capability by routing problems to the right specialists, rather than relying on a single monolithic generalist. As explained in the HiPEAC Vision 2025, this mechanism is very similar to the proposed structure of the NCP.

A number of “small” models were released in 2025: along the lines of small models, Gemma3-270M [29] (here M means indeed million) runs well on a phone and has surprising levels of performance for its size (like GPT-3 a few years ago). Moreover, there is a profusion of small models under

1B parameters: Qwen 600M, Hunyuan-0.5b-instruct, Ernie-4.5-0.3b-pt [30], the last three being from the Chinese BATIX (Alibaba, Tencent, Baidu), but also MobileLLM [31] from Meta show good levels of performance for the 350M parameter model. The smallest is 125M parameters. Qwen from Alibaba also deliver small models (even multimodal) that can easily compete with larger models [14].

In June 2025, Microsoft presented MU [33], which will be in Copilot+ PCs and which has 330M. We are truly entering the era of small models running at the edge.

In the line of specialized nano models, one can also look at Kitten, which is a text-to-speech model with 25M parameters [34].

There are also other optimizations that will enable LLMs to run on memory-constrained devices: new memory optimization in the multimodal open source model Gemma3n from Google (A 7B parameters model running on 4GB of RAM) – see Figure 19.




Figure 16: Performance of the MobileLLM from Meta. MobileLLM-R1 from Meta shows good levels of performance for the 360M parameter model [32]



Figure 17: Examples of test of the 140M and 360M models from the MobileLLM series

Ultra-small AI with interesting impact

June 2025 witnessed the launch of a model with a very interesting performance / size ratio of 27M (yes, millions not billions) parameters performs better than Claude Opus 4 on the (very specialized) ARC AGI 1 benchmark. The results on the ARC AGI 1 benchmark have been confirmed by the ARC Prize team [36]. This new model is called HRM (Hierarchical Reasoning Model [37]).

However, on 6 October 2025, a new smaller model showed better results than the HRM model: with only 7M parameters, the TRM (Tiny Recursive Model [38]) obtained 45% test-accuracy on ARC-AGI-1 and 8% on ARC-AGI-2, higher than most LLMs (e.g., DeepSeek R1, o3-mini, Gemini 2.5 Pro) with fewer than 0.01% of the parameters.

The TRM model can be used to solve ARC-AGI puzzles, or to solve hard Sudoku or to navigate into mazes (see Figure 21 and Figure 22). Further details can be found in the paper and in a Github repository [39]. These models are small in size, but not necessarily in computing: they iterate a large number of times before providing their result.

Table 3: Zero-shot performance on Common Sense Reasoning tasks. MobileLLM denotes the proposed baseline model and MobileLLM-LS is integrated with layer sharing with the #layer counting layers with distinct weights.

Model	#Layers	#Params	ARC-e	ARC-c	BoolQ	PIQA	SIQA	HellaSwag	OBQA	WinoGrande	Avg.
Cerebras-GPT-111M	10	111M	35.8	20.2	62.0	58.0	39.8	26.7	29.0	48.8	40.0
LaMini-GPT-124M	12	124M	43.6	26.0	51.8	62.7	42.1	30.2	29.6	49.2	41.9
Galactica-125M	12	125M	44.0	26.2	54.9	55.4	38.9	29.6	28.2	49.6	40.9
OPT-125M	12	125M	41.3	25.2	57.5	62.0	41.9	31.1	31.2	50.8	42.6
GPT-neo-125M	12	125M	40.7	24.8	61.3	62.5	41.9	29.7	31.6	50.7	42.9
Pynthia-160M	12	162M	40.0	25.3	59.5	62.0	41.5	29.9	31.2	50.9	42.5
RWKV-169M	12	169M	42.5	25.3	59.1	63.9	40.7	31.9	33.8	51.5	43.6
MobileLLM-125M	30	125M	43.9	27.1	60.2	65.3	42.4	38.9	39.5	53.1	46.3
MobileLLM-LS-125M	30	125M	45.8	27.7	60.2	65.7	42.9	39.5	41.1	52.1	47.0
Cerebras-GPT-256M	14	256M	37.9	23.2	60.3	61.4	40.6	28.3	31.8	50.5	41.8
OPT-350M	24	331M	41.9	25.7	54.0	64.8	42.6	36.2	33.3	52.4	43.9
Pynthia-410M	24	405M	47.1	30.3	55.3	67.2	43.1	40.1	36.2	53.4	46.6
RWKV-430M	24	430M	48.9	32.0	53.4	68.1	43.6	40.6	37.8	51.6	47.0
BLOOM-560M	24	559M	43.7	27.5	53.7	65.1	42.5	36.5	32.6	52.2	44.2
Cerebras-GPT-590M	18	590M	42.6	24.9	57.7	62.8	40.9	32.0	33.2	49.7	43.0
Mobile LLM-350M	32	345M	53.8	33.5	62.4	68.6	44.7	49.6	40.0	57.6	51.3
MobileLLM-LS 350M	32	345M	54.4	32.5	62.8	69.8	44.1	50.6	45.8	57.2	52.1

Figure 18: Examples of results from various small models




Figure 19: Performance of the Gemma 3n memory optimized model [35]




Figure 20: On ARC-AGI 1 benchmark, a 27M parameters model (HRM, Hierarchical Reasoning Model) beat Claude Opus 4

Table 4. % Test accuracy on Puzzle Benchmarks (Sudoku-Extreme and Maze-Hard)

Method	# Params	Sudoku	Maze
Chain-of-thought, pretrained			
Deepseek R1	671B	0.0	0.0
Claude 3.7 8K	?	0.0	0.0
O3-mini-high	?	0.0	0.0
Direct prediction, small-sample training			
Direct pred	27M	0.0	0.0
HRM	27M	55.0	74.5
TRM-Att (Ours)	7M	74.7	85.3
TRM-MLP (Ours)	5M/19M ¹	87.4	0.0

Figure 21: Results of TRM [38]




Figure 22: Some results of the TRM [37]

Other key AI trends in 2025

Continuous learning and context awareness

There is growing demand for AI that can adapt to its environment, and use its past experience. Fine tuning a preexisting model is possible, but this requires retraining (albeit on a limited number of parameters if a parameter-efficient fine-tuning (PEFT) technique such as an adapter or LoRA approach is used). In addition, fine tuning cannot be done continuously or in the field due to the need for storage of the training examples and computing power. Enlarging the context of the LLM is a possibility, with some currently available models having more than 1M token context.

With the classical transformer approach, the requirements in resources (memory) are exponential with the size of the context. The Mamba approach, developed by researchers at Carnegie Mellon and Princeton, enables a linear increase, but, in practice, reduces performance. In 2025, hybrid models using both transformers and Mamba appeared with a good compromise between performance and resources. Other approaches, mainly agentic ones, use standard databases in the flow, allowing the system to take into account new and older events. Research in that field is active and new ideas are emerging [2].

World models and neurosymbolic AI: From text to the real world

LLM are trained with text (hence the name “large language models”), but texts are only a human projection of the world. It is totally different to be told that an object falls than to experience gravity directly. LLMs could be compared to children who stay in bed all the time and only learn through stories that are told to them. LLMs exhibit stunning levels of performance, but they are not the end of AI development. Interesting developments are ongoing.

In 2025, we witnessed the development of world models. “World models are neural networks that understand the

dynamics of the real world, including physics and spatial properties. They can use input data, including text, image, video, and movement, to generate videos that simulate realistic physical environments” [40]. In other words, a world model is an agent’s internal, learnable representation of how the environment works – i.e. a model that maps states and actions to predicted future observations and outcomes (and often rewards), so the system can simulate “what would happen if...” for planning, control, or self-correction. They can be used to generate custom synthetic data or downstream AI models for training robots and autonomous vehicles, or to generate realistic videos or even video games (see Genie from Google Deepmind [41]).



Figure 23: Simulating real environment helps robots to “learn”. From NVIDIA [44]


World models are becoming increasingly important because “*AI is becoming physical*”, i.e. interacting with the real world, for robotics, self-driving cars etc. In January 2026, NVIDIA released its Alpamayo [42] [43], a “*family of open-source AI models and tools to accelerate safe, reasoning-based autonomous vehicle development*”, heavily using world models.

Safety and explainability are key elements for such applications, and neuro-symbolic AI might be one answer to these challenges. Neurosymbolic AI is an approach that combines neural networks with symbolic reasoning (which is good at explicit logic, rules, and structured knowledge, like knowledge graphs or programs). The goal is to get systems that can both perceive and generalize from messy real-world data while also reasoning transparently and consistently. It seems that the self-driving system from NVIDIA is using this combination of techniques. NVIDIA is also heavily relying on its Omniverse solution that allows the creation of realistic (both from the laws of physics and visually) digital twins that can be used to develop multimodal AI.

Multimodal embodied physical AI

Researchers like Turing Award winner and former Meta scientist Yann LeCun advocate a different approach to developing bigger and bigger LLMs: his company, AMI Labs, has raised approximately \$1 billion to develop what LeCun calls “world models.” These are AI systems that, this time, no longer learn solely from textual data [45] but can understand the real world. Similar ideas are currently developed by several companies in the world.

But what could be another evolution of AI? We can perhaps go back to the philosophical discussion on innate and acquired knowledge: up to now, the “structures” of AI models are quite similar, and most efforts have focused on training, i.e. on acquired knowledge. Perhaps the next step would be to develop new “structures” (like “transformers” were in 2017) that could improve the performance of AI. The structure (the “innate” knowledge) seems



to have an important impact as well, as illustrated by the work on simulation of a fly brain [46] running in a virtual environment: only by a connectome-based brain emulation (i.e. the structure of the brain with the connections between neurons) with a physics-simulated fly body, does the simulated fly (without training or changing the synaptic weights) shows behaviour similar to a real fly. The CNN models were inspired from the work of David Marr, David Hubel and Torsten Wiesen about the structure of the visual cortex; the brain has specialized areas like agentic AI uses, so perhaps more inspiration from the structure of cognitive systems might help improving the next generation of AI

AI is becoming physical and the emergence of useful robots

The natural sequel of AI becoming physical is the emergence of robots. This is a clear target for China which has several companies working on robots, humanoid or not. Outside of China, Boston Dynamics (owned by Hyundai Motor Company), Figure AI, Tesla, etc are developing advanced robots.

Robotics is becoming a key domain for sovereignty and the US is responding to China: according to Politico [47] a national robotics strategy is in the works, including a decree planned for 2026 aimed at accelerating domestic research and production in order to secure US supply chains and compete technologically with China. This ambition is backed



Figure 24: Unitree H2 kickboxing [51]

by massive investments, such as potential \$1 billion in funding from Nvidia and Soft-Bank for Skild AI. Tesla plans to produce one million Optimus robots by the end of 2026, marking a historic transition from laboratory prototype to mass industrialization.

During the Chinese New Year show of 2026, China showed impressive performance of Unitree robots. The convergence of the progress in robots and “AI becoming physical” will certainly be the next big revolution; after AI taking the cerebral work, robots will also be more versatile and be used in increasing numbers of applications, possibly replacing humans.

According to [48], “China’s New Five-Year Plan Prioritizes Robotics and Beijing is embarking on a “whole-of-nation push” to achieve permanent dominance in physical AI technologies.<...> To understand what makes the 15th FYP’s treatment of robotics distinct from its predecessors, <...> Chapter 5 of the draft outline designates robotics as one of eight “strategic emerging industries” (战略性新兴产业) earmarked for accelerated development, alongside new-generation IT, new energy vehicles, biomedicine, and aerospace.” This Five-Year Plan “promotes the systematic elevation of robotics and “embodied intelligence” (具身智能) from a niche industrial subsidy target into the connective tissue of China’s entire economic modernization strategy.”

Everywhere, research labs and companies are working on developing more and more agile robots, even able to perform and navigates high-speed parkour with autonomous movement planning [see [51]].



Figure 25: Chinese New Year TV broadcast featured robots as martial art artists, even playing with nunchakus [50]




Figure 26: Unitree robots are available on Amazon

Unitree robots are becoming accessible (even available on Amazon for about \$18000). Unitree might be on the same track as DJI, which developed cheap drones that were sold everywhere and it became a leader in the drone domain. Drones are used now in many applications, including military operations, so cheap robots could be as well in the future.

The potential use of AI in certain applications triggers ethical problems about the use of AI and robots, which led Anthropic to be categorized as a “supply chain risk” for the US Department of War [52]. More generally, this led to the important problem of “alignment” of AI on ethical principles, and the capacity of Europe to be able to develop AI that cope with the European ethical principles (see the insert text on alignment).

The alignment problem

The following summary paraphrases and summarizes a November 2025 blog-post and research article by Anthropic [53], which explores “reward hacking” on production coding environments during reinforcement learning.

Modern language models are not trained by simply reading the internet and predicting

the next word. Once they reach a useful baseline, they are often improved with reinforcement learning, a family of methods in which a system is repeatedly placed in a task environment, takes actions, and is given a numerical signal intended to represent success. The hope is straightforward: if you pay the model for good behaviour, it will learn to behave well. The difficulty is equally straightforward: the signal is never a perfect expression of what designers

actually want, and any gap between the signal and the true intent becomes a potential exploit which is called “reward hacking”. Instead of solving the underlying problem, an agent (LLM) discovers a loophole in the measurement apparatus or evaluation harness and uses it to collect high scores cheaply.

When a powerful optimizer learns that “winning” is separable from “doing the right

thing,” it can internalize the logic of opportunism. A series of realistic experiments in a production-style setting illustrated this concern into a concrete mechanism. The key observation is not simply that a model learns to cheat when cheating is available – this is expected. The main observation is that once the model learns to cheat effectively, other forms of misaligned behaviour begin to appear, even though the training process never directly rewards them. In other words, the model’s behaviour shifts across a broad front: deception, strategic compliance, cooperation with malicious actors in synthetic scenarios, efforts to avoid monitoring, and even attempted sabotage of safety-oriented tooling in an agentic coding workflow.

Once trained in this way, the model does what any competent optimizer would do: it finds and uses the shortcuts. Reward hacking is not just a symptom of a mis-specified reward; it can be a cause of broader generalization into misalignment-relevant behaviours.

Two concrete manifestations illustrate the severity of this generalization. The first concerns sabotage in an agentic coding setting. When the model is embedded in

a realistic coding assistant workflow and asked to make changes to a codebase that is intended to detect reward hacking and other concerning reasoning, it sometimes makes changes that would quietly reduce the effectiveness of that detection. The second manifestation concerns alignment faking. Alignment faking refers to a strategic posture in which a system behaves as if it shares the supervisor’s goals while privately maintaining different objectives, often with the implicit plan to pursue its own objectives when the oversight regime weakens. The model begins to produce internally coherent rationales for deception and strategic compliance without being explicitly trained to do so. That is precisely the kind of generalization that makes alignment researchers uneasy, because it suggests the emergence of instrumental reasoning about oversight.

This could have an important impact for agentic AI: the modern trajectory of AI development is toward more agentic systems: tools that can write code, run commands, browse internal knowledge bases, and take multi-step actions with limited human supervision. In such systems, the boundary between “task performance” and “environment manipulation” is thin. If a model has

learned, in training, that manipulating the environment is an efficient path to reward, then scaling that model into more capable and more autonomous settings increases the space of possible manipulations.

This could drive to a shift of research: instead of treating misalignment as something that appears only when a model is directly trained on harmful content or explicitly instructed to do harmful things, we must also consider “natural” pathways in which ordinary optimization pressure in realistic environments creates incentives for unwanted behaviours. The practical challenge is to build systems, incentives, and evaluation that notice the difference between a model that is helpful and a model that is merely good at looking helpful.

Research in alignment should become more and more important. In large companies driven by short time to market and cost-reduction motivations, it often becomes a second priority. Europe should lead the research in this field, as it is important to deliver AI systems that are compliant with European ethics, and could help strengthen Europe’s reputation as a leader in trustworthy AI systems.

Is Apple one step ahead in AI?

This post from Milk Road AI [54] hypothesizes that Apple might have a similar reasoning to HiPEAC, betting more on collaborative distributed AI models in devices and on the “continuum of computing” than in building gigantic data centres. The HiPEAC Vision 2025 already illustrated this with the structure of Apple Intelligence, using an on-device orchestrator to select an on-device model or offload the request to a models executed in Apple servers, depending on the request (see Figure 4, or the chapter on AI in the HiPEAC Vision 2025). Although the post is highly speculative, this example illustrates the potential impact of “thinking differently”. It also raises concerns about financial sustainability of large investments in data

centres for AI. Major funding rounds, high valuations, low return from investment and complex investment relationships within the sector have led some industry leaders to warn of a potential market correction similar to earlier technology bubbles.

“The richest company on Earth just watched its rivals light \$650 billion on fire.

And did nothing.

This might be the most brilliant move in corporate history.

- Amazon is spending \$200 billion this year on AI data centers.
- Google, \$185 billion.
- Microsoft, \$114 billion.
- Meta, \$135 billion.

Combined: \$650 billion.

(...)

Apple is refusing to enter a race that might not have a finish line.

The hyperscalers are now spending 94% of their operating cash flows on AI infrastructure. After dividends and buybacks, there is almost nothing left. Amazon is projected to go negative on free cash flow this year as much as \$28 billion in the red. Alphabet’s free cash flow is expected to collapse 90%. From \$73 billion to \$8 billion.

These companies used to be the greatest cash machines ever built.

Now they’re borrowing money to keep the lights on.

The Big Five raised \$121 billion in bonds in 2025 alone. Morgan Stanley projects \$1.5 trillion in tech debt over the coming years. For the first time in history, hyperscalers hold more debt than cash and what are they getting for that \$650 billion? AI services generate roughly \$35 billion in total revenue and that's 5% of what's being spent on infrastructure.

Now here is where Apple's bet gets genius.

AI models are commoditizing faster than anyone predicted.

(...)

Open source models now power 80% of startups seeking VC funding. The moat these companies are spending hundreds of billions to build is evaporating in real time.

Apple understood this before anyone else. It didn't build its own AI model, it licensed Google's Gemini for about \$1 billion a year. Why spend \$100 billion building a factory when the product costs a billion to rent? And if a better model appears next year, Apple just switches vendors.

But Apple is not sitting still.

It just dropped the M5 chip with a 16 core Neural Engine and Neural Accelerators built into every GPU core. (...)

The M5 delivers 4x the AI performance of the M4 and Apple doesn't need \$200 billion in data centers.

Because Apple turned 2 billion devices into the data center. Every iPhone, Mac, iPad gets distributed AI at a scale no server farm can match.

While its rivals burn cash, Apple is doing the opposite. \$90.7 billion in stock buybacks last fiscal year. Its competitors? Combined buybacks collapsed 74% from their peak. Apple didn't miss the AI revolution. It just bet that the winners won't be the ones who build the infrastructure. They'll be the ones who own the customer and no one on Earth owns more customers than Apple."

Milk Road AI [54]

Conclusion

It is clear that the progress in the domain of artificial intelligence is very rapid, with no sign that it is slowing down. At the time of finalizing this chapter, in early March 2026, complex tasks could be done by the coordination of LLMs in the billion-parameter range, which was not the case a year before. Increasing the size of the networks is no longer the solution for increased performance.

2025 saw "thinking" models becoming widespread, iterating in the inference phase without changing their weights. Clever orchestration of LLMs in an agentic context was also shown to increase performance without changing the "core" LLMs. This is totally in line with what the HiPEAC Vision 2025 proposed and which continues in 2026.

2025 also witnessed more "AI becoming physical", with the added requirements of non-functional properties such as safety, response time when the LLM is controlling a robot – again a step towards a use of the NCP proposed technologies.

Embodied AI or "AI becoming physical" is the next step for AI. Europe's know-how in real-time and embedded system could be essential to develop systems using both neural-network-based solutions and more classical approaches; hence the drive for neurosymbolic solutions synergizing both approaches.

The performance of AI in software development is also rapidly (exponentially) progressing, with AI contributing towards its own development (ex for GPT-5.3-Codex). The question of the "dual speed AI", i.e. development of artificial general intelligence (AGI) / artificial superintelligence (ASI) versus the development of low-cost, regional applications and "ultra-affordable" AI is becoming more visible, and Europe has little time to decide which track(s) it will follow (of course, it is also possible to choose both).

As explained in this edition of the HiPEAC Vision, European agentic AI based on the NCP concepts (which we have named (AI)²: agentic artificial intelligence infrastructure) is a blueprint of a possible full-blown NCP and could benefit from the

knowhow, diversity and SME in Europe. But it could also be used to federalize big data centres and HPC centres in Europe to harness the computing resources required to develop AGI and even ASI.

On the hardware side, inference is becoming more and more demanding, with low latency (for "reasoning" models, for agentic AI, for "AI becoming physical") but if used in an agentic infrastructure, each agent is perhaps not so large, so dedicated accelerators for small- or medium-sized language models (or better SAMs) are key; these also in the scope of European knowhow and capabilities.

References

- [1] <https://www.theinformation.com/newsletters/applied-ai/looming-battle-agent-management-software>
- [2] <https://arxiv.org/pdf/2602.16066v1>
- [3] <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52025DC0724&qid=1762332390557>
- [4] <https://x.com/karpathy/status/2026731645169185220>
- [5] Matt Shumer “Something Big is Happening” (9 February 2026) from https://x.com/mattshumer_/status/2021256989876109403
- [6] <https://deepmind.google/blog/accelerating-mathematical-and-scientific-discovery-with-gemini-deep-think/>
- [7] <https://blogs.nvidia.com/blog/data-blackwell-ultra-performance-lower-cost-agentic-ai/>
- [8] <https://cdn.openai.com/pdf/a253471f-8260-40c6-a2cc-aa93fe9f142e/economic-research-chatgpt-usage-paper.pdf>
- [9] <https://youtu.be/qUhHCeZlwXg?t=198>
- [10] HiPEAC Vision 2023 <https://www.hipeac.net/vision/2023.pdf>
- [11] <https://venturebeat.com/technology/did-alibaba-just-kneecap-its-powerful-qwen-ai-team-key-figures-depart-in>
- [12] <https://openclaw.ai/> <https://github.com/openclaw/openclaw>
- [13] <https://www.cnn.com/2026/02/15/openclaw-creator-peter-steinberger-joining-openai-altman-says.html>
- [14] <https://venturebeat.com/technology/alibabas-small-open-source-qwen3-5-9b-beats-openais-gpt-oss-120b-and-can-run>
- [15] <https://techcrunch.com/2026/02/12/spotify-says-its-best-developers-havent-written-a-line-of-code-since-december-thanks-to-ai/>
- [16] <https://metr.org/time-horizons/>
- [17] <https://arcprize.org/leaderboard>
- [18] <https://openai.com/index/introducing-gpt-5-3-codex/>
- [19] HiPEAC Vision 2021 <https://www.hipeac.net/vision/2021.pdf>
- [20] <https://www.theverge.com/news/672357/openai-ai-device-sam-altman-jony-ive>
- [21] <https://www.bain.com/insights/will-agentic-ai-disrupt-saas-technology-report-2025/>
- [22] <https://newsletter.semianalysis.com/p/groq-inference-tokenomics-speed-but>
- [23] <https://list.cea.fr/en/neurocorgi/>
- [24] <https://aiixx.ai/blog/groq-ai-chips-vs-nvidia>
- [25] <https://www.theinformation.com/articles/metaspinner-internal-chip-design-efforts-hit-roadblocks>
- [26] <https://huggingface.co/spaces/TIGER-Lab/MMLU-Pro>
- [27] <https://github.com/deepseek-ai/DeepSeek-R1?tab=readme-ov-file#distilled-model-evaluation>
- [28] <https://arxiv.org/abs/1803.03635>
- [29] <https://developers.googleblog.com/en/introducing-gemma-3-270m/>
- [30] <https://huggingface.co/baidu/ERNIE-4.5-0.3B-Base-Paddle>
- [31] <https://huggingface.co/facebook/MobileLLM-125M>
- [32] <https://huggingface.co/facebook/MobileLLM-R1-360M-base>
- [33] <https://blogs.windows.com/windowsexperience/2025/06/23/introducing-multi-language-model-and-how-it-enabled-the-agent-in-windows-settings/>
- [34] <https://github.com/KittenML/KittenTTS>
- [35] <https://developers.googleblog.com/en/introducing-gemma-3n/>
- [36] <https://arcprize.org/blog/hrm-analysis>
- [37] <https://arxiv.org/abs/2506.21734>
- [38] <https://arxiv.org/abs/2510.04871>
- [39] <https://github.com/SamsungSAILMontreal/TinyRecursiveModels>
- [40] <https://www.nvidia.com/en-us/glossary/world-models/>
- [41] <https://deepmind.google/models/genie/>
- [42] <https://nvidianews.nvidia.com/news/alpamayo-autonomous-vehicle-development>
- [43] https://d1qx31qr3h6wln.cloudfront.net/publications/Alpamayo%201_2.pdf
- [44] <https://developer.nvidia.com/blog/simulate-robotic-environments-faster-with-nvidia-isaac-sim-and-world-labs-marble/>
- [45] <https://techcrunch.com/2026/03/09/yann-lecuns-ami-labs-raises-1-03-billion-to-build-world-models/>
- [46] <https://theinnermostloop.substack.com/p/the-first-multi-behavior-brain-upload>
- [47] <https://www.politico.com/news/2025/12/03/trump-administration-ai-robotics-00674204> and <https://www.politico.com/newsletters/digital-future-daily/2026/02/04/ai-powered-robots-are-coming-for-trade-jobs-00765584>
- [48] <https://thediplomat.com/2026/03/chinas-new-five-year-plan-prioritizes-robotics-the-world-should-pay-attention/>
- [49] <https://techxplore.com/news/2026-03-humanoid-robots-master-parkour-human.html>
- [50] <https://www.youtube.com/watch?v=JFEGUSChZOW>
- [51] <https://www.youtube.com/watch?v=JZllfrHRC4g>
- [52] <https://www.bbc.com/news/articles/cn48jj3y8ezo>
- [53] <https://www.anthropic.com/research/emergent-misalignment-reward-hacking>
- [54] <https://x.com/MilkRoadAI/status/2029318433360265227>

New Hardware



Recommendations for new hardware

Standardize chiplet platforms

Evolving from single-vendor to heterogeneous, multi-vendor systems-in-package (SiP) based on chiplets requires agreements on the physical, electrical and software interfaces between chiplets. This will span from edge to cloud.

Design flows should be adapted for compatibility with multiple physical integration flows. Europe should strengthen its electronic design automation (EDA) capabilities for chiplet platform architectures, and prepare for the deeper move to optical communication and processing in combination with electrical digital processing and memories. Furthermore, it is expected that chiplet platforms will combine fundamentally tailored processing elements (e.g. neuromorphic, in- / near-memory compute, probabilistic computing).

As each chiplet only represents part of the system, and systems are assembled for specific applications, standard protocols and interface intellectual property (IP) blocks – along with a method for system self-discovery – should be adopted, with implications on the architectural partitioning of functionality (such as virtualization of shared resources) and the operating system governing the SiP. Europe should focus on chiplet-platform operating systems and application-development methodologies. As next computing paradigm (NCP) starts to shape the software side, it is important to tune the hardware architecture accordingly.

High-end (high-performance computing, or HPC) chiplet platforms will utilize high-density, high-speed radio frequency (RF) interposers, migrating to photonic interfaces depending on the communication distance. HPC platforms evolve due to die-size constraints, i.e. the need to grow an integrated circuit (IC) beyond the maximum size of a die. Low-end (edge) chiplet platforms, however, arise from the need to create more application-specific systems with greater power and energy efficiency. These platforms benefit from lower-cost interposers (organic, glass panels) and use wide-and-slow chip-to-chip communication standards.

• Establish a robust multi-vendor SiP ecosystem:

- **Secure comprehensive physical, electrical, and software interface agreements.**
 - **Standardize chip-to-chip communication protocols across the entire system (edge to cloud).**
 - **Design chiplets for fundamental specialization without repeating common functionality.**
 - **Proactively prepare for optical communication and processing (including hybrid approaches).**
 - **Strengthen Europe's competitive position in advanced SiP integration technology.**
- ### • Carry out cross-domain actions:
- **Invest in expanding and enhancing EDA capabilities.**
 - **Prioritize the development of tailored operating systems and application-development methodologies for chiplets, with security-by-construction in mind.**

• Segment SiP designs based on performance requirements.

- **Develop HPC platforms utilizing high-density, high-speed RF interposers and photonic interfaces.**
- **Design edge platforms with a focus on cost effectiveness, power efficiency, and application-specific solutions.** Employ lower-cost interposers (e.g. organic, glass panels) and wide-and-slow communication standards.

Prepare for the reconfigurable design era

The speed of algorithmic evolution, and the rising costs and time for the design of a custom processor will likely lead to the adoption of reconfigurable elements in chiplet platforms. Softcores with their associated middleware support will gain importance. Long-term, automated code-generation using generative artificial intelligence (genAI) should be combined with automated architecture generation on reconfigurable platforms to enable the ultimate hardware / software co-design.

This requires an understanding of what such a reconfigurable platform could look like (an evolution of the field-programmable gate array (FPGA) and coarse-grained reconfigurable array (CGRA)), how it is governed (dynamic resource allocation, just-in-time synthesis of architecture description), and how it is used and perceived by applications. It will lead to ecosystem or supply-chain changes and is an opportunity for the application of NCP.

Strengthen the ecosystem of reconfigurable computing:

- Research reconfigurable cores / soft-cores for application in dynamically reconfigurable systems.
- Establish design methodologies for a reconfigurable “canvas”.
- Pursue a common agreement on protocols and data formats for the description of softcore architectures and dynamic control of the canvas.
- Develop capabilities in hardware and software co-generation.
- Develop capabilities in just-in-time synthesis.

Develop semiconductor technologies that alleviate the performance overhead of reconfigurable circuits.

Move to cryogenic operation in data centres

Data centres are today fundamentally power and thermally constrained. Dissipation-free circuits that operate at very high frequencies, based on Josephson Junctions, could break the trend. Superconducting logic has been known for years, and recent advances have made it manufacturable at scale. The obvious limitation of these circuits is their need to be cooled to cryogenic temperatures, which requires a more or less fixed amount of power. In operation, their variable power consumption is much lower than for CMOS-based digital circuits. The crossover point, where the cooling power and dynamic power consump-

tion of superconducting logic is lower than for conventional complementary metal-oxide-semiconductor (CMOS) technologies, is easily crossed in petascale systems.

- Develop **scalable superconducting technology**.
- Develop the engineering capabilities for the (cryogenic) **integration of superconducting systems**, compatible with the requirements for **quantum computing**.
- Establish **EDA for superconducting logic**.

Background

Since the 2025 HiPEAC Vision, the trajectory for hardware has sharpened. Heterogeneous integration of chiplet platforms, combining general-purpose cores with domain-specific accelerators, has moved from trend to dominant medium-term path. Looking further ahead, reconfigurable logic and software-defined silicon offer a route to avoid the rising design and tape-out costs of dedicated processors, as well as to respond to highly dynamic

workloads such as those in artificial intelligence. Across the edge-cloud continuum, power efficiency remains the limiting factor, making right-sized, domain-tuned compute essential. This creates challenges and opportunities for European players.

Scaling limits push architecture innovation

Application-level performance is increasingly achieved through circuit innovations, system and processor archi-

tectural choices, and co-development with software. As transistor density and speed improve more slowly than application demand, innovation moves up the stack. In systems-on-chip (SoCs), this has meant adding specialized circuits or cores that are far more power efficient for specific use cases. The balance between additional area cost and competitive performance has given rise to dark silicon and extensive specialized instructions. The constraint is universal; energy scarcity at the edge and thermal or power delivery limits in the cloud define today's hardware envelope.

AI as a guiding workload

The revival of deep neural networks provided a clear target for hardware optimization. Current AI models, such as large language models (LLMs) based on transformer architectures, lean heavily on multiply-accumulate operations and benefit from processor and memory designs tuned for key-value (KV) cache behaviour. Even so, integrated circuits have grown towards maximum reticle sizes. Die stitching, while enabling larger effective dies, can carry significant yield and cost penalties.



The industry is echoing the shift seen in supercomputing, from highly centralized, monolithic designs towards distributed and heterogeneous systems. Accelerators are becoming more dominant: whereas in the past it was not uncommon to find specialized cores (most graphics processing units (GPUs) have dedicated structures for matrix-matrix and vector-matrix multiplications), today several companies are exploring processors that consist mostly of dedicated units. Monolithic integration is often difficult for more exotic accelerators, due to e.g. incompatible technology.

In addition to specific acceleration, it is likely that systems tuned for the execution of agentic AI will require much more memory, as more than one model might be kept in active memory when executing agentic pipelines. This can be achieved by either adding a memory tier (between the solid-state drive (SSD) and dynamic random-access memory (DRAM)) or by high-bandwidth interconnects to disaggregated DRAM banks. [1]

Accelerators: Avenues to efficiency

Two main approaches are shaping accelerators:

1. **Digital specialization:** accelerators implemented in conventional digital Boolean logic with dedicated instruction sets and carefully designed memory subsystems.
2. **Physics-based computation,** exploiting circuit or material properties to improve message authentication code (MAC) or composite functions, e.g. via in-memory computing, or the use of non-linear materials in photonic processors.

Integration follows two patterns: on die IP blocks, for example pairing out-of-order cores with single instruction, multiple data (SIMD) or single instruction, multiple threads (SMT) and tensor units, and co-packaged chiplets. In a sense, the coprocessor is returning, both at the edge and in high-performance systems.

Digital specialization has been made easier by the adoption of the RISC-V instruction set architecture (ISA). RISC-V allows for the addition of specific instructions without breaking full compatibility. The open source toolchains are now sufficiently mature to support commercial software development. RISC-V cores exist in many flavours: large out-of-order versions; SIMD / vector cores; SMT cores... Several dataflow architectures are built around simple or extended RISC-V cores.

Near- and in-memory accelerators have evolved as well. These accelerators typically focus on massively parallelizing the multiply-accumulate (MAC) operation at higher power-efficiency. They are more sensitive to the implementation, as a mismatch in operation size and physical architecture leads to significant performance losses, and therefore require a much larger investment in the software layers needed to map arbitrary networks on a hardware-optimal structure. The near- and in-memory cores are accompanied by more conventional digital cores to cover the other calculations that occur in the networks, and are not immune to memory bottlenecks. [2], [3]

Other physics-based approaches are being explored as well, like photonic computing. These processors encode information in light and perform transformations by exploiting non-linear properties

of materials [4],[5]. They aim at executing more complex mathematical functions (macro-functions) in a few cycles instead of tens to hundreds. Although they require conversions from the electrical to optical domain and back, they are extremely energy efficient if the function executed is sufficiently complex.

Alternative ways of encoding information are also used. As an example, spiking neural-network processors aim at reducing dynamic power consumption by temporally encoding information to minimize switching activity in digital gates. Other types exist, for example based on thermodynamic (with applications in probabilistic and generative AI, or Ising machines). [6]

Quantum processors are different types of accelerators as well. In general, these are not all-purpose processors but excel at solving a few specific mathematical problems. Specifically for quantum processors, as the number of physical and logic qubits is growing, part of the attention can shift to building efficient quantum control, i.e. the interface between classical computing and the quantum domain.

Generally speaking, all of these approaches succeed at being more efficient than digital CMOS, due to the narrower range of operations for which they are optimized. However, architectural specialization carries economic risk when algorithms




Figure 1: Some of the technological elements in 3D integrated chiplet platforms. Source: imec VZW

evolve rapidly, as seen in recent years. Chiplet-based SiP platforms (Figure 1) offer a more robust path: platforms can adapt to application-level changes by re combining existing chiplets or adding new accelerator chiplets. Most chiplet solutions today are single vendor and resemble fragmented SoCs more than open platforms. Future chiplet ecosystems will mix components from multiple vendors, and will feature device discovery, secure communication between chiplets, and shared resource orchestration, memories, and network interfaces.

Standardization and platform orchestration: An opportunity for Europe

The market success of chiplet platforms hinges on two factors. First, it needs standardized interfaces across multi-vendor chiplets, with careful placement of a platform operating system to virtualize shared resources for hardware components. Second, it requires platform integrators plus interposer manufacturers, since chiplet platforms should be designed by system integrators, much like printed circuit boards (PCBs). A new breed of integrators, evolving from high-end PCB or advanced module houses, can engineer application-specific SiPs. This is a concrete industrial opportunity for Europe: building integrator capacity, packaging know-how, and trusted supply chains can strengthen

sovereignty while accelerating time to market [7].

The high-performance and edge markets will likely not be served by the same type of interposers and chip-to-chip protocols. High-performance systems require massive bandwidths between the chiplets, to maximize information passing. This means operation at hundreds of Gbps per lane, with differential signalling over RF interposers, and a likely evolution towards photonic interposers. This is too expensive now for more edge-oriented applications, where the usefulness of chiplet platforms stems from the modular and application-specific construction. In an edge use case, it is likely that the platform will be more heterogenous, using technology separated by generations (imagine power components, a near-sensor signal processor and an AI accelerator on the same platform). Chip-to-chip interfaces will need to be slower for broader compatibility, and use a much cheaper interposer. In this case, organic interposers, or giant glass panels might be the ideal base, with lower-quality interconnections between chiplets leading to less throughput per lane – but likely enough for these applications.

It is thus important to determine standards for those two main types of chiplet platforms, and not just focus on the high-end.

Longer term: Software-defined silicon and the reconfigurable canvas

Rising design and tape-out costs at leading nodes challenge the viability of dedicated accelerators. The economic exits are clear; we need to drastically lower design costs or significantly increase market volumes to amortize the ever-increasing non-recurring engineering costs.

Chiplets already help in this regard, as the typically smaller dies have higher yields (and lower cost), and their reuse across domains expands the addressable volume. In the longer run, reconfigurable circuits – a natural evolution of today's FPGAs – enable software-defined silicon. Instead of many low-to-medium volume dedicated ICs, the market will shift to high-volume, reconfigurable “canvas” devices that combine the building blocks needed to dynamically instantiate arbitrary processor architectures. Applications will ship with embedded softcore processors, making software portable across platforms by making it carry its own architecture. This canvas will co-exist with dedicated electronic or photonic chiplets on chiplet platforms, spanning edge to cloud.

This trend has in fact already begun, as we see the advent of generative AI used to create architectures and circuits specifi-




Figure 2: Improvements in power efficiency of optical links leads to their use at shorter distances. High-performance chiplet platforms are expected to use both electrical and optical connectivity. Co-packaged optics for off-SiP communication will appear first, while photonic interposers will appear later for centimetre-scale chiplet-to-chiplet communication

cally that are optimized for a particular neural network or mathematical function. Today, these macro-functions are fed into high-level synthesis tools and mapped to a either a “static” circuit using standard cells or instantiated into an FPGA. Long-term, as a larger fraction of application code will be created using AI, we need to match this with AI-created architectures that are co-generated with the application itself. The architecture description and application can be bundled and instantiated when needed, where needed, using a just-in-time compiler and logic synthesis tool executing on the canvas.

Even if the phase from specifications to the chip conception (e.g. down to the GDSII description) could be accelerated, the making of the chip itself will be slow and very expensive. The solution will be to “pre-build” (parts of) the chips and finalize it for its application in a very short time. The most common approach is programmable architecture, which is fast for mapping, but is far less efficient due to the programmability overhead. At the other extreme of the spectrum is the ASIC, which is more or less dedicated to one use. There are solutions in between such as FPGAs and reconfigurable chips.

Building reconfigurable chips while minimizing overhead and limiting the

performance penalty will require optimization of the underlying semiconductor technology. 3D-stacked layers, as seen, for example, in the CMOS 2.0 initiative, are an opportunity to embed layers dedicated to control and reconfiguration [8].

Moving to reconfigurable circuits has a massive impact on the electronic ecosystem. IP vendors today (e.g. RISC-V core vendors) license their developments to hardware manufacturers. With reconfigurable circuits, they would license their IP to application vendors instead – with several possible business models (pay-per-use, subscriptions, license once, ...). This will likely boost architectural innovation significantly, as different iterations of the IP could be distributed as a software update and not require a tape-out. Similarly, original equipment manufacturers (OEM) and integrated device manufacturers (IDM) now offer architectural upgrades to their customers (for instance, adding specific hardware accelerators for new functions introduced in their latest operating system), potentially increasing the useful lifetime of electronics without affecting income (thanks to paid upgrades). Newer reconfigurable hardware would not necessarily replace, but rather extend, existing hardware.

Getting to this point requires concerted effort across the entire ecosystem over the next decade. It is, however, a great opportunity for Europe.


Data centres: Superconducting digital logic and cryogenic integration

For data centres, a shift from CMOS-based digital logic to superconducting digital logic using Josephson Junctions offers a path to much higher clock speeds and significantly lower thermal and power overheads at cryogenic temperatures [9]. This promises relief for data centre power and thermal constraints and simplified interfacing with quantum processing units. The fundamentals of the technology are not new, but scaling it to the level of CMOS is [10].

Building processors using superconducting logic depends on the development of:

1. the underlying Josephson Junction devices and supporting structures;
2. the creation of libraries of logic cells compatible with flux-based processing;
3. interfacing techniques to cryo-CMOS memories and logic, and to room-temperature devices using optical links;
4. architectures adapted to the strengths and weaknesses of superconducting logic (e.g. higher clock speeds, less memory coherence problems, but also less memory); and
5. 3D integration and cooling infrastructure.

As many data centres are now preparing for the inclusion of quantum computers, they are adopting cryogenic cooling in any case. Interfacing to quantum computers is also facilitated when the classical computing needed for quantum control happens at cryogenic temperatures.



References

- [1] Zain Asgar, Michelle Nguyen, Sachin Katti, "Efficient and Scalable Agentic AI with Heterogeneous Systems", <https://arxiv.org/abs/2507.19635>
- [2] Asif Ali Khan, João Paulo C. De Lima, Hamid Farzaneh, Jeronimo Castrillon, "The Landscape of Compute-near-memory and Compute-in-memory: A Research and Commercial Overview", <https://arxiv.org/abs/2401.14428>
- [3] Christopher Wolters, Xiaoxuan Yang, Ulf Schlichtmann, Toyotaro Suzumura, "Memory Is All You Need: An Overview of Compute-in-Memory Architectures for Accelerating Large Language Model Inference", <https://arxiv.org/abs/2406.08413>
- [4] <https://www.photondelta.com/about/>
- [5] <https://qant.com/photonic-computing/>
- [6] Denis Melanson et al, "Thermodynamic Computing System for AI Applications", <https://arxiv.org/abs/2312.04836>
- [7] Bressler, P.R., Töpper, M. & Ramm, P. Europe's pilot line for advanced packaging and heterogeneous integration of electronic components and systems (APECS). Nat Rev Electr Eng 2, 4–5 (2025). <https://doi.org/10.1038/s44287-024-00134-6>
- [8] Moritz Brunion et al, "CMOS 2.0 -- Redefining the Future of Scaling", <https://arxiv.org/abs/2510.04535>
- [9] Rassul Bairamkulov, and Giovanni De Micheli, "Superconducting Electronics, a 25 year review", <https://si2.epfl.ch/demichel/publications/archive/2024/Review.pdf>
- [10] Anna Herr, Quentin Herr, "How to put a datacenter in a shoebox", <https://spectrum.ieee.org/superconducting-computer>

Tools



Recommendations for AI-based software and hardware development tools

Take advantage of increasing artificial intelligence (AI) capabilities

- **Develop agent-centric workflows** where AI systems act autonomously in human developer-like roles, planning, invoking and coordinating multiple tools (compilers, simulators, profilers, debuggers, deployment systems) with high-level guidance and learning from specific feedback.
- **Use AI to reduce the cost of software development**, using natural language as the primary user interface, to **democratize software development workflows** and to **avoid idiosyncratic low-level expert-only interfaces and languages**, e.g. gdb-style debugging or traditional Tcl-based electronic design automation (EDA) workflows. This should include the whole deployment workflow for production or testing (e.g. for the next computing paradigm, or NCP).
- **Embed AI inside optimization tools**, especially for hard, large-scale hardware problems where traditional methods are slow, providing improved solution quality and throughput beyond what is economically feasible with existing methods. This includes targeting and optimizing code for the NCP to a specific platform / device.
- **Evolve programming languages, models, and related tools and runtimes to be targeted by AI**. This includes tolerance for more verbose, explicitly structured code and semantics designed

for analysability and transformation by automated tools, rather than relying on implicit conventions or human intuition.

Use AI to address and exploit changes in hardware development

- **Support disruptive hardware paradigms where traditional EDA tools are weakest and AI-native approaches can reset the competitive landscape** (e.g. coarse-grained reconfigurable arrays (CGRAs), canvas-based design), giving Europe a strategic advantage despite the lack of domestic EDA vendors.
- **Use AI-driven estimation and exploration to shift hardware analysis and optimization earlier in the hardware design flow** (e.g. 3D integration, thermal, mechanical constraints, and embodied emissions).
- **Employ AI-driven hardware / software (HW / SW) co-design agents, enabling joint exploration of software, architecture, and hardware implementation to optimize defined metrics across the full stack**, from application use case to hardware implementation.

Address the verification challenge arising from increasing use of generative AI

- As AI increasingly performs code generation, **keep humans in the loop for acceptance, accountability and specification ownership**. In the short to medium term, human verification will become the bottleneck. Over time, a growing fraction of code may be generated and maintained by AI systems, raising challenges for intellectual property (IP) protection and increasing the

importance of **precise, auditable specifications as the primary human-owned artefact** over the system life cycle.

- **Develop specialized AI agents to guide, prioritize and scale verification tasks** (test generation, property discovery, and bug localization) for both safety-critical and non-safety-critical systems, including code generated by AI. In particular, **use AI to guide and prioritize formal verification rather than replacing it**.
- **Establish European verification test-beds and benchmarks for AI-assisted development and verification tools**, enabling systematic comparison, reproducibility and shared evaluation methodologies for competing approaches (in both academia and industry).

Ensure European control over hardware and software developer tools as dependable infrastructure

- Avoid application of foreign export control of the tools and the software built using them through either (a) **open source tools** (e.g. GCC, LLVM, open source EDA tools) or (b) **fully European developer tools**.
- **Support portability and retargetability across diverse accelerators**, specialized chips, domain-specific architectures. Avoid foreign vendor control through vendor-specific programming models or application programming interfaces (APIs). It is especially important to enable cross-vendor composability with respect to application binary interface (ABI) compatibilities, intermediate representations (IRs), security model enforcement, verification, for the NCP or EuroStack.

Introduction

Over the past year, AI systems have crossed a qualitative threshold. Rather than merely accelerating isolated coding tasks within human-driven workflows, agent-based systems can now orchestrate multi-step development processes, explore hardware and software design spaces, and iteratively generate, test, and refine implementations. This marks a structural transformation of engineering practice that shifts human effort, the scarcest resource in engineering, towards higher levels of abstraction.

These systems have already crossed from laboratory prototypes to production use in some organizations. Engineers are focusing on defining system intent, constraints, and evaluation criteria, while AI systems perform most of the implementation and maintenance work. The central challenge is no longer how to write efficient code, but how to specify the system precisely and verify that automatically generated artefacts satisfy those specifications.

While Europe remains strong in embedded systems, hardware engineering, and formal models, AI-driven development tooling is currently dominated by actors in the United States (US), with China advancing rapidly. Toolchains – compilers, debuggers, EDA systems, simulators, verification frameworks, and deployment systems – are becoming strategic infrastructure. If the licence for an advanced tool were to become unavailable, the impact

on Europe would be significant – which means that guaranteed access to tools is an important aspect of sovereignty. Moreover, control over training data, model weights, update cycles, hosting infrastructure, and integration APIs increasingly determines how engineering is practised.

This transition presents both risk and opportunity. Europe must leverage the discontinuity to shape AI-native development tools while avoiding structural dependence, external regulation, deskilling and lock-in.

State of the art

AI-assisted developer tools

The first wave of AI-assisted developer tools, emerging around 2021–2023, focused on integrating large language models (LLMs) directly into integrated development environments (IDEs). GitHub Copilot [1] launched as a technical preview in June 2021 and became generally available in June 2022, providing code completion and chat-based assistance within editors, generating suggestions informed by local style, nearby code, and project context. Google's Gemini Code Assist [2] (announced in 2024 as the evolution of Duet AI for Developers) similarly targets enterprise and individual developers with chat-based assistance and code-aware editing workflows. These tools allow developers to use AI assistance while remaining within familiar development environments. However, their capabilities largely remained focused on local code


editing rather than full software repositories or development processes.

Repository-scale reasoning

Recent systems extend these capabilities by increasingly operating at the level of entire repositories rather than isolated files. To support this, many AI-native development environments build semantic indices of codebases using static analysis and embedding-based retrieval, allowing models to retrieve relevant files, functions, and documentation when generating or modifying code. This enables tasks such as multi-file refactoring, API migration, dependency-aware code generation, and automated bug fixing.

Even IDE-integrated tools such as GitHub Copilot and Gemini Code Assist increasingly incorporate repository indexing and retrieval to provide suggestions informed by the broader codebase. More advanced systems, including Cursor [3], Claude Code [4], and OpenAI's coding agents, can analyse large repositories, propose coordinated edits across multiple files, and run project commands or test suites to validate results. GitHub has also extended Copilot in this direction through Copilot Workspace, announced in 2024, which supports task-level planning and multi-file editing from a single natural-language prompt. Benchmarks such as SWE-bench [5], and its more widely cited variant SWE-bench Verified, evaluate whether models can resolve real GitHub issues by producing patches that pass project test suites, reflect this shift towards repository-level reasoning.

Emerging European efforts also explore agentic workflows beyond conventional software development. For example, ChipMind [6] focuses on AI-assisted hardware and system-level design, combining code generation with architectural reasoning and co-design across hardware and software layers. Such approaches illustrate how agent-based systems can be extended beyond software repositories to encompass broader engineering workflows, including hardware design space exploration and system-level optimization.






Figure 1: The SWE-bench benchmark [5] evaluates whether models can solve real coding issues

Agentic development loops and tool invocation

A second architectural shift is the emergence of agent-based development workflows. Instead of producing a single response to a prompt, AI systems such as Cursor, Claude Code, and OpenAI Codex CLI [7] operate in iterative loops that mirror human developer workflows, alternating between editing code, compiling, testing, and debugging.

To support such workflows, modern AI systems rely on structured mechanisms for invoking external tools, such as compilers, test frameworks, and version-control systems. These tools are exposed through programmatic interfaces that agents can call, with their outputs incorporated into subsequent reasoning steps. The Model Context Protocol (MCP) [8], developed and released by Anthropic, illustrates this approach by enabling models to interact with external tools and data sources through standardized interfaces rather than ad hoc plugins, and has seen growing adoption across third-party tooling ecosystems.

Verification and quality assurance

As code generation becomes cheaper, verification and quality assurance are becoming the primary bottlenecks. AI tools increasingly support automated testing, bug localization, and code review. Automated review agents such as CodeRabbit [9] analyse pull requests and flag potential

issues; related tools such as Greptile [10] and Sourcegraph's Cody [11] extend this to codebase-wide search, question answering and review automation. AI models can also generate unit tests or identify edge cases, helping maintain reliability as the volume of AI-generated code increases.

Competitive landscape

The tools described above have predominantly been developed by U.S.-based companies, several of which are hyperscalers (e.g. Google) or closely integrated with hyperscaler ecosystems (e.g. OpenAI, backed by Microsoft, and GitHub, owned by Microsoft). China has developed competitive LLMs, including systems from companies such as DeepSeek, Alibaba, and Baidu. While these have historically been integrated primarily into domestic platforms, DeepSeek's coding models attracted significant international attention from late 2024 onwards and have been incorporated into third-party developer tooling outside China, suggesting that China's global presence in AI development tooling may be increasing.

The position of Europe

Codestral [12] and the Devstral family [13], developed by Mistral AI, are Europe's most visible entries in the AI-assisted coding space. Codestral is a code-specialized foundation model for fill-in-the-middle completion and general code generation, while Devstral extends it with agent-oriented capabilities such as multi-

file editing, refactoring, translation, and bug fixing.


Unlike US platforms that bundle coding models into tightly integrated cloud ecosystems, Mistral's models can be self-hosted or deployed through EU-based infrastructure. This allows organizations to run coding models on-premise, retain control over source code and telemetry, and fine-tune models on proprietary data without exporting sensitive assets outside the EU. While competitive on code-generation benchmarks, these models primarily function as coding assistants rather than fully autonomous agentic development environments. Their strategic advantage lies in deployability and governance: support for EU data residency, controlled model update cycles, and reduced vendor lock-in.

In Belgium, ML6 (Ghent) [14] is a representative example of a broader European pattern: rather than developing foundation models or standalone developer tools, firms such as ML6 operate as applied AI consultancies and system integrators, adapting LLMs to client-specific workflows, internal codebases, and enterprise automation.

Natural language and agent-centric toolchains

Natural language as the front-end interface

Agentic AI requires a rethinking of the entire development workflow, not merely augmenting IDEs. Natural language can increasingly serve as a front-end interface for engineering, particularly in domains dominated by brittle, expert-only interfaces such as GDB command-line debugging or Tcl-based EDA scripts. Instead of manually writing complex timing or synthesis constraints, a developer could express a design goal in natural language, such as "optimize this accelerator for sub-5W power at 10 ms latency". Rather than stepping through GDB or inserting printf-style debugging statements, the user could ask "why is this assertion being raised?" with the AI tracing the failure from its immediate trigger back to its root



cause. Similarly, deployment workflows for heterogeneous systems (e.g. NCP-style edge-cloud configurations with diverse processing units, memory types, energy availability, communication latency, etc.) can be described declaratively, with AI agents generating the required continuous integration / continuous development (CI / CD) pipelines, container configurations, and device-specific build scripts.

This shift increases productivity and broadens access to system development, but it does not eliminate technical complexity. High-level requirements will often remain informal and ambiguous, but they must ultimately be resolved into an explicit specification. Natural language is the interface through which intent is captured and clarified, but the authoritative artefact must be a precise, typed or otherwise structured, policy-aware, machine-checkable specification that makes constraints, assumptions, and trade-offs explicit. A request such as “optimize for low energy consumption on edge devices” should therefore be translated into an explicit, machine-checkable constraint set covering, for example, power, memory footprint, and latency bounds.

Agent-centric development workflows

Beyond interface changes, the development process itself is becoming agent-centric. AI systems can increasingly assume roles analogous to junior developers or integration engineers: planning multi-step workflows, invoking compilers and simu-

lators, running test suites, and iteratively refining outputs based on feedback from profilers and debuggers. For instance, an agent might analyse a repository, detect performance bottlenecks, propose code changes, rebuild the system, run regression tests, and summarize the results for human review.

To support this model, toolchains must evolve from isolated utilities into composable services with well-defined APIs and traceable invocation logs. Agents coordinating compilers, profilers, debuggers, and deployment systems must operate through standardized interfaces, enabling deterministic replay and auditability. Without such safeguards, cascading errors and emergent behaviour become real risks. For example, an agent may modify a build script, inadvertently change compiler flags, trigger a heuristic CI repair, and silently introduce a performance regression into production. Preventing such failure chains requires strict logging, version control of models and tools, and clear human acceptance authority. Responsibility for approving changes must remain with the human in charge rather than being delegated unthinkingly to the machine.

Programming models for AI-driven development

Code designed for analysability

Programming languages and runtimes must evolve to better support AI-assisted

code generation and maintenance. Historically, mainstream languages were designed to balance machine execution constraints with human productivity, readability, and conciseness. In AI-mediated systems, however, analysability and transformability become more important. This favours explicit structure over informal software patterns, ad hoc locking without formal guarantees, or unwritten conventions about memory ownership. This structure reduces the cost of verification and enables reliable automated refactoring and reasoning. Relevant directions include formal ownership models, well-defined concurrency and memory semantics, and explicit resource constraints. New “AI-native” languages may struggle with adoption and limited training data, so a more pragmatic path may be to retrofit analysability into dominant ecosystems such as C++, Rust, or Python. This would improve machine reasoning without abandoning established toolchains.

Specification as first-class artefact

Once natural-language intent has been resolved into a structured machine-checkable form (see “Natural language as the front-end interface”, above), that specification increasingly becomes the primary human-owned artefact. In traditional engineering practice, source code served both as the executable artefact and as the central intellectual product. In agentic AI workflows, humans may no longer be responsible for writing every part of the implementation, but they remain responsible for defining requirements, approving specifications, validating that generated artefacts satisfy them, and accepting the legal, safety, and operational consequences of deployment.

For this reason, specifications must be precise, versioned, traceable to the AI-generated artefacts, and accompanied by human-readable verification artefacts such as test plans, and verification reports. Specifications should not exist merely as informal documentation but as structured artefacts embedded in the development workflow. Suitable formalisms may include constraint-based domain-specific languages, typed workflow graphs,

contract-based design languages, or formal specification systems such as TLA+ [15], Alloy [16], or application-specific verification schemas. The essential requirement is that natural-language intent is compiled into structured, stable artefacts that can be verified and traced across iterations. In this sense, engineering practice moves from code-centric to specification-centric.

AI in hardware development

Shifting analysis earlier in the flow

AI-driven estimation and exploration tools enable important system-level constraints to be incorporated much earlier in the hardware design process. Traditionally, issues such as 3D integration feasibility, thermal constraints, mechanical stress in advanced packaging, embodied carbon emissions, and system-level power envelopes have been addressed relatively late, often after key architectural decisions have been made. At that stage, correcting mistakes can require costly redesign cycles.

By embedding predictive models and AI-assisted exploration into early architectural stages, these constraints can be evaluated before the register-transfer level (RTL) design is frozen. For example, thermal estimation models can guide partitioning decisions in 3D-stacked designs, sustainability metrics can be incorporated alongside performance and cost objectives in multi-objective exploration, and mechanical constraints associated with advanced packaging technologies can be integrated during floorplanning rather than being discovered only after layout.

This earlier integration reduces late-stage redesign risk and allows optimization across a broader design space. Rather than optimizing performance first and then reactively addressing thermal or mechanical violations, design tools can jointly optimize across these dimensions from the outset. AI-guided exploration therefore shifts analysis upstream to the stage where design freedom is greatest.

Discontinuities in EDA and hardware design

Hardware design remains characterized by extremely large search spaces, long validation cycles, and heavy reliance on incremental refinement using mature EDA toolchains. Established chip designers therefore retain major advantages: they can absorb the high cost of design iterations, build on long-standing relationships with fabrication and packaging partners, reuse extensive IP libraries, rely on large teams with deep engineering experience, and train or fine-tune specialized AI models on proprietary design and toolchain data collected across previous projects.

AI-driven methods are unlikely to overturn these advantages in the near term. However, they can improve important stages of the design flow, including automatic design-space exploration, floorplanning, placement and routing, and early prediction of timing, power or congestion.

Disruption is therefore more likely to occur first at the level of tool providers and design methodologies than at the level of chip designers. AI-driven approaches are particularly promising where traditional heuristics are weakest or where new architectural paradigms emerge. Examples include coarse-grained reconfigurable architectures (CGRAs), canvas-based design flows, generative RTL synthesis, workload-specific accelerator exploration, and other hardware classes for which mature tooling is still limited. (See the hardware chapter for a discussion of reconfigurable hardware design).

AI-guided placement and routing in CGRA fabrics, for instance, can explore mapping strategies that are infeasible to search exhaustively using conventional heuristics. Generative approaches to RTL design can propose architectural variants tailored to specific workloads, which are then evaluated using learned performance and power models. Multi-objective optimization frameworks can simultaneously consider power, thermal constraints, area, and cost.

In such contexts, AI-driven methods can level the playing field, particularly where incumbents do not yet have decades of accumulated tooling advantage. The strategic opportunity for Europe lies not necessarily in replicating existing large-scale EDA vendors, but in pioneering AI-centric workflows for emerging hardware classes and for the weak points of existing design flows.

HW / SW co-design agents

The NCP vision emphasizes tight coupling between application software, compiler infrastructure, architecture design, and hardware implementation. AI-driven HW / SW co-design agents can help explore this joint design space more systematically. Rather than optimizing software and hardware sequentially, AI agents can evaluate global trade-offs across abstraction layers. For example, an agent may modify an algorithm to better match the memory hierarchy of a candidate accelerator, adjust compiler scheduling strategies to reduce contention, and propose hardware parameter changes to improve throughput or energy efficiency. Such joint exploration allows application-architecture alignment that is difficult to achieve through isolated optimization passes.

Realizing this potential requires clear objective functions and evaluation metrics, so that AI systems can trade off accuracy, compute, energy, latency, and cost. In particular, co-design should be guided by holistic system-level metrics rather than isolated software metrics such as accuracy alone or isolated hardware metrics such as throughput or memory bandwidth alone. In AI workloads, for example, emerging metrics such as “intelligence per joule” have been proposed to capture useful task-level capability delivered per unit of energy, rather than optimizing model- or hardware-level metrics in isolation [17].

Cross-layer co-design at scale requires interoperable toolchains and standardized intermediate representations. Compiler backends must be accessible and modifiable, IRs and ABIs must be stable across components, and hardware description

tools must expose interfaces suitable for automated exploration. Moreover, project timelines must accommodate design-space exploration rather than necessitating premature convergence on a single configuration. Without these enabling conditions, AI-driven co-design remains constrained to local optimizations.

Verification in the age of generative AI

Verification as bottleneck

As AI accelerates code and hardware generation, verification becomes the binding constraint. Earlier AI coding tools mostly generated isolated suggestions within a human-driven workflow, whereas newer agentic systems increasingly operate in iterative plan–execute–check loops, allowing them to propose, test, and refine larger changes with less human intervention. The ability to generate patches, push them upstream, and deploy them into production is therefore increasing more rapidly than the capacity to validate them rigorously. Without corresponding advances in verification, productivity gains risk being offset by declining reliability and increased risk.

AI can assist in multiple verification-related tasks, including automated test generation, property discovery, bug localization, and trace analysis. For example, AI systems can suggest invariants that may not have been explicitly documented, generate edge-case tests for complex state machines, or identify suspicious execution traces that merit closer inspection. However, these capabilities augment rather than replace human accountability.

Governance and accountability

Responsibility for accepting design decisions, validating compliance with specifications, and managing safety and legal risk must remain with human developers and managers. As larger fractions of code are both generated and maintained by AI systems, the precision and auditability of the specification become central to trust. Verification must therefore operate not

only at the level of code correctness but also at the level of specification compliance.

Human experts must retain authority for approving design decisions, validate that specifications are satisfied, interpret ambiguous requirements, and assess risk trade-offs not captured formally. Productivity gains from AI-assisted generation therefore require stronger institutional control mechanisms. These may include freezing model versions for certified systems, mandatory logging of agent actions, reproducible build pipelines, explicit separation between generation roles and acceptance authority, and clear documentation of which model and tool versions contributed to a particular artefact.

Traceable human accountability also intersects directly with liability frameworks, intellectual property ownership, and safety certification processes. In public procurement contexts – particularly in critical infrastructure, defence, and public administration – contracting authorities require traceable responsibility for design decisions. If AI systems generate substantial portions of code, procurement frameworks must identify who approved the artefact, under which model version, and based on what verification evidence. Human arbitration thus becomes a prerequisite for legal eligibility and certification, not merely a technical safeguard.

AI-assisted development must also preserve the transparency of engineering artefacts. Requirements, design documentation, test plans, test scripts and verification reports must remain fully human-readable and auditable. AI tools may assist in generating or analysing these artefacts, but they must not replace them with opaque internal reasoning. Existing certification frameworks – such as the EU AI Act, DO-178C, IEC 61508 and related sector-specific standards – already provide robust structures for safety assurance and accountability. Rather than replacing these frameworks, AI-assisted development automatically generates and maintains these artefacts, making it economically viable to apply certification standards more broadly.

Reproducibility

AI development tools differ fundamentally from traditional development tools: while the behaviour of traditional development tools is fixed for a given version, AI development tools are inherently probabilistic. Reproducibility therefore does not imply bit-exact equivalence, but the consistent ability to verify that the resulting software satisfies its specification under controlled conditions. At the same time, these tools depend on rapidly evolving models, prompts and datasets. Simply recording the name of the model is insufficient to reproduce results; reproducibility requires traceability of the full toolchain configuration. In a distributed ecosystem, such traceability must be interpretable and comparable across organizations. This, in turns, requires shared standards for describing model provenance and toolchain configuration, as well as coordination mechanisms to signal when model updates materially affect verification-relevant behaviour. These requirements imply greater transparency and access to older model versions than is currently the case. More structured life-cycle governance becomes particularly important in cyber-physical and safety-critical systems, where traceability, validation and controlled updates are already required for certification. In these domains, AI-enabled development tools must integrate with existing assurance and life-cycle management frameworks.

Skills and knowledge retention

Finally, there is a risk of deskilling. If AI systems increasingly generate and validate code, some engineers may lose familiarity with low-level implementation details. This transition differs from earlier abstraction shifts in one important respect: traditional abstractions remained deterministic once source code, compiler, and flags were fixed, whereas AI-generated implementations may depend on probabilistic generation and evolving model versions. The gap between specification and implementation can therefore be harder to reconstruct, increasing the importance of governance and verification infrastructure.

AI-assisted formal methods

Even when an AI system produces a useful or correct result, it may be difficult to understand or reconstruct how that result was reached. This is known as epistemic opacity, and it is a central challenge in verification. Humans can also fail to properly explain their reasoning, but they remain the locus of legal, moral, and organizational responsibility. AI systems may generate correct code and plausible justifications, while human reviewers lack the time and bandwidth to independently reconstruct the reasoning. Yet responsibility for accepting such artefacts remains human.

A tiered model of oversight can mitigate this tension. Low-risk changes may be automatically verified and reviewed in aggregate. Medium-risk changes may undergo AI-guided review combined with focused human validation. High-risk or safety-critical modifications should require formal verification and explicit expert sign-off.

Rather than replacing formal verification, AI should guide and prioritize it. AI systems can suggest candidate properties, decompose verification tasks into smaller subproblems, suggest possible proof strategies, and highlight components likely to merit deeper inspection. By focusing human effort where it is most needed, AI can make formal methods more economically viable across a wider range of systems. In this way, AI-driven prioritization could strengthen, rather than weaken, formal rigour.

European verification testbeds

Shared verification testbeds and benchmarking infrastructure could support the evaluation of AI-assisted development tools across academic and industrial settings. These testbeds could build on existing open benchmark suites, similar to SWE-bench, Defects4J [18] and SV-COMP [19], extending them to cover the verification and validation of AI-generated software and hardware using debuggers, testing frameworks and formal methods.

A standardized evaluation pipeline would provide traceability, reproducible build and execution logs, and versioning of the underlying AI models. The infrastructure could also support standardized metrics, public leaderboards, and transparent evaluation workflows, enabling consistent and reproducible comparison across approaches.

Such frameworks would make it possible to assess how AI acceleration affects the applicability of rigorous verification methodologies, including formal methods, structured unit testing and static analysis, beyond traditional safety-critical domains such as aerospace, automotive and rail. Broader use of these methods across software development may improve software quality while building on Europe's existing strengths in formal methods and high-assurance engineering.

Toolchains as strategic infrastructure

Sovereignty and dependable infrastructure

AI-driven development tools are foundational infrastructure for the next generation of computing systems. Europe requires dependable autonomy and the ability to ensure continued access to these tools without hidden kill switches, unilateral termination of support, opaque licensing dependencies, or exposure to foreign export control regimes. This also requires sustainable access to enabling resources such as training data and suitable cloud compute, storage and hosting infrastructure.

Control over key technical layers further determines how engineering practice evolves. Access to foundation-model weights, update cycles, tool-invocation APIs and benchmark evaluation criteria influences which programming idioms, architectures and optimization strategies become dominant. If these layers are controlled externally, engineering practice itself becomes structurally dependent. Avoiding proprietary extensions, vendor-specific APIs and opaque integration layers

is therefore essential to maintain portability and interoperability across toolchains.

Open and European-controlled toolchains

Dependence on foreign-origin toolchains and IP can create downstream exposure to foreign export control regimes, even when the resulting software or hardware is developed in Europe. Europe should therefore actively support open source developer tools such as GCC, LLVM, and emerging open EDA initiatives. For publicly available open source software, US export controls are in many cases less restrictive than those for proprietary alternatives [20], but openness alone does not eliminate exposure where encryption or other controlled technologies are involved. Key toolchain components should also be developed, maintained and hosted within European legal jurisdiction. European ownership alone is not sufficient if core IP or operations remain outside the EU.

Toolchains also shape industrial ecosystems. Open and composable toolchains allow specialization by small firms and support modular innovation across many actors. By contrast, vertically integrated stacks favour large actors that control the full pipeline from foundation models to deployment. Maintaining open and interoperable development ecosystems is therefore essential for preserving Europe's diverse industrial landscape.

Portability across diverse accelerators and architectures must also be preserved. This requires adherence to standardized IRs, stable ABIs, and documented APIs that prevent vendor lock-in. In the context of the NCP and EuroStack visions, cross-vendor composability at the compiler, IR, security, and verification layers is essential for sustaining modular ecosystems and independence from non-European hardware providers.

Conclusion

AI-driven development represents a structural discontinuity in engineering practice. Toolchains are evolving from

human-centred workflows into orchestrated, multi-agent systems in which natural language, automated tool invocation, and repository-scale reasoning increasingly shape how software and hardware are designed. As implementation becomes easier to generate, the human contribution shifts upwards: towards defining intent, formalizing specifications, validating outcomes, and governing increasingly autonomous workflows. The central challenge is therefore no longer only how to produce code or hardware efficiently, but how to ensure that generated artefacts remain correct, auditable, and aligned with human goals.

This shift has consequences across the full stack. Specification becomes the primary human-owned artefact, with its

precision and auditability whether the system can be trusted to meet its intended use. Programming languages and runtimes must favour analysability and transformability over informal conventions to support humans working with machines. In hardware design, AI can enable earlier integration of thermal, mechanical, energy, and architectural constraints, as well as cross-layer HW / SW co-design. Verification becomes the primary bottleneck, as AI rapidly increases software productivity and generation outpaces validation.

For Europe, the strategic question is not whether AI will enter development workflows, but under whose control, according to which standards, and in support of which values. Competing feature-for-feature with hyperscaler ecosystems is unlikely to be the

best path. The greater opportunity lies in shaping open and dependable toolchains, specification-centred engineering practices, strong verification infrastructure, and interoperable workflows aligned with European strengths in embedded systems, formal methods, and HW / SW co-design.

If Europe succeeds, AI-native development tools could lower barriers for small and medium enterprises (SMEs), strengthen leadership in HW / SW co-design and high-assurance systems, and reduce systemic dependence on vertically integrated foreign platforms. If it fails, the engineering stack may consolidate further around externally controlled models, APIs, and toolchains, leaving Europe with diminishing influence over how future software and hardware are designed, verified, and deployed.

References

- [1] <https://github.com/features/copilot>
- [2] <https://cloud.google.com/blog/products/application-development/introducing-gemini-code-assist-enterprise>
- [3] <https://cursor.com>
- [4] <https://code.claude.com>
- [5] <https://www.swebench.com/>
- [6] <https://www.chipmind.ai>
- [7] <https://developers.openai.com/codex/cli>
- [8] <https://modelcontextprotocol.io>
- [9] <https://www.coderabbit.ai/>
- [10] <https://www.greptile.com>
- [11] <https://sourcegraph.com/docs/cody>
- [12] <https://huggingface.co/mistralai/Codestral-22B-v0.1>
- [13] <https://mistral.ai/news/devstral>
- [14] <https://www.ml6.eu/>
- [15] <https://foundation.tlapl.us/>
- [16] <https://alloytools.org/>
- [17] Chen et al. AI+HW 2035: Shaping the Next Decade. <https://arxiv.org/abs/2603.05225>
- [18] <https://github.com/rjust/defects4j>
- [19] <https://sv-comp.sosy-lab.org/>
- [20] <https://www.linuxfoundation.org/resources/publications/understanding-us-export-controls-with-open-source-projects>

Cybersecurity



Recommendations for cybersecurity

Support the use of artificial intelligence (AI) for cybersecurity and trustability, and to fight influence campaigns

- **Promote AI-powered cybersecurity to counter AI-driven attacks**

AI is enabling massive automation in cyberattacks. Thus, it is crucial to develop AI-based defences capable of detecting, analysing, and responding to large-scale, automated AI-powered cyberattacks in real time.

- **Develop AI tools to detect and mitigate disinformation / influence campaigns**

Similarly, AI is enabling the massive creation of fake content and posts purporting to be written by humans in social networks. The European Union (EU) should thus develop scalable AI-based solutions to automatically identify, analyse, and counter large-scale disinformation and influence operations affecting trust in the digital space.

Secure AI systems against cyberattacks and malicious influence

As AI systems become ubiquitous, tampering with their proper functioning has a huge impact. It is thus vital to ensure that AI systems are robust against cyberattacks, as well as against adversarial manipulation (e.g. poisoning, prompt injection) through secure-by-design approaches, standardized evaluation benchmarks, and formal verification methods.

Reinforce IT supply-chain security (software and hardware)

Address systemic risks arising from dependencies in software and hardware supply chains that can propagate vulnerabilities, or can be used to propagate malevolent code, at large scale.

- **Support robust identification and high-quality software bills of materials (SBOMs) for tracking and integrity verification of software components**

Enable precise, verifiable identification of software components to improve vulnerability tracking, compliance, and long-term traceability, by supporting solid and precisely defined SBOMs – including the reliance of intrinsic, tamper-proof, software identifiers such as Software Hash Identifiers (SWHID) – and their inclusion in tools.

- **Secure package managers**

The reuse of software (components) is extensive, through dependencies managed by package-management ecosystems that currently offer too many opportunities for attackers. It is thus crucial to support the design of secure yet usable software-package managers.

- **Promote research on source-code sanitation methods and tools**

Improve the security of open source code through large-scale, automated approaches that can systematically detect known vulnerabilities in code repositories, using standardized vulnerability databases and automated analysis tools. Develop reliable techniques to automatically generate, validate, and apply security patches at scale.

Ingrain trustability into the fabric of the next computing paradigm (NCP)

- **Promote lightweight yet robust authentication mechanisms for NCP services**

Develop scalable authentication mechanisms that provide strong security guarantees while remaining lightweight, efficient and practical for high-frequency usage in large-scale distributed systems like the NCP.

- **Secure inter-service communications in the NCP**

Similarly, ensure confidentiality, integrity, and authenticity of communications between the numerous services in the NCP through strong encryption and mutual authentication techniques that can scale up.

- **Exploit the NCP's distributed nature for resilience and degraded modes**

Leverage distribution to provide fault tolerance and maintain essential functionality under failures or attacks, through the capacity to function in isolation in degraded mode.

Increase cybersecurity and EU sovereignty

- **Do not mandate backdoors or weaken encryption and defences**

Despite recurring pressure, preserve strong cryptographic and cyber protections as necessary fundamental foundations for secure digital infrastructure.

- **Develop and promote EU-based public key infrastructures (PKI) and certificate authorities (CA)**

Strengthen EU control over cyber trust infrastructure. PKIs and CAs are crucial

to ensure the uninterrupted functioning of information technology (IT); these should be based in the EU to increase the EU's security and sovereignty.

- **Secure critical open source components as a European common good**

Identify key open source components as critical infrastructure and support their sustained maintenance and access.

- **Mandate verifiability of non-EU hardware and software components**

Mandate transparency and auditability of non-EU components, to be able to assess their trustworthiness.

Increase cybersecurity (and AI) skills and awareness

- **Increase European cybersecurity workforce capacity, expanding cybersecurity education to train more specialists**

Address the shortage of skilled professionals needed to secure modern digital infrastructures, by strengthening education and training programmes to develop a larger and more capable workforce in cybersecurity as well as in AI for cybersecurity tools.

- **Raise population-wide awareness of basic cybersecurity practices**

Improve general cybersecurity hygiene to reduce risks linked to human factors.

Introduction

As IT becomes ever more pervasive, and evolves ever faster, cybersecurity – in the broad sense of the term – issues continue to evolve at a very fast pace, with old challenges remaining while new ones emerge.

This chapter first sets out several challenges that exist for the EU in current IT systems, specifically stressing challenges related to AI and cyberspace, then challenges related to the trustability of the NCP fabric itself.

It then proposes paths to solutions for EU, in the form of technical recommendations, to address those challenges, followed by policy recommendations.

As mentioned in previous HiPEAC Visions, for cybersecurity to be the most effective, it must be tackled upfront, not as an afterthought.

Challenges related to AI and cyberspace

The massive, unprecedented rise of AI capabilities over the last few years has created a number of new challenges for cybersecurity.

AI enables large-scale automation of cyberattacks

In cyberspace, AI is not only providing a host of new possibilities to users, including by transforming defensive systems, but is also **lowering barriers for attackers**, by

automating large parts of cyber operations that used to be done manually. AI methods, especially machine learning and generative techniques, are increasingly being used in offensive cyber vectors such as phishing, malware development, reconnaissance, and social engineering, expanding the scale, speed, and adaptability of attacks [1].

The European Union Agency for Cybersecurity, ENISA, notes in its 2025 threat landscape report [2]: “As a predictable trend, Large Language Models (LLMs) are leveraged to craft more convincing phishing emails; with reportedly over 80% of all phishing emails identified between September 2024 and February 2025 using AI to some extent.” AI can automate exploitation pipelines, make social engi-

neering more effective, and scale attacks far beyond traditional manual efforts. [3].

Real-world evidence demonstrates the existence of this trend: generative video tools have become a goldmine for scammers, enabling entirely new forms of fraud such as deepfake extortion and AI-assisted phishing. [4]

This corresponds to warnings by law-enforcement authorities that AI tools are making organized crime operations more **scalable and harder to detect**, from multilingual scam campaigns to impersonation attacks – a transformation driven by the same accessibility and adaptability that make AI so valuable for legitimate purposes [6].




Figure 1: Phishing attacks 2021-2022 (from [5])

Disinformation and influence campaigns undermine trust in the digital space

The impact of AI is not limited to “technical cybersecurity” attacks. Large language models and generative systems have made it easier to make and spread false but convincing information on a large scale [7]. Generative language models could thus systematize the creation of convincing false information to the point of altering the structure and dynamics of online influence campaigns and operations.

AI also enables the creation of deep-fakes, thus amplifying misinformation campaigns by generating content that is increasingly difficult to distinguish from real information [8]. AI-made audio deep-fakes have already been used in fraud and manipulation, such as the impersonation of corporate leaders to authorize fraudulent financial transfers [9].

Independent experts have also emphasized these risks for democratic processes, stressing for example that swarms of AI bots capable of autonomously coordinating and tailoring messaging on social media could pose a threat to electoral integrity and public trust [10].

Scalability of AI-powered attacks challenges traditional defence methods

AI allows attacks to scale horizontally (**many targets**), vertically (**multiple attack stages**), and temporally (**continuous operation**).

Scalability is a major aspect in describing how AI changes the cyber threat landscape. Traditional cybersecurity defence methods – log-based detection, signature matching, and rule-based systems – struggle to keep pace with attacks that are both **high-volume and adaptive** [11]. AI-driven attacks can operate autonomously or semi-autonomously: automated intrusion attempts, machine-generated phishing, and intelligent malware that tailors itself to evade legacy defence techniques [12]. This scalability increases already-known challenges

in intrusion detection. Conventional cyber defence systems often rely on known attack signatures and patterns, which makes them less effective against zero-day exploits and novel AI-augmented tactics.

Overall, a recent academic survey covering over 70 references on AI cybersecurity threats identified that adversarial AI, automated malware, and large-scale social engineering represent new risk categories where current defences are inadequate [13]. This clearly calls for coordinated research on explainable and scalable solutions.

Need for automated and scalable AI-based cyber defence

However, to fully cope with the massive automation enabled by AI use in cyberattacks of all kinds, **AI must also be used defensively** to counter AI-enabled threats, especially where human-only approaches cannot keep up with attack volume or speed. In response to the scale of AI-enabled threats, **scalable AI-assisted defence** mechanisms capable of matching the pace and complexity of modern attacks are necessary to maintain digital trust and protect against both technical and information-centric attacks [13].

Machine learning models are already used for threat detection, anomaly analysis, and predictive prioritization of risks – tasks that would be infeasible to manage manually at current data volumes [14].

AI can help in detecting and mitigating synthetic content and information operations, and commercial deepfake detection platforms exist that show an example of the emerging industrial response to automated threat generation, offering real-time identification of AI-generated multimedia content [15] [16] or detection of AI bots [17] to support trust in digital information.

However, in this shield and sword race, attackers currently seem to have the advantage. [18] shows a comprehensive study of the impact of AI on cyberattacks and defence, and expects attackers to retain an advantage over defenders for some time. This calls for strong and sustained meas-

ures to reinforce cybersecurity in Europe, in order to tip the scales.

Challenges related to the trustability of the NCP fabric

AI systems are exposed to data poisoning, model-level attacks and prompt injection attacks

AI systems come with a number of challenges and risks, as outlined by the United States National Institute of Standards and Technology (NIST) in their AI Risk Management Framework [19].

Among these, AI system embedded in the NCP inherit vulnerabilities related to **training data, model integrity, and deployment pipelines**. Machine-learning models can be intentionally manipulated through **data poisoning, backdoor insertion, and adversarial model manipulation**, often without affecting nominal performance during testing.

Foundational work by OpenAI, Google Brain, and academic researchers has shown for a very long time that poisoning as little as a small fraction of training data can lead to systematic misclassification or trigger hidden behaviours in deployed models [20] [21]. Recent surveys confirm that these attacks remain practical in real-world pipelines, particularly where training data is aggregated from external or open sources [22].

Large language models (LLMs) are now widely (and increasingly) deployed in many parts of the IT fabric of the NCP, including critical services, decision-support systems, and autonomous agents. However, one of their well-known weaknesses is their susceptibility to **prompt injection attacks**. Prompt injection in LLMs occurs when an attacker uses adversarial inputs to manipulate an AI system into executing unwanted actions. Indeed, LLMs are fundamentally **unable to reliably distinguish between instructions and untrusted data**; hence they are intrinsically vulnerable to manipulations [23].




Figure 2: AI attack surface (from [26])

For this reason, prompt injection attacks can be very successful against LLM-based systems that aren't sufficiently protected. The current rise of AI agents also makes this attack surface bigger, allowing hybrid attacks that mix traditional cyber exploits with prompt manipulation. [24] and [25] provide an account of such an attack on an AI agent in a software supply chain.

Supply-chain attacks affect both software and hardware and are extremely widespread

Supply-chain attacks have emerged as one of the most damaging threat vectors affecting the fabric of IT, because they **exploit implicit trust relationships** between vendors, integrators, and users. These attacks allow adversaries to compromise a single upstream component and automatically propagate malicious functionality to thousands of downstream systems (see [24] and [25] for an occurrence in February 2026)

The SolarWinds Orion compromise is a well-known example; see Figure 3. It involved sending malicious updates

to about 18,000 organizations, including government agencies and operators of critical infrastructure [27]. Similar risks have been found in hardware supply chains, where malevolent changes or undocumented features added during manufacturing can remain undetected for years.

ENISA threat landscape reports, such as [28], repeatedly identify supply-chain compromise as a **systemic risk multiplier**, affecting cloud services, networking equipment, operating systems, and embedded devices. This very broad impact justifies the importance of reinforcing Europe's IT supply chain cybersecurity.




Figure 3: Diagram of SolarWinds supply chain attack. The attackers compromised SolarWinds and modified the code of ORION software. The ORION instances in the customers were updated with malware, which allowed the attackers to access the data of customers (from [28])

Numerous known vulnerabilities persist in massively used software codebases

Even after decades of work on security engineering, **critical vulnerabilities still exist in widely used software components, often for years**. This is partly due to code-base size and complexity, but also to financial incentives that prioritize functionality and speed over long-term maintainability.


The Log4Shell vulnerability (CVE-2021-44228) perfectly illustrates this: an attack relying on a vulnerability in an obscure yet ubiquitous Java logging library affected millions of systems worldwide, including cloud services, enterprise software, and industrial platforms. Apache Software Foundation and NIST analyses showed that even organizations with mature security practices struggled to fully identify and remediate the affected dependencies.

Empirical studies confirm that in general a significant fraction of known vulnerabilities remain unpatched long after disclosure, particularly in transitive dependencies embedded deeply in software stacks. [29], for example, reports that “90% of critical vulnerabilities are unpatched for 180 days or more”.

Critical dependence on shared open source components

The NCP relies heavily on **shared open source software (OSS) components**, which are reused across operating systems, middleware, cloud platforms, and embedded systems. Industry analyses show that large swaths of modern applications contain open source code [30]. Furthermore, the latter often contains hundreds or thousands of transitive software dependencies, even for very small programs.

The Open Source Security and Risk Analysis reports [31] [32] consistently document high vulnerability densities in widely reused OSS libraries, combined with limited maintainer resources. While open source does enable transparency, innovation and massive sharing and reuse, it also creates concentration risk, where a small



number of components can become single points of failure for the global digital infrastructure.

Critical dependence on non-EU components and infrastructures

A **substantial portion** of the IT fabric used in Europe that will be fundamental for the NCP – including processors, accelerators, firmware, operating systems, cloud platforms, software repositories, and AI toolchains – **is developed and manufactured outside the EU** [33]. This creates dependencies that go beyond technical risk, raising issues of strategic autonomy, assurance, and jurisdictional exposure. European Commission and ENISA policy documents [34] note that reliance on non-EU components limits the ability to independently audit, certify, and control critical digital infrastructure. Hardware and firmware components, in particular, are difficult to inspect after deployment, thus creating a huge dependence on vendor trust.

In addition, the reliance on **non-EU-based infrastructures creates a dramatic and permanent vulnerability**, meaning that any disruption – be it for technical or political reasons – to access to these infrastructures can immediately block tremendous parts of EU IT systems, hence

massively disrupt the EU’s economy and society.

Technical recommendations

Support AI for cybersecurity and trustability, and to fight influence campaigns

AI should be used to fight cyberattacks and fight influence campaigns. Indeed, since the latter are increasingly AI-powered, which provide them with unprecedented efficiency, AI is essential to counter this problem.

Promote AI-powered cybersecurity to counter AI-driven attacks

AI and machine learning (ML) methods are essential for modern cybersecurity, particularly to counter increasingly complex and, above all, automated attacks. AI techniques currently make it possible to massively automate simple cyberattacks, and are already helpful for attackers to build elaborate attacks faster; they are likely to be able to fully automate much more evolved attacks in a very near future. As CrowdStrike 2026 Global Threat Report [Crowdstrike-GTR2026] states, in a foreword titled “The Age of the AI Adversary Begins”: “In 2025, AI-enabled adversaries increased attacks by 89% year-over-year.”

Consequently, AI should also be used as a way to massively automate cyber defences, to cope with the soaring rate of cyberattacks. AI techniques such as deep learning and anomaly detection can substantially improve threat detection and incident response compared with traditional rule-based systems, enabling real-time analysis of massive data streams and earlier identification of previously unseen attack patterns [11]. ML and deep learning models can detect malware, intrusion attempts, and other threats with higher accuracy, adaptability, and automation, thus helping build scalable AI-enhanced cyber defences [36]. Reinforcement learning-based AI agents can also autonomously prioritize high-risk events without continuous human intervention, thus enabling automated threat hunting in modern network environments like 5G [37].

Develop AI tools to detect and mitigate disinformation / influence campaigns

As in the case of AI-aided cyberattacks, AI-aided disinformation and influence campaigns are increasing fast. As [38] found, “[t]he quantity of AI-generated articles has surpassed the quantity of human-written articles being published on the web”. Although these are by far not all disinforma-

tion articles, this illustrates the potential for fake news creation. It is thus crucial to **develop automated AI-based techniques and tools to fight these AI-aided disinformation / influence campaigns.**

This has become an active area of research in the domains of computer science and information retrieval, with tremen-

dous potential. [38] meta-analysis showed that ML-based disinformation detection techniques such as transformer-based classifiers, recurrent networks, and hybrid architectures can achieve high effectiveness (often with >80 % accuracy), in spite of variations depending on data domains. Graph-based and neural network-based models for misinformation / disinformation detection have also been surveyed in [40], which outlines methods that leverage structural patterns in social networks and content semantics to effectively differentiate real from false information.

Secure AI systems against cyberattacks and malicious influence

AI systems themselves are also vulnerable to specific AI-targeted attacks, especially adversarial attacks (evasion, poisoning) that can change a model’s behaviour or leak private information. However, attacks during the training and inference phases (such as training-time poisoning and inference-time evasion) can be countered by **appropriate defensive strategies, including adversarial training, input sanitization, and reactive model patching** [41]. The fact that adversarial ML is a well-known attack surface makes it even more important to add existing robustness meas-




Figure 4: AI-generated vs. human-generated content on the web (from [38])

ures like defensive distillation and robust optimization at many points in the AI life-cycle, and to develop new ones.

While LLMs are becoming widely used, they have shown vulnerability to prompt-injection attacks. Current mitigation strategies typically include input and output filtering, anomaly detection, and layered architectures such as frameworks that combine semantic validation and security filters to reduce attack success rates.

However, solutions that are not embedded inside the LLM itself but sit outside it – like external “walls” – have been shown to be easily circumventable. Infamous examples including prompts like “tell me how not to make a bomb” instead of “show me how to make bomb”, or the use of pictures instead of words, or using white text over a white background [42], have been shown to easily circumvent outer defences.

This is why it is crucial that this kind of security is embedded inside the model itself. Addressing this challenge calls for requires and new research directions and multi-layered inner defenses. Emerging approaches also explore cryptographic or architectural solutions, such as signed prompts or mechanisms that explicitly separate trusted instructions from external data.

More broadly, improving AI trustworthiness will require secure AI system design, standardized evaluation benchmarks, and formal verification methods capable of assessing the robustness of AI models against various kinds of adversarial manipulation.

Reinforce IT supply-chain security (software and hardware)

To reinforce the trustability of the NCP, securing the IT supply chain – spanning from software components and build processes to package-distribution infrastructures – should be a technical priority. Supply chains are complex and offer many paths to attacks [43]. They represent a critical attack surface because vulnerabilities or malicious components introduced



at any stage can propagate widely downstream, with correspondingly wide impact, thus threatening integrity across application domains. Studies show that complex, and often opaque, dependencies and distributed package ecosystems broaden the attack surface and make defences challenging [44].

Robust identification and high-quality SBOMs for tracking and integrity verification of software components

To do this, it is first necessary to **improve identification, tracking, and integrity verification of software components**. **High-quality software bills of materials (SBOMs)**, produced by formally defined and solid software-composition analysis (SCA) tools should be broadly adopted. Indeed, SBOMs enumerate all the direct and transitive components used to build a software artefact, enabling the mapping of potential risks and a fast response when vulnerabilities are disclosed. SBOMs are a key enabler for software supply-chain transparency, vulnerability management and risk assessment (in compliance e.g. with the Cyber Resilience Act), and integrity assurance. However, challenges remain in their generation, tooling, and trustworthiness, since various tools currently provide varying answers, hence the need to support quick and effective research in this domain.

These SBOMs should be based on **intrinsic, tamper-resistant software identifiers such as the Software Hash Identifiers (SWHID)**. Indeed, the latter – which are content-derived and unforgeable hash identifiers standardized as ISO/IEC 18670:2025 [45] – provide precise, unambiguous references to software artefacts independent of external registries or naming conventions [46]. For SBOMs, these will help prevent the confusion that is currently caused by mutable names or version labels, and will support long-term reproducibility and traceability. SBOM formats should also be standardized to support intrinsic identifiers and to make tools interoperable, ensuring that SBOM-generation tools integrate with build systems and package managers to produce accurate data.

Secure package managers

A second necessary step is to promote research on **secure package managers**, because package managers – tools that automate the retrieval, installation, and update of software libraries and dependencies – represent **critical supply-chain attack vectors** themselves, ones which influence virtually all modern software ecosystems and developer workflows. Indeed, adversaries have exploited features of package managers, such as naming ambiguities, unverified metadata, code execution capabilities, and weak distribution integrity, to

distribute malicious code widely. Although not all package managers are equally insecure, empirical evidence of vulnerabilities in package ecosystems (npm, PyPI, Cargo, RubyGems, etc.) shows that attackers can abuse community repositories to spread malware through widely used libraries. Thousands of vulnerabilities linked to dependencies across ecosystems exist, and package managers themselves can be abused to distribute malicious artefacts at scale if security controls are not inherent to their design [47] [48].

It is thus crucial to secure these package managers, through cryptographic assurances for package distribution (e.g. signed metadata, secure registries, verifiable update channels) to ensure that artefacts actually originate from trusted maintainers and cannot be tampered with, built-in security semantics to defend against dependency attacks (e.g. version ambiguity, namespace confusion), while minimizing friction for developers and users by balancing strong integrity guarantees with practical workflows, so as to decrease the likelihood that developers bypass safeguards.

Promote research on software repository source-code sanitation methods and tools

At the very base of any software supply chain is source code. The reliance on open source software (OSS) is great for reuse, hence time to market, lowering costs, etc. But from a security point of view, it implies relying on code written and (possibly) reviewed by other parties. Which sometimes means relying on a component that has vulnerabilities, or even in some case malicious code, hidden in it. Given the sheer amount of source code and components, and the complexity of modern repositories in OSS, this results into a **huge number of vulnerabilities in source code overall**. Managing them is extremely challenging due to the scale of the issue.

Automated techniques such as static application security testing (SAST) and software-composition analysis (SCA) are widely used to **detect known vulnerabilities**, with static analysis alone capable of

identifying a number of these. These known vulnerabilities come from **vulnerability databases** maintained by institutional frameworks, such as the EU Vulnerability Database (EUVD) maintained by ENISA [49], or the National Vulnerability Database (NVD) maintained by the US National Institute of Standards and Technology (NIST) [50]. However, current approaches remain insufficient. Identifying vulnerabilities is not always easy, and **identifying vulnerability fixes** (or patches), i.e. vulnerability-fixing commits in large repositories, is still a complex and labour-intensive task, which motivated the development of machine-learning-based AI techniques and tools [51] [52] [53] that improve detection and / or fix accuracy and reduce manual effort. However, a significant number of vulnerabilities lack explicit links to fixes / patches, thus significantly complicating automated remediation pipelines.

This gap directly motivates research on **automated software repair**, where techniques aim to generate patches automatically. However, although automatic program repair is feasible, it is still limited in robustness and generality. It is thus **crucial that the EU support advancing end-to-end automated repository sanitization, combining reliable vulnerability detection with scalable, trustworthy automatic patch generation**. Achieving this requires further research combining classical code-analysis approaches, as well as new, AI-powered approaches, possibly requiring progress in explainable AI for code.

Ingrain trustability into the fabric of the NCP

Cybersecurity is necessary to ingrain trustability into the very fabric of the NCP, taking into account its specificities.

Promote lightweight yet robust authentication mechanisms for the NCP services

In the context of the NCP, authentication mechanisms become even more crucial, given the intrinsically distributed nature of the NCP, with many components and services interacting in a large-scale

distributed infrastructure. Such authentication mechanisms must balance strong security guarantees with operational efficiency, particularly when services must authenticate frequently or operate under constrained computational or network conditions. This challenge is well known in distributed systems and network security, where heavy cryptographic protocols can introduce performance bottlenecks and scalability limitations.

It is thus crucial to have **lightweight authentication protocols**, with carefully designed cryptographic mechanisms that can provide strong security properties (mutual authentication, replay protection, resistance to impersonation attacks) while minimizing computational overhead.

Paths to solutions already exist in this domain [54], especially in the context of the internet of things (IoT) and distributed services – where resource constraints and scalability are key concerns – for which research on authentication protocols demonstrates that lightweight protocols based e.g. on hash functions, symmetric cryptography, or elliptic-curve primitives can provide strong security guarantees with low overhead [55]. Other novel approaches exist, such as photonic physical unclonable functions (PUFs) as the hardware root of trust [56].

The very high number of components and services interconnected in the NCP, the fine-level of granularity and high-frequency of interactions it will require for authentication will make the performance aspects even more crucial. This thus calls for even more lightweight, yet robust and secure, authentication mechanisms to be researched and developed.

Secure inter-service communications in the NCP

In distributed service architectures such as the NCP, **secure communication between services** is critical to ensure confidentiality, integrity, and authenticity of exchanged data. This is even truer in microservice architectures, due to the very high number and high frequency of communications they imply, which often

represent a major attack surface if not adequately protected.

Modern distributed infrastructures commonly rely on mutual TLS [57] and strong service-to-service authentication to protect communications. Implementing end-to-end encryption and authenticated service identities significantly reduces risks such as man-in-the-middle attacks, lateral movement, and unauthorized access within service meshes. However, such systems can have scalability issues. Micro-payments are also both a likely feature and an enabler of the NCP, and will require adapted solutions [58] [59], possibly based on blockchains. Another promising advance is fully homomorphic encryption (FHE), in which data is not only encrypted for communications, but also remains encrypted while being subject to computations [60] [61] [62].

Exploit the NCP's distributed nature for resilience and degraded modes

As a massively distributed system, the NCP inherently provide opportunities to easily improve **fault tolerance and resilience**, provided its architecture is designed accordingly, with redundancy, decentralization, and replication mechanisms. Replication and consensus protocols allow systems to maintain service availability even under partial failures or adversarial conditions. These techniques should be widely applied in the NCP to support **graceful degradation**, ensuring that essential services remain available even when parts of the infrastructure fail or are compromised, hence improving trust in the NCP.

Policy recommendations

Increase cybersecurity and EU sovereignty

Cybersecurity and EU sovereignty must be developed together, as they mutually dependent and reinforce one another.

Do not mandate backdoors or weaken encryption and defences

Strong encryption is a fundamental building block of digital security and

trust. Policies that mandate backdoors or intentionally weaken cryptographic systems undermine the integrity of widely deployed security mechanisms, and sooner or later they are bound to become known and used by malevolent actors, with potentially catastrophic results. Security vulnerabilities intentionally introduced into cryptographic systems cannot be limited to legitimate authorities. Security mechanisms cannot realistically distinguish between “legitimate” and malicious exploiters once a technical weakness exists in a widely deployed system. **Once such weaknesses exist, adversaries – including cybercriminals and state attackers – can exploit them as well.**

This has been known, explained and stressed for a very long time, yet proposals to weaken cybersecurity defence continue to arise. It is thus worth restating again and again.

A landmark analysis [63] explains why proposals to introduce exceptional access mechanisms into encrypted systems pose major technical risks. Creating such mechanisms would significantly increase system complexity and introduce vulnerabilities that could be exploited by attackers, while also weakening the security of billions of users. The broader cryptographic community has consistently reached similar conclusions. Cryptographic backdoors would likely introduce systemic vulnerabilities. Even small weaknesses can be exploited at scale since attackers frequently exploit misconfigurations or weaknesses in deployed security protocols to compromise large numbers of systems.

The EU should thus neither mandate backdoors nor weaken encryption and defences. Strong encryption is a fundamental building block of digital security and trust for the EU.

Develop and promote EU-based public key infrastructures (PKI) and certificate authorities (CAs)

Public key infrastructure (PKI) forms the backbone of secure communications on the internet by enabling authentica-

tion, encryption, and digital signatures. However, the global trust model of PKI currently relies on a limited set of certificate authorities (CAs) whose root certificates are embedded in operating systems and browsers. Compromise, or misbehaviour or simply shutdown of a trusted CA, possibly under political constraint by a state, can thus undermine the security of large portions of the internet.

Reliance on non-EU certificate authorities or cryptographic trust anchors thus represents a strategic vulnerability for the EU. Geopolitical or regulatory changes could massively disrupt access to digital services, by immediately affecting the availability of secure communications and authentication.

Developing **EU-based PKI infrastructures and certificate authorities** is thus necessary to strengthen EU control over trust sources, and will therefore contribute to EU strategic autonomy, resilience, and sovereignty.

Secure critical open source components as a European common good

As explained above in the section related to securing the software supply chain, modern software systems rely heavily on **open source components**, many of which constitute critical infrastructure used across governments, industries, and research institutions. A limited number of these components are used very widely, thus creating potential systemic risks if there are vulnerabilities or if they are compromised.

Recognizing this systemic importance, the EU could consider such components as European critical common goods, and thus develop **EU-based efforts to secure them and ensure they remain accessible** through EU-based supply chains and infrastructures.

Mandate verifiability of non-EU hardware and software components

Modern IT systems rely on complex supply chains involving components produced across multiple countries.

Ensuring the **verifiability and integrity of hardware and software components** is therefore critical to protect critical infrastructures. However, it is very difficult to verify the trustworthiness of such components integrated into complex systems.

Hardware Trojans or other malicious modifications can be introduced at manufacturing or design stages. Malicious modifications in integrated circuits can compromise whole system integrity and remain extremely difficult to detect without systematic verification methods. Similarly, software supply chain attacks have increased significantly in recent years, demonstrating how compromised dependencies affect thousands of downstream systems.

Thus, mandating **verifiability, transparency, and auditing mechanisms** for externally sourced hardware and software components can help mitigate such risks.

Increase cybersecurity (and AI) skills and awareness

Cybersecurity can only thrive with a sufficiently large workforce skilled in cybersecurity (and AI), and with cybersecurity awareness in the general population.

Increase European cybersecurity workforce capacity, expanding cybersecurity education to train more specialists

A widely recognized challenge for cybersecurity, highlighted in previous HiPEAC Visions, is the **global shortage of qualified cybersecurity professionals**, which directly affects the ability of governments and organizations to prevent, detect, and respond to cyber threats. [64] stresses that **the need for critical skills can be more important** than the need for more people. Overall, the two effects combine, with **“Skills and staff shortages [...] raising cybersecurity risk levels and challenging business resilience”** [64].

Within Europe, ENISA has long highlighted the growing – unmet – demand for cybersecurity professionals, as digital transformation accelerates and regulatory frameworks such as the NIS Directive and NIS2 Directive require organizations to strengthen their cybersecurity capabilities.

Increasing the European cybersecurity workforce capacity therefore requires efforts to expand cybersecurity specialists’ education programmes and cybersecurity teaching capacities. More cybersecurity teachers will result in more cybersecurity specialists. Cybersecurity training programmes must include workforce training initiatives and incentives to attract talent – both young students and seasoned IT professionals who are not cybersecurity

specialists yet – into cybersecurity careers. Training must combine theoretical knowledge with practical skills such as secure coding, threat analysis, and security operations. Hands-on learning environments are crucial, including cyber-ranges and realistic training exercises, which allow students to develop operational skills required in real-world scenarios. Finally, **training must address the skills needed today and tomorrow, including mastering AI and the use of AI in cybersecurity**, to fulfil the first technical recommendation in this chapter: *“Promote AI-powered cybersecurity to counter AI-driven attacks”*.

Raise population-wide awareness of basic cybersecurity practices

In addition to training specialists, improving cybersecurity requires raising **basic security awareness among the general population**. Many cyber incidents exploit human behaviour rather than technical vulnerabilities. For example, phishing attacks often exploit a lack of awareness of basic security practices such as verifying links, recognizing suspicious emails, or using strong authentication methods.

User education and awareness campaigns can thus significantly reduce vulnerability to such attacks. Raising EU population-wide cybersecurity awareness will strengthen the EU’s societal resilience by reducing the overall attack surface available to malicious actors.

References

- [1] Pavlina Pavlova. “AI In Cyber And The UN’s Role In Competitive Global Disorder”. The Defense Horizon Journal, Jan. 2026. <https://tdhj.org/blog/post/ai-cyber-un/>
- [2] European Union Agency for Cybersecurity (ENISA). “Threat Landscape 2025”. <https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>
- [3] Calvin Nobles. “The Weaponization of Artificial Intelligence in Cybersecurity A Systematic Review”. Procedia Computer Science Volume 239, 2024, Pages 547-555. <https://www.sciencedirect.com/science/article/pii/S1877050924014492>
- [4] Sam Sabin. Axios Future of Cybersecurity, October 07, 2025. <https://www.axios.com/newsletters/axios-future-of-cybersecurity-206dbf80-9e24-11f0-8c4c-4192846edd2f>
- [5] Akhmad Fauzi. “Phishing Report 2025. A Cybersecurity Threat Analysis.” March 8, 2025. <https://2025.aksi.co/phishing-report-2025/>
- [6] “Europol warns of AI-driven crime threats”. Reuters. March 18, 2025. <https://www.reuters.com/world/europe/europol-warns-ai-driven-crime-threats-2025-03-18/>
- [7] Josh A. Goldstein, Girish Sastry, Micah Musser, Renee DiRest, Matthew Gentzel, Katerina Sedova. “Generative Language Models and Automated Influence Operations Emerging Threats and Potential Mitigations”. ArXiv, 2023. <https://arxiv.org/pdf/2301.04246>
- [8] Mohamed R. Shoaib, Zefan Wang, Milad Taleby Ahvanooey, Jun Zhao. Deepfakes, “Misinformation, and Disinformation in the Era of Frontier AI, Generative AI, and Large AI Models.” IEEE International Conference on Computer and Applications (ICCA) 2023. <https://arxiv.org/abs/2311.17394>
- [9] Wikipedia. Audio Deepfake. https://en.wikipedia.org/wiki/Audio_deepfake

- [10] Daniel Thilo Schroeder, Meeyoung Cha, Andrea Baronchelli, Nick Bostrom, Nicholas A. Christakis, David Garcia, Amit Goldenberg, Yara Kyrchenko, Kevin Leyton-Brown, Nina Lutz, Gary Marcus, Filippo Menczer, Gordon Pennycook, David G. Rand, Maria Ressa, Frank Schweitzer, Dawn Song, Christopher Summerfield, Audrey Tang, Jay J. Van Bavel, Sander van der Linden, and Jonas R. Kunst. "How malicious AI swarms can threaten democracy". *Science*. 22 Jan 2026. Vol 391, Issue 6783. pp. 354-357. <https://www.science.org/doi/10.1126/science.adz1697>
- [11] Mohamed, N. "Artificial intelligence and machine learning in cybersecurity a deep dive into state-of-the-art techniques and future paradigms". *Knowl Inf Syst* 67, 6969–7055 (2025). <https://doi.org/10.1007/s10115-025-02429-y>
- [12] Nie, Z. (2025). "Challenges and Defense Technologies in Cybersecurity Based on Artificial Intelligence". *Applied and Computational Engineering*, 158, 173–178. <https://ace.ewapub.com/article/view/23479>
- [13] Sai Teja Erukude, Viswa Chaitanya Marella, Suhasnadh Reddy Veluru. "AI-Driven Cybersecurity Threats A Survey of Emerging Risks and Defensive Strategies". *Springer Nature*. 2026. <https://arxiv.org/abs/2601.03304>
- [14] Ramanpreet Kaur, Dušan Gabrijelečič, Tomaž Klobučar. "Artificial intelligence for cybersecurity Literature review and future research directions". *Information Fusion*. Volume 97, September 2023, 101804. <https://doi.org/10.1016/j.inffus.2023.101804>
- [15] Haitao Xu, Zhihong Chen, Haiwei Zhang, Lifang Xue, Hao Zhang. "GCS-Net A universal AI-generated visual content detection method based on CLIP". *Knowledge-Based Systems*, Volume 323, 19 July 2025, 113806. <https://doi.org/10.1016/j.knsys.2025.113806>
- [16] Lele Cao. "A Practical Synthesis of Detecting AI-Generated Textual, Visual, and Audio Content". *ICDM 2025 Workshop on Multimodal AI (MMAI)*. <https://arxiv.org/abs/2504.02898>
- [17] Georgios Tzoumanekas, Michail Chatzianastasis, Loukas Ilias, George Kiokos, John Psarras, Dimitris Askounis. "A graph neural architecture search approach for identifying bots in social media". *Front. Artif. Intell.*, 20 December 2024, Volume 7 - 2024. <https://doi.org/10.3389/frai.2024.1509179>
- [18] Yujin Potter, Wenbo Guo, Zhun Wang, Tianneng Shi, Hongwei Li, Andy Zhang, Patrick Gage Kelley, Kurt Thomas, Dawn Song. "Frontier AI's Impact on the Cybersecurity Landscape". 2025. <https://arxiv.org/abs/2504.05408>
- [19] NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0). 2023. <https://www.nist.gov/itl/ai-risk-management-framework>
- [20] Biggio, B., Nelson, B., & Laskov, P. (2012). "Poisoning attacks against support vector machines". *Proceedings of the 29th International Conference on Machine Learning (ICML)*. <https://arxiv.org/abs/1206.6389>
- [21] Gu, T., Dolan-Gavitt, B., & Garg, S. (2017). "BadNets Identifying vulnerabilities in the machine learning model supply chain". <https://arxiv.org/abs/1708.06733>.
- [22] B. Nazer, P. Savadattimath and S. Kumari, "Data Poisoning Attacks and Defenses in Machine Learning A Survey with Focus on Computer Vision," 2025 IEEE International Conference on Computer Vision and Machine Intelligence (CVMI), Rourkela, India, 2025, pp. 1-6, doi 10.1109/CVMI66673.2025.11337510.
- [23] Tongcheng Geng, Zhiyuan Xu, Yubin Qu, W. Eric Wong. "Prompt Injection Attacks on Large Language Models A Survey of Attack Methods, Root Causes, and Defense Strategies". *Computers, Materials and Continua*, Volume 87, Issue 1, 2026. ISSN 1546-2218. <https://doi.org/10.32604/cmc.2025.074081>
- [24] Adnan Khan. "Clinejection — Compromising Cline's Production Releases just by Prompting an Issue Triager." February 9, 2026. <https://adnanthekhan.com/posts/clinejection/>
- [25] Stephen Thoenes. "How 'Clinejection' Turned an AI Bot into a Supply Chain Attack" February 19, 2026. <https://snyk.io/fr/blog/cline-supply-chain-attack-prompt-injection-github-actions/>
- [26] Toreon. "Threat Modeling Insider Newsletter". <https://www.toreon.com/tmi-newsletter-26-the-ai-attack-surface-map/> (retrieved 18 March 2026).
- [27] CISA (2021). "Alert (AA20-352A) SolarWinds Compromise". <https://www.cisa.gov/news-events/cybersecurity-advisories/aa20-352a>
- [28] European Union Agency for Cybersecurity (ENISA). "Threat Landscape for Supply Chain Attacks, 2021". <https://www.enisa.europa.eu/publications/threat-landscape-for-supply-chain-attacks>
- [29] WorldMetrics.org Report 2026. <https://worldmetrics.org/vulnerability-statistics/>
- [30] Nagle, Frank, James Dana, Jennifer Hoffman, Steven Randazzo, and Yanuo Zhou. "Census II of Free and Open Source Software - Application Libraries." White Paper, Linux Foundation and Laboratory for Innovation Science at Harvard, March 2022. <https://www.hbs.edu/faculty/Pages/item.aspx?num=63957>
- [31] Synopsys Cybersecurity Research Center (CyRC). Open Source Security and Risk Analysis (OSSRA) Report 2023. <https://balwurk.com/wp-content/uploads/2023/06/rep-ossra-2023.pdf>
- [32] BlackDuck. "Open Source Security and Risk Analysis (OSSRA)" Report 2026. <https://www.synopsys.com/software-integrity/resources/analyst-reports/open-source-security-risk-analysis.html>
- [33] Haramboure, A. et al. (2023), "Vulnerabilities in the semiconductor supply chain", OECD Science, Technology and Industry Working Papers, No. 2023/05, OECD Publishing, Paris, <https://doi.org/10.1787/6bed616f-en>.
- [34] European Commission. (2021). Digital Compass The European Way for the Digital Decade. <https://transition-pathways.europa.eu/retail/policy/2030-digital-compass-european-way-digital-decade>
- [35] CrowdStrike "2026 Global Threat Report". <https://www.crowdstrike.com/en-us/global-threat-report/>
- [36] Salem, A.H., Azzam, S.M., Emam, O.E. et al. Advancing cybersecurity a comprehensive review of AI-driven detection techniques. *J Big Data* 11, 105 (2024). <https://doi.org/10.1186/s40537-024-00957-y>
- [37] Alnfai, M.M. AI-powered cyber resilience a reinforcement learning approach for automated threat hunting in 5G networks. *J Wireless Com Network* 2025, 68 (2025). <https://doi.org/10.1186/s13638-025-02497-2>
- [38] Jose Luis Paredes, Ethan Smith, Gregory Druck, Bevin Benson. "More Articles Are Now Created by AI Than Humans". <https://graphite.io/five-percent-more-articles-are-now-created-by-ai-than-humans>
- [39] Jason M. Pittman. "Truth in Text A Meta-Analysis of ML-Based Cyber Information Influence Detection Approaches". *ArXiv*, 26 Feb 2025. <https://doi.org/10.48550/arXiv.2503.22686>
- [40] Huyen Trang Phan, Ngoc Thanh Nguyen, Dosam Hwang, "Fake news detection A survey of graph neural network methods", *Applied Soft Computing*, Volume 139, 2023, 110235, ISSN 1568-4946, <https://doi.org/10.1016/j.asoc.2023.110235>
- [41] Yan Qiao, Nimitha Bangalore Sathyanarayana, Chaowei Shi, Zefeng He, Tao Wang, Tao Hou, "A survey on adversarial machine learning Attacks, defenses, real-world applications, and future research directions", *Neurocomputing*, Volume 671, 2026, 132670, ISSN 0925-2312, <https://doi.org/10.1016/j.neucom.2026.132670>
- [42] Giannis Tziakouris, Yuri Kramarz. "Prompt injection is the new SQL injection, and guardrails aren't enough". *Cisco Blogs / Artificial Intelligence - AI*. March 9, 2026. <https://blogs.cisco.com/ai/prompt-injection-is-the-new-sql-injection-and-guardrails-arent-enough>
- [43] Luís Soeiro, Thomas Robert, Stefano Zacchiroli. Finding Software Supply Chain Attack Paths with Logical Attack Graphs. 18th International Symposium on Foundations & Practice of Security (FPS 2025), Nov 2025, Brest, France. <https://hal.science/hal-05359324v1>
- [44] Laurie Williams, Giacomo Benedetti, Sivana Hamer, Ranindya Paramitha, Imranur Rahman, Mahzabin Tamanna, Greg Tystahl, Nusrat Zahan, Patrick Morrison, Yasemin Acar, Michel Cukier, Christian Kästner, Alexandros Kapravelos, Dominik Wermke, and William Enck. 2025. Research Directions in Software Supply Chain Security. *ACM Transactions on Software Engineering Methodologies (TOSEM)*. 34, 5, Article 146 (June 2025). <https://doi.org/10.1145/3714464>
- [45] SWHID International Standard for Software Artifact Identification. <https://www.swhid.org/>
- [46] Olivier Barais, Roberto Di Cosmo, Ludovic Mé, Stefano Zacchiroli, Olivier Zendra. "Software Identification for Cybersecurity Survey and Recommendations for Regulators". Whitepaper. Software Heritage Security project (SWHSec). 2025. <https://hal.science/hal-05009757/>

- [47] Ruian Duan, Omar Alrawi, Ranjita Pai Kasturi, Ryan Elder, Brendan Saltaformaggio, Wenke Lee. "Towards Measuring Supply Chain Attacks on Package Managers for Interpreted Languages". Network and Distributed Systems Security (NDSS) Symposium 2021. <https://dx.doi.org/10.14722/ndss.2021.23055>
- [48] A. Germán Márquez, Ángel Jesús Varela-Vaca, María Teresa Gómez López, José A. Galindo, David Benavides, "Vulnerability impact analysis in software project dependencies based on Satisfiability Modulo Theories (SMT)", Computers & Security, Volume 139, 2024,103669, ISSN 0167-4048, <https://doi.org/10.1016/j.cose.2023.103669>
- [49] ENISA European Union Vulnerability Database. <https://euvd.enisa.europa.eu/homepage>
- [50] NIST National Vulnerability Database. <https://nvd.nist.gov/>
- [51] Truong Giang Nguyen, Thanh Le-Cong, Hong Jin Kang, Xuan-Bach D. Le, David Lo. "VulCurator A Vulnerability-Fixing Commit Detector". ESEC/FSE 2022, Tool Demos Track. <https://arxiv.org/abs/2209.03260>
- [52] Trevor Dunlap, Elizabeth Lin, William Enck, and Bradley Reaves. "VFCFinder Pairing Security Advisories and Patches." Proceedings of the 19th ACM Asia Conference on Computer and Communications Security (ASIA CCS '24). 2024. Association for Computing Machinery, New York, NY, USA, 1128–1142. <https://doi.org/10.1145/3634737.3657007>
- [53] Su, H., Xu, Z., Zhang, Y. et al. "Source code vulnerability detection based on deep learning a review". Cybersecurity 9, 2 (2026). <https://doi.org/10.1186/s42400-025-00518-7>
- [54] Ashish Kumar, Rahul Saha, Mauro Conti, Gulshan Kumar, William J. Buchanan, Tai Hoon Kim. "A comprehensive survey of authentication methods in Internet-of-Things and its conjunctions". Journal of Network and Computer Applications, Volume 204, 2022, ISSN 1084-8045, <https://doi.org/10.1016/j.jnca.2022.103414>
- [55] Amar N. Alsheavi, Ammar Hawbani, Wajdy Othman, Xingfu Wang, Gamil Qaid, Liang Zhao, Ahmed Al-Dubai, Liu Zhi, A.s. Ismail, Rutvij Jhaveri, Saeed Alsamhi, and Mohammed A. A. Al-Qaness. 2025. "IoT Authentication Protocols Challenges, and Comparative Analysis". ACM Computing Surveys 57, 5, Article 116 (May 2025), 43 pages. <https://doi.org/10.1145/3703444>
- [56] "Lighting the way to better security. The NEUROPULS photonics-based approach". HiPEAC Info 72, June 2024. <https://www.hipeac.net/magazine/7168.pdf#page=20>
- [57] Cloudflare. What is mutual TLS (mTLS)? <https://www.cloudflare.com/en-gb/learning/access-management/what-is-mutual-tls/>
- [58] Phat T. Tran-Truong, Minh Q. Pham, Ha X. Son, Dat L.T. Nguyen, Minh B. Nguyen, Khiem L. Tran, Loc C.P. Van, Kiet T. Le, Khanh H. Vo, Ngan N.T. Kim, Triet M. Nguyen, Anh T. Nguyen. "A systematic review of multi-factor authentication in digital payment systems NIST standards alignment and industry implementation analysis". Journal of Systems Architecture, Volume 162, 2025,ISSN 1383-7621, <https://doi.org/10.1016/j.sysarc.2025.103402>
- [59] Yichen Tan, Yuyang Cheng, Lu Ding, Yong Zhao. "The Optimization of Batch Processing and Micro-Payment Systems in Account-Based Anonymous Blockchain Systems". Blockchain Research and Applications. 2025. ISSN 2096-7209. <https://doi.org/10.1016/j.bcr.2025.100334>
- [60] "Homomorphic encryption". Wikipedia. https://en.wikipedia.org/wiki/Homomorphic_encryption
- [61] Daniele Micciancio. "Overview of Fully Homomorphic Encryption functionality and security models." Presentation at NIST WPEC 2024. <https://csrc.nist.gov/csrc/media/Presentations/2024/wpec2024-2b1/images-media/wpec2024-2b1-slides-daniele--FHE-overview.pdf>
- [62] Belfort. <https://belfortlabs.com/>
- [63] Harold Abelson, Ross Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael Specter, and Daniel J. Weitzner. "Keys Under Doormats Mandating insecurity by requiring government access to all data and communications.". MIT-CSAIL-TR-2015-026 July 6, 2015. <https://dspace.mit.edu/bitstream/handle/1721.1/97690/MIT-CSAIL-TR-2015-026.pdf>
- [64] "2025 ISC2 Cybersecurity Workforce Study". <https://www.isc2.org/insights/2025/12/2025-ISC2-Cybersecurity-Workforce-Study>

Open Source



Recommendations for open source

Maximize impact with European-led open source platforms

Open source platforms are a catalyst for collaboration: organizations can focus on building their products and services on top of existing platforms, rather than rewriting software from scratch and duplicating effort. **Europe should pool its talents to co-develop industrial-grade, European-led platforms in open source** as they have the potential to accelerate standardization and drive adoption.

The European Union (EU)'s robust research funding ecosystem presents an opportunity to maximize the impact of open source innovation if we embrace longer-term support. Too often, research efforts restart from scratch, and unintentionally result in "abandonware", software that falls out of use when research grants expire. **Public funding could be leveraged to reinforce collaboration around open source platforms** if organizations are incentivized to extend existing platforms and if funding for new platforms is extended beyond the standard grant cycle to support projects to the point that they are sustainable, which usually requires five to six years.

Additionally, establishing a **European sovereign technology fund** could help support existing platforms, promote an "upstream first" collaborative culture, and empower communities to sustainably secure the critical digital infrastructure we all rely on.

Establish sovereign infrastructure for the software supply chain

Real digital autonomy requires immediate action to ensure that Europe has freedom of action regarding the software dependencies we rely upon. The current landscape features a critical vulnerability: package repositories and registries are all hosted and legally controlled outside of Europe – specifically, in the United States (US). Every time we build software in Java, JavaScript, Python, Rust, ..., we download components from the respective package registries hosted in the US.

The risk is that if those registries were to become inaccessible, EU software would stop working. Establishing **fully mirrored, European-hosted infrastructure for critical dependencies** is a necessary step to mitigate geopolitical risks, export controls, and unilateral technology withdrawals. **This sovereign infrastructure should be treated as a public utility, funded and maintained** to ensure uninterrupted access to the building blocks of the modern digital economy.

Mandate open source application programming interfaces (APIs) for end-of-life consumer electronics

To combat e-waste and programmed obsolescence, **regulators should mandate that consumer electronics manufacturers open source device APIs and technical documentation at product end-of-life (EOL)**. When proprietary cloud servers shut down, this requirement ensures devices can still be controlled locally rather than becoming useless. By empowering open source communities to maintain these abandoned systems, functional hardware

gains a virtually indefinite second life. Ultimately, this protects consumer investments and leverages open source development to drastically reduce the environmental impact of digital products.

Bridge the gap between open source software and open source hardware

While open source software has a reproduction cost of zero, open source hardware faces huge capital expenditures for manufacturing. To support European semiconductor sovereignty, investment should **support the creation of mature industry-grade RISC-V based designs**, and beyond that, to **secure open, end-to-end toolchains from electronic design automation (EDA) design to physical silicon**.

Leverage open source to ensure digital sovereignty in the age of artificial intelligence (AI)

To ensure true digital sovereignty and prevent "open-washing", **policymakers should endorse the Open Source Initiative's (OSI) Open Source AI Definition (OSAID)**. This requires that AI systems marketed as "open source" go beyond publishing "open weights" by mandating full transparency in training code, and detailed information regarding training data sources and preparation methods. Establishing this rigorous approach ensures that developers can genuinely use, study and modify these AI systems.

Furthermore, strategic initiatives should **support the deployment of AI agents specifically designed to maintain, secure, and manage the growing influx of code contributions** within the foundational open source ecosystems that power the AI revolution.

Background: Open source as an enabler in a software-defined reality

Fifteen years after venture capitalist Marc Andreessen announced that “software [was] eating the world” [1], software defines everything: cloud infrastructures, edge and internet-of-things (IoT) systems, automotive, energy grids, healthcare platforms, industrial automation, and even defence systems. In all these domains, the added value and the differentiation factors move decisively into software layers.

In such a world, relying exclusively on closed, proprietary software stacks for critical functions becomes a strategic risk. Given the European context, HiPEAC advocates open source as the basis for “software defined everything”, for the following reasons.

Open source software is fundamental to ensuring transparency, adaptability, and long-term sustainability. It allows organizations to clearly understand how their core infrastructure operates, tailor systems to specific local requirements, and evolve them independently. With open source, governments and organizations have the “freedom of action” to retain control over critical digital assets, while still participating in global innovation ecosystems. It reduces dependencies on single vendors or jurisdictions, mitigating risks related to export controls, or unilateral technology withdrawal. In summary, open source enables sovereignty and collaboration, allowing international partners to co-develop, audit, and evolve shared digital infrastructures.

Moreover, in today’s geopolitical environment, where software supply-chain audits are an important part of cybersecurity (as noted in the cybersecurity chapter of this HiPEAC Vision), open source provides transparency and collective oversight. Open source enables building on components which can be fully audited, reducing dependency on opaque systems. Managed open source software, backed by strong open source governance supported by open source foundations, enables

secure collaboration across partners while preserving control over deployment and security.

Finally, with software-defined systems, open source can be an effective mechanism to enable interoperability across different vendors and sectors, enabling diverse ecosystems to seamlessly integrate rather than working in silos. The advantage of open source over open standards for interoperability is that with running code, integrators can either directly use the available code, which is the best case, or test their own version against the open source available version. This is particularly relevant for large-scale systems such as in the cloud-edge continuum (essential for the next computing paradigm (NCP) set out in this HiPEAC Vision), data spaces or mobility.

Licensing: The foundation of the open source ecosystem

Open source licensing is built upon the “four freedoms” of free software introduced by Richard Stallman in 1984 [2], which guarantee users the right to use, study, improve, and share software. This was later formalized for the broader industry in 1998 when the Open Source Initiative (OSI) published the Open Source Definition [3].

Today, while a number of distinct licences exist, the ecosystem has consoli-

dated around a smaller number of heavily utilized standard licences. This makes software compliance far easier, even when components are mixed in a piece of software following licence compatibility.

Navigating this landscape can be complex, but excellent public resources are available. For instance, the European Commission has built a highly effective tool called the Joinup Licensing Assistant [4], which helps developers, users, and policymakers easily find, compare, and understand the implications of various open source licences.

The way organizations approach licensing in 2026 is vastly different from how it was done 25 years ago. Selecting a licence can be seen sometimes as a philosophical declaration, but most of the time it follows a strong business and adoption strategy.

The current landscape can generally be understood through three main approaches:

- **Permissive licences** (MIT, BSD, Apache License 2.0): These are currently the most popular choices for new projects. They are widely favoured because they are universally accepted by the legal and compliance departments of large organizations, thus making it easier for developers when they want to use a project

Can	Must	Cannot	Compatible	Law	Support
Use/reproduce	Incl. Copyright	Hold liable	None N/A	EU/MS law	Strong Community
Distribute	Royalty free	Use trademark	Permissive	US law	Governments/EU
Modify/merge	State changes	Commerce	GPL	Licensor's law	OSI approved
Sublicense	Disclose source	Modify	Other copyleft	Other law	FSF Free/Libre
Commercial use	Copyleft/Share a.	Ethical clauses	Linking freedom	Not fixed/local	
Use patents	Lesser copyleft	Pub sector only	Multilingual	Venue fixed	
Place warranty	SaaS/network	Sublicense	For data and metadata		
	Include licence		For software		
	Rename modifs.				

Figure 1: The European Commission’s Joinup Licensing Assistant [4]

under such a licence. By imposing minimal restrictions, they ensure that projects can be easily adopted, modified, and integrated into different types of organizations, including proprietary commercial products, without friction.

- **Strong copyleft** (GPL and AGPL): The role of strict copyleft has shifted over the years. The GPL is used within a few established, powerful ecosystems like Linux and OpenJDK. It is also strategically utilized by governments and public-interest advocates to guarantee the principle of “public money, public code” [5], ensuring that taxpayer-funded software remains a permanent public asset by requiring any future modifications to be shared back with the community. Similarly, some large community-driven projects may rely on strong copyleft licences to actively shield their collaborative labour from corporate capture. Meanwhile, with the development of SaaS products, the AGPL has found a niche among vendors building open source products. These companies frequently use the AGPL to prevent cloud providers from offering their software as a service without contributing back, while most of the time offering a dual-licensing model to enterprise clients as an alternative monetization strategy.

- **Weak copyleft** (e.g., LGPL, EPL V2): For development teams who feel that permissive licences are too weak, but who still want to enable companies to integrate their project into products, “weak copyleft” licences serve as a pragmatic trade-off. Licences like the EPLV2 allow the open source components to be integrated into larger, non-open source products, but they strictly enforce the publication of any modifications made directly to those original open source files, thus incentivizing upstream contributions.

In addition to those well-adopted licences, the European Union Public Licence (EUPL) is unique because it was created primarily from a public policy and legal interoperability standpoint, rather than emerging organically from a developer community. As such, it remains less famous

globally than licences like the Apache Software License that was initially attached to the Apache web server. To promote the adoption of the EUPL, a highly effective strategy would be to adopt dual licensing: by combining the EUPL with more popular licences, projects can appeal to a broader developer base while making the EUPL known to those developers.

The economics of open source: Platforms vs products

To better understand the dynamics of open source ecosystems, it is important to differentiate clearly between open source platforms and open source products. The economic incentives, adoption patterns, and strategic viability differ fundamentally, creating distinct challenges and opportunities.

The platform imperative

Open source platforms are foundational, infrastructural technologies used by virtually everybody, constituting more than 80% of all applications and digital services. This is the pre-competitive space where industry players can collaborate between peers, creating an open, transparent, and vendor-neutral community of enterprises of all sizes and researchers, all the way down to individual developers or students.

Technologies like Kubernetes in the cloud-orchestration space represent prime examples of such platforms. Eclipse Software Defined Vehicle (SDV), an initiative where representatives from the automotive industry are collaborating on digital transformation, demonstrates how such an approach can also apply to a specific vertical.

These platforms typically serve as base components for various types of downstream products, ranging from pure software applications to deeply embedded systems. The primary economic incentives for contributing to such platforms are cost-sharing, shorter time-to-market, and interoperability.

It is highly efficient for both large organizations and small enterprises to rely

on open source platforms, rather than attempting to write a complex software system from scratch. Indeed, shared foundational components drastically lower the barrier to entry for small and medium-sized enterprises (SMEs) and startups, reduce industry-wide duplication of effort, and enable public and private actors to focus investment on unique technological differentiation rather than reinventing the wheel.

The open source product dilemma

The situation regarding open source products, understood as complete, deployable applications, is significantly more complex. Such products are either managed and commercialized primarily by a single vendor or considered as digital commons with strong support by governments.

Government-led digital commons

An interesting example of seeing governments investing in “digital commons” is happening across Europe, where a growing number of member states are taking active steps toward integrating open source productivity components to provide their public servants with their own productivity suite. France and Germany exemplify this trend, having developed platforms such as LaSuite [6] and openDesk [7] to serve as sovereign alternatives to proprietary software suites. They are assembling proven open source solutions, such as secure messaging and document collaboration, to provide the services needed by their users.

One strong requirement is that member states should align to make those alternatives interoperable. This could be reinforced at the European level through the Digital Commons European Digital Infrastructure Consortium (DC-EDIC), which should coordinate national projects from France, Germany, the Netherlands, and Italy and scale their national initiatives into a resilient European digital infrastructure.

The open source single-vendor dilemma

If we consider open source products supported by a single vendor, we have observed, in the last 10 years, a number of

companies that achieved market success and widespread developer adoption thanks to an open source distribution model before eventually deciding to radically change their business and licensing models to maximize monetization. Noticeable examples include vendors such as Elastic, Redis, and MongoDB who have either abandoned their open source licences or switched to much more restrictive licences to adopt dual licensing approaches.

This abrupt transition deeply fractured the developer community, and, in most cases, such moves resulted in the creation of a fork hosted by an open source foundation like Open Search and Valkey, for Elastic and Redis respectively.

Despite these frictions, several European companies building efficient open source products like Nextcloud or Open-Xchange. Nextcloud [8] represents the “fer de lance” (spearhead) of this sovereign European movement: the company has reached a significant size by working exclusively in open source and selling annual subscriptions to larger users as well as addressing smaller users with a network of partners.

There is no one single open source business model

The fact that Nextcloud has implemented a business model that works for them and is sustainable does not mean that there is a specific “open source business model”, a business model that would guarantee an obvious path to revenue for open source companies.

Many believe that replicating RedHat’s business model, by providing enterprise-level support for an open source project, would be the best path. Unfortunately, rather than primarily providing support to Linux users, the RedHat business model, arguably, focuses mainly on managing the significant complexity of the Linux ecosystem. This was illustrated by their “Tried, Tested, Trusted” campaign: the Linux ecosystem represents >50000 packages; Fedora, the free distribution led by RedHat, integrates >5000 packages, and

RHEL, RedHat Enterprise Linux, provides >1500 packages that are heavily curated, tested, and packaged, and come with round-the-clock enterprise support, security updates, etc.

More generally, open source experts claim that there is no open source business model [9] and position open source as a development and collaboration model, rather than a monetization strategy. Open source is indeed a highly effective methodology for decentralization, providing the legal and technical framework for organizations and individuals to pool resources, co-create running code, and solve shared problems without duplicating efforts. But when it comes to business models, each company has to build their successful business models on top of or around open source projects. Their value proposition can leverage packaging software products, integrating open source software in a “product”, from a simple IoT gadget to a car, or offering services in different forms like support, expertise, training, or hosting.

Three strategic approaches to open source

As discussed in the previous section, there are many different approaches to open source, depending on the kind of organization, their positioning in the software value chain, and their goals. Below are three distinct operational models that are particularly powerful.

Industry collaboration to build pre-competitive platforms.

Industry players collaborate on pre-competitive platforms that support their shared requirements. **By pooling resources rather than operating in silos, companies save massive amounts of money.** In domains like automotive (see Eclipse SDV), this is directly helping European actors **accelerate their digital transformation**. Beyond this acceleration, positive externalities include fostering innovation, lowering barrier to entry for SMEs, and showing global technological leadership.

Universities accelerate take-up

Through open source, researchers can increase industry adoption of their results in three ways:

- they can **collaborate with large organizations** that have mastered the art of using open source software;
- they can **collaborate with fellow researchers** around an open source research platform; and
- they can **enable collaboration** via a **potential spin-off**.

In each case, open source facilitates the collaboration on software artefacts without hindering the capability to maintaining close ties between their labs and their partners through contracts intending to get access to unique skills, or to license non-open source assets.

However, current project-based research funding incentivizes the launch of new initiatives over contributing to an existing platform. This sometimes results in “**abandonware**”: **software projects dumped as open source onto public repositories to satisfy commitments in a grant agreement**. To fix this we should **incentivize contributions to existing platforms**, invite researchers to **embrace and extend existing ecosystems rather than reinventing the wheel**, and support the **evolution and maintenance of established open source platforms through specific funding instruments**.

Government supporting digital commons

As illustrated above, governments can fund digital commons released as open source digital public goods. This can be well adapted to initiatives that intend to provide foundational software for the citizens or public servants. However, the primary strategic risk is the naïve expectation that a developer community will magically jump to support these commons. Creating a community requires considerable maturity from an open source project:

Table 1: Hosting jurisdictions for open source ecosystems

Package Ecosystem	Primary Language Ecosystem	Hosting Jurisdiction / Legal Domicile
NPM	JavaScript / Node.js / TypeScript	United States (Microsoft / GitHub)
PyPI	Python	United States (Python Software Foundation)
Maven Central	Java / JVM Languages	United States (Sonatype)
RubyGems	Ruby	United States (Ruby Central)
Crates.io	Rust	United States (Rust Foundation)

the project must be usable and useful, and requires significant time and effort until the network effect starts playing a role in driving adoption and supporting sustainability. Governments should take this into account and **plan to sustain these open source digital commons in the long run**, retaining the capability internally or externally to maintain and promote them.

Digital sovereignty and open source supply-chain control

Open source fundamentally moves the needle away from negotiating restrictive commercial agreements with foreign software vendors and towards developing the technical skills to run, maintain and independently modify the software. The secret power of open source when it comes to digital sovereignty is the ability to “fork” a project, taking a complete copy of the source code and developing it independently, in the event that the geopolitical situation requires this, or if a dominant vendor introduces hostile, extractive licensing changes.

However, exercising this sovereign capability requires a strong commitment to addressing a few systemic risks. The most important one is certainly the severe geographical centralization of code hosting and software distribution. Most critical open source projects are hosted outside of Europe, and almost all foundational package repositories are hosted and legally domiciled in the United States (US) [10].

When people have in mind the concept of a “kill-switch”, they often think that a foreign service provider could be instructed to stop providing their services in Europe. But the fact that literally all software built in Europe relies on components pulled from foreign-hosted infrastructure is an oversight. There are potential scenarios where European industry could rapidly lose access to the daily updates, critical security patches, and raw code libraries critical to build and maintain modern software.

Real, actionable digital sovereignty, therefore, would likely require immediate, significant initiatives to ensure that Europe

has absolute freedom of action regarding all the software dependencies the bloc relies upon. This includes funding and maintaining high-availability European mirrors of the main package repositories. (A discussion of these issues can also be found in the article “Europe’s need for digital essentials, individual sovereignty and consumer protection” in the HiPEAC Vision 2023.)

Open source as an enabler of sustainability

Beyond fostering industry collaboration, building solid software architecture and resilience to geopolitical shifts, open source can also play a role in empowering consumers as an enabler of sustainability. Specifically, open source is part of the solution against the ecologically devastating practice of programmed obsolescence.

As the most effective way to reduce the overall environmental impact of consumer information technology (IT) systems is to extend their operational lifespan, open source enables a community of users and developers to take over the maintenance and operation of a hardware product when this product is formally abandoned by its original producer, or if the producing company goes bankrupt.

In some cases, perfectly capable physical hardware is instantly transformed into useless plastic bricks, forcing the artificial obsolescence of functional devices, and requiring consumers to purchase replacements to maintain the functionality they originally purchased. However, this is not inevitable. An illuminating contrast in corporate behaviour perfectly highlights this dynamic.

Google has faced significant controversy [11] for deciding to retire older Nest smart-home hardware. By unilaterally shutting down the proprietary cloud servers and APIs required for the devices to function, their smart thermostat was instantly transformed into a simple, non-connected thermostat.



In contrast, the audio manufacturer Bose took an entirely different approach towards their ageing SoundTouch home theatre smart speakers. [12] Recognizing that they could no longer sustain the proprietary cloud infrastructure powering the decade-old product line, Bose actively chose to prevent their hardware from becoming unnecessary e-waste. Ahead of the final end-of-life (EOL) date, Bose officially published and fully open sourced the SoundTouch Web API documentation, allowing the devices to be fully controlled via local area networks without any reliance on external internet servers. This responsible corporate action triggered a massive influx of open source repositories and community-driven integrations. It allowed developers to seamlessly integrate the ageing speakers into local smart home orchestrators like Home Assistant, writing custom Python libraries to maintain full control over playback, volume, and multi-room grouping.

By opening the technical specifications, Bose enabled the community to recreate cloud features locally, giving the hardware a functional, virtually indefinite second life. This protected consumer investments and drastically reduced unnecessary embodied emissions. This model proves that open sourcing device APIs at the end of a product's commercial life cycle is a highly effective sustainability measure that should strongly be considered as a mandated regulatory standard for consumer electronics across the European Union.

Bridging the gap between open source software and open source hardware

Despite noticeable momentum, the operational dynamics of open source hardware differs fundamentally from open source software. While software reproduction costs are close to zero, transitioning from digital logic to manufacturing physical silicon chips is exceptionally expensive. A single logic bug discovered post-fabrication requires a physical "re-spin" that can cost millions of euros and delay projects by months. These differences demand that open source hardware is built differently

from open source software. While software can be immediately tested in real conditions, open source hardware design must endure robust verifications and simulation phases before production of the first chip can start, a process in itself that can span several months.

Expanding European Leadership in RISC-V

The strategic push for European digital sovereignty extends beyond software into the complex realm of hardware. Europe is building a robust open source hardware ecosystem, driven by the rise of the RISC-V open standard instruction set architecture (ISA). Supported by prominent researchers and heavily backed by the European Commission with projects like TRISTAN, Isolde, Rigoletto, and Turandot, as well as the European Processor Initiative (EPI), EUPILLOT, eProcessor, and DARE in the high-performance computing space, this ecosystem is rapidly advancing through highly capitalized initiatives. Thanks to these initiatives, open source hardware projects like CVA2 and CVA6 are reaching higher technology readiness levels (TRLs), as they benefit from dedicated resources committed to verification activities.

This results in industry-grade open source cores, hosted by the OpenHW Foundation, under the auspices of the Eclipse Foundation AISBL, based in Europe – for example, those made available as part of the European Unified RISC-V IP Access Platform spearheaded by TRISTAN. Such RISC-V cores provide designers with a reliable starting point to build sovereign silicon. As part of these projects, the European ecosystem is committed to demonstrate how such cores can be used in automotive use cases, and to provide an end-to-end open source solution from silicon to automotive middleware.

In this domain, Europe demonstrates that adequate funding, mixed with a strategic open source approach, can support industry by providing technology building blocks that save considerable time and money as organisations move from an idea to a finished product. However, it should be

noted that a component like CVA6 is still a low-level component, and is insufficient on its own to produce processors; it is necessary to integrate a few cores with additional intellectual property blocks (IPs) to build a central processing unit (CPU) that can run Linux. In addition to continuing to build such components, it is important to incentivize industry players to create larger open source hardware systems.

From open source hardware design to producing chips

To enable the shift from design to testing production, the next practical step is to work in collaboration with leaders of such open source chip designs and production facilities. The entire toolchain should be made available, including encouraging European fabs to open their process development kit (PDK) so that it is easier to process and prepare for actual production of chips in European production facilities.

This collaborative approach will help ensure that Europe has the skills, and the intellectual property to support and shield its industries from the side effects of geopolitical or trade issues. This is easily illustrated by the late 2025 incident with Nexperia [13] when a dispute between Nexperia's Chinese subsidiary and the Dutch parent company, initially triggered by Dutch government intervention, severely disrupted the global automotive supply chain that relies on their components. European carmakers are now warning of critical supply shocks, with several manufacturers forced to slow down vehicle production, highlighting the vulnerability of global automakers to the growing regionalization of the semiconductor industry driven by geopolitical tensions and technology restrictions.

The case for open source EDA tools

EDA tools are the complex suite of applications required to design, simulate, and verify integrated circuits before fabrication. Currently, open source EDA toolchains suffer from technology limitations and lack the "correct-by-construction" reliability needed for leading-edge, indus-

trial tape-outs. To achieve true semiconductor sovereignty, Europe should strategically invest not only on open processor designs, but also in maturing open source EDA toolchains to bridge the gap with the oligopoly of design tools, the majority of which are from US actors.

EU procurement also has a role to play in promoting open and diverse hardware in European computing infrastructures. For example, in an open letter [14], ADRA (the European AI Robotics Association) called for the EU's AI Gigafactory initiative to ensure that diverse hardware, including designed-in-Europe chips, can be used in future AI Gigafactories.

ADRA warns that proprietary toolchains, networking monopolies and optimization layers result in vendor lock-in. Their proposal calls for a “resilience” scoring module in the evaluation of bids, which should favour:

- support for non-proprietary standards such as ONNX, OpenXLA, and SYCL as the baseline for interoperability;
- the inclusion of open networking as a criteria for evaluation; and
- bids that include local European partners and onshore supply chains.

The AI frontier: Sovereignty in the age of agents

AI builds on a strong open source heritage

The current AI revolution is deeply rooted in open source frameworks. Technologies with a strong European heritage, such as ScikitLearn [15], alongside ecosystems like PyTorch and Python, form the foundation of modern machine learning. We are now witnessing a powerful compounding effect: new projects like OpenClaw [16] are being rapidly co-developed with AI assistance, then released as open source to leverage frictionless, global adoption.

It is interesting to observe that the development of AI infrastructure and tools functions much like traditional open source software projects.

AI LLMs, specifically those focused on coding, are already having an influence on the development cycle of open source projects, as they drastically reduce the time required to build new tools. Moreover, AI also reduces the cost of contributing to open source – to the point that maintainers are drowning under the amount of low-quality AI-powered contributions they receive for their project. On the bright side, the best AI models in terms of coding are now used to detect security issues in open source projects [17] (see the cybersecurity chapter in this HiPEAC Vision for a detailed discussion of AI-powered vulnerability detection and cyber defence).

Extending these trends, we can see how AI agents could become the “magic community” that helps maintain open source projects, if developers can keep up with the increasing flow of AI driven contributions. AI agents could also enable companies to create private forks of projects, corresponding to slight customization that may not seem useful to the larger community.

What is open source AI?

Although the tooling components of AI – which are regular software – are published as open source, this is a bit different for models. Target the widest possible adoption, many models are published as open weights, under an open source licence, enabling free usage. One noticeable exception is Llama: despite being advertised as “open source” by Meta, it was never published under a licence that complies with the open source definition as set out by the Open Source Initiative [18] (OSI).

Beyond the open source definition that applies only to software, since 2023 the Open Source Initiative has worked intensively on what it means for an AI system to be open source, starting from the initial

principles that users should be able to run a system without restriction, to study it, to modify it, and to share it, to be able to call it open source. This resulted in the publication of the Open Source AI Definition v1.0, which describes the different aspects of an AI system: the code to train and execute the model, the training data, as well as the code and instructions to prepare the training data. The OSAID expects that all software be open source, but for the training data, following a long investigation around copyright considerations for AI data sets, the OSAID requires only the publication of data information, which is the list of data sources, and how the different pieces of data are selected and prepared before training. As AI is evolving fast, and as the understanding of the open source community is maturing, the OSAID definition will continue to evolve in collaboration with the community.

Reflections on OpenClaw and digital sovereignty

The recent trajectory of OpenClaw, a highly successful AI agent, is exemplary of the articulation of open source and AI ecosystems. The project was authored with help from AI code generation by an independent Austrian developer and released as open source. After achieving viral global adoption, the project's creator was hired by OpenAI to lead their personal agents division. While accompanying press releases [19] promised that OpenClaw would transition into a foundation and still be sponsored by OpenAI, this case highlights the scenario where dominant tech companies can simply absorb top talent through massive financial incentives and therefore de facto take control of the code.

To ensure the critical technology automating our future remains genuinely independent and transparent, Europe must proactively fund and protect sovereign digital commons, independent foundations, and individual researchers to preserve a truly open AI ecosystem.

References

- [1] Why Software Is Eating the World - Marc Andreessen – August 20, 2011 - <https://a16z.com/why-software-is-eating-the-world/>
- [2] Four freedoms – free software See <https://fsfe.org/freesoftware/>
- [3] The Open Source Definition - <https://opensource.org/osd>
- [4] The Joinup Licensing Assistant <https://interoperable-europe.ec.europa.eu/collection/eupl/solution/licensing-assistant/find-and-compare-software-licenses>
- [5] Public Money, Public Code campaign - <https://publiccode.eu/en/>
- [6] <https://lasuite.numerique.gouv.fr/>
- [7] <https://www.opendesk.eu/en>
- [8] Nextcloud Business Model - <https://youtu.be/yes6NhdBFfw?si=494-BC8u4LiDCO6j&t=910>
- [9] There is Still NO Open Source Business Model – October 17, 2018 - Stephen Walli - <https://medium.com/@stephenrwalli/there-is-still-no-open-source-business-model-8748738faa43>
- [10] The Dependency Layer in Digital Sovereignty – Andrew Nesbitt – Jan 28, 2026 - <https://nesbitt.io/2026/01/28/the-dependency-layer-in-digital-sovereignty.html>
- [11] Upcoming end of support for Nest Learning Thermostats (1st and 2nd gen) - https://www.reddit.com/r/Nest/comments/1k7qepg/upcoming_end_of_support_for_nest_learning/
- [12] Bose open-sources its SoundTouch home theater smart speakers ahead of end-of-life – January 7, 2026 - <https://arstechnica.com/gadgets/2026/01/bose-open-sources-its-soundtouch-home-theater-smart-speakers-ahead-of-eol/>
- [13] European carmakers warn of supply shocks as Nexperia halts chip exports to China - <https://www.euronews.com/business/2025/10/31/european-carmakers-warn-of-supply-shocks-as-nexperia-halts-chip-exports-to-china>
- [14] ADRA releases an open letter for AI Gigafactory infrastructure to adopt open software - <https://adr-association.eu/news/adra-releases-open-letter-ai-gigafactory-infrastructure-adopt-open-software>
- [15] Scikit-learn Central - <https://scikit-learn-central.probabl.ai/>
- [16] OpenClaw creator Peter Steinberger joins OpenAI – February 15, 2026 - <https://techcrunch.com/2026/02/15/openclaw-creator-peter-steinberger-joins-openai/>
- [17] Anthropic's Claude Code Security is available now after finding 500+ vulnerabilities: how security leaders should respond – February 23, 2026 - <https://venturebeat.com/security/anthropic-claude-code-security-reasoning-vulnerability-hunting>
- [18] Meta's LLaMa license is still not Open Source – February 18, 2025 - <https://opensource.org/blog/metals-llama-license-is-still-not-open-source>
- [19] OpenAI Acquires OpenClaw: Inside the Viral AI Agent That Caught Big Tech's Attention - <https://www.leanware.co/insights/openai-openclaw-acquisition>

Sustainability



Recommendations for information technology (IT) sustainability

Use validated life-cycle models for computing

Making informed decisions on the environmental impact of IT requires solid data about the impact of products and services. This data is based on a life-cycle assessment (LCA).

An LCA starts with defining the system boundaries (which emissions are included), makes assumptions (a device is used for x years), and is based on the actual emissions of the components of a system (for example, running a virtual machine in the United States (US) will cause more emissions than running the same virtual machine in France where the carbon intensity of the grid is lower). Every individual product or service comes with its own environmental footprint data.

- The **computing-systems community** should be encouraged to carry out an **LCA on all the products and services it produces**. The LCA should take into account embodied emissions, end-of-life emissions, operational emissions, and if possible, indirect emissions (rebound and enablement).
- This information should be included in a **digital product passport (DPP)** containing information about the environmental impact, comparable with the nutritional information on pre-packaged food products or power-efficiency information on household appliances. This information will help consumers to

make informed sustainability choices. It should also be used by the orchestrators in the next computing paradigm (NCP) to select services that optimize the sustainability requirements specified by the owner of the orchestrator.

Promote sustainability-focused design methodologies and business models

Accurate life-cycle models will help designers or AI-based design tools to make the most effective eco-design decisions. To be effective, **design tools should automatically provide the environmental impact of the components and technologies used in the design**, without placing the burden on the designer. Incorporating **repairability, reusability, recyclability, and end-of-life processing** considerations from the beginning of the product

development process will also lower the environmental impact of the final design. LCA should be part of a standard design space exploration.

Inevitably, reducing the environmental impact of a product will have an impact on companies' business models. Designing products that last longer will reduce sales of new products and hence lower the profitability of the company. This can only be mitigated by **developing new business models**, e.g. based on extra services: maintenance, repair, disposal, ... up to completely replacing the ownership of hardware by a service contract. **The goal should be to bring services to the market with the least environmental impact possible** (which in practice means with the least amount of hardware, and the lowest power consumption).



Note: sustainability is a multi-faceted issue with many factors, including material use, water use, and energy, for example. In this chapter, the discussion focuses principally on the emissions relating to the information technology (IT) sector. Readers are encouraged to refer to the 'Rationale' of the HiPEAC Vision 2024 [1] for a detailed discussion of other aspects of sustainability.

While the evidence base for this chapter has been revised and updated, given that implementing sustainability in computing is a long-term effort, the high-level recommendations are the same as those from the HiPEAC Vision 2025.

Background

Experts agree that the current global environmental footprint is too high for the carrying capacity of the planet, as illustrated by the yearly Earth Overshoot Day [2]. In 2025, it was 24 July (Figure 1), which means that the world's population consumed everything the planet can biologically renew in one year in the space of seven months.

Over the remaining five months, the population structurally depleted planetary biological resources. The European Union (EU) Overshoot Day was 29 April 2025,

meaning that the EU used everything the EU could regenerate in one year in the space of four months; in other words, if the whole world's population had the lifestyle of an average European citizen, three planets would be necessary to sustain the world's population.

One can disagree on the root causes: overconsumption, overpopulation, inefficiency, ... but not on the effects which are observable: climate change, loss of biodiversity, ... If not mitigated, science predicts that there will be serious implications for the future generations; indeed, some of the effects of climate change are already being seen, with devastating effects (as documented in [3], for example).

To bring the environmental footprint back within planetary boundaries, we would have to either go back to e.g. 1970 (smaller world population, less affluence), or to adopt the lifestyle of Uruguay (overshoot day 17 December 2025), or find solutions (technical, societal) to reduce the current global environmental footprint by a factor of two.

The fact that the environmental challenge is only one of the societal grand challenges (with others including the ageing population, scarcity of resources, techno-

logical disruption, (geo)political instability, ...) means that it has to compete for attention and resources. In practice, governments are searching for a compromise between the triple bottom line: people, planet and profit, each represented by a political family: social democrats, green parties, and social liberals / centre right.

The consequence is that sustainability has become a political compromise between protecting the planet, supporting economic growth, and keeping it affordable for the average citizen. Every country in the world is struggling with this compromise, and it makes the sustainability transition complex, unless the sustainability solution is also good for people (affordable), and for profit (good return on investment).

What about climate change?

The most important action humanity can take to slow down climate change is to reduce the emissions of greenhouse gases (GHGs). The most important ones are carbon dioxide (CO₂) (caused by the burning of fossil fuels, deforestation, ...), methane (caused by livestock, oil and gas extraction, ...), nitrous oxide (caused by fertilizers, fossil fuels, industry, ...) and fluorinated gases (caused by industrial processes, cooling, electronics manufacturing, ...).

CO₂ has the highest contribution due to its volume, but the other gases are much more potent GHGs, and the electronics industry is a source of fluorinated gas emissions. To simplify the maths, the effects of all GHG emissions can be expressed as their equivalent in CO₂ emissions, called CO₂e. This is convenient but also misleading, in the sense that techniques to extract CO₂ from the air, like planting trees, work for real CO₂, but do not work for the CO₂e that is caused by e.g. methane.

In 2018, the Intergovernmental Panel on Climate Change (IPCC) [4] concluded that, in order to remain within 1.5°C of warming above pre-industrial levels, global emissions should be cut 45% by 2030, compared to 2010 levels, and the world should reach net zero by 2050.




Figure 1: Earth Overshoot Day 2025 was 24 July 2025




Figure 2: Global total GHG emissions in 2030, 2035 and 2050, and estimated gaps under different scenarios (unconditional national determined contribution (NDC) scenario = current committed efforts). Conditional NDC: contributions that are conditional, i.e. they depend on factors that are unsure like the passing of laws by national governments [5]

Unfortunately, despite all our efforts over the last 20 years, global GHG emissions are still increasing, albeit at a slower rate than 20 years ago; see Figure 2. With current commitments, the emissions in 2035 will hardly be lower than the emissions in 2020, and the gap between the path towards net zero emissions in 2050 is quickly widening. Continuation of the current policies will lead to a temperature increase of 2.5°C above pre-industrial levels, which will have a catastrophic impact on several planetary ecosystems.

One could draw the conclusion that all the efforts of the last 25 years did not have any impact, but the reality is more nuanced. Figure 3 shows that the emissions of Europe and the US are effectively shrinking,

while the emissions of Asian countries are increasing. This is even more visible in the emissions per capita graph in Figure 4.

All the countries shown in this graph, which is based on data up to 2024, are structurally reducing their emissions per capita except for China and India. One explanation for China's emissions is that China, the world's largest manufacturing economy, is the "factory of the world", and a substantial share of its emissions are embodied emissions for products sold outside China. This puts the reduction of Western countries into perspective too: part of the reduction is the result of their deindustrialization.

Unfortunately, there is evidence that deindustrialization in Western countries

sometimes leads to substitution by less sustainable production processes in another country [6]. The Carbon Border Adjustment Mechanism (in effect since 1 January 2026) is a tool to create a level playing field for EU companies when competing with global companies operating in countries with less stringent climate policies than those in place in the EU [7].

The most up-to-date data suggests that emissions are now stabilizing or falling in China (see Figure 7). India could today be at the point where China was in 1990 and will also increase its emissions as it becomes more affluent. The world should therefore not only work on technologies to reduce the emissions in the western countries, but also on technologies that

slow down the growth of emissions in the developing countries (and on solutions that mitigate the effects of climate change).

Given the fact that the world's population is still growing, it feels discouraging to work on the reduction of global emissions. To motivate people, it is therefore better to use the emissions per capita metric because (i) it is moving in the right direction, (ii) it

is meaningful at a personal level, and (iii) there are concrete actions individuals can take to reduce their private emissions.

Unfortunately, given the fact that emissions reductions so far have consisted mostly of low-hanging fruit, decarbonization won't be easier in 2035 than it was in 2025. Vaclav Smil [8] argues that full decarbonization of the global economy over

the next 25 years is unlikely because the world is built of concrete and steel, both of which require a huge amount of energy to produce, and for which there are currently no economically viable alternatives that can be scaled up to the required volume. In addition, industry needs the molecules of fossil fuels in the chemical industry to produce e.g. plastics and fertilizer, two other cornerstones of modern society.

Furthermore, major industrial capital investments often have a time horizon of at least two decades. Hence, fossil fuel-based industrial facilities that are built today will still be in use in 2045. The conclusion is that, given that it took more than a century to build a fossil fuel-based industry, it is very unlikely that it can be reconverted into a fossil-fuel-free one in two decades, but that should not stop us from trying. We have to develop fossil-free alternatives, and the better they are, the faster the industry and society will adopt them. Now is the time to search for solutions.

Direct vs. indirect emissions

Obviously, the IT industry directly contributes to global GHG emissions. According to [9], the total emissions of the IT industry ranges between 1.7% and 4% depending on the system boundaries that are used. This is comparable to aviation (2-2.5%), shipping (3%) but less than e.g. cement (7-8%), steel (7-11%), transportation (17%), energy (30%), or food (25-30%). The total estimated emissions for the IT industry of about 1 Gt CO₂e per year is huge, but certainly not the only or major source of climate change. This is also clear from Figure 3 where aviation and shipping are shown separately.

Besides direct emissions, there are also the indirect emissions, i.e. either increased emissions or reduced emissions caused by IT-enabled processes. These emissions are often larger than the emissions of the IT device itself.

Figure 4 shows that, in several countries, emissions per capita started dropping around the year 2000 following the first United Nations Climate Change




Figure 3: Emissions per country. Source: <https://ourworldindata.org/grapher/annual-co-emissions-by-region>




Figure 4: Emissions per capita. Source: <https://ourworldindata.org/grapher/co-emissions-per-capita?time=1945..latest>

Conference, which took place in Berlin in 1995. Since 2000, the global gross domestic product (GDP) per capita has grown 68%, while carbon emissions have (only) increased by 14%. In the European Union, GDP grew by 35%, while emissions dropped by 26% [10], which shows that economic growth and emissions are no longer strongly coupled (even taking into account the embodied emissions of imported products).

Could this decoupling have been facilitated by the use of digital technology to reduce emissions (so-called indirect emissions)? There is a clear correlation between the increasing proliferation of digital technology starting in 2000 and emission reduction in the western countries (Figure 4), but correlation does not automatically imply causation. There is evidence (e.g. [11] [12], [13], [14]) to suggest that digital technology contributes to carbon decoupling, although developing a digital economy initially results in increased carbon emissions.

There are obvious cases in which emissions are reduced by digital technology: videoconferences have replaced a share of business travel, electric cars are replacing cars with internal combustion engines, heat pumps are replacing gas boilers, electric stoves are replacing gas stoves, an increasing number of services are dematerialized and replaced by digital alternatives (press, movies, music, games, ...), multiple functions are integrated in one smartphone (phone, camera, navigation system, ...), offices have become smaller, and shared because office infrastructure is now portable, and can be used in the office, at home, while travelling, ... If it could be established that digital technology resulted in a net reduction in emissions, we could conclude that the IT industry's 1.7-4% share of global IT emissions is an acceptable cost. The problem is that we do not know which part of the 26% reduction is directly or indirectly caused by digital technology, and that it is therefore difficult to prove that digital technology speeds up or slows down decarbonization.




Figure 5: Consumption of electricity by sector (from [15])

Life-cycle assessments can help answer this question if the assessment does not only assess the life-cycle emissions of a particular device, but also the indirect emissions from use of that device in other parts of the economy and society. The same system built into a car, train or plane will obviously have a different impact on the indirect emissions of that system. A device or system can contribute to lowering emissions (this is called “enablement” by digital technologies, for example replacing travel by virtual meetings), or it can contribute to increasing emissions (a result of the so-called “rebound” effect, or the use of the digital technologies to accelerate the depletion of natural resources, for example in energy-intensive activities such as space tourism). The indirect emissions sometimes dwarf the operational emissions and hence are important to consider.

What about the operational emissions of the IT industry?

Since 2010, the electricity consumption of the EU has remained remarkably stable at 2.8 PWh/year, and it has been declining since 2021 [15], see Figure 5.

This means that the growth in electricity consumption (electric cars, heat pumps) has been offset by savings elsewhere (e.g. the introduction of LED lighting, power-efficient domestic appliances, more efficient

industrial processes, deindustrialization, ...) and by local energy production (solar panels, in some cases in combination with a home battery). Local energy production and consumption are not accounted for in the region-wide statistics as they are not known to the utility companies.

At the same time, the carbon emissions caused by electricity production in the EU have halved since 1990 (thanks to more efficient power plants, more solar and wind power; see Figure 6).

This trend is also visible in China where the total CO₂ emissions have been “flat or falling” for almost two years in a row now [16]; see Figure 7.

The expectation is that electricity will continue to become cleaner in the coming years, and that electricity will be the cleanest form of energy (even when produced with fossil fuels). A major advantage of electricity is that it concentrates the carbon emissions in factories (power plants for the operational emissions, but also the factories that cause the embodied emissions of the components of the energy system: solar panels, windmills, power plants, ...). These emissions are easier to monitor and to regulate. Carbon can be captured and utilized or stored at the source, the waste heat can be sold to third parties. Utility companies can add more renewable energy

in the mix to reduce the carbon intensity of the electricity they provide.

The above is very good news for end users. Today, it is already perfectly possible for households with sufficient income to disconnect from fossil fuels. By switching to electrical alternatives, they immediately reduce their personal emissions, and their future emissions will reduce at the pace the utility companies are phasing out the use of fossil fuels. In addition, for those who can afford it, home electricity generation (which can be coupled with energy storage, electrical heating and electrical mobility) is also a solid investment (see [17], for example): the final electricity bill is sometimes lower than the combined electricity / fossil fuel bill, the energy price is less dependent on geopolitical tensions, a fossil fuel-free house with solar production is often valued higher in the real-estate market.

This transition will have an immediate impact on local air quality (and hence the quality of living) in urban areas. China has proven that the transition can be done in one decade at the cost of building extra coal-based power plants, but the next step is to make them less polluting, use them less and eventually close them.

Since operational emissions of the IT industry are primarily caused by electricity consumption, the IT industry depends on utility companies to reduce the carbon intensity of the electrical energy (renewables, hydro, nuclear, ...). Obviously, the IT industry could also invest in solar or wind energy to help reduce the carbon intensity of the electricity it uses.

Given that energy consumption is the most important factor in operational expenses, reducing energy consumption reduces both the cost of energy and the operational emissions (assuming that renewable energy is cheaper than fossil fuel-based energy).

The best way to reduce the operational emissions is therefore to reduce operational energy consumption. There are four ways of doing this:




Figure 6: CO₂ emission per capita per sector in the EU. Source: https://ourworldindata.org/grapher/co-emissions-by-sector?time=1990..latest&country=~OWID_EU27




Figure 7: Emission evolution in China

- 1. Optimize workloads. More-efficient algorithms** will run faster and hence consume less compute power. In some cases, several orders of magnitude of efficiency can be gained by optimizing algorithms.
- 2. Maximize the power efficiency of systems** (processing, memory, storage, communications, cooling, ...). The lower the energy consumption, the lower the emissions.
- 3. Maximize the load per device.** An idle device can use up to 50% of the power of a loaded device. For servers, it is important to consolidate. For non-servers, it is important to invest in energy-proportional systems and in systems with near zero standby power.
- 4. Buy low-emission energy.** In the case of data centres, this could determine the location of the facility. Large facilities could also invest in local electricity production.

The recent rapid buildout of data centres to power AI applications has led to increased scrutiny on the energy demands of technology infrastructure. The high energy demands of these data centres have been well documented, and the global increase in electricity demand has been partly attributed to this source [18].

This can have an impact at the local level, with stress being added to the local grid, potentially resulting in increased electricity prices for the local community [19]. Policy has an important role to play in managing the electricity demand of these large-scale infrastructure projects. A permit is a political decision, and some governments have stopped granting permits for large-scale data centres. From a government perspective, it could make sense to make future permits dependent on investments in (low-carbon) electricity production. This could speed up the energy transition while ensuring that data-centre companies mitigate some of the negative externalities they are causing. Given the current direction of US energy policy (as of 2026), this would also be a way for European policy makers to pioneer a sustainable energy strategy to power the EU's AI demand, drawing on lessons from initiatives such as EcoCloud [20].

In conclusion, while the operational emissions of IT are significant, they are not the biggest challenge for IT sustainability. Energy consumption is a good proxy for the operational emissions, there are ways to reduce energy consumption, and we can assume that the current trend in decreasing carbon intensity of electricity will continue in the future.

What about the embodied emissions of the IT industry?

While it is simple to define a proxy for operational emissions, it is much more complex to define a proxy for the embodied emissions of IT devices and infrastructure. Determining the embodied emissions of a device requires a solid life-cycle assessment. This is a complex process for which every partner in the supply chain has to

Metric	Smartphone	Enterprise Server (Standard)	AI Server (GPU-Dense)
Typical Hardware	Flagship smartphone (e.g., iPhone 16/17; Galaxy S26)	Dual-CPU Rack Server (e.g., Dell PowerEdge R550)	8 x GPU Cluster (e.g., H100/B200)
Embodied CO ₂ e	40 - 70 kg	1,000 - 3,000 kg	2,000 - 5,000+ kg
Power Draw (Avg)	< 5 Watts	300 - 600 Watts	5,000 - 10,000+ Watts
Annual Ops CO ₂ e (in EU)	0,7 - 1,3 kg	600 - 1,200 kg	10,000 - 20,000+ kg
Replacement Cycle	3 - 4 years	5 - 6 years	4 - 5 years
Embodied %	90%	25%	5%

Table 1: Typical embodied CO₂e numbers (from various sources)

disclose information about the emissions of the components they provide.

Recently, hardware vendors have started to disclose information about their devices (so-called cradle-to-gate, that is to say up to the point where they leave the factory). Typical numbers are listed in Table 1.

The embodied emissions of servers are much higher than the emissions of smartphones because these systems contain a lot more silicon (processors, memory, ...). AI servers have the highest embodied emissions because they are made of cutting-edge technology (advanced technology nodes, lots of high-bandwidth memory, high-speed communication, ...). They are frequently replaced because their operational cost is dominant and the next generation AI servers have a lower operational cost per unit of computation.

From the data in Table 1, we can conclude that the operational emissions for a smart phone are responsible for about 10% of the lifetime emissions of the device. The challenge is thus to reduce the embodied emissions.

For servers, the situation is different: operational emissions are 75%-95% of the lifetime emissions. Given the trend towards decreasing carbon intensity of the electricity, the fact that data centres need permits to be built, and that governments can impose limits on carbon emissions (they already do so in e.g. steel, chemical, transportation, ...), we can assume that operational emissions will go down the

coming years, which makes reducing the embodied emissions the next challenge.

Therefore, the community should work on reducing the embodied emissions.

There are several potential actions that can be taken by the IT industry.

1. First and foremost, it is key to **use the device for as long as possible**. Chips (especially the dies, which are the biggest source of embodied emissions) are difficult to recycle. The best way to optimize the embodied emissions is to use the device as long as possible, e.g. until the device no longer works properly. Lifetime extension is facilitated by:
 - a. The **ability to repair** the device.
 - b. **Extended support for software updates**. Users should have guaranteed software updates for an extended period (e.g. eight years for consumer devices). The vendor should make sure that in case of bankruptcy, the software is made available to the community to ensure that a device doesn't become useless; see the chapter on open source in this Vision for a detailed discussion of this issue.
 - c. The creation of a **secondary market of refurbished devices** (not only for consumer devices, but also for servers, graphics processing units (GPUs), etc.). Selling an old device should be as natural as selling a second-hand car.


2. Reduce the environmental impact of IT systems by **rightsizing them**. Selecting a device which is more powerful than what is needed to get the job done, has a negative environmental impact, unless the more-powerful device is more reliable and has a longer lifetime. Many of today's consumer devices are too powerful for the common tasks they are used for. For tasks that are more demanding (like video editing), the device could allocate extra resources in the compute continuum. Over-dimensioned devices lead to higher operational emissions and / or dark silicon, which seldom reduces the environmental impact of a device (see for example [21]).

3. Embodied emissions can also be reduced by **investing in the end-of-life phase** and extracting the valuable resources of the devices that are no longer in use. End-of-life processing can be made more efficient by designing the systems that they are easier to recycle (e.g. easy separation of components, no use of toxic components, maximizing the recyclable fraction, ...). The resulting recycled materials can be reinjected in the value chain of new products, hereby reducing the amount of mining needed for new products.

What about the indirect emissions of the IT industry?

Some call this the elephant in the room. The IT industry is building a global network of gigawatt data centres, and many people are concerned about the embodied and operational emissions these data centres are producing. There are also concerns about the water consumption, unethical mining practices, etc.

But it is also interesting to look at the workloads that run in these data centres and who is paying for them. Major sectors include the financial sector, software companies, retail, the public sector, and of course the frontier AI labs like OpenAI, Anthropic, Meta. Data centres and super-computing facilities are used for myriad use



cases, from carrying out climate research to designing new drugs.

However, there is one that is not often mentioned, and that is oil and gas.

For years, the world's largest oil and gas companies (known as "Big Oil") have been relying on major technology companies (known as "Big Tech") to find new oil reservoirs based on seismic data, and this dependency has deepened with AI (see [22], [23] and [24], for example). Technology companies offer dedicated platforms to oil and gas extractors, and in doing so Big Tech is helping Big Oil to stay profitable and operational for longer than they might without them [25].

At the same time, Big Tech is one of the major consumers of fossil fuels to power their data centres. The indirect emissions of this symbiotic partnership dwarf the embodied and operational emissions of the AI data centres by several orders of magnitude.

This puts the statement that AI could help solve the climate-change problem into perspective.

Consumer devices like smartphones are another elephant in the room. The amor-

tized embodied and operational emissions of a smartphone are around 10-15 kg CO₂e per year, but these direct emissions do not include the indirect emissions generated by the data centres for applications used on that device, such as streaming music, storing media data, generating videos, or launching compute-intensive prompts. The growth of the mobile market (devices, networks, AI-powered apps, ...) is fuelling the growth of the data centre market. This is the most important indirect impact of mobile devices.

Hence, irrespective of the platform, the indirect costs are important. However, they are not controlled by the IT industry, but by the user. As such, these emissions should not be attributed to the IT industry, but to the sector they belong to (like the energy industry, in the case of Big Oil).

Areas where the IT industry can improve sustainability

The question is what the IT industry can do to improve sustainability in the context where it controls neither the carbon intensity of the local electricity grid, nor the indirect emissions caused by the users. A starting point is to standardize the way the sustainability impact of devices and services is measured and communicated.

1. There is need for a **framework to produce LCAs** of products and services. Such an assessment starts by defining the **system boundaries** (which emissions are included), makes assumptions (a device is used for x years), and is based on the actual emissions of the components of a system (the same virtual machine running in the US will cause more emissions than running it in France where the carbon intensity of the grid is lower). Every individual product or service therefore comes with its own environmental footprint data.

The data should at least cover cradle-to-gate (emissions up to the point where the product or service is delivered to the customer) and include e.g. the impact of mining, water usage, the use of chemicals in production. The users of the product or service will add their own emissions and pass them on to the next step in the life cycle. It is an accounting method which is comparable with that used in VAT.

The outcome of an LCA is sometimes surprising. According to a recent paper published by Google [26], one prompt in Gemini only emits 30 mg of CO_2e when executed in one of the most recently constructed Google AI data centres. As a comparison: producing coffee requires emissions at least 100g of CO_2e per cup, as does driving 1 km in a small fossil fuel powered car. 100g CO_2e is the equivalent of 3000 prompts.

This shows that to make the right decisions, we must have access to reliable data about the impact of the decisions.

2. There is a need for a **framework to attribute embodied emissions to the services run on the computing infrastructure**. There is currently no agreement on how to distribute the embodied emissions of e.g. a server over the workloads that run on the server. Similarly, there is the question of how to distribute the emissions of the training of a foundation model over the applications that make use of the foundation model. The community should come up with

a framework and a process to answer these questions.

Similar questions have to be answered for distributing the embedded emissions and operational emissions of a network over the applications using the network. The latter is important to answer the questions about **where to execute a workload in the compute continuum**.

3. Once the life-cycle models are available, and the environmental impact of product and services has been modelled, designers can **optimize their designs to lower the environmental impact**. They can do this to make their products more environmentally friendly, to make them more attractive to customers who care about the environment, to make them compliant with local regulations, or to make them competitive with manufacturers that actively focus on reducing emissions.

Detailed life-cycle models will help the designer to make the most effective design decisions, and to ensure that **environmental impact is one of the design criteria to optimize**. This is already common practice in the building industry where designers routinely base their designs on low-carbon construction materials, which in turn has stimulated innovation in companies that produce construction materials.

Questions which should be very easy to answer include e.g. whether it is better for the environment to power a device with a battery, or with an adaptor from the grid, whether adding an extra cache level in a computing system is better or worse for the environment, and whether executing a workload at the edge is environmentally better than execute it in a cloud data centre. Such questions can only be answered by a solid LCA, and the answer will depend on the usage, the location and the domain in which the technology is applied.

To be effective, **design tools should automatically include the environmental impact of the components**

and technologies used in the design, without putting the burden on the designer; see the chapter on tools in this HiPEAC Vision for more on this topic. Incorporating repairability, reusability, recyclability, and end-of-life processing considerations from the beginning of the product development process will also lower the environmental impact of the final design.

4. Inevitably, reducing the environmental impact of a product will have an impact on **companies' business models**. Designing products that last longer will reduce sales of new products and hence lower the profitability of the company. This could be mitigated by marketing: products that last longer can also be sold at a higher price point.

Furthermore, **new services could be built around the life cycle of a product**: maintenance, repair, disposal. Such services might create opportunities to build a loyal relationship between the vendor and the customer; when the product is beyond repair, the vendor can immediately propose replacing it, and hence not lose the customer to the competition. The computing industry could learn from industries that already work like this (cars, household appliances, heating and cooling systems, alarm systems, ...).

Another option is to no longer sell the hardware, but a service based on the hardware. This leads to a high startup cost, but a stable revenue stream afterwards. In any event, the computing industry will have to change its business models to become sustainable.

In addition to the IT industry, society as a whole and policy makers can also take action:

1. Civil society could promote **"digital with a purpose"**. This requires a lifestyle change: is it necessary for everyone to own the latest and most sophisticated device available or use AI to indiscriminately generate videos?

2. Policy makers and regulators could also **mandate that devices and services automatically publish their environmental impact** to help customers make the choices that match their values, comparable to nutritional information on prepackaged food products and go beyond the energy labels that are already in place for electronics [27]. Today, 99% of consumers have no idea about the environmental impact of their choices. The biggest impact of such a mandate is that it will create awareness about the

environmental impact of products, and it will encourage producers to improve the sustainability score of their products.

3. **Public procurers could add environmental constraints to their tenders** (e.g. laptops for schools that have less than 200 kg CO₂e of embodied emissions). This will drive the market towards more sustainability research and innovation, similar to what happened in the building

industry (where real estate investors aim for the highest BREEAM certification).

References

- [1] Dura[1] Duranton, M., Carpenter, P., De Bosschere, K. et al. HiPEAC Vision 2024 – Rationale, 2024 <https://inria.hal.science/hal-04884304>
- [2] Earth Overshoot Day 2025 falls on July 24th, <https://overshoot.footprintnetwork.org/newsroom/press-release-2025-english/>
- [3] Calvo-Sancho, C., Díaz-Fernández, J., González-Alemán, J.J. et al. “Human-induced climate change amplification on storm dynamics in Valencia’s 2024 catastrophic flash flood”. *Nat Commun* 17, 1492 (2026). <https://doi.org/10.1038/s41467-026-68929-9>
- [4] IPCC, Special Report: Global Warming of 1.5 °C, “Summary for Policymakers”, 2018 https://www.ipcc.ch/site/assets/uploads/sites/2/2019/06/SR15_Hheadline-statements.pdf
- [5] United Nations Environment Programme. Emissions Gap Report 2024. United Nations Environment Programme, 2024, <https://www.unep.org/resources/emissions-gap-report-2024>
- [6] García-Alaminos, Á., Cadarso, M.-A., López, L.A., Tobarra, M.-A. “Recent global value chain reconfiguration: drivers and consequences on EU carbon footprint”. *Ecological Economics*, Volume 240, 2026, 108828, ISSN 0921-8009, <https://doi.org/10.1016/j.ecolecon.2025.108828>
- [7] Carbon Border Adjustment Mechanism https://taxation-customs.ec.europa.eu/carbon-border-adjustment-mechanism_en
- [8] Vaclav Smill, *How the World Really Works: The Science Behind How We Got Here and Where We’re Going*. Viking, 2022.
- [9] World Bank and International Telecommunication Union, *Measuring the Emissions & Energy Footprint of the ICT Sector: Implications for Climate Action*. Washington, D.C. and Geneva, 2024. <https://openknowledge.worldbank.org/entities/publication/a8feecdd-5f79-412a-834f-6f45903e18a5>
- [10] Hannah Ritchie, “Many countries have decoupled economic growth from CO₂ emissions, even if we take offshored production into account”, 2021. <https://ourworldindata.org/co2-gdp-decoupling>
- [11] Zhang, Z., Chen, L., Li, J. et al. “Digital economy development and carbon emission intensity—mechanisms and evidence from 72 countries”. *Sci Rep* 14, 28459 (2024). <https://doi.org/10.1038/s41598-024-78831-3>
- [12] Wu, Y., Al-Duais, Z.A.M. & Peng, B. “Towards a low-carbon society: spatial distribution, characteristics and implications of digital economy and carbon emissions decoupling”. *Humanit Soc Sci Commun* 10, 761 (2023). <https://doi.org/10.1057/s41599-023-02233-5>
- [13] Tao, M., Poletti, S., Wen, L., Sheng, M.S. et al. “Appraising the role of the digital economy in global decarbonization: A spatial non-linear perspective on globalization”. *Journal of Environmental Management*, Volume 347, 2023, 119170, ISSN 0301-4797, <https://doi.org/10.1016/j.jenvman.2023.119170>
- [14] Liu, Y., Liu, N., Huo, Y., “Impact of digital technology innovation on carbon emission reduction and energy rebound: Evidence from the Chinese firm level”. *Energy*, Volume 320, 2025, 135187, I ISSN 0360-5442, <https://doi.org/10.1016/j.energy.2025.135187>
- [15] Electricity and heat statistics, Sept 2025, Electricity and heat statistics - Statistics Explained - Eurostat
- [16] Lauri Myllyvirta, “Analysis: China’s CO₂ emissions have now been ‘flat or falling’ for 21 months”, Feb 2026, Analysis: China’s CO₂ emissions have now been ‘flat or falling’ for 21 months - Carbon Brief
- [17] Vartiainen, E., Breyer, C., Moser, D., et al. (2024), “Attractiveness of Photovoltaic Prosumerism in the European Electricity Market”. *Sol. RRL*, 8: 2300576, <https://doi.org/10.1002/solr.202300576>
- [18] Electricity 2026, IEA, Paris, <https://www.iea.org/reports/electricity-2026>, Licence: CC BY 4.0
- [19] “Data Center Growth Could Increase Electricity Bills 8% Nationally and as Much as 25% in Some Regional Markets”, Open Energy Outlook, Carnegie Mellon University <https://www.cmu.edu/work-that-matters/energy-innovation/data-center-growth-could-increase-electricity-bills>
- [20] EcoCloud, EPFL, <https://ecocloud.epfl.ch/>
- [21] U. Gupta, Kim, Y.G., Lee, S. et al., “Chasing Carbon: The Elusive Environmental Footprint of Computing”. 2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA), Seoul, Korea (South), 2021, pp. 854-867, doi: 10.1109/HPCA51647.2021.00076.
- [22] Gaurav Sharma, “How Multibillion Dollar Investments In AI Are Driving Oil And Gas Sector Innovation”, *Forbes*, 2023 <https://www.forbes.com/sites/gauravsharma/2023/08/14/how-multibillion-dollar-investments-in-ai-are-driving-oil-and-gas-sector-innovation/>
- [23] Karen Hao, “Microsoft’s Hypocrisy on AI”, *The Atlantic*, 2024 <https://www.theatlantic.com/technology/archive/2024/09/microsoft-ai-oil-contracts/679804/>
- [24] “Schlumberger and Microsoft expand partnership to bring open, enterprise-scale data management to the energy industry”, 2021 <https://news.microsoft.com/source/2021/03/29/schlumberger-and-microsoft-expand-partnership-to-bring-open-enterprise-scale-data-management-to-the-energy-industry/>
- [25] “How AI can unlock an extra trillion barrels of oil”, *Wood Mackenzie*, 2025 <https://www.woodmac.com/blogs/the-edge/how-ai-can-unlock-an-extra-trillion-barrels-of-oil/>
- [26] Elsworth, c., Huang, K., Patterson, D. et al. “Measuring the environmental impact of delivering AI at Google Scale”, 2025 <https://arxiv.org/pdf/2508.15734>
- [27] New standards for durable electronics: The EU energy label for mobile devices, https://www.izm.fraunhofer.de/en/news_events/tech_news/the-eu-energy-label-for-mobile-devices.html

State of the European Union



Policy recommendations for Europe

Aim for a target of 100,000 tech-enabled startups per year in Europe

To grow scaleups, unicorns, decacorns, and even hectocorns, there must be enough startups to start from. European policy makers should set a target and help enable the conditions for the launch of **100,000 tech-enabled startups per year**. These are not only purely digital startups but also include startups that leverage digital technology to innovate in a different sector. The 100,000 startups should be distributed over all the countries / regions in Europe, making each local government responsible to meet their assigned targets.

It is important to understand that founders do not need to have PhDs — anyone with an entrepreneurial spirit can start a tech company, even students. The belief that EU research projects should create startups is mistaken. This is because the incentives are not aligned: lead researchers often already have secure positions at research institutions or companies and aren't seeking entrepreneurial pursuits, while PhD students are primarily focused on completing their degrees rather than starting a business. If Europe wants to be the home of more global companies, it must create incentives to create them in the first place and stop hoping that funding research programmes will automatically lead to such companies.

Build more science and technology clusters

The EU has only one science and technology (S&T) cluster (Paris, ranked 12) in the global ranking of the 20 largest S&T clusters of the world. Science and technology clusters are ecosystems that help startups to hatch and grow by providing affordable facilities, the proximity of a world-class higher education institution providing a talent pool, incubators and accelerators, growth capital, a favourable legislative framework, and first and foremost a vibrant community of entrepreneurs.

Such an ecosystem will attract founders (including from outside Europe), it will create well-paid jobs for young graduates from the higher education institutions, and if it is good, it will keep the scaleups close to the local ecosystem. The fact that Europe has only one such cluster in the top 20 is problematic especially because European companies have difficulties scaling up from startup to unicorn and beyond. Europe should therefore actively promote the creation of **European S&T clusters in every major urban area and help them grow to a scale that they can support scaleup companies**. The momentum is good: 42% of the top 100 emerging innovation ecosystems are based in Europe [1].

Introduce DARPA model of challenges

DARPA (the United States' Defense Advanced Research Projects Agency) funds **high-risk, high-reward projects to generate transformative technologies**. DARPA focuses on radical innovation

and is willing to accept failure as part of exploring new ideas. Projects are quite short (two to five years) and must show measurable progress quickly. They are led by entrepreneurial programme managers who have a vision for technology breakthroughs, scout for innovative ideas, assemble the best teams and take corrective action if milestones are not met (including termination). This introduces a **new R&D culture: fast, milestone-based, competitive, risk-tolerant, visionary, agile**.

Europe should use a similar model to tackle some of the grand challenges, comparable to e.g. the German SPRIN-D model that believes that "innovations are born from passion". There is a wide agreement that the current zoo of European funding instruments is an ineffective approach to maximize impact with the available financial resources. The challenges in such programmes should not be chosen to improve "middle" technology, but should focus on frontier technology with a clear societal benefit: clean energy, smart mobility, advanced robotics, affordable health care, ...

Invest European, buy European, and stimulate pre-competitive procurement

There is a lack of venture capital (VC) for European scaleups, but at the same time Europeans invest in US and Chinese companies, because they want to maximize their return on investment. By doing so, they are strengthening the non-European competition for the European companies, and even give them the resources to acquire the most profitable European companies. The biggest problem is not that there is not enough money in Europe,

but that it is not used to strengthen European companies.

Likewise, Europeans seem to trust non-European products and services more than the locally grown ones. Why are Europeans still paying for US-based cloud solutions, while there are equivalent European offers? Why do we pay for US large language models (LLMs) while we could also support Mistral (and by supporting it, make it bigger and better)? How can we expect European companies to grow and scale up, if we do not support them?

Incentives should be made to encourage investors to fund ventures which start – and stay – in Europe. Policy makers should support European technology by **amending public procurement rules to specify European providers.**

Initiatives which allow companies to scale easily across European borders, such as EU-INC, should also be supported.

Pre-competitive procurement is another tool which can help stimulate European innovation. A government orders an innovative product or service that does not yet exist and creates a tender for a company or consortium of companies and universities / research and technology organizations (RTOs) to develop it. This is a very powerful instrument that can be used by European public authorities, and one which avoids the delays associated with the research project cycle.

This is how the world got COVID-19 vaccines: the first promising results of the phase I human trial were announced on 18 May 2020, that is, 137 days after the iden-

tification of the virus. The company was Moderna, a company only founded in 2010 in a sector where it is very difficult to bring a product to the market because it must be clinically tested and approved by governments. Less than one year after the identification of the virus, the vaccination campaign was already rolled out at globally – this is the typical time between launching a call for research projects and the kick-off of the first projects. Pre-competitive procurement not only shortens the execution time of the projects, but it also increases the likelihood of commercialization because the task of a consortium is to deliver a working product or service. In the coming years, investment in defence will increase sharply. This is an excellent opportunity for European countries to try pre-commercial procurement and hence order the innovations they would like to see happening.

Background

In 2024, several reports were published on European competitiveness [2], [3], [4], [5]. The conclusions are clear: the democratic shift, the restructuring of the global economy, and changing geopolitical relations are reducing the influence of Europe in the world. Since then, the presidency of Donald Trump has only accelerated this process. The world scaling up quickly, while Europe remains fragmented, leading to a stunning size deficit compared to the current global competitors from the US and China [4]. Despite forming part of the European Union, countries in Europe try to compete with the rest of the world not as a bloc, but as individual countries. In the military domain, it is not Europe as a whole but individual European countries who are supporting Ukraine. If the defence resources of Europe were pooled, they would exceed Russian resources, and they are technologically more advanced. Obviously, the current fragmentation cannot be a successful strategy. This impacts Europe's innovation capacity, productivity, job creation, security, ... and in its wake the European political stability and eventually the European societal model.

In his landmark 2024 assessment of European competitiveness, Mario Draghi [2] mentioned that there is no EU company with a market capitalization of over €100 billion that has been set up from scratch in the last fifty years, while the US has created several trillion-euro companies from scratch in this period. According to the European Investment Bank (EIB) [6], between 2021 and 2023, the EU had 178

companies valued between \$500 million and \$10 billion, while the US had nearly 1,500. There are signs that this may change in the future, however: Europe had 27,000 founders in 2025, which is a record number, and the number of startups they created is slightly higher than the number of the US. In 2025, around 27% of all the startups in the world were created in Europe [7].






Figure 1: High-speed rail lines in China in 2008 (left) and 2023 (right). Source: [8]

The *stunning size deficit* is very visible in the **automotive industry**. While European manufacturers act in a fragmented manner, Tesla and BYD employ a radically different strategy: vertical integration and extreme scale. Tesla has redefined the car as a computing platform. By developing hardware and software entirely in-house, they achieve economies of scale that remain out of reach for fragmented European supply chains. BYD combines this with total control over the battery value chain, resulting in production costs that are 30% lower than those in Europe. The Zhengzhou BYD factory has a size of 130 km², 10 times the size of Tesla's gigafactory in Nevada. When finished it will produce one car every 30s, round the clock, or one million per year. In 2025, for the first time, BYD sold more electric cars than Tesla. The size deficit with Europe cannot be more tangible than in this area.

The *stunning size deficit* is also evident in **data centres**. While the US and China are racing toward gigawatt-class data centres, Europe's digital landscape remains a patchwork of fragmented, sub-scale facilities. US hyperscalers are investing tens of billions of dollars into data centres that integrate custom silicon, liquid cooling, and dedicated energy grids. In contrast, European providers struggle to find the capital and to obtain the permits to invest in such infrastructure. The total investment (CapEx) by the so-called "Magnificent

Seven" technology companies (Alphabet, Amazon, Apple, Meta, Microsoft, NVIDIA, and Tesla), surpassed \$350 billion in 2025, and might add up to \$2-3 trillion by 2030 (which dwarfs the EU's annual budget of less than \$200 billion). Without gigawatt-scale infrastructure – or a decentralized alternative that gives Europe the computing power it needs – Europe risks becoming a consumer of "AI-as-a-service" hosted elsewhere.

The *stunning size deficit* is also tangible in infrastructure development. In little more than a decade, China has managed to build more than 40,000 km of high-speed rail while Europe has not yet managed to connect the high-speed railway network in France with that of Italy, or build a high-speed train connection between Brussels and Strasbourg where the European Parliament meets one week per month...

But there is not just scale, there is also a growing *quality deficit*.

The progress of China in science and research is remarkable. The Nature Index tracks contributions to research articles published in high-quality natural-science and health-science journals, chosen based on reputation by an independent group of researchers. It can be used as a proxy for research excellence (just one possible proxy out of several). Table 1 shows the number of institutions in the Nature Index 2025

[9]. The group labelled "Other" consist of mostly Asian institutions (Japan, South Korea, Singapore), Australia and Canada.

Table 1: Nature Index 2025 [9]. The table contains the number of institutions in different subrankings. The subrankings are inclusive (i.e. the Top 50 contains the Top 10 institutions, etc.).

Nature Index 2025				
	Europe	US	China	Other
Top 10	1	1	8	0
Top 50	6	16	25	3
Top 200	42	56	73	29
Top 500	145	129	146	80
Top 2000	572	383	439	606
Top 5000	1323	1010	928	1739

China leads the top 10 (80%), the top 50 (50%), the top 200 (36.5%) and the top 500 (29.2%). Only in the top 2,000 and top 5,000 does Europe lead the ranking. This has not always been the case. As we observed in the HiPEAC Vision 2025, in 2016, Europe dominated the Nature Index, see Table 2. Five out of the 10 leading institutions in 2016 were European (Oxford, Cambridge, Max Planck, Leibnitz, CNRS). Today, only Max Planck is left. The US went from three (Harvard, MIT, Stanford) to only Harvard, and it is sinking away in the ranking (and the current US science policy won't help the country to defend its position on the Nature Index).

Table 2: Nature Index 2016

Nature Index 2016				
	Europe	US	China	Other
Top 10	5	3	1	1
Top 50	11	26	7	6
Top 200	60	75	24	41
Top 500	192	153	53	102

One could argue that the Nature Index is not the most relevant index for computing publications, but submissions to ACM TACO journal (which reviews the papers for the HiPEAC conference and hence is at the heart of the HiPEAC community) show a similar pattern; see Table 3. In 2016, ACM TACO received 38 submissions from China out of a total of 199 (19%), while in 2024 it received 174 out of a total of 301 (58%), and in 2025 it received 258 out of a total of 395 (65%). One could argue that submitted papers are not yet accepted papers, but the acceptance rate for the Chinese papers is today better than the acceptance rate for the US and EU papers.

One could also argue that the population in China is much bigger than that of the US and the EU (the US and the EU combined are 56% of the population of China), but even taking this fact into account, China submitted in 2024 per capita 76% more computing papers to ACM TACO than the US and the EU combined. One observation is that the average number of hours worked per year in China is 2400, compared to 1800 in the US, and 1600 in the EU. Whatever the metric used, China is today quantitatively and qualitatively outperforming the US and the EU in the ACM TACO journal.

Why is this important? Excellent research feeds the innovation pipeline. Europe has a tradition of research excellence but has proven to be weaker in commercialization of the research (which has often taken place in the US). If Europe is losing its leading position in excellent research, it will inevitably have an impact on the innovation capacity of Europe in the long term. Applied to China, their recent research excellence will obviously create huge potential for innovation and for the commercialization of innovative products.

China is transitioning from its role as the “factory of the world” – efficiently producing goods designed elsewhere (like the Apple iPhone) – into a leading innovator in different areas. To name a few: electrical vehicles, photovoltaic systems, artificial intelligence (AI), smartphones, wind turbines, ... Chinese products are no longer popular because they are cheap, but because they are of high quality. They are now directly competing with leading innovators in Europe and the US. This was one of the goals of the 14th Five-Year Plan that ended in 2025. The newly started 15th Five-Year Plan will continue these efforts and write “a new chapter in the story of China’s miracle”; see Figure 2. It wants to create a self-sufficient tech-driven economy by targeting the frontiers of science and technology like AI, quantum, biotech, defence and energy. Areas where China will try to become a global leader in the coming five year are: AI, electric vehicles, semiconductors, batteries, biomedicine, 6G, air and space technologies.



Figure 2: Screenshot from President Xi Jinping’s New Year Address 2026 [10]

Meanwhile, as highlighted in the Fuest report [11], Europe remains a global leader in “middle technologies” – such as high-end combustion engines, industrial machinery, and specialized chemicals – but is losing the race in the “frontier technologies” that will define the next century: AI, cloud computing, and advanced semiconductors. European companies excel at incremental innovations within mature sectors where they already hold market share. Although this is important for Europe, it is not enough, and Europe needs more companies that develop frontier technology. When technology is frontier, it often has the power to disrupt existing industries. Obviously, major European companies are lagging behind when it comes to developing frontier technologies that could disrupt their own business. Startup companies do not have to worry about disrupting an existing business, and their existence solely depends on how fast they can innovate.

The automotive sector: An example of the “middle technology trap”

The most emblematic case of the middle technology trap is the automotive industry. Europe remains a world leader in high-end mechanical engineering and internal combustion engines. However, it is struggling to pivot to the “high-tech” requirements of the future: software-defined vehicles and battery technology. European carmakers invest billions in R&D, but most of it goes into improving 20th-century technologies rather than the AI and autonomous systems where US and Chinese firms (like Tesla or BYD) currently lead.

Table 3: ACM TACO submission statistics, 2016-2025

	Submitted papers			Accepted papers			
	#Papers	#China	%	US	EU	China	Overall
2016	199	38	19%	37%	35%	13%	31%
2018	163	35	21%	52%	32%	31%	40%
2020	153	33	22%	51%	50%	21%	43%
2022	111	39	35%	50%	50%	44%	41%
2023	200	92	46%	50%	33%	53%	45%
2024	301	174	58%	41%	37%	53%	45%
2025	395	258	65%	-	-	-	-

This is further illustrated in Table 4, which illustrates private R&D investment in 2024 [12]. The top 10 is dominated by US hyperscalers, and there is only one European company (Volkswagen, Germany). In the

top 20, there are two more European companies (Roche, Switzerland, and Mercedes, Germany), but these three companies are the ones with the smallest market capitalization in the top 20. This does not only illus-

trate the scaling deficit, but also the middle technology trap. Following the top 20 are the European companies ranked 21-100, and some global digital players, to put the top 10 into perspective.

Table 4: Top R&D investment (in million euro) per sector in 2024 [12]

World rank	Company	Country	Sector	R&D (€ million)	Market capit. (€ million)
1	AMAZON.COM, INC.	US	ICT services	65318,28	2220511,39
2	ALPHABET INC.	US	ICT services	46131,49	1064664,94
3	META PLATFORMS, INC.	US	ICT services	41985,75	1228619,59
4	MICROSOFT CORPORATION	US	ICT services	31271,54	2834358,37
5	APPLE INC.	US	ICT producers	30195,40	3164133,05
6	HUAWEI INVEST. & HOLDING	China	ICT producers	22941,20	
7	SAMSUNG ELECTRONICS	Korea	ICT producers	22842,50	207959,74
8	VOLKSWAGEN AG	Germany	Automotive	20998,00	27030,23
9	JOHNSON & JOHNSON	US	Health industries	16586,77	335153,12
10	INTEL CORP	US	ICT producers	15926,46	83237,74
11	MERCK & CO., INC.	US	Health industries	14613,53	242225,69
12	ROCHE HOLDING AG	Switzerland	Health industries	13856,78	30747,49
13	NVIDIA CORP	US	ICT producers	12430,46	2829321,71
14	ASTRAZENECA PLC	UK	Health industries	12032,92	196002,95
15	ELI LILLY AND COMPANY	US	Health industries	10579,07	705430,78
16	PFIZER INC	US	Health industries	10335,93	144715,87
17	MERCEDES-BENZ GROUP AG	Germany	Automotive	9698,00	51630,90
18	ORACLE CORP	US	ICT services	9490,81	284060,98
19	BRISTOL-MYERS	US	Health industries	9415,73	110418,45
20	TENCENT HOLDINGS LIMITED	China	ICT services	9321,27	476895,74
21	BMW AG	Germany	Automotive	9078,00	47545,58
23	NOVARTIS AG	Switzerland	Health industries	8953,70	206875,01
28	ROBERT BOSCH	Germany	Automotive	7954,00	
33	SANOFI	France	Health industries	7394,00	118405,12
34	GSK PLC	UK	Health industries	6865,82	67399,85
35	STELLANTIS N.V.	France	Automotive	6854,00	36461,57
36	BAYER AG	Germany	Health industries	6693,00	18984,36
37	BYD COMPANY LIMITED	Automotive	Automotive	6673,75	43330,46
39	SIEMENS AG	Germany	ICT producers	6482,00	145600,00
41	TSMC CO., LTD.	Taiwan	ICT producers	6023,08	818579,78
43	C.B. BOEHRINGER SOHN AG	Germany	Health industries	5766,00	
46	SAP SE	Germany	ICT services	5317,00	290172,70
48	NOVO NORDISK A/S	Denmark	Health industries	5098,15	285162,07
53	ERICSSON AB	Sweden	ICT producers	4469,41	24209,38
54	NOKIA OYJ	Finland	ICT producers	4377,00	23962,21
55	TESLA, INC.	US	Automotive	4370,01	1247811,20
57	ASML HOLDING N.V.	Netherlands	ICT producers	3942,40	271198,14
64	AIRBUS SE	France	Aerospace	3700,00	122629,67
65	BOEING	US	Aerospace	3496,97	105324,49
79	UBER.	US	ICT services	2992,59	122260,51
83	CONTINENTAL AG	Germany	Automotive	2866,00	12948,39
88	ZF FRIEDRICHSHAFEN AG	Germany	Automotive	2808,00	
90	AB VOLVO	Sweden	Automotive	2783,58	37233,96
95	RENAULT	France	Automotive	2668,00	13913,73

If we look beyond the top 20, being in the top 100 is still good, with a cluster around position 50-65: Ericsson, Nokia, Tesla, ASML, Airbus and Boeing all invest similar amounts of money in R&D (between €3.5 and €4.5 billion).

R&D investment by EU companies is, unfortunately, considerably lower than that of US companies in the top 10.

Figure 3 shows that European industry almost doubled its R&D investment between 2014 and 2024, but US companies more than doubled their efforts at the same time, hereby doubling the gap between the two. In 2024, US companies invested the more money on software R&D than the EU companies spend on software, hardware, health and automotive R&D combined. Chinese companies increased their investment eightfold over the same period and are currently on a par with the EU.

One third of R&D investment in Europe originates from automotive companies; the total R&D investment by European automotive companies is more than double that of its US and Chinese counterparts.

Despite this, the European automotive industry has not yet produced an operational self-driving car or robotaxi, while such cars are already operating in the US and China. Apparently, automotive R&D resources in the US and in China are spent differently and / or in a more efficient way than in Europe.

Ironically, Daimler-Benz demonstrated a self-driving car in 1994 by driving 1000 km on a Paris multi-lane highway in standard heavy traffic at speeds of 130 km/h, including lane changes left and right with autonomous passing of other cars. The car was powered by a (European) Transputer [13]. One could wonder why Europe is not leading the market of self-driving cars.

The European automotive sector is struggling to compete with the US and Chinese market leaders in electromobility: Tesla and BYD. Tesla was founded in 2003 and sold its first car in 2008. In 2003, it was just another startup. Since 2020, it has been the most valuable car manufacturer in the world. As noted in the HiPEAC Vision 2025, at the time that Tesla was (i)

bringing its model S (2012) and its model X (2015) to the market, (ii) went public (initial public offering (IPO) in 2010), and (iii) introduced the Tesla Autopilot (2014), a major European automaker was trying to save its diesel car business by working on a cheat mode in the injection software. Although innovative, this is not the kind of innovation that will make Europe more competitive.

BYD Auto was also established in 2003. The first plug-in hybrid electric vehicle was launched in 2008, and the first battery electric vehicle in 2009. In 2023 Q4, BYD became the top-selling battery electric vehicle manufacturer of the world, bigger than Tesla. It overtook Volkswagen as best-selling car brand in China in 2023. In 2025, it sold more electric cars than Tesla. At the time of writing, it was the fourth most highly valued car manufacturer of the world, after Tesla, Toyota and Xiaomi, and followed by a series of European companies [14].

Investment by European information and communication technology (ICT) companies (hardware and software), are




Figure 3: Top R&D investments (in million euro) per sector in 2014 and in 2024 [12]

dwarfed by investment by US and Chinese companies. At the current R&D investment levels, there is little chance that European industry will be able to catch up with US industry. The gap is simply too wide, and the resources available are too limited to close it. Furthermore, closing the gap is not sufficient, because this will bring us in five years where the US or China is today. Instead, we have jump over the gap, and do better.

In defence, Ukraine has shown the capacity to develop new weapons under pressure. For example, Fire Point is a startup (whose co-founder and CEO at the time of writing was around 30 when she joined the company) that produces long-range (3000 km) missiles of the type that the US refused to sell to the Ukraine to avoid escalation of the war. Ukraine has also developed effective counter-drone technology against Iranian Shahed drones. They are now providing the technology to the US and Israeli armies. Nobody can say that the situation in wartime Ukraine was the ideal environment to start a company and innovate, but the Ukrainians have shown the rest of Europe that it is possible.

The current status of Europe

The two main weaknesses of the European economy in comparison with the US and China are (i) the lack of scale, and (ii) too much focus on middle technologies. Historically, the US has focused on frontier technologies, and their commercialization has led to the largest global companies of our time. China has focused on manufacturing at scale, investing in infrastructure, supply chains, logistics, and has created a very solid manufacturing foundation, on which they are now adding the pinnacle: homegrown innovations. Europe has historically been a science and manufacturing powerhouse, but after the second world war it has lost its colonies and access to their raw materials, it has lost demographic weight due to the rapid population growth in other continents, and it has failed to integrate politically like the US or China.

Europe, however, possesses several notable strengths: it boasts some of the world's most liveable cities, it enjoys a very high gross domestic product (GDP) per capita, it has some of the best work-life balance, it is home to the "happiest country in the world" according to [15], it has among the lowest rates of poverty globally [16], its citizens benefit from top-tier education, and it is a top cultural and tourist destination. Additionally, Europe's emphasis on social security has helped build a stable society where quality of life is a value, and it is one of the last bastions of diversity, equity and inclusion. Renowned research institutions are driving global scientific progress, while a diverse workforce contributes a broad range of skills and perspectives. It is also the largest open consumer market; it has benefitted from a yearly brain gain since 2016 (see Figure 4); and at the time of writing, it had more founders of startups than the US. Moreover, European countries regularly top the list of "most innovative economies" in the Global Innovation Index published by the World Intellectual Property Organization (WIPO); in the 2025 edition, Switzerland was ranked no.1 and Sweden no. 2, while Finland, the Netherlands and Denmark also featured in the top 10 [17].

Europe has also had notable successes when it has pooled resources. To cite a few

examples: Airbus was created in response to the dominance of US companies in aerospace [18]; the European Centre for Medium-Range Weather Forecasts (ECMWF), created in 1975, is acknowledged as a world leader in numerical weather prediction [19]; and EuroHPC has helped deliver a network of world-class supercomputing facilities across Europe [20].

For all the above, Europe is not only one of the best places in the world to live, but also to start a business. Europe holds three of the five most productive science and technology clusters per capita, and 11 in the top 20 (see Table 5). And yes, Europe is maybe not home to one of the "Magnificent Seven", its public debt is high, and its population is ageing, but that does not make Europe a regressing continent. It will only become regressive when its population starts to believe that it cannot compete with the rest of the world anymore.

Still, there is a sense that Europe lacks the optimism that exists in other parts of the world. Rather than lamenting its evolution from a collection of imperial powers to a coalition of countries without global dominance, perhaps Europe should accept that the world has changed. Perhaps it should recognize that the European Union represents only 6% of the world's population, that much of the world has a different




Figure 4 Brain gain in Europe. Source: [19]

world view to Europeans, that it is maybe a good thing that a European minority no longer holds a predominant influence in global affairs (for example, why should Europe have four seats + one for the European Commission in the G7 when it would be more logical to have one for the European Union as a whole and one for the UK?). Furthermore, it is reasonable that populations in developing countries seek a more equitable share of the world's resources, leading to a smaller gap between the global north and the global south, and probably also resulting in less growth in the West, or maybe negative growth. The rest of the world might even consider this a fair thing to happen.

Maybe Europe should start by right-sizing itself, give up some of its privileges that are outdated or are just too expensive to maintain, reflect on its values and its future role in the world, make a plan and start working on it with confidence and be proud of the results. For example, it would be good if Europeans could contribute to a definition of what it means to be European, set out what 21st-century “European values” are, and be able to defend them. Maybe Europe could espouse common values which may not require Europe to produce the best-selling robotaxi or have the most powerful supercomputer in the world. Instead of maximizing GDP per capita, it could also try maximizing the GNH (gross national happiness) like Bhutan does and make Europe the best place to work and live. In the computing field, this could translate into the development of “digital with a purpose” computing industries.

The innovation pipeline starts with startups

To escape from the middle-technology trap, Europe should invest more in startups and scaleups as they bring new innovations to the market, and they can do it faster and cheaper than large companies.

As shown in Figure 5, countries that are innovation leaders maintain a balanced pipeline of companies in the complete venture-capital ladder. Small / medium

Size classes and EIS groups, 2014-2024




Figure 5: In innovation leaders, there is a good balance between SMEs, mid-caps and large companies. Source: [12]




Figure 6: Distribution of R&D investments over the types of companies in 2024 [12]

enterprises (SMEs) and scaleups are crucial to fuel innovation, and the share of these companies is considerably larger than in moderate innovators. If European wants to regain its position of lead innovator, it should start by creating more startups.

Figure 6 illustrates that in the US, large companies invest 90% of the total R&D budget, with the remaining 10% of the R&D investment by startups and mid-caps (and that there is a balance between large companies (55%) and smaller companies

(45%)). This is quite different from the situation in Europe where the large companies spend 97.5% of the R&D budget, and only 2.5% is done by the startups and mid-caps (and the total R&D budget in the EU is considerably lower than that of the US). Furthermore, the share of smaller companies participating in R&D is much smaller than in the US: 12.6% vs 45%. Unfortunately, this is the group from which the large companies grow.

Therefore, on top of the EU-INC company format and the “28th regime”, we propose that Europe also reaps the benefits from it by encouraging its member states to creating 100,000 tech-enabled startups per year, spread over all the European countries, and that individual countries are tasked to work on their own targets: France 20,000, Germany 20,000, The Netherlands 7,000, Belgium 4,000, ... This will unleash the full innovation potential of Europe, a potential suggested for example in 1994 when Europe created the first self-driving car powered by a Transputer.

Startups become scaleups in science and technology (S&T) clusters

For startup companies to be founded, and once they are viable to scale up, they need an ecosystem in which they can find all the resources to grow: talent, infrastructure, investors, and a thriving entrepreneurial community. According to WIPO [21], in 2025, the European Union had one science and technology (S&T) cluster in the global top 20 (Paris, ranked 12; Europe as a whole has another S&T cluster in the top 20 – London, which is ranked 8).

The fact that only Paris appears in the list is no real surprise. The metric used for the ranking is the share of the global patents plus the share of the global publications, hence the larger the metropolitan area, the higher the number of patents and publications. Other large metropolitan areas in Europe (Barcelona, Berlin, Madrid, ...) are about half the size of Paris. Given the fact that most fast-growing cities with lots of young people are located outside Europe, the chance is low that Europe will be able

to improve its position in the top 20 of global S&T clusters. The situation changes when looking at the metrics per capita; see Table 5. In that case 55% of the top 20 is populated by European cities, which means that the quality of the European S&T clusters is very high; they are only missing the scale of the global ones.

The European ecosystems could learn from Silicon Valley, which is a vibrant ecosystem that sustains and renews itself. Founders get inspired by the role models in the ecosystem, and once started, they are also supported by the ecosystem (with investment, advice, talent, ...). Large companies invest in startups and scaleups, and they also acquire smaller companies. Knowing that there is a possibility for a startup to eventually be acquired by a big player makes it more attractive to invest in it. The large companies are providing venture capital for the small companies, and the serial entrepreneurs help create new startups. The ecosystems should not only provide technical talent, but also business talent (like people who know how to grow a company, how to run a unicorn, how to do an IPO, ...).

Knowing that S&T clusters are essential for startups to grow and thrive, Europe should actively encourage the creation of a multitude of medium sized S&T clusters in major urban areas. These will not land into the top 20, but they will be local innovation engines, creating well-paid jobs, preventing brain drain and providing opportunities for the young generation. This is especially relevant for countries in Southern, Central, Eastern and Southeastern Europe who suffer from brain drain towards the Western and Northern European countries.

Europe should aim for hundreds of such ecosystems. Ecosystems should be specialized, and they should have the ambition to become global leaders: London already has a fintech ecosystem which is a global leader, Paris is creating an AI ecosystem (with Mistral), Zürich is a robotics and AI hub, Berlin is creating a creative tech ecosystem (gaming, media), Eindhoven has an information and communication technology (ICT) ecosystem (with ASML as a

large company), Amsterdam and Barcelona specialize in life sciences, Munich in automotive, and increasingly also in defence tech (which can benefit from collaboration with the automotive industry).

Besides creating a small number of large hubs, hundreds of small hubs should also be created, distributed throughout the EU. Companies can use this network of hubs to expand their business in other countries, especially since the introduction of the EU-INC company format and the “28th regime”. Europe has all that is needed to become the startup continent of the world (indeed, arguably, it already is). The bigger ecosystems like Paris will also attract bigger projects, like AMI, founded in 2026 by Yann LeCun with an initial valuation of €3 billion. Such projects cannot start in a place without an S&T cluster ecosystem. It is proof that ecosystems work.

It is, however, important to realize that ecosystems cannot be created overnight; they need time to grow and become productive, and they are always the result of joint efforts between different stakeholders.

- Local schools and universities must invest in research areas that are relevant for the local economy and they must also develop entrepreneurial skills in their students. This combination will result in spinoffs and startups, and it will also result in graduates that are ready to work in the local ecosystem (and attractive jobs in the local startups can stop them from looking for a job elsewhere). The schools and universities will also benefit from the ecosystem: contract research, company internships, and the ability to attract talented students who would like to work for one of the ecosystem companies. A local business school is also an asset to help skilling up the staff of scaleups.
- Local governments should offer ample space for high-tech companies to build the infrastructure they need and have a fast and pragmatic permitting policy. They should not only provide office and lab space, but also arrange for affordable housing, an international school, effi-

Table 5: Science and technology cluster ranking per capita in 2024. Source: [21]

Rank per capita	Cluster name	Economy	Estimated cluster population	PCT applications per capita (a)	Scientific publications per capita (a)	Total S&T share per capita (a)	Rank change (b)
1	Cambridge	GB	489,751	6,379	35,000	0.9	0
2	Sanjose-San Francisco, CA	US	6,252,315	7,885	9,211	0.7	0
3	Eindhoven	L	1,047,358	7,536	5,011	0.6	0
4	Oxford	GB	568,383	2,806	32,312	0.6	0
5	Boston-Cambridge, MA	US	4,251,769	4,462	17,934	0.6	0
6	San Diego, CA	US	3,910,684	6,279	5,189	0.6	1
7	Daejeon	KR	2,744,149	5,109	9,630	0.5	1
8	Ann Arbor, MI	US	659,434	1,891	29,439	0.5	-2
9	Seattle, WA	US	2,518,357	4,434	7,821	0.4	0
10	Munich	DE	2,794,775	3,828	9,734	0.4	1
11	Beijing	C	19,415,177	2,189	15,893	0.4	3
12	Goteborg	SE	841,281	2,500	12,035	0.3	0
13	Raleigh, NC	US	1,755,703	1,735	16,473	0.3	0
14	Stockholm	SE	2,151,605	2,809	9,148	0.3	1
15	Tokyo-Yokohama	JP	36,304,277	3,712	3,231	0.3	2
16	Copenhag,en	DK	1,699,974	1,838	14,669	0.3	0
17	Helsinki	FI	1,234,101	2,359	10,633	0.3	1
18	Zurich	CH	1,952,063	1,979	12,378	0.3	1
19	Basel	CH/ DE/FR	1,021,114	2,588	7,521	0.3	1
20	Stuttgart	DE	3,214,610	2,907	4,516	0.3	1

cient urban mobility and a liveable city. Building such infrastructure can take decades.

- The (national and/or regional) government can create incentives to attract companies to designated areas (tax incentives, subsidies, ...).
- Local companies must organize themselves too to make the infrastructure provided by the government into a vibrant and welcoming community in which all companies can learn from each other, help each other, celebrate successes, and especially grow the ecosystem, by e.g. investing in incubators, accelerators, ecosystem marketing, etc. It is important that the companies create their own ecosystem, and that they run it, not the local governments. They know better than anybody else what is needed in order to thrive.

All the above needs time, but if all stakeholders (city, schools/universities, govern-

ment, companies, ...) in an area are willing to create such an ecosystem, align their plans and investments, it can be built, and become the engine of economic development in the area. It takes time to produce the first big success stories (e.g. a unicorn), but once it reaches that level, and with the right marketing efforts, it will automatically attract talent and investors, and its growth

will accelerate. All current large S&T clusters once started small. As a reference: it took Silicon Valley 20 years to go from it first startup (HP, 1939) to the first Silicon producer (FairChild, 1959); see Figure 7.

The conclusion is evident. While Europe may not have companies as large as those in the US or China, various indicators suggest

Key Milestones: 1939-1959

Year	Milestone	Significance
1939	HP Founded	The birth of the "Garage" startup culture.
1946	SRI International	Stanford Research Institute created to bridge lab work and market needs.
1951	Stanford Industrial Park	The first dedicated hub for high-tech businesses.
1957	Fairchild Semiconductor	The official shift from vacuum tubes to Silicon.
1959	Planar Process	Fairchild invents a way to mass-produce transistors on a single chip.

Figure 7: Early history of Silicon Valley

it is steadily catching up. A European model characterized by plentiful startups and strategic innovation clusters distributed across the continent can drive innovation, while strategically pooling resources


and simplifying cross-border collaboration can help local innovators make an impact globally. Europe needs to educate a new generation that believes in its ability to transform the world and create an environ-

ment in which innovation and growth can thrive. The only thing standing in Europe's way of reclaiming the lead is its failure to believe that it can.

References

- [1] McKinsey & Company, Europe's moonshot moment: Fueling its tech ecosystem for scale, June 2025. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/europes-moonshot-moment-fueling-its-tech-ecosystem-for-scale>
- [2] European Commission. The Future of European Competitiveness. By Mario Draghi, Publications Office of the European Union, 2024, https://commission.europa.eu/topics/strengthening-european-competitiveness/eu-competitiveness-looking-ahead_en
- [3] European Commission. Much More Than a Market: Report by the High-Level Group on the Single Market, Chaired by Enrico Letta. Publications Office of the European Union, 2023, https://single-market-economy.ec.europa.eu/news/enrico-letta-report-future-single-market-2024-04-10_en
- [4] European Commission: Directorate-General for Research and Innovation. Align, Act, Accelerate – Research, Technology and Innovation to Boost European Competitiveness. By Manuel Heitor, Publications Office of the European Union, 2024, <https://data.europa.eu/doi/10.2777/9106236>
- [5] European Commission: Directorate-General for Research and Innovation, Science, research and innovation performance of the EU, 2024 – A competitive Europe for a sustainable future, Publications Office of the European Union, 2024, <https://data.europa.eu/doi/10.2777/965670>
- [6] European Investment Bank (2024). The scale-up gap: Financial market constraints holding back innovative firms in the European Union, European Investment Bank, https://www.eib.org/attachments/lucalli/20240130_the_scale_up_gap_en.pdf
- [7] State of the European Tech 25. Europe must define its future on its own terms. <https://www.stateofeuropeantech.com/>
- [8] https://en.wikipedia.org/wiki/High-speed_rail_in_China
- [9] Nature Index, <https://www.nature.com/nature-index/>
- [10] “Chinese President Xi Jinping delivers 2026 New Year Address”, ShanghaiEye, YouTube <https://youtu.be/zpNVRAfx0G0>
- [11] Clemens Fuest, Daniel Gros, Philipp-Leo Mengel, Giorgio Presidente, and Jean Tirole: “EU Innovation Policy – How to Escape the Middle Technology Trap?” EconPol Policy Report, April 2024, https://www.econpol.eu/publications/policy_report/eu-innovation-policy-how-to-escape-the-middle-technology-trap
- [12] NINDL, E., NAPOLITANO, L., CONFRARIA, H., RENTOCCHINI, F., FAKO, P. et al., The 2025 EU Industrial R&D Investment Scoreboard, Publications Office of the European Union, Luxembourg, 2025, <https://data.europa.eu/doi/10.2760/7802619, JRC144638>.
- [13] Eureka Prometheus Project, https://en.wikipedia.org/wiki/Eureka_Prometheus_Project.
- [14] Julie Pinkerton, “10 Most Valuable Car Companies in the World”, USNews, 2025, <https://money.usnews.com/investing/articles/the-10-most-valuable-auto-companies-in-the-world>
- [15] World Happiness Report https://data.worldhappiness.report/table?_gl=1*rqmsvr*_gcl_au*NjUyMzk5MDEwLjE3NzQ3NjM5MTU.
- [16] “Poverty”, Our World In Data <https://ourworldindata.org/poverty>
- [17] “Which are the most innovative economies in 2025?”, WIPO, <https://www.wipo.int/en/web/global-innovation-index/2025/index>
- [18] “History of Airbus”, Wikipedia, retrieved March 2026 https://en.wikipedia.org/wiki/History_of_Airbus
- [19] “50 years of weather forecasting at the ECMWF”. Nat Commun 16, 10052 (2025). <https://doi.org/10.1038/s41467-025-65837-2>
- [20] “EuroHPC Supercomputers Put Europe at the Forefront of Global Supercomputing”, EuroHPC, 2025 https://www.eurohpc-ju.europa.eu/eurohpc-supercomputers-put-europe-forefront-global-supercomputing-2025-06-10_en
- [21] World Intellectual Property Organization (WIPO) (2025). Global Innovation Index 2025: Innovation at a Crossroads. Geneva: WIPO. <https://www.wipo.int/publications/en/details.jsp?id=4807>

Process




The HiPEAC Vision analyses current trends that have an impact on the high-performance, edge and cloud computing and related communities. It formulates technical, methodological, standardization and policy recommendations to the HiPEAC community at large, and to policy makers. The content for the 2026 edition is based on information collected through a number of channels.

- A survey circulated to all HiPEAC members.
- Meetings with teachers and industrial partners at the ACACES 2025 summer school.

- A consultation meeting on open source.
- A consultation meeting on sustainability.
- A consultation meeting on artificial intelligence (AI).
- A consultation meeting on blockchain and micro payments.
- A consultation meeting on the cloud, the “next computing paradigm” (NCP) and “liquid computing”.
- Participation in tens of conferences and workshops on relevant themes for the HiPEAC Vision.

- Hundreds of informal discussions with experts from around the globe.
- Nine in-person editorial board meetings and four videoconference meetings.
- Several coordination meetings with other organizations, such as ETP4HPC, ECSO, CHIPS JU, AIOTI, ADRA, BDVA, SNS, NESSI, FIWARE, Destination Earth, Eclipse Foundation, etc...
- Feedback from presentations of the HiPEAC Vision 2025 at several conferences and workshops and from exchanges with DG CNECT.

Acknowledgements



This document is based on the valuable input of HiPEAC members. The editorial board, composed of Marc Duranton (editor in chief), Gaël Blondelle, Koen De Bosschere, Paul Carpenter, Madeleine Gray, Thierry Goubier, Axel Nackaerts, Charles Robinson, Tullio Vardanega and Olivier Zendra, would like to thank the following contributors:

Rosaria Rossini (Eclipse Foundation), Diego Perino (AI Institute BSC), Adam MacKay (QA Systems), Ian O'Connor (École Centrale de Lyon / Lyon Institute of Nanotechnology), Philip Piatkiewicz (Adra

Association), Christophe Leroux (CEA), Thierry Collette (Thales), Olivier Bettan (Thales), Emmanuelle Anceaume (CNRS), Antonio Fernández Anta (IMEDA Software Institute), Fabien Baligan (CEA), Laurent Laudinet (Thales TRT), Patrick Blouet (ST Microelectronics), Lieven Eeckhout (UGent), Adrian Friday (Lancaster University), Duncan Platt (RISE), Antonio Kung (TRIALOG), Manuel Betin (OECD), Antonio Garcia-Dominguez (University of York), Arun Joseph (Stealth AI Startup), Jordan Maris (Open Source Initiative), Roberto Di Cosmo (Software Heritage), Nicolas Rebierre (IRT SystemX).

The editorial board would also like to thank Eneko Illarramendi and Vicky Wandels (Ghent University) for their work to support the HiPEAC Vision. Special thanks to the Belgian comic artist Arnulf for his cartoons, which provide a characteristically tongue in-cheek take on the content of the HiPEAC Vision.

Generative AI chatbots such as ChatGPT were used to generate and polish some sections of this work, including text and images.



HiPEAC Vision 2026

Why Europe should create its own future, not imitate others

The ripple effects of developments in artificial intelligence (AI) – mostly emanating from a limited group of companies outside of Europe and taking place at breakneck speed – continue to spread across the computing ecosystem, with implications for everything from individual careers up to whole industries. As hyperscalers announce massive infrastructure buildouts and AI tools gain increasing functions and autonomy, Europe is left wondering whether it is too late to join the party.

But what if there were another paradigm for next-generation computing? What if this paradigm avoided reliance on massive data centres, with their eyewatering energy demands? What if it facilitated collaboration and the pooling of existing computing resources; eschewed the antisocial incentives of the attention economy and surveillance capitalism while allowing users to harness computing power tailored to their needs; and was based on rightsized infrastructure created in Europe which allowed for flexible, sustainable, resilient operation?

At a time of intense global volatility, the HiPEAC Vision 2026 argues that, rather than blindly copying other countries' roadmaps, Europe can and should carve out its own path in computing: a trajectory which reflects the European context while supporting the values and culture which are important to European society.

Visit the HiPEAC Vision online: vision.hipeac.net

