

The EU's digital transition and AI



JOURNAL



AI is already shaping (...) the public record on which democratic accountability depends.



EUROPEAN
COURT
OF AUDITORS

The Architecture of Intelligence

By **Matt Stroud**, author and technologist



7

17



Artificial intelligence and protecting knowledge systems in the European Union

By **Professor Raluca Csernaton**, Centre for Security, Diplomacy and Strategy, Brussels School of Governance, Vrije Universiteit Brussel, and Fellow at Carnegie Europe

INTERVIEW

with **Henna Virkkunen**, European Commission Executive Vice-President for Tech Sovereignty, Security and Democracy

'I want this continent to become an AI Continent'



24

34



Advancing EU public administration readiness for AI Act compliance

By **Wojciech Wiewiorowski**, European Data Protection Supervisor (EDPS)

INTERVIEW


with **Jorg Petrovic**, ECA Member and Reporting Member for the ECA's AI strategy

'I have no doubt: AI's advantages far outweigh all the négatives!'



38

48




The ECA's audit of the EU's microchips ambitions – examining an industrial ecosystem

By **Rafal Gorajski**, Investment for cohesion, growth and inclusion directorate, and **Austin Maloney**, private office of Annemie Turtelboom, ECA Member

From experimentation to integration: advancing AI in European public audit

By **Alvar Nouakas**, National Audit Office of Estonia



75

84



When the rulebook arrives before the road is built

By **Mustafa Elmoslhey**, space engineer and entrepreneur

The curse of Big Tech: how digital giants are eroding the constitutional foundations of liberal democracy

By **Reijer Passchier**, Professor of Digitalisation and Constitutional Democracy, Open University of the Netherlands, and Assistant Professor of Constitutional Law, Leiden University



96

101



INTERVIEW

with **Michael McNamara**, Member of the European Parliament and European Parliament rapporteur on AI

'For democracies legitimacy comes from protecting the rights of their citizens, also on AI...'

Contents

- 5** **Editorial**
Preserving the human component in AI
By Gaston Moonen
- 7** **The Architecture of Intelligence**
By Matt Stroud, author and technologist
- 13** **Parallel tracks : AI and platforms in the EU**
By Ian Gauci, GTG Legal, Malta
- 17** **ECA Journal Long read**
Artificial intelligence and protecting knowledge systems in the European Union
By Professor Raluca Csernatosi, Centre for Security, Diplomacy and Strategy, Brussels School of Governance, Vrije Universiteit Brussel, and Fellow at Carnegie Europe
- 24** **'I want this continent to become an AI Continent'**
Interview with Henna Virkkunen, European Commission Executive Vice-President for Tech Sovereignty, Security and Democracy
- 28** **The European Commission's Cloud Sovereignty Framework: a new brick for building digital sovereignty**
By Vincent Coudrin and Philippe Merle, Directorate-General for Digital Services, European Commission
- 34** **Advancing EU public administration readiness for AI Act compliance**
By Wojciech Wiewiorowski, European Data Protection Supervisor (EDPS)
- 38** **'I have no doubt: AI's advantages far outweigh all the négatives!'**
Interview with Jorg Petrovic, ECA Member and Reporting Member for the ECA's AI strategy
- 43** **AI – is the EU picking up pace?**
By Mihails Kozlovs, ECA Member
- 48** **The ECA's audit of the EU's microchips ambitions – examining an industrial ecosystem**
By Rafal Gorajski, Investement for cohesion, growth and inclusion directorate, and Austin Mloney, private office of Annemie Turtelboom, ECA Member
- 53** **From promise to practice: how AI is shaping ECA audits**
By Emanuele Fossati, ECA Directorate of the Audit Quality Control Committee
- 58** **Directors' cut**
'Artificial intelligence as a thinking partner'
Interview with ECA directors Ioanna Metaxopoulou and Marco Barros Lourenço
- 65** **Digitalisation of the Czech public administration: between ambition and reality**
By Stefan Kabatek, Supreme Audit Office of the Czech Republic
- 69** **Digital transformation and the changing role of supreme audit institutions**
By Burhan Gün, Rudi Turksema and Colin van Noordt, Netherlands Court of Audit
- 75** **From experimentation to integration: advancing AI in European public audit**
By Alvar Nouakas, National Audit Office of Estonia

- 80** **Governing AI while building it – accountability in the EU’s digital**
By Andreas Braun and Hazal Kantarci, PwC Luxembourg
- 84** **When the rulebook arrives before the road is built**
By Mustafa Elmoslhey, space engineer and entrepreneur
- 89** **AI for access to justice – an important opportunity**
By Assistant Professor Hannes Westermann, Maastricht University
- 92** **‘Safeguarding technology is a shared responsibility that begins with those who develop and deploy it’**
Interview with Katarina Wallin Bureau, Microsoft
- 96** **The curse of Big Tech: how digital giants are eroding the constitutional foundations of liberal democracy**
By Professor Reijer Passchier, Professor of Digitalisation and Constitutional Democracy, Open University of the Netherlands, and Assistant Professor of Constitutional Law, Leiden University
- 101** **‘For democracies legitimacy comes from protecting the rights of their citizens, also on AI...’**
Interview with Michael McNamara, Member of the European Parliament and European Parliament rapporteur on AI
- 107** **Foresight and audit**
The future of AI in audit – opportunities and challenges
By Anna Zygierewicz, Directorate of the Presidency
- 111** **Foresight and audit**
What does AI say about its impact on the future of audit?
AI-assisted article, reviewed by Daniela Hristova, Directorate of the Presidency
- 116** **New ECA Members**
‘Supreme audit institutions are pillars of democracy!’
Interview with Pierre Moscovici, ECA Member since 1 January 2026
- 121** **New ECA Members**
‘Professionalism is the key trademark of the ECA’
Interview with Daniel Caspary, ECA Member since 1 March 2026
- 126** **Reaching out**
Maintaining forward momentum on gender equality
By Annemie Turtelboom, ECA Member and Dean of the Investment for cohesion, growth and inclusion Audit Chamber
- 129** **Reaching out**
Our outreach work on the EU’s next seven-year budget – 12 ECA opinions published this winter
By Damijan Fiser, Directorate of the Presidency
- 134** **Focus**
ECA publications from January to June 2026
- 144** **Next edition**
EU efforts to combat fraud and corruption

Preserving the human component in AI

By Gaston Moonen

In an ECA Journal on the theme of *The EU's digital transition and AI*, this editorial might just be generated by AI, or at least with the help of AI. Using input from others, actually from all of you, putting data into large language models – LLMs – to generate my own unique product. Well, I did not use AI for this editorial, you have my word. But how can you be sure? How do you know, as the reader, that what I write is really mine and not 'borrowed' from others, or worse, based on 'non-existent content'?

This uncertainty has not only existed since AI rose to the level of its current users. Before the introduction of AI, I could plagiarise or publish fake news. But a thorough review could catch me out on that, and reviews are one of the key challenges that AI poses today. AI enables us to generate a lot of information, handed to us on a digital silver platter – information that would otherwise require a lot of effort to gather, if we found it at all. AI provides us with new ideas, new insights and potentially huge efficiency gains, provided that what is presented is correct. And that is the first challenge: LLM checks on information accuracy are far from 'error-proof' and the task of reviewing for correctness is daunting. First, because of the 'hallucinations' – ranging from small mistakes to outright 'fake news' – that this AI-generated information may contain. And secondly, because of the enormous scale of the information provided, and hence of the related review work. Which brings me to the second challenge: who can be trusted to do that review? Which organisations, which institutions have the reputation, the capacity and the ability to execute a review that subsequently inspires trust that what is presented is true and reliable?

This is how society works: trust is generated through reliable review. This goes for commercial relations, government, democratic processes, and equally for personal relationships. Fake news is not something recent; it has always existed, even in democracies, for manipulation purposes or similar. The big difference now with AI is that, consciously and unconsciously, we as users generate data which, out of our sight or control, creates news and information. This may be consciously manipulated by certain selection criteria or by sheer business models going for profits. This manipulation, and hence dependency, dawned upon me for the first time about ten years ago, when I heard that a certain US EV company could remotely extend the battery range of its newest car model in an emergency. How kind! And what kind of power this gave those controlling the parameters! All the more so if those in control on AI are now in fact very few in number and huge in financial and societal power, with more financial leverage than many companies active in other areas, more than various governments or international organisations.

Any concentration of power raises the question of checks and balances, to ensure that the interests of all citizens are respected and not those of a given few. This is true whether it concerns power concentration in politics, business, an executive authority or even sports. This ECA Journal does not primarily focus on the techniques related to the digital transition and the multiple AI possibilities out there – it would risk being outdated tomorrow. Somehow, though, almost all the contributors to this edition, whether specifically requested or not, bring up the issue of dependence on a few big players and how to create a governance structure preventing this dependency. How can we ensure a minimum level of autonomy for other businesses, for governments, for institutions, ultimately for citizens themselves? Eventually, it is the governance model and its application that will determine autonomy. And while that governance discussion may appear far removed from our practical use of AI, it is the prime focus of the contributions by Professor Raluca Csernatonu ([page 17](#)), who terms it epistemic security, and Professor Reijer Passchier ([page 96](#)), who calls for constitutional resilience against Big-Tech power concentration and the incursion of their tentacles into politics. Matt Stroud, in his opening article on what AI entails, reformulates AI as the Architecture of Intelligence, highlighting the EU's choice for a third path that keeps markets and states accountable ([page 7](#)).

To its credit, the EU has not shied away from taking this path. Despite the dominance of US Big Tech (over 70% of all EU data – personal, commercial and government – is stored in US-controlled clouds), the EU has carved out its own governance model, as explained by Ian Gauci, who dives into the design and operational logic of the EU's Digital Services Act and the AI Act ([page 13](#)). But does the EU have sufficient leverage in this area? Perhaps with the example of the EU's General Data Protection Regulation (GDPR) – which set a world standard on personal data protection – in mind, Henna Virkkunen, European Commission Executive Vice-



President for Tech Sovereignty, Security and Democracy, remains optimistic not only about having and applying a thorough regulatory framework on AI, but also about increasing Europe's AI independence in a rapidly changing landscape where scale, capital and access are key ([page 24](#)). She advocates for simplification to stimulate EU-grown AI solutions, yet sufficient regulation – including financial enforcement powers – to protect citizens' rights from AI abuse.

The EU employs a carrot-and-stick approach to reach these objectives. 'Carrots' include promoting the European development of 'enablers' for digital transition and AI capacity, such as cloud storage capacity and microchip development and production. European Commission experts Vincent Coudrin and Philippe Merle explain how the launch of an EU Cloud Sovereignty Framework is a crucial element of a functioning and competitive European digital ecosystem ([page 28](#)). ECA auditors Rafal Gorajski and Austin Maloney look into the outcomes of the European Commission's support for European microchip development, an essential enabler for increasing the EU's autonomy, including in AI ([page 48](#)). The 'stick' aspect surfaces in the contribution by Wojciech Wiewiorowski, the European Data Protection Supervisor (EDPS), who can impose administrative fines under the AI Act, while at the same time stimulating close and structured cooperation between various levels of governance ([page 34](#)).

With the EU's AI ambitions set high, reality checks are all the more important, as explained by Mustafa Elmoslhey, an entrepreneur and practitioner trying to apply the EU's AI regulation and experiencing compliance asymmetry followed by implementation headaches ([page 84](#)). Reality checks are needed regularly, as advocated by ECA Member Mihails Kozlovs ([page 43](#)), reiterating findings from an ECA audit report that called for stronger governance and more focused investments. Two years after publication, the findings still appear to be highly relevant. Our counterparts from member-state audit institutions find similar gaps. Štefan Kabatek reports on an overall failure to achieve the intended digitalisation of the Czech public administration ([page 65](#)). Buran Gün, Rudi Turksema and Colin van Noordt share their audit experiences when auditing digitalisation in the Dutch government, including identifying how AI can empower external auditors ([page 69](#)).

The ECA's own AI ambitions are also high, as explained by Jorg Petrovič, ECA Member and reporting Member for the ECA's AI strategy. He explains that for the ECA to be able to examine auditees' use of AI, ECA staff need to be AI users too ([page 38](#)). Emanuele Fossati, AI expert at the ECA, highlights, with concrete examples, how AI contributes to more efficient use of ECA resources while preserving high-quality audit work ([page 53](#)). This last aspect, delivering high-quality output with the aid of AI, is also the key ambition and concern of ECA director Ioanna Metaxopoulou. Together with her fellow ECA director Marco Lourenco, she embraces AI technology while staying within the ECA's guardrails when it comes to quality, accountability and reliability ([page 58](#)). This positive attitude is also encouraged by Alvar Nõuakas of Estonia's National Audit Office, who sees AI explicitly as a tool for enhancing accountability and public trust ([page 75](#)). Assistant Professor Hannes Westermann reaches the same conclusion in relation to AI and access to justice: AI tools offer great potential for making legal systems more accessible and hence contributing to trust in the system ([page 89](#)). ECA staff members Anna Zygierevicz ([page 107](#)) and Daniela Hristova ([page 111](#)), the latter presenting a specimen of an AI outcome, elaborate on what AI may entail for the future of external audit – a driver for change in audit work and in the audit profession.

But how can we govern AI while building it? This question is raised by Andreas Braun and Hazal Kantarci of PwC Luxembourg. For them, the human factor remains indispensable ([page 80](#)). We also tried to get feedback on this question from both US and European Big-Tech companies, but there the human factor was most often absent; it was almost impossible to get hold of a human to respond! In the end, Katarina Wallin Bureau of Microsoft was kind enough to respond to some questions, including on Big Tech's responsibility to users to embed safeguards in design and throughout the product lifecycle ([page 92](#)). That human factor is also the guiding element for Michael McNamara, Member of the European Parliament and its rapporteur for AI. For him, this separates the EU framework from anything else and is what the EU was created for: protecting the rights of its citizens, including when it comes to digitalisation and AI.

As things stand now, in digitalisation and the use of AI, the human factor cannot and should not be ruled out. How could this ever be otherwise if the overall human interest is the key objective? To prevent hegemonic AI influence or excessive enrichment of tech barons from prevailing, the EU needs to lead by example, as it has in various other areas when it comes to protecting the fundamental rights of its citizens. Through its Digital Services Act and AI Act, the EU makes its own agenda for AI and places the burden for reliability, for proper review, for accountability regarding how AI works, on those developing and deploying it. Just like in any other product area, creating trust in its functioning through responsible monitoring. Employing resources for review, being transparent on the how and what, engaging the human factor in all this, may decrease profit margins for those offering AI solutions, but will increase public trust in these solutions. Ultimately, public trust is a key ingredient for continued interest and use, whether it relates to an AI tool or a fact-finding report, hopefully even more so when produced by an external audit institution seeking facts instead of opinions.

The Architecture of Intelligence

AI, data and Europe's fork in the road

By Matt Stroud, investigative journalist



© GPT / Matt Stroud

For most of its history, AI advanced quietly. The technology did useful work in pockets, in fraud detection, image classification, and recommendation engines, but it remained largely invisible to the public. According to Matt Stroud that changed when three things converged: new model architectures, enormous datasets and industrial-scale compute. He argues that together they turned a specialist research field into a general-purpose technology. The narrow systems that once classified, predicted and recommended have been joined by generative models that write, code and reason across domains, and by agents that use digital tools autonomously to complete tasks. Matt Stroud is a specialist in AI governance, decentralised personal data, and the commercial and societal implications of AI-driven transformation, and author of the book *Digital Liberty* and finalising two new books – *Behavioural Physics* and *Irreversibility*. He guides us into the vast realm of what AI entails and influences, including a short historical exploration. He shows that that the opportunities are real, whether it relates to healthcare, finance or corporate processes. AI can spot patterns humans miss, accelerate expert work, personalise services at scale and compress the distance between decision and action. But this value does not come from the model alone. Matt Stroud argues it depends on the data the model has access to, the systems it can act through, and the rules that govern how its outputs become consequences. Each is a decision someone has made, and together they form a structure that almost nobody is invited to examine. For Matt Stroud the question that matters most is who controls that structure. The answer, for Europe, will shape not only competitiveness but the balance of power between platforms, states and individuals.

Origins

You probably encountered AI three times before lunch. The autocomplete that finished your sentence, the spam filter that decided which messages deserved your attention, the route your phone suggested without your asking. None of it announced itself. The most consequential AI in your day is, more often than not, the AI you barely notice.

The technology behind those quiet interventions is older than most people think. Alan Turing asked whether machines could think in 1950¹, and six years later a workshop at Dartmouth College gave the field its name: artificial intelligence². For most of the seven decades that followed, progress was steady, if occasionally glacial, punctuated by two long winters in which both the funding and the faith ran out. Then, in 2017, something

How we got here

The first generation of AI researchers, working through the late 1950s and 1960s, believed that intelligence could be captured in rules. You wrote down everything a doctor knew about infectious diseases, encoded it as if-then logic, and the system would diagnose patients. The approach worked just well enough to be funded for decades, and just badly enough to remind everyone, periodically, that intelligence might be more than a sufficiently long list of rules.

By the 1980s and 1990s, researchers had shifted strategy. Rather than trying only to hand-code intelligence as rules, they increasingly built systems that learned from data, drawing on probabilistic approaches such as Bayesian networks, statistical methods such as support vector machines, and long-running work on neural networks. These methods were less elegant than logic, but they worked on the messy real-world problems where the rules could never be fully written down. Then, in 2012, a deep neural network called AlexNet won the ImageNet image-recognition competition by a margin so large it stunned the field⁴. The era of deep learning had begun. Five years later, a team at Google published a paper called 'Attention Is All You Need', introducing the transformer architecture that now sits behind many of the most prominent generative AI systems⁵.

What sits inside today's AI systems is, at root, three ingredients. The first is data: the corpus on which a

shifted. A new architecture called the transformer, combined with enough data and enough compute, began doing things that had previously seemed years away. Five years later, the public noticed. We are now, plausibly, at an inflection point.

Father John Culkin's observation that 'We shape our tools, and thereafter our tools shape us' has never been more prescient³. AI may become one of the most consequential tools humans have ever shaped, which is another way of saying that it may become one of the most consequential shapers of us. In this article I try to do two things: to explain how the technology actually works, and to keep one slightly uncomfortable question in view throughout: how do we want our societies to be reshaped?

model is trained, often drawn from the internet, books, code, scientific literature and other curated datasets. The second is the model itself: a neural network with billions or trillions of adjustable weights, often built on the transformer architecture. The third is compute: the brute force required to do the training and run the system, supplied largely by specialised chips, global supply chains and vast data centres.

The intuition for how a neural network learns is simpler than the mathematics behind it. Imagine a system with billions of tiny dials. Show the system a million examples, and each time it produces the wrong answer, nudge the dials slightly in the direction that would have made the answer less wrong. Repeat often enough and what emerges, eventually, is a system that has internalised the statistical regularities buried in the examples. That internalisation is the AI model learning, which is then retrieved in response to future inputs or queries. The system does not understand in the way that you and I do; it is a sophisticated pattern-matching machine operating at enormous scale.

A consequence of this is that an AI model is not shaped by architecture alone. It is also shaped, deeply and durably, by the data on which it was trained. The same training method, run on different data, produces dramatically different systems. Whoever curates the data, in other words, shapes the model.

¹ Turing, A. M., Computing Machinery and Intelligence. in: *Mind*, 1950.

² McCarthy et al. / Dartmouth 1956 background, via LLNL or Dartmouth archive.

³ John M. Culkin Jr. was an American academic and former priest who was a leading [media scholar](#) and [critic](#).

⁴ Krizhevsky, A., Sutskever, I., and Hinton, G., *ImageNet Classification with Deep Convolutional Neural Networks*, 2012.

⁵ Vaswani et al., 'Attention Is All You Need', 2017.

Three kinds of AI system

Policy conversations about AI often mix together three different technologies. Each has different use cases, fails differently and needs different governance.

The first category is narrow predictive AI. It does one thing, in one domain, and usually does it without anyone calling it AI: the model that flags a fraudulent card transaction, the radiology tool that highlights a suspicious mass, the recommender that decides what appears on your screen this evening. This is by far the most economically productive AI in the world today, and most of it is invisible. It fails by being overconfident on edge cases, by inheriting bias from its training data, and by losing accuracy silently when deployed on a population the training data did not represent.

The second category is generative foundation models, including systems such as ChatGPT, Claude, Mistral, Gemini and Copilot. The word 'foundation' matters: each model is trained once on a broad corpus of data and then adapted for many specific uses. This is the category that drives both the public conversation and much of the regulatory anxiety around AI. It fails by hallucinating, by giving different answers to the same question, and by being almost-but-not-quite right in ways that are genuinely hard to audit.

The third category is agentic systems. Here a foundation model is given a goal, given tools, and allowed to chain

actions over time: a scheduler that actually books the meeting, a research assistant that searches across documents, a coding agent that writes and runs its own code. These systems are the category for which existing regulation is least developed. They fail by acquiring sub-goals nobody specified, by acting faster than humans can supervise, and by compounding errors as actions chain.

For Europe, this distinction has industrial as well as regulatory consequences. The AI most useful to European manufacturing, healthcare and agriculture is not necessarily the consumer-facing chatbot trained on broad internet text. It is the domain-specific model trained on European industrial, clinical and operational data. Such data exists in abundance. The problem is that it sits in silos, kept apart by walls that are partly legal and partly competitive. Whether Europe builds the AI it needs or imports it depends less on developing algorithms than on whether those silos can be safely connected. The unglamorous work of connecting them is already underway. [Gaia-X](#) is developing federated data-infrastructure models; the European Health Data Space has entered its implementation phase; and the Common European Data Spaces, supported by the Data Governance Act, are intended to enable trusted data sharing across strategic sectors. These initiatives matter more than the headlines they fail to make.

AI in the wild

Theory is one thing, deployment another. The clearest picture of what this technology is already capable of comes from looking at four places where it has unmistakably arrived. Each show something different: where the technology delivers, where it fails predictably, and how much of its real impact turns on the architecture surrounding the model rather than on the model itself.

Healthcare is where AI is most clearly useful and most dangerous. DeepMind's work with Moorfields Eye Hospital showed AI diagnosing age-related macular degeneration⁶ from retinal scans at expert level, and AlphaFold cracked a fifty-year-old problem in protein-structure prediction⁷. The failure cases are equally instructive. Models trained on one population and deployed on another lose accuracy in ways that nobody notices until somebody investigates. Dermatology classifiers trained on lighter skin perform worse on darker skin⁸. Mortality-prediction algorithms that use historical healthcare spending as a proxy for medical need systematically under-prioritise patients

who have been historically under-served⁹. The bias is not a moral failing of the engineers; it is an architectural property of the data they were given.

Finance is where the same kind of data dependence reaches your wallet. A modern credit decision is increasingly informed not only by what you wrote on a form, but by wider behavioural and financial data trails accumulated over years, much of it provided without any conscious moment of consent. The result is a decision taken by a machine that has read you more carefully than you have read it, against benchmarks calibrated for people you will never meet. Open Banking under the [EU's PSD2 directive](#)¹⁰ is a partial response: when you change your financial services provider, your banking history can go with you. The principle is right, but confined to a single sector its scope is narrow.

Public services can be transformed by AI technology, but only if the right data infrastructure is built first. Estonia has become the European exemplar. The country has a

⁶ De Fauw et al. / [DeepMind-Moorfields work on retinal disease, plus Moorfields summary](#).

⁷ Jumper et al., 'Highly accurate protein structure prediction with AlphaFold', in: *Nature*, 2021.

⁸ Daneshjou et al., 'Disparities in dermatology AI performance', in: *Science Advances*, 2022.

⁹ Obermeyer et al., 'Dissecting racial bias', in: *Science*, 2019.

¹⁰ The EU PSD2 directive (Directive 2015/2366) modernizes payment services across the EU, enhancing security, consumer protection, and competition.

national digital identity that underpins an interoperability layer called X-Road, which lets agencies share data securely¹¹. A 'once-only' principle spares citizens from submitting the same information twice, and an access log lets each citizen see who has looked at their records, and why. The contrast with other European cases is bracing. In the Netherlands, an algorithm in the childcare benefits programme wrongly accused thousands of families of fraud, leading to significant political fallout in 2021¹². In the United Kingdom, an algorithmic adjustment of A-level grades in 2020 was withdrawn within days under public pressure¹³. What went wrong in both cases was not only the algorithm but the architecture around it: no transparency, no appeal, no recourse.

The fourth case is the corporation itself. Most firms are still pouring new wine into old bottles, using AI to do the old job slightly better; the more interesting ones are redesigning the bottle, and what they are mostly after is speed. The military strategist John Boyd, who trained fighter pilots after the Korean War, described a decision cycle he called the OODA loop: observe, orient, decide, act. The pilot who runs the loop faster, he argued, beats the pilot flying the better plane. The same is now becoming true of firms. A company whose observe-decide-act cycle runs in days will routinely outmanoeuvre one whose cycle runs in quarters, irrespective of who started larger. For European business, that turns AI from an efficiency question into one about organisational design.

Box 1 - Three innocuous facts identify you

1990s
In the 1990s, the Harvard researcher Latanya Sweeney showed that

3 INNOCUOUS FACTS
were sufficient to uniquely identify **87%** of the United States population.

ZIP CODE (the American equivalent of a postcode) + DATE OF BIRTH + PERSON'S SEX

2019
In 2019, a study in *Nature Communications* updated the result with modern inference techniques:

99.98% of Americans could be re-identified using **15 DEMOGRAPHIC ATTRIBUTES** in any anonymised dataset.

AGE, GENDER, RACE, EDUCATION, INCOME, MARITAL STATUS, OCCUPATION, ZIP CODE, HEALTH STATUS, HOME VALUE, VEHICLE OWNERSHIP, VOTING HISTORY, CREDIT RATING, PURCHASE BEHAVIOUR, AND MORE

The regulation guards the front door...
...while inference climbs in through the windows.

Data may be anonymous. You are not.

© Matt Stroud

The studies referenced in **Box 1** are:

- “87% identified by ZIP code + date of birth + sex - Latanya Sweeney’s paper: ‘Simple Demographics Often Identify People Uniquely’ Carnegie Mellon / Data Privacy Lab, originally based on 1990 U.S. Census data. <https://dataprivacylab.org/projects/identifiability/paper1.pdf>
- 99.98% re-identifiable using 15 demographic attributes - Luc Rocher, Julien M. Hendrickx, and Yves-Alexandre de Montjoye: ‘Estimating the success of re-identifications in incomplete datasets using generative models’ in: *Nature Communications*, 2019. <https://www.nature.com/articles/s41467-019-10933-3>

¹¹ e-Estonia / X-Road sources on interoperability and once-only principle.
¹² Amnesty International, ‘Xenophobic Machines’, on the Dutch childcare benefits scandal.
¹³ House of Commons Library / Ofqual on the 2020 A-level algorithm.

The data foundations of AI

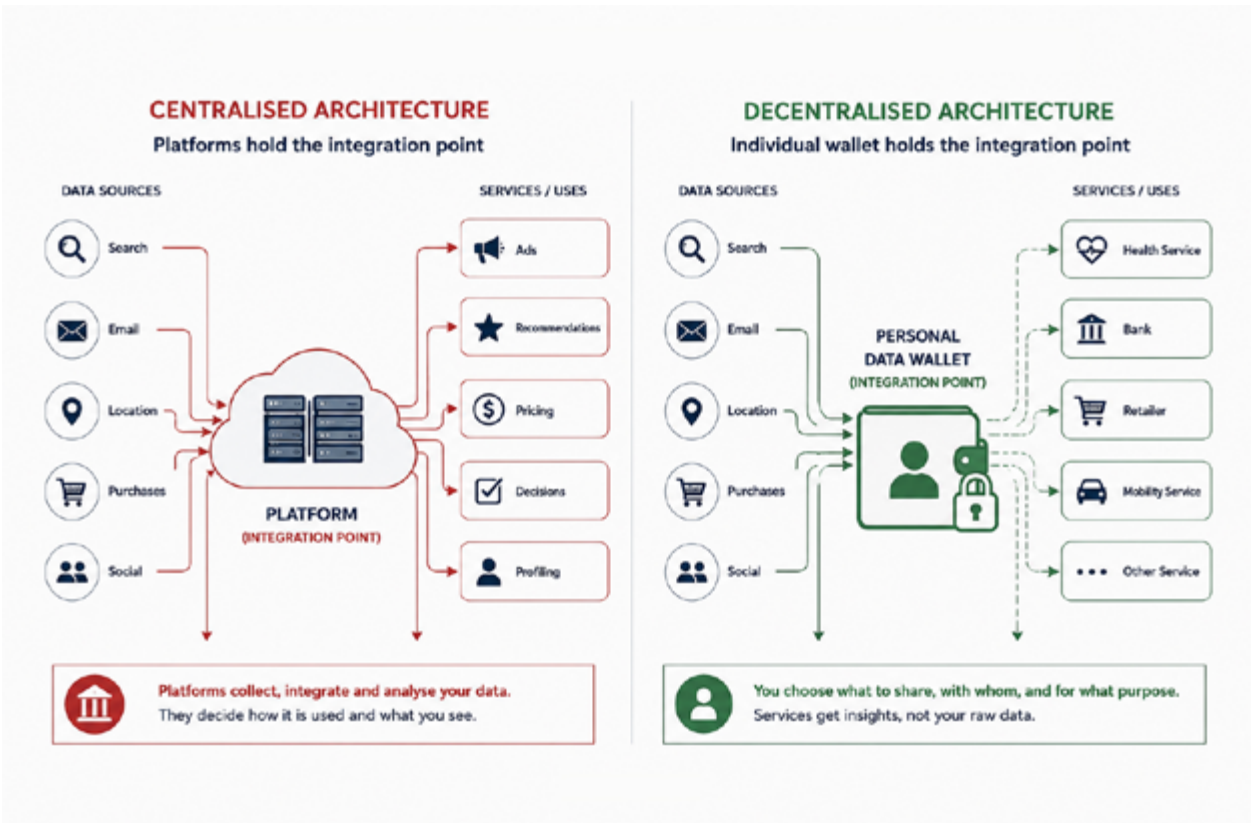
Mention AI and most people picture a chatbot. The chatbot is the visible tip of a much larger structure. Below the waterline sit the model itself, the data it was trained on, the platforms that supplied that data, and the infrastructure that runs every query and logs every answer. Most of what determines a chatbot’s behaviour sits quietly out of view. Many of the consequential questions about an AI system live in that hidden part. Who trained the model, and on whose data? Who owns the system that runs and logs your every interaction with it? These are questions about architecture, and the architecture is invisible by design. The choices being made about it now will outlast the applications they enable.

The architecture beneath European AI is largely shaped by the General Data Protection Regulation (GDPR). The GDPR is the most ambitious data-protection regime any jurisdiction has built, and it has rightly become the global grammar for much of what has followed. In short, it has

been globally influential and, in many respects, successful. The problem is that it was designed for a world where data types had stable boundaries, where collecting your shoe size was clearly different from collecting your HIV status, and where one could draw a sensible line around things called ‘special categories’. With the arrival of modern AI, that world no longer exists (see also **Box 1**).

The reason it no longer exists is inference. Modern machine learning can reconstruct sensitive attributes from data that looks utterly innocuous. Your phone’s accelerometer can predict Parkinson’s disease years before a clinician would. Your typing rhythm can flag the onset of cognitive decline. Your supermarket shopping can identify a pregnancy before you have told your family. None of these inputs sits in the GDPR’s special category list, yet through inference they become medical data, psychiatric data, data that most people would be horrified to learn could be derived from their digital exhaust¹⁴.

Figure 1 - Two architectures for personal data



© Matt Stroud

¹⁴ 12.Sweeney, Simple Demographics Often Identify People Uniquely, 2000. On re-identification; Rocher et al in: *Nature communications*, 2019.

In response, European data and AI policy will need to evolve. It must stop asking only who is allowed to hold which categories of data, and start asking what is allowed to be inferred from data, by whom, under what conditions, and with what disclosure to the subject. The shift is hard, because it requires regulating an action rather than an object; but that is the inescapable challenge we now face.

Fortunately, new architectures are emerging that enable a response. Useful analytics generally require the integration of data across domains, especially when those analytics are about people, because people do not live in silos. Our lives span them: we orchestrate planes, trains, hotels and restaurants to assemble a holiday; we move fluidly between healthcare, work and family. The silos are pieces of a single picture, and an engine that can see the whole picture produces better and timelier predictions than one that sees only a slice. In the AI services of the next decade, context will be king. The question is who holds the integration point at which that context is assembled.

Today the platforms hold it. Google integrates your search, your email and your location. Meta integrates

your social graph, attention patterns and behavioural responses. The insights produced give these firms both commercial power and a measure of power over the individuals concerned. The new architectures move the integration point back to the individual. In the strongest version of this model, a personal data wallet holds the full picture. Services request specific insights, computed locally or under user-controlled conditions, and receive answers rather than raw data (see **Figure 1**).

What changes here is not the AI itself but the data layer beneath it. That layer, however, shapes almost everything the model on top of it can do. Europe is well placed to lead this transition in personal data architecture, with projects from [DECODE](#) in Barcelona to [Inrupt](#) and [Dataswyft](#) in the United Kingdom, and the European Digital Identity Wallet now rolling out under [eIDAS 2.0](#). The components of such an architecture already exist; what they need is political backing to grow from pilot projects into default infrastructure. That shift would change what kind of AI Europe ends up with. What AI does to us, in the end, will be shaped less by the algorithms in the headlines than by the architecture sitting quietly beneath the surface.

The choice in front of us

The question in front of European policymakers is not whether AI will reshape society; it will. The question is what shape of society we want it to fashion. Two paths are visible from here. One ends with algorithmic gatekeepers concentrating data and AI-powered inference into a handful of firms, with citizens as raw material. The other ends with individuals holding their own data and using that control to shape their relationships with AI services. Neither path is inevitable.

The AI ecosystem is still in flux, but it will not remain so for long. Architectural choices made in the next few years may lock in a trajectory for decades. Between the

American model that lets the market decide and the Chinese model that lets the state decide, there is space for a third path, one that lets individuals decide while keeping markets and states accountable. Europe has the regulatory weight to choose that path, with wide-ranging consequences for our economies, societies and lives.

Key questions that need to be addressed relate to governance, sovereignty, audit, regulation and the practical realities of industry. We have shaped a tool powerful enough to shape us back. The hard work now is to design the architecture beneath it, so that what it shapes us into is something we have chosen.

Parallel tracks: AI and platforms in the EU

By Ian Gauci, GTG Legal, Malta



© mixmagic / depositphotos.com

The Digital Services Act and the Artificial Intelligence Act were each designed with a different regulatory object in mind. Yet they increasingly have a bearing on the same operators, the same services, and sometimes the same technical functions. In this article, Ian Gauci argues that the central weakness of the current framework is not the overlap between the two regimes as such, but the absence of any integrated mechanism for managing that overlap where it matters most. He concludes that the EU should confront this directly rather than treat it as a scheduling problem. Ian Gauci is Managing Partner at GTG Legal, Malta, specialising in financial services regulation, fintech, and technology law. He chairs Malta's AI in Fintech Strategy Group and also writes regularly on EU digital regulation. Below he explains how the physical and legal boundaries of digital services and AI are converging and that a key focal point for those designing and implementing digital governance should be operational coherence.

When convergence outpaces coordination

The European Union's digital regulatory architecture now turns on a practical rather than a conceptual question. The Digital Services Act (DSA) and the Artificial Intelligence Act (AI Act) were designed around different regulatory objects, yet they increasingly have a bearing on the same services, the same operators, and sometimes

the same technical functions¹. Overlap was foreseeable. The more difficult question is whether the EU has built a supervisory architecture capable of managing that overlap in a disciplined way. Despite recent Commission work on the interaction between the DSA and other EU legislation, and its 2025 simplification package for the

¹ Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act) [2022] OJ L 277/1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 2024/1689.

AI Act, there is still no dedicated cross-regime procedure for cases in which platform-level DSA supervision and obligations under the AI Act converge².

The AI Act entered into force on 1 August 2024. Its prohibitions took effect on 2 February 2025. The rules on general-purpose AI models, together with key governance provisions, took effect on 2 August 2025. The DSA has applied generally since 17 February 2024, and the Commission has already opened proceedings and taken enforcement steps against major platforms such as X³, Meta and TikTok⁴ under its powers for regulating very large online platforms (VLOPs). These proceedings and preliminary findings have focused, inter alia, on transparency obligations, advertising repositories, data access and systemic risk governance. The January

2026 investigation into X's recommender systems and Grok illustrates the point: an online platform's risk management, recommender architecture, and embedded generative functionality can now be scrutinised in the same enforcement environment⁵.

In this article I argue that the central weakness is not overlap as such. It is the absence of an integrated mechanism for dealing with overlap when it matters most. The EU has built two serious and effective regimes; however, the connections between them remain underdeveloped. This is the problem which the current debate on simplification should consider more directly.

The design logic and its limits

The DSA and the AI Act were each designed with a different regulatory subject in mind. The DSA addresses intermediary services and, in its enhanced due diligence layer, very large online platforms and very large online search engines that reach at least 45 million average monthly active users in the EU. Its central concerns are transparency, content governance, systemic risk, and user safeguards. The AI Act, by contrast, regulates AI systems and general-purpose AI models through a risk-based structure that ranges from prohibited practices to limited transparency obligations and, at the stricter end, high-risk systems and general-purpose AI models with systemic risk⁶.

Precision matters in this context. The DSA's rules governing recommender systems do not mirror the provisions of the AI Act. Article 27 of the DSA requires providers of online platforms to set out in their terms and conditions the main parameters of their recommender systems and the options available to recipients to modify or influence those parameters. Article 38 then adds a specific obligation for VLOPs and very large online search engines⁷ (VLOSEs) to offer at least one recommender option that is not based on profiling. The AI Act addresses different legal objects. Article 26 governs deployers of high-risk AI systems, requiring use in accordance with the provider's instructions, human oversight, monitoring, and certain requirements pertaining to context-specific information. The instruments therefore converge in practice without creating identical legal tests⁸.

The AI Act goes to the design of governance itself. The AI Act's conformity assessment architecture and lifecycle requirements apply to the system or model. The DSA's systemic risk obligations under Article 34 pertain to the service and to its effects across the EU. A provider may satisfy the requirements of one regime and still encounter difficulties under the other. Article 2(4) of the AI Act makes clear that the Regulation operates without prejudice to other EU law. That is an important coexistence clause. It is not, however, a conflict rule⁹.

The same point becomes sharper in relation to general-purpose AI. Article 3(63) of the AI Act defines a general-purpose AI model by reference to its generality and capability to perform across a wide range of tasks. Articles 51 to 55 then impose a more demanding regime on providers of general-purpose AI GPAI models with systemic risk. Where a platform relies on such a model, platform-level requirements under the DSA duties and model-level requirements under the AI Act duties may both bear have a bearing on the same service environment at once. The Commission's November 2025 interaction report and the Digital Package proposals show that this problem is now formally recognised. What still does not exist is a binding cross-regime procedure telling operators or enforcers how concurrent determinations should be aligned in concrete cases.

² European Commission, [Commission evaluates the Digital Services Act's interaction with other EU laws and its designation threshold for VLOPs and VLOSEs](#), 17 November 2025, accessed 11 March 2026; European Commission, [Digital Package](#), 20 November 2025, accessed 11 March 2026; see also H. Graux and others, *Interplay between the AI Act and the EU digital legislative framework*, European Parliament 2025.

³ European Commission, Commission opens formal proceedings against X under the Digital Services Act, 17 December 2023, accessed 8 April 2026.

⁴ European Commission, Commission preliminarily finds TikTok and Meta in breach of their transparency obligations under the Digital Services Act, 24 October 2025, accessed 8 April 2026.

⁵ European Commission, Commission investigates Grok and X's recommender systems under the Digital Services Act, 26 January 2026, accessed 11 March 2026.

⁶ Article 33(1) of the DSA; European Commission, *Commission evaluates the Digital Services Act's interaction with other EU laws and its designation threshold for VLOPs and VLOSEs* (fn 2). A

⁷ Under Article 33(1) of the DSA, very large online search engine providers are online search engines reaching an average of at least 45 million monthly active users in the EU. See also recital 77 to the DSA.

⁸ Articles 27 and 38 of the DSA; Articles 14 and 26 of the AI Act; European Commission, *Navigating the AI Act* (fn 4), under 'What are obligations of deployers of high-risk AI systems?'

⁹ Article 34 of the DSA; Articles 2(4), 8 to 15, and 43 of the AI Act.

Two enforcement systems; limited cross-regime coordination

The deepest structural problem is institutional. This is not because no cooperation exists at all. Both instruments contain internal coordination structures. Under the DSA, Digital Services Coordinators, the European Board for Digital Services, and the Commission operate within an established cooperation framework. Under the AI Act, the AI Office, the AI Board, national competent authorities and market-surveillance authorities perform parallel governance functions. The difficulty lies elsewhere: these are two separate systems of cooperation, not a single shared procedure encompassing both the DSA and the AI Act¹⁰.

That gap has gained in importance as more active cases have accumulated. The Commission's January 2026 action against X and its Grok-related functionalities is instructive precisely because it shows how platform supervision and AI concerns can converge within the same factual setting. The point is not that every recommender system is a general-purpose AI model; this would be far too crude. It is that some large platform functions now sit at the boundary between service governance and AI governance, with no settled institutional rule for joint handling¹¹.

A narrower and sharper criticism of the current framework therefore emerges. The problem is not the total absence of cooperation. The problem is the absence of a dedicated joint procedure for cases where the same deployment or service function is material to both regimes. A platform should not have to determine, for itself, how a DSA systemic risk assessment will sit alongside a high-risk AI system review, a general-purpose AI model inquiry, or a

supervisory intervention by a sector-specific authority.

Sectoral expertise complicates the picture further. A financial services supervisor assessing an AI-based creditworthiness tool according to the requirements of annex III, point 5(b) of the AI Act¹² brings an understanding of consumer credit, prudential context, and market practices that a horizontal AI authority may not possess. Yet the case for horizontal coordination remains strong. General-purpose models do not respect sectoral boundaries, and inconsistencies across sectoral silos can produce misunderstandings as readily as they produce nuance. The issue is not whether expertise should be centralised entirely. It is whether the system has a disciplined way of combining horizontal and sectoral expertise when both are needed¹².

The cumulative regulatory burden adds to the concern. A firm might now have to consider the provisions of the AI Act, the DSA, the Digital Operational Resilience Act (DORA), the NIS2 Directive (cybersecurity), the Data Act, and sector-specific rules in the same operating environment. Each instrument pursues a specific objective. The problem lies in the aggregate design. As Mario Draghi's 2024 report on European competitiveness emphasised, fragmentation and administrative burden are not neutral frictions. They affect the capacity to scale. Larger firms can absorb the burden of cumulative compliance through the work of their in-house legal, technical, and governance teams. Smaller and mid-sized firms cannot do so as easily. This is not a sweeping argument against regulation. It is a reason to take regulatory interaction seriously as part of market design¹³.

The simplification debate

Seen in this light, the Commission's simplification agenda is plainly relevant. The November 2025 Digital Package does more than simply adjust dates. It also proposes to centralise oversight of AI systems built on general-purpose models in the AI Office, to concentrate oversight of AI embedded in VLOPs and VLOSEs at Commission level, and to clarify the interplay between the AI Act and other EU legislation. That is a significant step because it acknowledges that implementation difficulty is partly

a governance problem rather than merely a scheduling problem¹⁴.

Timing nevertheless remains part of the story. The Commission has itself accepted that the delayed availability of harmonised standards and support tools complicates compliance. The third draft of the General-Purpose AI Code of Practice was published on 11 March 2025, the final version was received on 10 July 2025, and

¹⁰ Articles 56 to 66 of the DSA; Articles 64 to 70 of the AI Act; European Commission, *The cooperation framework under the Digital Services Act*, accessed 11 March 2026; European Commission, *Governance and enforcement of the AI Act*, 14 November 2025, accessed 11 March 2026.

¹¹ European Commission, *Commission investigates Grok and X's recommender systems under the Digital Services Act* (fn 4).

¹² Annex III to the AI Act, point 5(b); European Commission, *Navigating the AI Act* (fn 4), under 'Any examples of high-risk use cases as defined in Annex III?'

¹³ Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA) [2022] OJ L 333/1; Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS2) [2022] OJ L 333/80; Regulation (EU) 2023/2854 on harmonised rules on fair access to and use of data (Data Act) [2023] OJ L 2023/2854; Mario Draghi, *The future of European competitiveness* (Report to the European Council, September 2024).

¹⁴ European Commission, *Digital Package* (fn 2).

the Commission and the AI Board approved it on 1 August 2025. Even so, the standardisation work for high-risk AI systems has continued beyond the original timetable. Operators were therefore correct to say that parts of the framework were formally in force before the supporting compliance infrastructure was fully mature¹⁵.

But simplification should not become a synonym for dilution. The practices prohibited by Article 5 of the AI Act and the systemic risk obligations set out in Article 34 of

the DSA should not be treated like ornamental burdens. They embody legislative judgments about fundamental rights, public safety, and the structural power of very large services. The more pertinent simplification question is different. How can the EU reduce duplication, uncertainty, and inconsistent supervisory pathways without weakening the substantive choices already made in primary law? Framed that way, the case for interaction focused reform becomes much stronger¹⁶.

Towards a more coherent architecture

Three practical implications follow.

First, the EU needs an expressly articulated cross-regime coordination process for recurring intersection scenarios. Content ranking, recommender systems, platform-based generative tools, advertising systems, and AI-enabled moderation functions are obvious candidates. The aim should not be to collapse the DSA and the AI Act into a single instrument. It should be to require a common handling method where one fact pattern engages both.

Second, the next phase of implementation should distinguish more clearly between three levels of analysis: the model, the system built on the model, and the service environment in which the system operates. Many present debates become confused because those categories are merged. A binding coordination framework should require regulators to specify which of the three is under review, and why.

Third, the Commission's simplification agenda should be extended into a joint compliance-mapping exercise, by sector and by use case, that identifies where obligations run in parallel, where one regime adds a distinct layer of review, and where genuine legal tension remains. Some of that work can be done through guidance, coordinated templates, and administrative arrangements. But where primary law produces a real conflict, the honest answer is that guidance will not suffice. A conflict rule or priority rule with real legal force may require legislative amendments. The Commission's own 2025 package suggests as much by pursuing change at proposal level¹⁷.

None of this requires the risk-based logic of the AI Act or the systemic logic of the DSA to be abandoned. On the contrary, both logics are worth preserving. What is needed is an architecture that treats them as components of one regulatory ecosystem rather than as neighbouring projects that happen to affect the same operator.

Operational coherence as a key part of digital governance

The DSA and the AI Act are not failures of legislative design. Each represents a serious attempt to govern a difficult digital environment. The weakness lies in the space between them. As implementation advances, the question is no longer whether the EU can legislate ambitiously. It is whether it can make ambitious legislation work coherently when multiple regimes converge on the same technical and commercial reality.

That is why the real challenge is institutional integration.

The simplification debate is useful only if it focuses not simply on lighter obligations or later dates, but on operational coherence. The EU's credibility in digital governance will depend less on the number of rules it has enacted than on whether those rules can be applied together, predictably and fairly.

Building that coherence is now the harder task. It is also the more important one.

¹⁵ European Commission, [Drawing up a General-Purpose AI Code of Practice](#), accessed 11 March 2026; European Commission, [The General-Purpose AI Code of Practice](#), 10 July 2025, accessed 11 March 2026; European Commission, Navigating the AI Act (fn 4), under 'What would be the role of standardisation in the AI Act?.'

¹⁶ Article 5 of the AI Act; Article 34 of the DSA; Charter of Fundamental Rights of the European Union [2012] OJ C 326/391.

¹⁷ European Commission, Digital Package (fn 2).

Artificial intelligence and protecting knowledge systems in the European Union

Providing a strategic context for the EU's digital transition

By Professor Raluca Csernatonu, Centre for Security, Diplomacy and Strategy, Brussels School of Governance, Vrije Universiteit Brussel, and Fellow at Carnegie Europe



© titima157/ depositphotos.com

Artificial intelligence is fundamentally changing how people across the world produce, share and assess the reliability of knowledge. Professor Raluca Csernatonu argues that “epistemic security”, i.e. protecting and ensuring reliable knowledge systems, should be a strategic objective in the European Union’s digital agenda, covering regulation, industrial policy, democratic governance and public audit. She argues that the EU’s move towards deregulation in relation to AI risks compounding regulatory weaknesses in certain areas, just when trusted and secure knowledge systems are ever more vital to underpin EU democracies. Europe’s AI debate is still too often presented as a trade-off between innovation and regulation but this misses the point. AI in all its forms is reshaping the publicly available systems where people go to seek information and knowledge. Making these systems secure and reliable is therefore a test of Europe’s digital sovereignty, democratic resilience and technological might.

From regulatory dilemmas to questions about the reliability of public knowledge and information

Europe's debate over artificial intelligence (AI) reflects the familiar opposing views of many policy arguments. Is it better to have innovation or regulation? What about strategic autonomy versus transatlantic alignment? How do we balance competitiveness versus fundamental rights? All these questions capture something important, yet each also focuses insufficiently on a deeper transformation underway across European democracies. Generative AI (which creates new content by learning from existing data) and agentic AI (which can act autonomously with minimal human input) are altering how we produce, distribute and validate knowledge. Thus the political question of how the EU regulates AI systems is inseparable from considering the security of those systems. Whose models select and organise what information for whom, and can they be trusted?

In this article, I trace how the EU's regulatory architecture, industrial strategy and audit practice address this challenge, and where they fall short. I argue that the move towards deregulation that began in late 2024 (and was further signalled at the informal European Council retreat at Alden Biesen in early 2026, with input from the [Draghi report on European competitiveness](#), the [Letta report on the single market](#), and the European Commission's [Competitiveness Compass](#)) risks compounding existing regulatory weaknesses. My analysis closes with five priorities to embed epistemic security in all aspects of the

European digital agenda. My arguments build on earlier work on [the EU's AI power play between deregulation and innovation](#) and I place that debate within a broader view of how AI is reshaping the European public sphere and democratic debate.

Not only is AI creating new regulatory problems for the EU, it is also changing the conditions under which democratic societies find out about, decide on, and hold those in power to account. Put differently, epistemic security could be considered part of democratic survival in the AI transition. Can democracies still produce reliable, pluralistic, and decision-guiding knowledge for elections, public authorities, in relation to crisis response, and to support societal trust, given that synthetic content, automated recommendations and searches using AI models increasingly shape what the general public sees as real?

For public audit institutions, this raises a new question. Auditing AI cannot stop at legality, procurement, or technical performance. It must also ask whether systems where AI actively shapes outcomes rather than just passively processing information uphold traceability, contestability, evidence quality and public accountability. Auditing epistemic security could provide a valuable review of many aspects of the AI transition by first breaking them down into an audit task plan.

ECA JOURNAL SHORT READ

AI systems **are altering** how people, public authorities and democratic institutions **produce, distribute and validate knowledge**. Just one example is how AI systems are influencing the legitimacy of those standing for public office. "Epistemic security" refers to protecting the **integrity of all the systems** through which society produces, shares and assesses the reliability of knowledge. Providing epistemic security would mean **regulating how public information is produced**, i.e. requiring traceable sources, contestable evidence and accountable intermediaries. The regulatory challenge arises because ways of checking evidence and accountability are being left behind as AI develops rapidly. Epistemic security offers a practical framework to **align technological innovation with the ability of democratic systems to continue operating effectively**.

The EU's regulatory architecture on AI is very detailed but has two main weaknesses: **fragmentation and a high risk of obsolete rules**. There is also a tendency towards **simplification**, but this carries a **risk for epistemic security**. Coordinated action based on shared rules and principles is missing. The aim should be to have control ("sovereignty") not just over "chips and clouds" but over the systems providing public information and knowledge. A number of EU initiatives exist in this area but they are not joined up. Addressing these issues and plugging governance gaps has become even more urgent in view of emerging "agentic AI" systems, which can act autonomously with minimal human input. Against the background of rapid developments in AI, protecting knowledge systems and ensuring they are reliable should focus on four key areas: **traceability, contestability, evidence quality and accessibility to test systemic effects**. If public auditors are to assess the overall performance of AI systems this **places demands on them** in terms of methods, capabilities, accessibility of information and cross-border cooperation.

There are **five priorities** which should be part of a European epistemic security strategy, ranging from making such security a central objective, to developing a suitable AI framework with proper guidance on using AI agents. AI models are reshaping the systems through which European society produces, shares and evaluates information. **Epistemic security** is the **strategic context** through which **Europe's autonomy, prosperity and democracy will be tested in the AI age**.

What is epistemic security?

Epistemic security refers to protecting the integrity of all social and technical systems through which societies produce, circulate and validate knowledge (see **Box 1**). Before the digital age, these systems were found in a relatively stable group made up of public broadcasters, peer-reviewed journals, statistical agencies, archives, electoral commissions, and a published press whose business model rewarded accuracy. Each contributed to what the general public experienced as a shared and contestable factual baseline. The concept of epistemic security does not presuppose a hostile adversary and also applies even when the erosion of the factual baseline stems from market forces, a wider diffusion of technology, or institutional decay. This is why AI poses a distinct epistemic challenge: it does not simply add new channels for misinformation, it is reorganising the systems through which knowledge gains visibility, authority and credibility.

[Generative AI](#) reduces the cost of producing plausible content to nearly zero, meaning that all sorts of content floods into the public arena. Recommendation systems concentrate on emotionally charged material, hollowing out the intervention of editors historically provided by public broadcasters. [Agentic AI](#) introduces autonomous and synthetic actors capable of sustained interaction with humans, collecting large amounts of information from the public arena and placing a lot of crafted material into

the public sphere. The cumulative effect is weakened audit trails for information “provenance”, displacement of primary sources by outputs from AI models and blurring of the boundaries between real speech and orchestrated influence.

AI is already shaping electoral legitimacy and the public record on which democratic accountability depends. Coordinated efforts to influence what people see, think, or believe using information enhanced by AI tools featured prominently in [Slovakia's parliamentary elections in 2023](#) and the [presidential campaigns in Romania in 2024 and 2025](#). The 2025 India-Pakistan crisis around attacks and killings of opponents showed how AI-generated and recycled media could contaminate the information environment during a period of acute geopolitical tension, producing false evidence of attacks, losses, and political declarations well before journalists, officials, or the general public could check whether the information was actually true. These examples share a common feature: AI systems are now embedded in the production of information which itself sometimes has serious societal consequences. There is thus an ongoing challenge to secure knowledge systems in both civilian and military contexts, and appropriate regulation cannot be neatly split between them.

Box 1 – Epistemic security and adjacent concepts

Epistemic security refers to protecting the integrity of all the systems through which a society produces, circulates and validates knowledge, and ensuring the resilience of those systems under technological, economic or political pressure. It does not require a hostile actor.

Hybrid threats denote the integrated use of military and non-military instruments by hostile state actors below the threshold of armed conflict. Using or mis-using information is just one component of such behaviour. The context is actor-centred and adversarial.

Cognitive warfare, a more recent NATO concept, designates adversarial efforts to shape perception, judgement and decision-making in target populations through coordinated use of information, including AI, and behavioural science. This concept also presupposes hostile intent.

Psychological warfare is the classical military discipline of using communication to influence the emotions and behaviour of audiences in conflict situations.

Knowledge systems are where knowledge is produced and stored, such as on computers, using AI models, in archives, via broadcasters, or in peer reviews, libraries or schools.

Hybrid, cognitive, and psychological warfare reflect hostile intent. Knowledge systems underpin knowledge and information production. Epistemic security cuts across all these areas and reflects what is at stake in the AI transition.

The point is not to securitise truth or to create a state monopoly over legitimate knowledge. Rather it is to make regulating the production of knowledge possible, with traceable sources, contestable evidence, accountable intermediaries and institutions capable of explaining how information at least partly produced by AI has been used to support decisions that affect the general public.

These distinctions matter because epistemic security is broader than counter-disinformation, platform regulation, or cyber defence. It concerns protecting the integrity of all the systems through which societies produce, circulate, verify and contest knowledge. It is not about policing opinion, imposing official truth, or treating all information

disorder as foreign interference. The regulatory challenge is that AI systems are being deployed more quickly than public institutions, regulators, courts, media organisations and the general public can adapt how they check facts and accountability. Epistemic security offers a practical framework to align technological innovation with the ability of democratic systems to continue operating effectively (“democratic resilience”). It directs attention to information provenance, auditability, transparency, data access, the ability of institutions to work effectively and the trusted public knowledge systems needed to sustain a shared factual baseline.

The EU’s regulatory architecture and its weaknesses

The EU has constructed an unusually detailed regulatory framework for digital governance. The [AI Act](#), in force from August 2024, establishes a risk-based regime with prohibitions, high-risk obligations, and transparency duties for general-purpose AI models. The [Digital Services Act](#) (DSA) and the [Digital Markets Act](#) (DMA) govern intermediary liability, systemic risk assessments for very large platforms, and competition in digital markets. The [Regulation on the transparency and targeting of political advertising](#) addresses paid political messaging. The [European Media Freedom Act](#) protects media pluralism and editorial independence. The [European Democracy Shield](#), presented by the Commission in November 2025, sets out around fifty action points across information integrity, elections and media, and societal resilience, although its instruments remain largely voluntary and spread across existing frameworks.

Assessed against the epistemic security challenge, there are two structural weaknesses. The first is fragmentation. Each legislative instrument addresses one segment of the information ecosystem, with distinct competent authorities, reporting cycles and methods of enforcement. Coordination among the [European AI Office](#), the [European Board for Digital Services](#), national audiovisual regulators, data protection authorities,

and the [European Centre for Algorithmic Transparency](#) remains underdeveloped. The Commission’s [AI Liability Directive](#) was intended to close some of these gaps, but it has been withdrawn.

The second weakness is related to timing. Rulemaking cycles spanning several years are colliding with model release cycles measured in months. By the time secondary legislation is adopted, the underlying technology has already shifted. This mismatch produces what could be termed “premature normative closure”, i.e. locking in rules before the technology has stabilised. Premature closure is doubly costly. It generates rules that quickly become obsolete and also signals to industry that European frameworks are lagging behind operational reality. The epistemic security challenge requires regulation that can develop alongside the technology, anchored in flexible standards, robust evaluation regimes and continuous oversight.

Epistemic harm is difficult to regulate after the fact. Once synthetic narratives, manipulated evidence, or model-generated errors circulate through public debate, the democratic cost is not only individual misinformation but delayed correction, contaminated archives and weakened trust in any subsequent verification.

Alden Biesen and the deregulatory drift

The [informal European Council retreat at Alden Biesen](#) in early 2026 underlined the fact that the bloc was at a political crossroads. Concerns about European competitiveness and the advance of American and Chinese AI ecosystems began to crowd out previous emphasis on guardrails in policy discussions. The Draghi and Letta reports, and the Commission’s Competitiveness Compass increased calls for simplification, revisions of EU legislation in single “omnibus” packages and selective relaxation of digital obligations. Several member states, along with parts of the Commission, are now even

treating the AI Act’s transparency requirements and the DSA’s risk assessment provisions as costs to be cut.

This adds another layer to the autonomy debate. A Europe dependent on external models, cloud infrastructures, ranking systems and provenance standards that all come from outside the EU is not just dependent on foreign industrial capacity. It is also reliant on foreign presentations of reality, i.e. how knowledge in the European public sphere is classified, retrieved, summarised and checked.

This deregulatory drift, however, carries risks to the integrity of knowledge that are seldom acknowledged in the competitiveness debate:

- loosening obligations on general-purpose AI providers reduces the incident reporting and transparency information that auditors, researchers and oversight bodies need to assess systemic risk;
- diluting access to platform data under [DSA Article 40](#) constrains the public-interest research that supports evidence-based policymaking;
- withdrawing the AI Liability Directive removes one of the few channels through which injured parties could seek redress for damage resulting from the use of AI.

Each of these changes can be justified from a competitiveness point of view. However, each also could be seen as undermining the reliability of systems where knowledge is created and organised and on which democratic governance depends.

The missing link is coordination. A strong system for protecting knowledge would connect these initiatives through shared rules and principles, securing the information supply chain from production and provenance through distribution, evaluation, audit and democratic use. This would make EU digital sovereignty not only a means of building capacity in this area but also of protecting democratic knowledge systems.

Strategic autonomy and the integrity of European knowledge

European strategic autonomy has been successively described as industrial autonomy, defence technological autonomy and digital sovereignty. Each of these descriptions brings into focus a different layer of technology and the need to build greater independence from foreign suppliers. An epistemic security framework would require a further breakdown of the components of European knowledge: computers, models, evaluation regimes, content-provenance standards and trusted information channels through which the general public in Europe encounters reality. Sovereignty over chips and clouds is necessary; sovereignty over how and where the public seek knowledge is indispensable.

Several EU initiatives point in this direction. The [EuroStack proposal](#) aims to assemble a connected system of computers, clouds, general AI models and applications across Europe. The [AI Factories programme](#) channels supercomputing capacity toward European AI startups and researchers. [GAIA X](#) and the [Common European Data Spaces](#) try to address data infrastructure needs. [The European Centre for Algorithmic Transparency](#) provides specialist support for DSA enforcement. The [open technical standard for content provenance](#), provided by the Coalition for Content Provenance and Authenticity (C2PA), and which several European actors are actively promoting, offers a technical baseline to authenticate media. These initiatives are not yet joined up and have still to be reworked as components of a coherent epistemic

security architecture connecting systems, regulation, evaluation and audit.

This is also a competition problem. The risks from contaminated, manipulated, or unverifiable information are spread across society, while the systems that produce this remain highly concentrated with just a few providers. Journalists, regulators, researchers, public authorities, and the general public bear the burden of verification, correction, audit and repairing trust, while dominant platforms continue to control visibility, amplification and monetisation of content. Hence, a coherent epistemic security architecture would have to connect digital sovereignty with competition policy. It would need to reduce dependence on a small number of private actors, widen access to trusted data and audit capabilities, and ensure that the burden of protecting public knowledge does not fall on already overstretched democratic institutions.

Overall, this shows a clear gap between intent and outcome. Europe is investing in the technical foundations of digital sovereignty but has not yet defined how those foundations will protect the integrity of public knowledge. Computing capacity, trusted data spaces, algorithmic oversight and content provenance are treated as separate policy areas rather than as mutually reinforcing safeguards for democratic resilience.

Agentic AI and governance gaps

A more demanding regulatory challenge is now emerging. Agentic AI systems combine large language models with tools, memory and the ability to act in digital environments. They are being deployed across business, public administration and consumer contexts. Their implications for knowledge differ qualitatively from those of generative AI models. An agentic system can

browse, summarise, transact, communicate and continue working across different interactions with the model. It can compile a person's regulatory exposure, draft and file a complaint and follow-up on outcomes. It can also be instructed by malicious or irresponsible actors to engage in large-scale and coordinated information production.

Current EU AI legislation was drafted against a background of general AI models and recommendation systems. However, agentic AI systems sit awkwardly across the AI Act's categories, the DSA's intermediary regime, the [EU's data protection framework](#), and the cybersecurity baseline in the [second Network and Security Directive \(NIS2\)](#). Liability allocation among the model developer, agent operator, deployer and end user has not been settled. Verifying the identity and authority of synthetic agents acting on behalf of humans has also not been fully developed. Signals to indicate the source of content generated by AI agents are inconsistently implemented. Each of these gaps is also an epistemic security gap. Closing them will require coordinated policy work across the European AI Office, the European Union Agency for Cybersecurity ([ENISA](#)), the European Data Protection

Board ([EDPB](#)), the [European Board for Digital Services](#), and the developing jurisprudence of the Court of Justice of the European Union ([CJEU](#)) on damage resulting from AI.

A practical epistemic security audit task plan would therefore ask four questions:

1. Can public services delivered or supported by AI be traced to original sources?
2. Can the general public contest automated outputs?
3. Can public bodies provide the evidence for AI-supported decisions?
4. Do external researchers and auditors have sufficient data access to test systemic effects?

Auditing the epistemic ecosystem

Public audit institutions have a specific contribution to make in this area. The European Court of Auditors examined the EU's AI ambitions and governance in [special report 08/2024](#) and the semiconductor strategy in [special report 12/2025](#). The ECA's work programme for 2026 and beyond focuses on digitalisation across many different topics. National supreme audit institutions in the EU have begun to develop ways of auditing AI systems used in public administration and assessing readiness to comply with the AI Act.

Auditing the effect of AI on knowledge and information has specific methodological demands. It requires the ability to assess the effects of system deployment, including concentration of content, displacement of sources and information imbalances between platforms and users. It calls for robust data-access provisions that withstand deregulatory pressure. It depends on cross-border cooperation between audit institutions, as the relevant data flows and platform operations go beyond national boundaries. The European Court of Auditors, working alongside its partners in the Contact Committee and the EUROSAI Working Group on the Audit of IT, is well placed to develop a shared methodological framework to audit public services delivered or supported by AI. Such work would extend the remit of public audit from verifying compliance to assessing the performance of the whole AI ecosystem, a shift that the technology now demands.

This shared work should avoid the familiar EU trap of treating all aspects in the same way. It should identify specific overlaps in the system:

- who has access to platform and model data,
- who validates provenance standards,
- who audits public-sector deployment,
- who maintains crisis-communication channels, and
- who is accountable when information failures resulting from the use of AI affect democratic processes.

Resilient knowledge systems depend on a broader group of independent actors: investigative journalists, academic researchers, fact-checking organisations, communities in civil society working with technology, and the public service media that the European Media Freedom Act seeks to protect. Ongoing funding through Horizon Europe, the Creative Europe programme, the European Democracy Shield and the Digital Europe Programme should be organised to support these groups and treat them as critical democratic structures.

This becomes even more urgent in a political context where regulatory simplification, including through the [Digital Omnibus](#) agenda, may weaken or dilute obligations for powerful digital actors. An epistemic security approach would also require counterbalancing measures. Where regulation falls short, the EU should strengthen public interest audit, independent research access, media pluralism, and ways in which the general public themselves can verify facts. This should make the EU better able to sustain trustworthy public knowledge despite uneven regulation across the world.

Five priorities for a European epistemic security strategy

First, treat epistemic security as a strategic objective in its own right, coordinated across the EU institutions. Without a horizontal framework, the existing instruments will continue to operate separately and fail to come together sufficiently.

Second, defend and operationalise data access for public-interest research under DSA Article 40 and the equivalent provisions of the AI Act. Without being able to view platform and AI model behaviour, regulation operates blindly. The European Centre for Algorithmic

Transparency, if properly resourced, can become the analytical backbone of European epistemic security.

Third, accelerate the deployment of content-provenance standards across European public communication. EU institutions, public broadcasters, electoral authorities and statistical agencies should, as a minimum, use C2PA-compliant digital signatures to prove the source and authenticity of official communications. This would build trust through technical standards as much as through institutional reputation.

Fourth, develop an agentic AI framework that addresses identity verification, action authorisation, audit logging and cross-system interoperability. ENISA, the AI Office, and the EDPB should jointly issue guidance on using AI agents in sensitive sectors, beginning with public administration, healthcare, and electoral processes. [European standardisation work in the three](#)

[EU standards organisations](#) should be prioritised. The strategic implication is that epistemic security is not a soft democratic add-on to the AI agenda but is needed so that competitiveness, autonomy, and a rights-based regulatory framework remain politically meaningful, providing a trustworthy basis for the general public to make decisions affecting their everyday lives.

Fifth, however improbable, sustain transatlantic and transregional coordination on epistemic security through successor arrangements to the Trade and Technology Council, as well as through the [Council of Europe Framework Convention on AI](#), the [G7 Hiroshima AI Process](#), and the [OECD's AI governance work](#). European positions are most credible when working in coalition with like-minded partners and expressed through shared standards.

Regulating the integrity of European knowledge

To conclude: the EU's digital transition stands at a crossroads. The familiar trade-offs between innovation and regulation will continue to dominate political debate. The deeper transformation, however, lies in the shifting conditions of public knowledge itself. AI models are restructuring the systems through which European society produces, shares and evaluates information. Whether this restructuring strengthens or erodes democratic debate will depend on choices made over the next legislative cycle.

It is important to recognise epistemic security as a strategic priority that should be embedded across the

EU's regulatory and industrial agenda and supported by renewed public audits, which would equip European institutions to treat and regulate the technology with the seriousness it now demands. The alternative is a continent that builds capacity in computing and chips while forfeiting capacity in the more fundamental resource on which any democracy ultimately depends, namely a shared, contestable and trustworthy public reality. Epistemic security is the strategic context through which Europe's autonomy, prosperity and democracy will be tested in the AI age.

‘I want this continent to become an AI Continent’

Interview with Henna Virkkunen, European Commission Executive Vice-President for Tech Sovereignty, Security and Democracy

By Gaston Moonen



© Mikko Mäntyniemi

When the first Von der Leyen Commission started work in December 2019, its digital ambitions were reflected in the package known as ‘A Europe fit for the digital age’. The EU’s digital transformation has remained a top priority in the second Von der Leyen Commission, as was confirmed when Henna Virkkunen was made Executive Vice-President responsible for the Digital and Frontier Technologies portfolio. Her role was clearly linked to security and democracy, with the aim of ensuring freedoms, justice and democracy. This connection is reflected in her portfolio title: ‘Tech sovereignty, Security and Democracy’. In this interview, she explains what the Commission is doing to create a resilient and independent Europe on the digital front, particularly in the area of AI.

Europe’s independence is our motto

At the European Commission, where you are responsible for Tech Sovereignty, Security and Democracy, you cover three foundational pillars of the EU. How do you view the relationship between these three areas, and what is your biggest concern for your portfolio?

In today’s world, technology and security are inseparable.

”

Henna Virkkunen: In today’s world, technology and security are inseparable. Those who master technology also master the world. In recent years, we have experienced many shocks that have shaped the security landscape in Europe.

However, geopolitical shocks can also act as a catalyst and an opportunity – and this is exactly what has happened in Europe. We are now determinedly building the capacity to defend our sovereignty across several critical sectors: technology, energy, raw materials, security, and competitiveness.

This means reducing our high-risk economic and security dependencies. We must ensure that we have the capability and sufficient capacity to act independently in any situation, while also maintaining reliable partnerships. The good news is that we are doing a lot right now, and in my view, the right things.

According to some researchers, at least 85% of the high-tech systems used by the authorities in EU member states depend on US Big Tech technologies and support. You have stressed the importance for the EU to create its own AI and tech capacity. What do you think the key impediments to building that capacity are; and what are the Commission's key actions in the short and medium terms to reduce this dependency and create European Big Tech players?

Henna Virkkunen: Europe's independence moment is this Commission's motto, and this spring we will present the Technological Sovereignty package to accelerate our autonomy in the digital space. Not because we want to do everything by ourselves, but because we want to retain our status as an innovative, competitive market that is relevant for doing business with our trading partners worldwide.

Europe's independence moment is this Commission's motto... ”

As I have already said, I want this continent to become an AI continent, ready to flourish in the decades to come. To achieve this, we need to take action, and we need to do so decisively as we are at a 'now or never' moment. With the Technological Sovereignty package, we are addressing the most important foundations of the AI continent: chips and compute. In the upcoming package, we do so by addressing the supply of and demand for semiconductors and cloud services, and by creating a virtuous ecosystem that will stimulate market dynamics in all sectors of the economy, but without depending on grants or subsidies. This is being supplemented by EU programmes such as the AI Gigafactories project, which will build very large-scale supercomputers to power the AI age.

With this strategy, we will boost our internal capacities and reduce dependence on outside players in sensitive areas, while of course remaining open to the world. In doing so, we will consolidate not just our competitiveness, but also our security.

Protecting citizens through regulatory dialogue

Besides creating a regulatory framework for all things digital, the European Commission has the power to fine those who break its rules. The Commission has issued numerous fines, particularly against US and Chinese Big Tech companies, such as Meta, Microsoft, X and AliExpress, fines that were often subsequently challenged by Big Tech companies with deep pockets. How effective is it to fine these companies when their subsequent actions do not demonstrate any intention to comply, and can you do anything to change this situation?

...our intention is not to sanction companies, but to engage in regulatory dialogue ... ”

Henna Virkkunen: First of all, our intention is not to sanction companies, but to engage in regulatory dialogue to ensure that EU law, as adopted by the Parliament and the Council, is effectively complied with.

...the broader approach underpinning the DMA and the DSA: promoting behavioural changes and fostering a culture of compliance... ”

Under the Digital Services Act (DSA) and the Digital Markets Act (DMA), the Commission has opened investigations and taken decisions concerning companies such as Meta, Apple and X, always with the objective of securing compliance and ensuring fair and safe digital markets. The Commission continuously engages with platforms and gatekeepers through regulatory dialogue, and closely monitors implementation of the rules. In several cases, including investigations into TikTok and AliExpress under the DSA, proceedings were concluded with commitments, demonstrating the companies' willingness to cooperate with the Commission and adapt their practices in line with EU requirements. This reflects the broader approach underpinning the DMA and the DSA: promoting behavioural changes and fostering a culture of compliance, while ensuring effective enforcement – including issuing relevant fines, where necessary.

The ultimate objective of DSA and DMA enforcement is to make compliance with the law the preferred and sustainable outcome, for the benefit of citizens' safety and rights online.

What is the Commission itself – including individual Commissioners – doing to promote alternative social media based in Europe?

“
...the broader approach underpinning the DMA and the DSA: promoting behavioural changes and fostering a culture of compliance...

Henna Virkkunen: Through the [European Democracy Shield](#), the Commission is stepping up support for innovative media projects that develop new formats and content for audiences across the Union and beyond, while also exploring future pathways for Europe's tech environment.

The initial focus is on next generation social networking and social media platforms, with the aim of strengthening EU digital sovereignty. This work also looks to support the development of pan-European platforms that can make real-time news and information from professional media outlets available to wider audiences across the EU in multiple languages. Today, we are encouraged to see European social media players being launched.

Various companies in the EU face the risk of being hacked, and SMEs in particular are grappling with issues of cybersecurity and possible blackmail or coercion by hostile regimes aiming for destabilisation. At the same time, data protection and privacy authorities are increasingly concerned about disinformation and fake news, sometimes aimed at influencing democratic processes. What do you think entrepreneurs and citizens can do about these issues?

“
Improving media literacy is one of the Commission's priorities...

Henna Virkkunen: Improving cybersecurity remains a top priority for the Commission. Ransomware is a cybersecurity threat which, in the worst cases, can put companies out of business. It poses a threat to our critical infrastructure, and also imposes a cost on our entire economy. For companies, taking the essential cybersecurity measures to protect against this threat is worth the investment. I am talking here about common-sense measures such as ensuring back-up management, training staff on cyber hygiene practices, and using authentication solutions such as multi-factor authentication.

In cyberspace, each one of us can contribute to our shared security by exercising good cyber hygiene. Similarly, individual citizens can contribute to the resilience of our democracy by taking a more critical approach to disinformation and fake news. Improving media literacy is one of the Commission's priorities, and is reflected in the Democracy Shield.

Focusing on simplification, regulatory delivery, and enforcement

Echoing various US Big Tech companies, several EU member states have voiced the need for less regulation in the digital realm so as to avoid EU rules stifling innovation and competitive growth in digital tech. Some have even complained that certain elements of EU legislation detract from free speech. Is this why the Commission has launched its Digital Omnibus Regulation proposal, and how do you balance EU citizens' digital rights with tech entrepreneurs' calls for less bureaucracy?

“
Regulatory simplification is one of my biggest goals...

Henna Virkkunen: Regulatory simplification is one of my biggest goals in my current role. Our rules and laws are very important, but too often their accumulation and interaction have resulted in a heavy administrative burden for EU businesses. This leads to exactly the opposite of our objective to make Europe more innovative and competitive, because businesses should spend their time on what they do best: their business.

This is why I have presented the AI Omnibus proposal and the Digital Omnibus proposal, both of which contain real simplification measures that could lead to at least €5 billion of cost savings for European businesses by 2029.

...regulatory simplification is not the same as deregulation.

”

Two things are important here. First, we are doing this out of necessity for our own start-ups, SMEs and businesses, who are clearly calling for it. We are not doing it because we are coerced to do so by third countries or foreign corporations. Second, regulatory simplification is not the same as deregulation. Innovation will not thrive in a lawless space. And companies will not develop world-class products and services in a fragmented European single market, where rules are at times unclear.

Balance has been key to the omnibus process...

”

I am convinced that regulatory burdens and lack of clarity can be fixed, without lowering any of our fundamental rights protections. This formed the basis for the Digital Omnibus proposal, and I am proud of what we have put forward. For example, with the proposal for a single-entry point for cybersecurity incident notifications, we will unlock an estimated €41.5 million of savings for companies each year. We are doing this simply by offering a solution so companies can report an incident once, while ensuring the information reaches all the necessary authorities.

Balance has been key to the omnibus process: with the Digital Omnibus on AI, the AI Act continues to safeguard safety and fundamental rights while enabling companies to become competitive. I am committed to simplifying the digital rulebook and delivering on its application and enforcement.

With the adoption of the DSA, the DSM and the AI Act, a whole regulatory framework has been created, but the proof of the pudding often lies in implementation or – as some would call it – enforcement. How could the European Court of Auditors help you in your work, i.e. what kinds of audit in which areas would you like to see from the ECA in the coming year?

Implementation is an issue that is often neglected.

”

Henna Virkkunen: Implementation is an issue that is often neglected. Public debate often dies down swiftly once a piece of legislation has been adopted. However, there usually remains quite a lot to be done, especially at EU level, because implementation typically involves several layers. For example, for the AI Act, the entire governance structure which the Act envisages actually has to be created and made to function.

Similarly, for the DMA legislation, the Commission acts as the sole enforcer. The Commission has very recently also conducted its first DMA review, which points out that in the first two years of application the DMA has opened up new opportunities for businesses and developers, while giving users more control over their experiences and devices. The Commission is committed to continuing to implement the DMA effectively.

...enforcement powers are split between the Commission and national Digital Services Coordinators.

”

Under the DSA, enforcement powers are split between the Commission and national Digital Services Coordinators. The Commission is in charge of enforcement regarding very large online platforms (VLOPs) and search engines. And we take this very seriously: just this year, we have opened proceedings against X – for Grok’s production of potential child sexual abuse material; Shein – for selling sex dolls that look like children; and Snapchat – for age assurance, account settings, and the grooming/recruitment of minors for criminal purposes. We have also preliminarily found several VLOPs in breach of the DSA: TikTok for its addictive design; four very large porn platforms for failing to carry out proper age verification; and Meta for not doing enough to keep children under 13 off Facebook and Instagram. Last December, we fined X €120 million for breaching its transparency obligations under the DSA. On the European Board for Digital Services, we work towards consistent application of the DSA and effective cooperation between national authorities and the Commission. We even join forces on specific cases. For example, the Netherlands Authority for Consumers and Markets [ACM] started the investigation into Snapchat, which we are now carrying out in close cooperation with the Authority. We really are a ‘DSA Enforcement Team Europe’.

As for the audits that the European Court of Auditors intends to carry out, I would always defer to the ECA’s institutional independence.

The European Commission's Cloud Sovereignty Framework: a new brick for building digital sovereignty

By Vincent Coudrin and Phillippe Merle, Directorate-General for Digital Services, European Commission



© European Union, source: European Commission.

The EU's digital transformation and AI capacity development need enablers such as data centres and microchip production to create an AI ecosystem that addresses its strategic autonomy goals. This is all the more important given the rapid geopolitical changes the EU is facing. The continuing rise of AI has also created a surge in data capacity, cloud services and microchip production. So how can the European Commission stimulate the growth of an EU cloud market, essential for EU digital independence? Philippe Merle, Head of Unit, Cloud Services & Software Broker, and Vincent Coudrin, Independent Cloud Expert, both with the European Commission's Directorate-General for Digital Services, elaborate on the Sovereignty Framework developed by the Commission. This framework applied to the Sovereign Cloud call for tender, launched to procure sovereign cloud services for EU institutions, bodies, offices and agencies worth up to €180 million over 6 years. Through this call for tender, the Commission aims to redefine what 'sovereign' means in practice for cloud services.

From cloud adoption to sovereignty imperative

The European Commission's Cloud Sovereignty Framework represents a significant evolution in the European Union's approach to digital infrastructure. Initially centred on promoting cloud adoption, the EU's strategy has gradually shifted towards securing strategic autonomy in a context defined by geopolitical instability, technological concentration, and growing dependency on non-European providers.

The European Commission's 2019 cloud strategy was part of a broader global wave of 'cloud-first' public policies aimed at accelerating digital transformation. At that time, the central policy challenge consisted of balancing the

benefits of cloud computing – such as scalability, agility, cost-efficiency, and innovation – with its risks, including vendor lock-in, cybersecurity vulnerabilities, legal exposure, and operational dependency.

While the benefits of cloud adoption clearly outweighed the risks in 2019, the global context has since undergone profound transformations that have reshaped the risk landscape and rendered the original balance obsolete. The Cloud Sovereignty Framework emerges as a response to these changes, seeking to redefine the criteria by which cloud services are evaluated and procured.

A changing risk landscape (2019–2025)

Several major global developments have exposed new vulnerabilities and highlighted the strategic nature of cloud infrastructure.

The COVID-19 crisis

The COVID-19 pandemic demonstrated both the necessity and fragility of cloud infrastructure. Governments across Europe relied heavily on cloud services to deploy emergency digital solutions, including remote education platforms and vaccination systems. However, the crisis also revealed supply chain vulnerabilities, particularly when disruptions in semiconductor production constrained cloud capacity¹.

This dual lesson – critical dependency combined with limited control – highlighted the risks of relying on globally distributed and externally controlled infrastructure.

Russia's war against Ukraine

Russia's war of aggression against Ukraine further emphasised the strategic dimension of cloud infrastructure. The Ukrainian government's ability to continue operating was largely enabled by migrating its systems to US-based cloud providers². While this demonstrated the resilience and capabilities of cloud technologies, it also underscored Europe's reliance on non-European providers for critical digital capabilities, including cybersecurity and advanced data processing³.

The rise of generative AI

The emergence of generative AI since late 2022 has significantly increased the strategic importance of cloud computing. AI development is deeply dependent on massive computational resources, which are largely controlled by a small number of US hyperscalers. The exponential growth of their investments highlights a widening technological gap between Europe and global leaders⁴. This raises concerns that Europe may lose not only infrastructure sovereignty but also the ability to develop and control next-generation AI technologies.

Data protection volatility

The instability of transatlantic data transfer frameworks, exemplified by repeated legal challenges and governance uncertainties, has further increased the risk of relying on external legal systems. The governance of personal data remains vulnerable to political and institutional instability, creating uncertainty for European stakeholders⁵.

Geopolitical signalling

Statements made at international forums, such as the Paris AI Summit in 2025, reinforced the perception that Europe risks being relegated to a consumer role in global technology ecosystems dominated by non-European players⁶. The sanctioning of the judges of the International Criminal Court tolled the death knell for European illusions about the real level of tension.

Strategic autonomy as a policy objective

In response to these developments, the EU has elevated digital sovereignty to a strategic priority⁷. The concept refers to Europe's ability to control its digital

infrastructure, technologies, and data in accordance with its values, laws, and interests.

¹ Microsoft confirmed as early as March 2020 facing such difficulties due to the closure of semiconductor factories <https://azure.microsoft.com/en-us/blog/update-2-on-microsoft-cloud-services-continuity/>

² Kaido Einama, *The Cloud Saved a Nation: How Ukraine Backed Up an Entire Country During the War*, in: *the Baltic Sentinel*, 26 October 2025.

³ Rita Konaev, *How Tech Giants Turned Ukraine Into an AI War Lab | Center for Security and Emerging Technology*, in: *Time*, 8 February 2024.

⁴ Charles Fitzgerald, *Platformonomics - Follow the CAPEX: The Clown Car Race Checkered Flag*, in: *Platformonomics*, 11 April 2024.

⁵ Silvia Lorenzo Perez, *What the PClOB Firings Mean for the EU-US Data Privacy Framework*, *Center for Democracy and Technology*, 2025.

⁶ *In Paris, JD Vance skewers EU AI rules, lauds US tech supremacy | TechCrunch*

⁷ European Commission, *European Commission digital strategy Next generation digital Commission, C(2022) 4388 final*.

The idea of a 'EuroStack' illustrates this ambition: a fully integrated technological ecosystem covering the entire value chain – from raw materials and hardware to software and expertise⁸. The objective is not necessarily total independence, which may be unrealistic, but rather the reduction of critical dependencies.

This shift requires rethinking risk assessment frameworks. Traditional risks such as data breaches or contractual limitations must now be complemented with new scenarios, including:

- loss of service availability due to geopolitical disruptions;
- restricted access to advanced technologies; and
- structural dependency on foreign providers.

Public procurement as a strategic lever

Public procurement plays a central role in advancing digital sovereignty. The European Commission is leveraging its purchasing power to shape the cloud market and incentivise the development of sovereign solutions.

In 2025, the Commission launched a €180 million procurement procedure focused on sovereign cloud services. While modest compared to the overall market⁹, this initiative is strategically significant. It creates demand for sovereign solutions, supports emerging European providers, and establishes standards for evaluating cloud sovereignty.

The procurement process incorporates the following key principles, aligned with the Commission's sovereignty objectives:

- enhancing resilience and service continuity;
- reducing reliance on proprietary systems through open-source technologies;
- diversifying supply chains;
- limiting exposure to non-European geopolitical control;
- promoting interoperability; and
- strengthening data protection through advanced encryption and security practices.

This approach ensures that sovereignty is not only a conceptual goal but also an operational criterion embedded in procurement processes.

Evolution of the cloud market

Procurement occurs in a market that has evolved significantly since 2020, when sovereignty was primarily assessed through data protection – specifically, providers' ability to resist foreign access requests. At that stage, legal jurisdiction was sufficient to differentiate operators.

Today, the cloud market is dominated by three US hyperscalers – Amazon Web Services (AWS), Microsoft, and Google – holding around 70 % of global market share. While offering unmatched technical capabilities, they introduce complexity and dependency. In parallel, a growing number of European providers have emerged, focusing on sovereignty, data protection, and sustainability. However, they remain much smaller: OVHcloud is about 100 times smaller than AWS, and other European players are considerably smaller still.

This has led to the emergence of a European sovereign

cloud segment, offering diverse solutions tailored to compliance and sovereignty needs. These providers often rely on open-source technologies and position themselves as alternatives to hyperscalers, particularly in response to extraterritorial laws, such as the CLOUD Act¹⁰. However, the market remains fragmented and may consolidate or stay niche without scaling beyond national boundaries.

Public initiatives, including national certification schemes such as SecNumCloud and the upcoming EU Cybersecurity Certification Scheme for Cloud Services (EUCS), aim to support this ecosystem, though the EUCS is not yet operational¹¹. Meanwhile, hyperscalers have adapted through EU-based subsidiaries or partnerships with local operators, providing partial sovereignty guarantees while still relying on proprietary technologies. Additional models include local deployment of EU cloud

⁸ European Commission, [How the DIGITAL Building Blocks can help bring EuroStacks vision of European digital sovereignty to life](#), May 2025

⁹ €167 billion in 2025 according to IDC ([Worldwide Spending on Public Cloud Services is Forecast to Double Between 2024 and 2028](#))

¹⁰ The Clarifying Lawful Overseas Use of Data Act or CLOUD Act (H.R. 4943) is a United States federal law enacted in 2018. The Act allows federal law enforcement to compel U.S.-based technology companies via warrant or subpoena to provide requested data stored on servers regardless of whether the data are stored in the U.S. or on foreign soil.

¹¹ Certification schemes, including the EUCS, are included in the new cybersecurity package proposed in January 2026.

https://ec.europa.eu/commission/presscorner/detail/en/ip_26_105

stacks, which can reduce vendor lock-in but may limit scalability.

Overall, the diversity of models reflects varying

sovereignty levels (**Table 1**), highlighting the need for a structured framework to compare providers in procurement.

Table 1 – Cloud service provider typologies

Typology	Description
<i>EU Full Stack</i>	Independent companies based in the EU, contributing to the development and implementation of a complete and open cloud stack.
<i>EU Pure Player using non-EU cloud stack</i>	Independent companies based in the EU whose offering is primarily based on a non-EU proprietary solution.
<i>Hyperscaler technology operated by local EU firm</i>	Cloud services operated by EU providers with their own staff and assets, based on hyperscaler technology.
<i>Hyperscaler cloud region (EU)</i>	Isolated region operated by EU subsidiaries of hyperscalers, employing EU residents and using EU-located hardware, based on hyperscaler technology.
<i>Hyperscaler standard offers</i>	Global hyperscaler offering; ubiquitous and generally operated from outside the EU.

The need for a sovereignty framework

Given the diversity of cloud offerings and sovereignty requirements, contracting authorities need a structured method to evaluate and compare solutions. The Cloud Sovereignty Framework responds to this need by introducing measurable criteria that allow organisations

to assess the level of sovereignty provided by different cloud services. It recognises that sovereignty is not binary but exists on a spectrum, depending on legal, technical, and operational factors.

Structure of the Sovereignty Framework

Sovereignty Objectives

The framework defines eight key dimensions of sovereignty:

1. *strategic sovereignty* – alignment with EU interests and governance control;
2. *legal & jurisdictional sovereignty* – protection from foreign legal interference;
3. *data & AI sovereignty* – control over data and AI systems;
4. *operational sovereignty* – ability to operate independently;
5. *supply chain sovereignty* – control over hardware and software dependencies;

6. *technology sovereignty* – openness and independence of technological stack;

7. *security & compliance sovereignty* – adherence to EU standards and control of security processes; and

8. *environmental sustainability* – long-term resilience in terms of resources and energy.

These objectives collectively define what ‘sovereignty’ means in the context of cloud services.

Sovereignty Effectiveness Assurance Levels (SEAL)

The framework introduces a maturity scale from SEAL-0 to SEAL-4, (**Table 2**).

Table 2 – Sovereignty Effectiveness Assurance Levels (SEAL)

Sovereignty Effectiveness Assurance Levels	Sovereignty Objective Descriptions
SEAL-0	No sovereignty. Service, technology or operations under <i>exclusive control of non-EU third parties</i> , governed entirely by <i>non-EU jurisdictions</i> .
SEAL-1	<i>Jurisdictional sovereignty</i> . EU law formally applies with <i>limited practical enforceability</i> ; service, technology or operations under exclusive control of non-EU third parties.
SEAL-2	<i>Data sovereignty</i> . EU jurisdiction, but <i>material dependencies remain</i> ; service, technology or operations under indirect control of non-EU third parties.
SEAL-3	<i>Technological sovereignty</i> . EU jurisdiction; EU stakeholders exercise <i>meaningful but not full influence</i> ; service, technology or operations under marginal control of non-EU third parties.
SEAL-4	<i>Full digital sovereignty</i> . Technology and operations under complete EU control, subject only to EU jurisdiction, with no critical non-EU dependencies.

This scale allows for granular assessment of providers and ensures that minimum requirements are met. In the context of the Commission’s call for tenders, the providers had to reach the minimum threshold, SEAL-2.

Sovereignty Score

In addition to SEAL levels, the framework uses a weighted

scoring system across the eight objectives. Strategic sovereignty carries the highest weight, reflecting its importance in long-term autonomy. The score is used to differentiate between providers within the same assurance level, ensuring that excellence in one area does not compensate for critical weaknesses in others.

Application in procurement

The framework is operationalised through procurement processes. Providers are evaluated based on their responses to technical questionnaires and supporting evidence.

One key principle is that minimum SEAL levels must be

met across all objectives. This ensures that providers meet baseline sovereignty requirements before being considered on comparative scoring. This approach enables contracting authorities to align procurement decisions with strategic objectives while maintaining transparency and fairness.

Current and future scope of use

In April 2026, the Commission awarded a framework contract that allows EU institutions, bodies, offices and agencies (Union entities) to procure sovereign

cloud services for up to €180 million over six years¹². The implementation of the framework in this €180 million procurement process demonstrated its practical

¹² [Commission advances cloud sovereignty through strategic procurement.](#)

effectiveness. It helped identify strengths and weaknesses of competing solutions and supported informed decision-making.

The Commission enforced some of the framework's criteria as contract terms. The most stringent forbids any data transfer outside the EU. These transfers cover customers' data as in all Commission cloud contracts, but also any data resulting from the usage of the cloud service known as meta-data or telemetry. The market responded positively to these constraints. Out of the four winning consortiums, the three who reached SEAL-3 are European players: a Luxembourgish-French partnership led by Post Telecom with OVHcloud and CleverCloud; the German company STACKIT (Schwarz Group); and the French company Scaleway (Iliad Group). The fourth

brings together Proximus and the joint venture of the defence company Thales and the hyperscaler Google. By obtaining the SEAL-2 level, the Proximus-led partnership shows that even hyperscalers can comply with European requirements, and that the European institutions do not exclude them, but take into account their lower resilience to geopolitical events involving non-EU entities.

Beyond this initial use case, the framework is gaining interest from both public and private players across Europe. It has the potential to become a standard tool for evaluating digital sovereignty. The European Commission is also exploring broader applications, including its use for other types of digital service, as an analysis grid to audit internal systems, and in supporting policy development and market regulation.

Towards a European digital ecosystem

The European Commission's Cloud Sovereignty Framework marks an important step in Europe's transition from cloud adoption to digital autonomy. It reflects a growing recognition that cloud infrastructure is not merely a technological issue but a strategic asset with implications for security, competitiveness, and sovereignty. By redefining risk, introducing measurable criteria, and integrating sovereignty into procurement, the framework provides a practical tool for navigating a complex and evolving landscape.

While challenges remain – particularly in achieving full independence from global supply chains – the

framework sets a clear direction. It empowers European institutions and organisations to make informed choices and contributes to the broader objective of building a resilient, sovereign digital ecosystem.

Ultimately, the framework does not offer a definitive solution but a structured way to measure progress and guide decision-making. Its success will depend on continued innovation, policy support, and collaboration across the European digital ecosystem.

internal systems, and in supporting policy development and market regulation.

Advancing EU public administration readiness for AI Act compliance

By Wojciech Wiewiórowski, European Data Protection Supervisor (EDPS)



© denisismagilov / depositphotos.com

With the entering into force of the EU AI Act on 1 August 2024, also the tasks and responsibilities of Wojciech Wiewiórowski, the European Data Protection Supervisor (EDPS), changed tremendously. In addition to the existing tasks and competences the EDPS has regarding applicable data protection legislation, he and his staff became tasked with the monitoring of the implementation and respect of the AI Act in the EU institutions, bodies, offices and agencies (EUIs). As EDPS he is entitled to impose administrative fines, under certain conditions, on those failing to comply with the regulation. Similarly, the EDPS acts as the market surveillance authority. Below he explains which key actions he has undertaken to advance AI Act compliance in the EU public administration.

AI transforming EDPS's work in multiple dimensions

On the occasion of the 38th International Privacy Conference at Marrakech, back in October 2016, the European Data Protection Supervisor (EDPS) presented a background paper on the implications of artificial intelligence, machine learning and robotics to steer

the discussion towards a matter that was, at the time, a forward-looking issue. Artificial intelligence, said the paper, could create endless possibilities for the best and for the worst. However, technology should not dictate but rather uphold our fundamental rights and values.

Shortly after, my predecessor, our dearly missed Giovanni Buttarelli, discussing on machine learning and bias, warned us of the need to make the most of the window of opportunity we had to embed the right values into these technologies prior to their widespread adoption.

The thing we could not have imagined at that time was what the future held for the EDPS in that area. In the past few years, both the adoption of AI technologies and systems and the EDPS mandate have been significantly transformed. As far as AI adoption is concerned, rapid and unprecedented advances in the AI domain have been driving change across both private and public sectors, with applications ranging from generative AI chatbots and agents to AI systems assisting public sector with administrative decision making, automating routine processes and optimising public service delivery. A recent EU report highlights increased AI adoption across European governments, indicating that 70% of public administrations have already deployed—or plan to deploy—generative AI within the next three years.

Similarly, AI has been extensively shaping the EU public administration, as evidenced by the high number of AI systems (predominantly generative AI) reported by European institutions in our recent mapping report. Concerning the EDPS mandate, with the entry into force of the AI Act, the EDPS as organisation has been transformed by adding an important new task of supervising the AI systems provided and deployed by EU institutions, offices, bodies and agencies (EUIs). In addition to our well-established role in data protection, we are now a market surveillance authority (MSA) for AI systems of EUIs and a notified body for conformity assessments of high-risk AI systems of EUIs in the areas of remote biometric identification, biometric categorisation and emotion recognition, under certain conditions.

In view of its duties envisaged by the AI Act, as the EDPS I have undertaken a number of key initiatives to advance AI Act compliance in the EU public administration. Below I present a summary of these initiatives.

Preparing for new roles - the EDPS Compass

Since August 2024, the EDPS as organisation is undergoing a profound transformation with the aim to prepare for our monitoring and supervisory duties under the AI Act. To guide this transformation, we have recently published the for our new role under the AI Act¹. The Compass outlines how we envision to fulfil our role as MSA and notified body to ensure the safe, compliant and human-centric development and use of AI systems by the EUIs based on four mutually reinforcing pillars: the supervision of EUIs' AI systems, the contribution to the AI Act governance and regulatory coordination, the institutional empowerment for trustworthy AI and the international engagement and exchange of best practices.

At the heart of this new mandate is the proactive supervision of EUIs AI systems. Rather than adopting a passive stance, we will engage in continuous oversight to gain early insights into how these systems perform in real-world scenarios. This approach will help to understand where risks manifest before they escalate, enabling the delivery of timely guidance that helps EUIs to address challenges as they arise.

Beyond individual institution supervision, my organisation is stepping into a new role as a hub for governance and regulatory coordination within the broader AI ecosystem. This involves developing robust mechanisms for cooperation with other competent authorities, facilitating the seamless exchange and dissemination of best practices. This to ensure that the regulatory response remains dynamic in managing emerging risks.

A critical aspect of our new role is the direct empowerment of public institutions to adopt and maintain trustworthy AI practices. By acting as both a market surveillance authority and a notified body for EUIs, we aim at providing resources and expertise to help these bodies build internal resilience against AI risks. This support is crucial given the sheer scale of the task, which involves monitoring hundreds of systems across approximately 80 different institutions.

While the global AI supervisory landscape remains fragmented, it reveals a trend of national authorities approaching AI supervision from diverse regulatory approaches, stressing the need for strong international cooperation in this area. As the EDPS I will also actively engage in international exchanges, aiming to ensuring our presence in international networks and collaborate with international organisations working on AI matters, to contribute with its expertise and know-how and to share best practices. I consider that international engagement is not optional but central to proactive oversight and fundamental for supervisory readiness.

From a practical perspective, I am taking concrete technical and procedural preparatory steps, including targeted adjustments to our organisational structure and internal processes, while building both technical capabilities and human expertise. Specifically, we have created the unit responsible for our duties according to the AI Act which is organisationally and functionally independent from the structure responsible for the EDPS's data protection duties.

¹ Towards trustworthy AI in the EU public administration: The EDPS Compass for its new role under the AI Act, available at: https://www.edps.europa.eu/data-protection/our-work/publications/ai-act/2026-03-17-towards-trustworthy-ai-eu-public-administration-edps-compass-its-new-role-under-ai-act_en

Leveraging data protection expertise

In my role as a data protection authority, I remain responsible for overseeing any processing of personal data carried out in the context of the development or deployment of AI systems by EUIs, even when these activities are also governed by the AI Act.

Even before the AI Act came into force, data protection authorities were actively regulating AI under EU data

protection law, both by issuing guidance and taking enforcement actions. In recent years, we also offered guidance to EUIs on AI from a data protection perspective. For example, we published and recently updated, to help EUI ensure compliance with data protection rules when using Generative AI systems² and to help³ EUIs detect and mitigate common technical risks associated with AI systems.

Building cooperation

My organisation is engaging closely with the supervised EUIs to ensure they possess the necessary skills, organisational structures, and tools to confidently manage compliance of their AI systems with the AI Act. In recent years, we have taken consistent and deliberate steps to strengthen the AI governance in EUIs through the establishment of a dedicated network of AI experts - the 'AI Act Correspondents Network'. The Network provides a central platform for compliance support allowing EUIs to identify synergies, consolidate cross-institutional

expertise on AI systems and share best practices.

As the EDPS I also play an important role of observer in the AI Board, a key expert group which plays a fundamental role in the multistakeholder governance established by the AI Act. This is complemented with our work as a full member of the Administrative Cooperation Group (also known as ADCO) that brings together all MSAs to address matters related to market surveillance and sector specific issues.

The EDPS AI Sandbox

In October 2025, we launched the EDPS AI Sandbox pilot project to explore practical ways to apply innovative regulatory approaches under Article 57(3) of the AI Act - a step that symbolises our ongoing commitment to learning by doing. The sandbox provides a collaborative space for EUIs to test and validate AI systems under regulatory guidance pursuant to the AI Act and EUDPR before they are deployed. Participation in the sandbox is

voluntary, and the pilot is designed to help institutions identify and address compliance problems at an early stage, avoiding the need for costly adjustments once the systems are live. The pilot project attracted overwhelming interest, reflecting high expectations from EUIs for this initiative and a strong need for support in fostering trustworthy AI innovation.

Mapping the 'market'

The AI Act creates a comprehensive governance framework for the use of AI systems combining risk-based, anticipatory measures, with monitoring measures to assess the operation of AI systems following their deployment as well as enforcement measures allowing supervisory authorities to intervene when AI systems fail to meet regulatory requirements. In this context, the EU public administration is currently preparing its own AI systems for the implementation of the AI Act. The strategy⁴, adopted on October 2025, recognises the public sector as a key domain for the effective adoption of AI. In this context, EU public administrations are encouraged to introduce AI into their operations, moving from policy

to on-the-ground implementation of safe, secure and trustworthy AI systems.

In this scenario, it is essential to understand how AI is already used in practice across EUIs. Therefore, in 2025, we conducted an initial voluntary⁵. The exercise offered a preliminary overview of the AI landscape within EUIs, helping us to identify priority areas for supervision, as well as informing both the technical capacity and required allocation of resources essential to fulfil our tasks under the AI Act.

The mapping identified around 186 AI systems currently in use, as well as those planned for future deployment

² Generative AI and the EUDPR. Orientations for ensuring data protection compliance when using Generative AI systems, available at: https://www.edps.europa.eu/system/files/2025-10/25-10_28_revised_genai_orientations_en.pdf

³ Guidance for Risk Management of Artificial Intelligence systems, 11 November 2025, available at: https://www.edps.europa.eu/data-protection/our-work/publications/guidelines/2025-11-11-guidance-risk-management-artificial-intelligence-systems_en

⁴ Communication from the Commission to the European Parliament and the Council Apply AI Strategy COM/2025/723 final, available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0723>

⁵ High-Risk AI Systems Mapping Report in European Institutions, Agencies and Bodies, available at: https://www.edps.europa.eu/data-protection/our-work/publications/ai-act/2025-12-04-high-risk-ai-systems-mapping-report-european-institutions-agencies-and-bodies_en

across EUIs. Many AI systems reported by EUIs were still in pilot or development phases. This shows that EU public administration is actively exploring AI, while progressively building internal experience and governance structures.

The highest concentration of high-risk use cases was found in the (AFSJ), and employment and recruitment. By providing greater clarity on where such systems may emerge, the mapping helped to demystify the concept of 'high-risk' AI and to underline that these systems are not prohibited but require strong governance and oversight.

The exercise also revealed that the AI tools used by EUIs form a highly heterogeneous AI ecosystem with different

system techniques, various hosting environments and roles by EUIs. Systems based on generative AI have emerged as a dominant AI technology. The hosting environments of EUIs' AI systems vary: some applications run on premises, others rely on public or private cloud infrastructure, and many use hybrid setups combining both environments. The EUIs are more often users than developers of AI systems. Most institutions rely on externally developed, off-the-shelf tools. This finding confirms the importance of ensuring that EUIs remain accountable for how AI systems are deployed and used, even when these systems are procured from third parties.

The road ahead

As EDPS I have a new, multifaceted role in enforcing the AI Act – as an MSA and a notified body for selected high-risk systems –, and still also as a data protection authority acting under the EUDPR. This presents challenges that require careful planning and the provision of technical and organisational resources that are up to the task. Against the backdrop of a rapidly evolving technological environment and a changing legislative and supervisory landscape, this reality makes close and structured cooperation between my organisation, national MSAs and

the [European Commission's AI Office](#) not merely desirable, but indispensable. It also calls for stronger and more coherent cooperation with fundamental rights authorities and bodies. As progress on frameworks and governance structures advance, the coming months and years will be pivotal in establishing a European AI ecosystem that upholds fundamental rights and puts at its core the health and the safety of its citizens. We will fulfil our obligations in this area, aiming to establish dialogue and cooperation as the foundation of close and effective oversight.



'I have no doubt: AI's advantages far outweigh all the negatives!'

Interview with Jorg Petrovič, ECA Member and Reporting Member for the ECA's AI strategy

By Gaston Moonen



© European Union, 2026. Source: ECA.

From left to right: Jorg Petrovič, ECA Member; Mirko Iaconisi, Attaché; Emanuele Fossati, Head of AI Strategy and Adoption – ECA DATA team.

Artificial intelligence (AI) affects almost every organisation, especially those involved with data, analysis and reporting. The possibilities of AI are huge for a fact-finding organisation such as the ECA. In 2023, the ECA carried out a deeper analysis of AI's potential, when the ECA presented its first AI strategy. Jorg Petrovič is the ECA Member whose task it is to report and lead the ECA on AI matters. Below, he not only explains the key issues of the AI strategy, but gives specific examples of how it affects our audit work, explaining what the ECA needs to do to make both optimal – and safe – AI tools available.

Identifying AI's potential at the right time

Jorg Petrovič has been the Reporting Member for the ECA's AI strategy from the outset – from the very first time the idea of developing this strategy was raised. He explains that the reason he rolled up his sleeves and got involved was not because he is particularly AI-savvy. 'I'm not much of an IT person, but I am someone who keeps their eyes open. I look at everything, ask questions about whatever I see, and if I see something that has potential, I'll discuss it. It turns out that one can often be wrong, but AI and its development came onto my radar at just the right time. I think that in the future it will be as integrated just like all the other information systems we use currently.' He adds that Mirko Iaconisi, a member of his private office, raised the topic at a morning briefing back in the spring of 2023. 'There was always something in the news about AI... ways to generate texts, photos, audio, music. We realised that something big was happening.'

In terms of text capabilities, the ECA Member explains that this was what caught his attention and made him keen to explore the potential uses of AI for the ECA's work. He launched the drafting of a reflection paper, which Mirko Iaconisi and Emanuele Fossati, the ECA DATA team expert, worked on together. 'We wanted to use this reflection paper to raise the issue at ECA College of Members level to get all the Members on the same page – and demonstrate the potential of this new technology.' He highlights that the

As public auditors, we cannot remain observers of AI. We must become users of it. And examiners of it.

”

paper also included warnings. 'We had to take control if we were going to use AI at the ECA. In particular, we had to be cautious when dealing with sensitive documents. Some information needs to be kept within the ECA, such as personal data, financial data and other sensitive data. This sort of data should not be put into AI systems, because no one knows where it goes. AI providers are private companies, and are often not even registered in the EU.'

Jorg Petrovič looks back on the content and impact of this reflection paper with great satisfaction. 'Those who drafted the paper, particularly Mirko and Emanuele, did a great job. It provided examples of the potential use of AI for us as auditors. It went through all the fields of compiling an audit text, all the way through to analysing audit evidence. The ECA tested AI with files that included deliberate mistakes, which the tool then detected. This testing basically showed that in the future, an AI agent could be used for some of the tasks that auditors currently do.' Adding, with a warning: '... which does not mean that it will replace auditors, but it will enhance their capabilities.'

He explains that after being reviewed by the ECA College, the reflection paper led to establishing an initial AI strategy in 2023, which was successfully implemented in 2024 and 2025, and has been now updated with more ambitious actions to cover the period 2026-2030. 'As public auditors, we cannot remain observers of AI. We must become users of it. And examiners of it. We need to develop in-house knowledge and skills to be able to audit any work carried out by our auditees involving AI.' He further develops this, adding that the type of work can range from short risk-analysis audits, to broader issues such as where data goes, who owns it and who uses it.

Going well beyond training in AI

”

One of the immediate needs was training. The ECA Member refers to the guidelines and training invitations for ECA staff members, which have been taking place since 2023. 'It usually starts with basic notions and a few prompting exercises. We quickly established a training path to cater for different needs and job profiles. This is now developing to the kind of level we never imagined at the beginning.' He explains that the ECA looked at what was happening in other institutions too, for example in the European Commission and the Court of Justice of the European Union. 'We also visited the OECD, which had a team was working on AI-related research in anti-corruption and integrity in government. It was very interesting, and they appreciated our reflection paper. Our structured approach to training was mentioned in one of their publications.' This spring, the plan is to distribute AI tools to all ECA staff in a controlled way. 'We rolled out Copilot, starting with 120 pilot users in January 2025. By this summer everybody will have it.'

We quickly established a training path to cater for different needs and job profiles.

Training is one thing; ownership is another. 'With new skills, there's always the risk of a "typical" reaction. Perhaps older colleagues might not be as keen as younger ones. But you also have people among the older colleagues who are very aware of the potential offered by AI, and who invest in a lot of training to master it. It may be a personal preference, but sooner or later everybody will use it, just like with computers. There isn't a single auditor nowadays who doesn't have a computer!'

He underlines that training also includes establishing a framework to listen to people's concerns and tackle them. 'We approach this in a multi-layered way: we have awareness sessions, where people can come and ask questions; we have more targeted training sessions; and we have informal meetings, such as 'GPT cafés', where people come along with a coffee to share their experience – these pop-up sessions are even sometimes led by early adopters.' According to the ECA Member, those testing Copilot and who are satisfied with it are often the ones to spread the news 'from person to person, highlighting (and advertising) the benefits. We were also concerned that staff members might use AI tools for purposes other than what they were intended.'

The more you move towards a customised, locally managed system, the less likely it is to be as up-to-date as external systems ...

”

Another element triggered by the strategy is a budgetary one. 'The strategy enabled our Secretary-General to organise the financial side of this challenge.' Jorg Petrovič goes on to address another topic: data ownership. 'Then there's the question about which tools we should use. There are several providers, models and different ways of using them. If the data we feed these models is open data or already in the public domain, like our published reports, we can use an external AI system. All we need to do is include the data in the tool, if it's not there already. The commercial tools are the most advanced. The more you move towards a customised, locally managed system, the less likely it is to be as up-to-date as external systems – these external systems are backed by huge investments, not only in infrastructure, but anything related to it. The size of the data, the way it's programmed, the energy efficiency.' He's not overly worried in terms of alarmist signals about AI's excessive energy consumption though. 'At one point, it was predicted that these machines would use more energy than humankind could produce. In my view this is nonsense, and creates a false dilemma. These nay-sayers didn't predict that technologies could evolve. Using AI can actually be more economical. I have no doubt: AI's advantages far outweigh all the negatives!'

In terms of the consequences for the ECA's work, Jorg Petrovič cites a specific example. 'In February and March 2026, we published several opinions related to the European Commission's proposals for the new multiannual financial framework. The Commission published their proposals in June and July 2025. Before the first requests had even arrived for the ECA to provide opinions on them, I had already asked Mirko whether it might be possible, with AI's help, to analyse some of these proposals. He produced a very interesting analysis, having initially fed the tool the recommendations from our published reports related to the Recovery and Resilience Facility and other, relevant special reports. He then compared them with the Commission's proposals to establish whether our recommendations had been taken on board. Of course, an individual can also do this task: take a recommendation, read all the Commission's legislative drafts, and then decide whether or not our recommendation features in the proposal. Using AI, the task was completed – for all the recommendations – in less than 48 hours.'

If everything is digitalised and in an appropriate format for computers, then we can use a digital agent, essentially a digital auditor.

”

He goes on to underline that this was a preliminary analysis, which was subsequently followed up by doing the actual spade-work, but using the AI tool to help and flag some potential issues. 'Like in any other audit, you need to go into it in detail, explore, and be sure about the evidence. This is what we did for the opinions we produced, which have been very well received by the European Parliament.'

Whether AI has the potential to change the ECA's audit methodology, for example in its statement of assurance (SoA) work, is something Jorg Petrovič cannot predict. 'What we currently use for that sort of work, namely monetary unit sampling – MUS – is well-tested and established. If AI can provide something that delivers better results than the MUS methodology, it will undergo critical analysis and testing in the big, global audit companies, particularly those that perform financial audits. If it is successful, it will land on our desks as well.' He sees a lot of potential in terms of how to run tests in an MUS context. 'We discussed using agents for repetitive tasks: this can make a huge difference – it's a real game changer – and is already now presenting us with the potential to change how we carry out our work. Until now, an auditor would delegate simple checks to a trainee or a junior auditor in their team. If everything is digitalised and in an appropriate format for computers, then we can use a digital agent, essentially a digital auditor. This AI "junior auditor" can work 24 hours, seven days a week. This is what we see as potential for change.'

He explains that at the moment, not only these programmes can serve as agents. 'As Mirko indicated, there is another recent development: AI "orchestrators". They have more of a director-type role. They organise all the

digital “junior auditor” agents.’ He reiterates that so much is developing and changing, with so many new ideas, every month – or even every week. ‘It’s a fast, moving target. We’re following these developments closely so that we can stay very much up to date.’

In this context Jorg Petrovič also refers to the ECA’s contact with other audit institutions. ‘We’re clearly not the only institution to be looking at AI. For example, the US Government Accountability Office has more than one hundred people in its research department. We have little contact with the big private audit firms because of substantial differences in our respective audit work, such as our compliance and performance audits.’ He underlines that generally speaking, the ECA’s AI strategy encompasses three goals. ‘First, to improve our operational efficiency in audit; second, to build our ability to audit AI-based projects, systems and processes; and third, to add value and contribute to EU-wide discussions on AI.’

Treading carefully within the EU’s established guardrails

...we made the [AI] goals broad enough so that we won’t have to make many changes at a high level. ”

With all the rapid AI developments, the elements in the ECA’s strategy relating to AI risks may soon be out of date. Jorg Petrovič is well aware that at some point during the strategy period it will need to be updated. ‘But we made the goals broad enough so that we won’t have to make many changes at a high level. The strategy covers several aspects, touching on training, tools and financial aspects. Sooner or later, staff questions will also emerge.’ He explains that two approaches were selected: a top-down approach and a bottom-up approach. ‘The top-down approach was driven by senior management. It was focused on setting an overall strategy, defining a roadmap and establishing a formal framework to align initiatives across the ECA. The bottom-up approach builds on staff initiatives. We’ve put in place an AI training pathway, bringing in ideas, contributing to the AI competence centre, etc. Promising ideas may be further developed to create wider organisational tools.’

Sometimes you don’t want to share everything you’re doing with your auditee... ”

In this respect there is also a link to in-house development or using something that others, for example, the Commission, are using. ‘We explored this option. Using the Commission’s ecosystem feels safe for us, because we already use other products which are a part of the same ecosystem. In addition, we’ve never had any serious IT breaches. On the other hand, the Commission is our main auditee. Sometimes you don’t want to share everything you’re doing with your auditee, at least not too early on.’

A key concern in using generative AI for audit is the limited repeatability of processes and results. ”

The ECA has chosen to work with Microsoft Copilot, which comes with certain contractual guarantees. ‘The ECA also has a local tool for more sensitive cases.’ However, opting for Copilot was not our first choice. ‘One of the questions was whether we could find something more European, such as Mistral AI, a French company, or Cradle, which is Dutch. European companies have a problem in that when they have to upgrade, they realise that they need more capital, and they are not in the same financing situation as American companies. Both Mistral and Cradle explained this issue at a public conference. Another consideration is: how long will these companies remain European? The risk is that they will be eaten up by others.’

Using Microsoft services raises questions related to dependency, in addition to the sort of guardrails the ECA has put in place to protect privacy. There are also accountability issues regarding the algorithms used. The ECA Member explains that it is essential that data remains in the EU. ‘We consider this privacy element to be crucial, in line with the EU GDPR and other considerations. In terms of accountability issues, the prompts, the step-by-step reasoning, and the subsequent results should all form part of your audit evidence. This is something that we have written in our strategy and about

It is inevitable that AI is and will be used for military purposes. (...) Any tool that we make can be used against us.

”

produce results are not fully understood. Jorg Petrovič explains: 'If you have a prompt that you have used to obtain results from an AI "black box" there will be a risk. The black box is likely to change substantially over time because AI tools are evolving so quickly, and the prompt may lead to different results if it is used repeatedly. But this could be mitigated, for example by documenting AI's step-by-step reasoning and the exact version of the black box used.' He adds that this is why links in reports and scientific literature always indicate the date on which they was used. This is 'because a webpage can change from one day to the next. I think a similar approach should be used with audit evidence.'

For Jorg Petrovič, AI tools come with risks attached, but this is the same for any tool. 'It is inevitable that AI is and will be used for military purposes. You can even use your table knife to kill someone. Any tool that we make can be used against us.'

Changing not only the *how*, but also the *what* in audit

One of the focal points of EU support has been to boost the enablers that are key to the EU's AI capacities, such as data storage capacity and promoting microchip developments. When discussing the ECA's capacity to audit these enablers for AI, there is a difference between policy and systems. Jorg Petrovič states that 'in terms of enablers, you can audit at policy level, with auditors who are knowledgeable and experienced in a specific policy area. And auditors can resort to experts if needed. If we have to audit the systems themselves, it could become really technical. We also have IT auditors, who need regular upskilling and training to ensure they have a good understanding of AI.'

We already speak to computers, just like we're speaking to each other right now. This is a big change.

”

Another area where he sees great potential is for presenting the ECA's audit results. 'It's already the case that when I go to the European Parliament to give a presentation, our team prepares potential questions and answers, with the support of AI. It can really help to be better prepared in this sort of situation.'

Another task that AI can help with is the interaction between human beings and computers. 'We already speak to computers, just like we're speaking to each other right now. This is a big change. And instead of reading a text, you can listen to it. I use this myself sometimes, for example if I have a long text. I change it into a voice file and I can listen to it - I'm also a big podcast fan. For me, this change in the way we communicate is important. I remember when the first touchscreens came out... how natural it seemed, how intuitive it was to use. Before that, you always had to type, for example, on a Blackberry. I vividly recall how quickly my children adapted to using tablets. Only a few years ago, we were taking our baby steps and just beginning to use AI for day-to-day purposes. Now, just three years later, it's really taken off - and there's a lot more to come. We, meanwhile, will continue to explore the best way to use it to improve our work.'

AI – is the EU picking up pace?

By Mihails Kozlovs, ECA Member



© European Union, source: ECA/Evelina Gaisonoka - Gemini created visuals

Artificial intelligence (AI) is booming, with its developers' promise, backed up already by a number of significant results, to transform economies and societies. While AI has a great deal of private money invested in it, AI also attracts public money from the EU, with billions of euros in EU support since 2018. The ECA, with Mihails Kozlovs as reporting ECA Member, published its first audit to assess the effectiveness of the Commission's contribution to developing the EU's AI ecosystem. Below, he highlights the key initiatives, the ECA's audit findings regarding these initiatives, and also outlines what the Commission has done since the findings were published in 2024. He identifies the next multiannual financial framework as being an opportunity to address a number of fundamental weaknesses.

The EU's ambition to catch up in the AI race

Artificial intelligence is a breakthrough technology. It has perhaps unprecedented potential to reshape our economies – and it already is doing so. There is therefore fierce international competition for leadership in this area; there is a high chance that the leader takes it all. In this context, sizeable and focused AI investment and clear deployment strategies are game changers in setting the speed of EU economic growth for years to come. We should also not forget the numerous applications of AI

in sensitive areas, such as internal security or the use of AI for military purposes. Indeed, without adequate AI-powered defence technology, how could Europe defend itself against a swarm of AI-managed drones?

Public policymakers have an important role to play in organising the AI ecosystem. The recommendation of the Organisation for Economic Cooperation and Development (OECD) on AI promotes principles for

responsible stewardship of trustworthy AI, and identifies five public action dimensions to encourage AI innovation and uptake:

1. investing in AI research and development, and open datasets, and encouraging private investment;
2. fostering a digital ecosystem for AI, including the development of and access to appropriate digital infrastructure and AI knowledge sharing;
3. shaping an enabling policy environment that encourages innovation and competition for trustworthy AI and supports the transition from research to deployment;
4. building human capacity and preparing for labour-market transformation; and
5. encouraging international cooperation for trustworthy AI.

All of the major world economies are currently racing to take the lead in the development and deployment of AI technologies. Europe, despite its scientific excellence, has been lagging behind on key aspects such as financing, scaling and creating new ecosystems. Some experts consider that catching up in some areas, such as developing new large language models (LLMs), is more or less impossible. According to estimates, the US significantly outpaces the EU in public and private AI investment, with roughly \$68 billion in US venture capital compared to only \$8 billion in the EU (and \$15 billion in China) in 2023. Companies creating generative AI models in Europe, such as Aleph Alpha and Mistral, need large investments to avoid losing the race to US firms. However, European markets are unable to meet this need, pushing European firms to look further afield for funding¹.

There are several reasons for why the EU is lagging behind. Europe started scaling up its AI capabilities much later than other regions, such as China and the US, which

Clear underuse of AI's potential

Let's take a closer look at the results of our work. On the positive side, we found that the EU had taken different measures to support an integrated and innovation-friendly ecosystem for AI. The ecosystem's scope corresponds to international best practice set by the OECD and has three main dimensions:

- coordination of investment and regulation;
- putting in place infrastructure and financial enablers; and
- direct EU investment in AI research projects.

surged in the post-2010s. The EU's starting point was in 2018, when the European Commission adopted the first plan on the development and use of European-made AI. It was co-designed with the member states and updated in 2021. The plan outlined actions to enhance AI research and create trustworthy regulations. The ambition was for the EU to become a global leader in cutting-edge, ethical AI.

Since 2018, the Commission has also tried to incorporate AI as a key component in the EU's strategies on the digital transition, industrial policy and strategic autonomy for this decade² recognising that AI can bring considerable societal and economic benefits to a wide range of sectors. Nevertheless, because one of the EU's core principles is that we need to ensure that AI products are safe and secure³, the EU has focused its efforts more on regulating rather than on boosting AI. There are therefore differences in how AI ecosystems are established around the world. Other jurisdictions, primarily the US, apparently chose a different path, which allowed an unrestricted flow of innovation and deployment – and only later defined the boundaries. While information on China is somewhat limited, the state strategy is clearly to catch up with the US, and eventually surpass it, whatever it takes.

This is why from our point of view as the EU's external auditor, it was so important to look at the established ecosystem. We focused our audit on what the EU had tangibly been doing to support the development and uptake of AI solutions and applications. We published the results of our work in 2024, in [special report 08/2024: 'EU Artificial intelligence ambition – Stronger governance and increased, more focused investment essential going forward'](#). Two years have passed, but our audit results are highly relevant, especially at a time where the ongoing discussions relating to new multiannual financial framework could address most of AI's significant design and funding weaknesses of the EU.

But we also found a number of areas where the EU should do better.

Firstly, we found that EU action so far had had a limited effect in developing the EU's own, single AI ecosystem. Member states differ significantly in the way they implement their AI policies. Moreover, not all member states have adopted their AI strategies, which leaves certain loopholes in terms of the possibility of having a comprehensive overview of the overall situation in the EU. Our audit revealed that there are currently 27 AI ecosystems in the EU, with a few member states such as Germany and France that have very developed AI policies. However, many member states are also lagging behind

¹ <https://ec.europa.eu/newsroom/eisma/items/864247/en>.

² Commission staff working document, [Strategic dependencies and capacities](#), SWD(2021) 352 final.


³ European Council, [A digital future for Europe](#), 2025.

and are only at the very early stages of considering AI's potential.

The latest statistical data (**Figure 1**) demonstrates this fragmentation. For example, the proportion of businesses using AI varies widely across member states, bringing the

EU average to 20 %, meaning that only one in every five companies in the EU uses AI. In the US and China, these numbers double – or even triple – with around 50-70 % actively using AI.

Figure 1 – European enterprises using AI technologies, 2024 and 2025 (% of enterprises)



(*) 2025: Break in the time series.
Source: Eurostat (online data code: isoc_eb_ai)

In the markets leading the way in AI, the lion's share of investment (and trials and errors) come from large private sector companies that invest themselves and constantly search for prospective startups. Thus, at any stage of development, a promising technology company has access to abundant funding, including risk capital. In Europe, large companies seem to underinvest in AI (**Figure 2**).

Secondly, at the time of the audit we found that EU investment targets were outdated and moreover were not sufficiently clear about the expected results. They

had not been updated since the first plan in 2018, despite very dynamic international developments in the area at the time. The EU set a target of mobilising €20 billion in AI investment between 2018 and 2020, and €20 billion annually in annual public private investments by 2030. We recommended that the Commission reassess by 2025, and justify investment targets, bearing in mind international and technological developments and the national investment needs of both the public and the private sector.


Ambitious EU objectives on AI, but results lag behind

Following the publication of our audit report, various measures to address those weaknesses were taken. The European Commission that took office in 2024 brought AI policy, yet again, to the centre of the debate – and rightly so. Most importantly, in 2025 the Commission launched the [InvestAI initiative](#), an initiative to mobilise €200 billion for investment in AI from 2025 to 2030, including a new

European fund of €20 billion for AI gigafactories⁴.

In April 2025, the Commission also approved the [AI Continent Action Plan](#) to accelerate EU AI policies. The plan focuses on investing in large-scale AI infrastructure, improving data access, boosting AI adoption in strategic sectors, strengthening skills and talent, and streamlining

⁴ [AI Gigafactories](#) are large-scale facilities dedicated to the development and training of next-generation AI models.

Figure 2 – Global private investment in AI by geographic area, 2013-2024

© Stanford University, [The 2025 AI Index Report](#), (2025)

regulation. It promotes strong collaboration between governments, companies, researchers and developers. The Commission also established the [European AI Office](#), which is mandated to work closely with member states through the AI Board to ensure a coordinated policy approach aligned with rapid technological progress.

Time will tell if all of these elements taken together will work, but known design weaknesses could hamper success. For example, from the design perspective, InvestAI faces the same difficulty – ambitious political statement at EU level regarding the €200 billion investment target, but with only €20 billion contributed directly through EU funding. The remainder will depend on private investments and national funding. This is the same as from the very beginning – the EU mainly relies on member states to fill the funding gap, and, given the size of the EU budget, there is no other way. However, it is not clear if, how, or when the funds will be made available by member states and, more importantly, if the funds will be deployed in a way that will enhance synergies and boost their leverage effect.

At this point, it is not possible to track national and private AI funding levels in member states because the implementation of national AI policies (including funding) is spread across several sectoral ministries or specific sectoral programmes. This makes it difficult to isolate individual measures or expenditure that support AI. Our audit showed that less than half of the EU member states (9 out of the 20 that responded to our survey) had specific budgets for national AI strategies. Furthermore, the Commission had few monitoring tools available to check national action and review their effect on EU targets. This leads to blind spots in the overall picture of investment volumes – and it is not clear how these ambitious

investment targets would be met.

In such a fragmented ecosystem, we might expect that EU measures would be the tool to supplement national measures, providing EU added value and filling the gaps. Unfortunately, this is not the case.

First, the EU plan envisaged measures to financially support AI innovators. However, despite multiple initiatives, our audit showed that in the past the Commission had only provided modest capital support to SMEs. In comparison, one of the main instruments available to the EU, the European Innovation Council, had a budget of €256 million in 2024, while the US allocated more than \$6 billion for this purpose⁵.

Second, the Commission allocated direct funds to numerous AI innovation projects. The fragmentation of funding and management resulted in a lack of coordination or monitoring of their collective contribution to the development of an EU AI ecosystem. Moreover, the mechanisms to ensure that research results will be used by academia or the market were weak. Despite the Commission's objective, AI research projects under Horizon 2020 did not trigger significantly higher private cofinancing beyond the average for the whole programme. In addition, the share of projects with patent filings was below what had been planned.

These findings confirm the weaknesses of the Commission's controls to ensure that the research results are used in the EU – in particular, that they translate efficiently into innovation. The recently proposed European competitiveness fund has earmarked €51 billion under the digital leadership policy window, and will priorities AI in addressing the current weakness in funding and use in the next multiannual financial

⁵ See: <https://ec.europa.eu/newsroom/eisma/items/864247/en>.

framework. The newly established [European AI Office](#) could be helpful in addressing these issues, but its design and its place in the institutional hierarchy raises doubts about its ability and powers to address them effectively. Indeed, the question is whether an administrative unit within a large public organisation can be agile enough to change the rules of the game in a highly dynamic market.


The EU was not overly successful in terms of boosting AI investments to the extent that would allow the EU to compete globally. But the EU is the frontrunner in regulating AI. This is an important feature of the European institutions' AI ecosystem, but is something that was outside the scope of the audit that produced our special report 08/2024. In approving the EU AI Act, the EU has now taken unprecedented steps to regulate AI and ensure its safety, leading the way at a global level. The AI Act entered into force on 1 August 2024 and will be fully applicable on 2 August 2026. Time will tell whether the strategy to pioneer regulation (or as some put it to regulate into a leadership position) will pay off. In any case, without serious work on enablers there may not be much leverage to regulate.

To conclude, more than eight years after the adoption of the first EU AI policy document – the EU AI Plan – weaknesses in implementation and performance monitoring demonstrate the need for the EU to focus more on delivering results and better coordination with member states. The ultimate goal is to build an effective and efficient policy framework to develop the European AI ecosystem.

It is high time to speed up the work to address the key success factors for AI development and deployment in the EU. AI is already a game changer in military strategy, so we have no time to waste. Considering the most recent international developments, the forthcoming implementation of the AI Act and the ongoing discussions on the next financing framework, there is a window of opportunity for the Commission, the Council and the Parliament to address the weaknesses. Given the importance of AI and the level of public funds deployed in the sector, the supreme audit institutions (as external auditors) should ensure that developments in this area remain firmly on their risk analysis radars, analyse the effectiveness of AI strategies in the EU, provide general insight, oversight, and recommend improvements.

The ECA's audit of the EU's microchips ambitions – examining an industrial ecosystem

By Rafal Gorajski, Investment for cohesion, growth and inclusion directorate, and Austin Maloney, private office of Annemie Turtelboom, ECA Member



© Edelweiss / stock.adobe.com

As the steam engine was to one phase of industrial development, the microchip is to the modern world. In a digitalised society, few things can operate without them, and the recent explosion of artificial intelligence has only reinforced their importance, as advanced AI systems depend on ever more powerful microchips, which is why securing reliable supplies and the race to the next generation of chips is a topic that is foremost in the minds of today's governments. The EU has long suffered from a declining share of the global market and increasing dependencies on non-domestic supplies. In response, it has launched multiple strategies over the past decades, including the 2023 Chips Act. In 2025, the ECA published its [special report 12/2025](#) on the EU's Microchips Strategy, assessing the early implementation of the Chips Act and the EU's general framework on microchips. In this article, that special report's Head of Task Rafal Gorajski gives us an insight into the challenges of putting together a performance audit on such a highly publicised and fast-moving topic, assisted by additional background info and analysis from team member Austin Maloney.

The Chips Act landscape


In 2024, Mario Draghi's [landmark report on the future of European competitiveness](#) stated that capabilities to innovate cutting-edge microchips were 'virtually non-existent in the EU' and investments were 'below the scale needed to sustain expected demand'. The report also noted the 'conflicting, uncoordinated postures of Member States'. These messages wouldn't have surprised many at the ECA.

At the time of publication, the ECA was in the middle of an audit on that very subject. And while the EU had strengths in some areas of the microchip value chain, the situation was a worrying one. In 2021, the EU's trade deficit for microchips was almost €20 billion. Its share of global revenue had crashed from more than 20% in the 1990s to around 10% in 2022. Shortages during the COVID-19 crisis had highlighted how dependent the EU was on external supplies – with supply-chain breakdowns causing German car manufacturing to collapse to 1975 levels¹.

In short, chips were more important than ever, and losing access to them could cause the production of everything from smartphones to cars, satellites and advanced military equipment to grind to a halt. The EU had ended up in a position where this was a real risk.

This wasn't a new situation. Back in 2013, the European Commission had [launched a strategy](#) aimed at reversing the decline in the EU's share of the global chips market (see **Figure 1**). However, it failed to meet its main objective. While the overall capacity of the EU increased, its share of the global market, which is ever-growing, continued to decline, both in terms of production capacity and capital expenditure. Recognising this failure, the Commission returned in 2022 with the [Chips Act package](#), which [passed in 2023](#). The Chips Act launched a new sweep of funding with the total envelope estimated at around €86 billion, and came with a new goal – 20% (by value) of global production of cutting-edge and sustainable microchips by 2030².

Figure 1 – Share of global chip capacity by region 2010-2030



Note: All values shown in 200 mm wafer size e equivalents; the chart excludes capacity below 5 000 wafer starts per month or less than 200 mm. This reflects the modern semiconductor manufacturing facilities' capacity, where wafer diameter is greater than or equal to 200 mm.

Source: ECA, based on BCG and SIA study [Emerging resilience in the semiconductor supply chain](#), 2024.

The importance of the area meant that it was a natural selection for an ECA performance audit. Annemie Turtelboom, the reporting ECA Member for the audit, said 'Microchips are the brains of modern systems - nothing works without microchips. They are present in everything, and are becoming only more so with time. In a modern car, there are around 1 500 microchips; by 2030, it's

expected to rise to 3000. There is a global competition to be the first to develop smaller, more powerful, and more efficient microchips. If we don't secure our supplies of chips, all our other plans – in industry, in defence, AI, everything, will be brought to a standstill. The Chips Act is therefore vital to the EU's interest, and that's why we picked it up for a performance audit.'

¹ [ECA Special Report 12/2025](#), The EU's strategy for microchips and European Commission [Communication \(2022\) : A Chips Act for Europe](#).

² Adopted from the EU's 2030 [Digital Decade Strategy](#).

Auditing an entire ecosystem

By the time the European Chips Act entered into force, microchips had become one of the most politically visible and strategically sensitive topics in Europe. For the audit team, however, the challenge was not understanding why chips mattered. The real challenge was figuring out how to audit an industry that spans continents, relies on highly confidential information and evolves faster than most public policies.

In fact, by the time the audit finally started, it already had a history. Within the ECA, microchips had been identified as a strategic topic several years, earlier. We had attempted to launch the audit twice before, but organisational impediments delayed this launching. Looking back, those delays may have been a blessing in disguise. By waiting until the Chips Act entered into force, we were able to assess not just an idea, but the beginning of its implementation.

Even before the audit formally started, the ECA invested considerable effort in understanding the sector. Fairly early in the planning phase, we realised that a traditional project-based audit would not get us very far. The microchips industry is simply too interconnected. The value chain stretches across multiple continents and involves manufacturers, equipment suppliers, designers, research organisations and public authorities. Looking just at a sample of projects would have provided only a very partial picture.

And so we started with the ecosystem itself. We mapped the main sources of public support and identified the companies receiving the largest share of funding. The results were striking. A relatively small number of major players accounted for a substantial share of investments across the European microchips value chain. That finding shaped the rest of the audit.

Not everyone was convinced this approach would work. Some colleagues questioned whether the headquarters

of global technology companies would be willing to spend time talking to public auditors. In practice, they were. The major companies we approached responded to our questionnaires, participated in interviews and shared their views on both the Chips Act and the broader competitive environment. We complemented these discussions with meetings involving leading research and technology organisations, including the Interuniversity Microelectronics Centre (IMEC), Fraunhofer, the Consiglio Nazionale delle Ricerche (CNR) and several organisations involved in developing pilot lines supported under the Chips Act.

One visit that remains particularly vivid in my memory was to ASML in Eindhoven. The campus resembles a small city dedicated entirely to advanced technology. Its lithography machines are indispensable for producing the world's most advanced chips, representing one of the few areas where Europe holds a truly dominant global position.

Another moment that challenged some of our assumptions came during a visit to a fabrication facility in Germany. One fact immediately stood out. Contrary to what many people imagine, microchip production is not a rapid process. A wafer may spend several months moving through hundreds of highly specialised production steps before becoming a finished and tested chip.

During discussions on the Chips Act, considerable attention was given to mechanisms intended to strengthen security of supply during future crises. Yet semiconductor manufacturing is not like producing consumer goods. Particularly in sectors such as automotives, qualification and validation processes can take months or even years before a chip can be integrated into a final product. Production cannot simply be redirected overnight from one customer or application to another.

Between ambition and industrial reality

Another recurring theme emerged as we spoke with companies, researchers and policymakers. The Chips Act placed strong emphasis on cutting-edge semiconductor manufacturing. At the time, this reflected a broader ambition to strengthen Europe's position in the global technology race. However, many stakeholders pointed out that Europe's traditional strengths often lie elsewhere: automotive, industrial and power microchips. The tension between ambition and industrial reality surfaced repeatedly throughout the audit and ultimately influenced some of our conclusions.

Perhaps surprisingly, the most difficult aspect of the audit was not the technology itself. The greater challenge was the political visibility of the topic. Microchips had become a flagship issue for Europe and the policy environment was evolving rapidly as the Chips Act moved from political vision to implementation.

Not every ambition we had at the start could be realised. We had hoped to conduct a much deeper quantitative analysis of the microchips market using specialised industry datasets. We even explored acquiring commercial data sources. In practice, obtaining suitable

data within the available timeframe proved difficult. In retrospect, this limitation did not materially affect our conclusions, which were supported by extensive evidence gathered from industry, research organisations, member states and the Commission.

The audit also coincided with the emergence of generative AI, which provided an efficient means of searching for information and acquiring background knowledge on microchip technologies, manufacturing processes and market dynamics. At the same time, the rapid development of AI was itself reinforcing the strategic importance of microchips, as advanced chips had become an essential prerequisite for training and deploying increasingly powerful AI models. While audit interviews and documentary evidence remained the primary sources of audit evidence, these tools helped the team navigate the sector's technical complexity more efficiently.

If there is one key takeaway from this audit, it is that

ambition and realism need to go hand in hand. Nobody questioned the importance of strengthening Europe's position in microchips. The more difficult question was whether some of the assumptions underpinning the 2030 objective remained valid as market conditions evolved. As highlighted in our special report, even if all announced projects were implemented successfully, the EU's share of global production would increase only modestly and remain significantly below the 20% EU target. Our recommendations reflected that concern, calling for better monitoring of progress and a reassessment of whether the strategy remained aligned with developments in the global market.

As Head of Task, examining an entire industrial ecosystem was fascinating and this task remains the most fascinating audit I have worked on. In a world where technological leadership and economic resilience are becoming increasingly important, that perspective may prove valuable far beyond the microchips sector.

The aftermath – the launch of Chips Act 2.0

The ECA published [special report 12/2025](#) (see **Box 1**) in April 2025, and it immediately attracted significant attention from policymakers, industry stakeholders and the media. All recommendations were accepted by the Commission. The report was presented to the Council of the European Union, where Member States [expressed strong support for its findings](#). It received extensive coverage in leading European publications such as [The Guardian](#), [Politico](#), [the Financial Times](#), [Le Figaro](#) and [Frankfurter Allgemeine Zeitung](#). According to the ECA's media metrics, it was one of the ECA's best-performing reports of the year.

One of the report's recommendations was that the Commission begin preparing the next-generation of Europe's semiconductor strategy. The debate moved quickly. In September 2025, [Member States referred explicitly](#) to the ECA's findings when discussing the future of the Chips Act, arguing that the 20% global market share target was both unrealistic and insufficiently focused on Europe's strategic priorities within the semiconductor value chain.

In June 2026, the Commission launched the Chips Act 2.0 as part of its broader [Tech Sovereignty Package](#)³. The new initiative places⁴ greater emphasis on resilience, supply-chain security, technological leadership and crisis preparedness. The Commission noted the EU's reliance on '[non-EU countries for over 80% of key digital products, services, infrastructure, and intellectual property](#)', at the launch, and presented the new framework as a response to those vulnerabilities.

The effectiveness of these new measures remains to be seen, and the ECA has not audited the new framework. Nevertheless, the evolution from Chips Act 1.0 to Chips Act 2.0 illustrates how rapidly policy assumptions can change in a strategic sector. It also suggests that the questions raised during the audit — about realism, monitoring and the long-term direction of Europe's semiconductor strategy — are likely to remain relevant for years to come.

³ [Commission Communication](#), on European Tech Sovereignty, accompanied by an EU Open Source Strategy, 2026.

⁴ [Proposal for a Regulation on a framework of measures for strengthening the Union's semiconductor ecosystem \(Chips Act 2.0\)](#).

Box 1 – ECA special report 12/2025 – main findings

The Commission has no mandate to coordinate national investments at EU level to ensure they align with the Act's objectives and is responsible for a small part of the Chips Act funding announced. The Chips Act provided new impetus, although it lacks clarity in its targets and monitoring, and it is difficult to know whether it takes proper account of the industry's current levels of demand for mainstream microchips. Nevertheless, the publicly funded projects we examined were aligned with the EU's strategic objectives.

Several other key factors affect the EU's competitiveness in the field, and the chances for successful implementation of the Chips Act. These include dependency on imports of raw materials, high energy costs, environmental concerns, geopolitical tensions and export controls, and a shortage of skilled workers. Furthermore, the EU microchip industry consists of a few large enterprises focused on high-value projects, meaning that funding is concentrated. The cancellation, delay or failure of a single project can therefore have a significant impact on the whole sector.

Overall, the auditors found that the Chips Act is highly unlikely to significantly increase the EU's share of the microchips market, or to meet the objective of 20% of global output. Implementation of the Chips Act is progressing, but too slowly to meet the Digital Decade 20% target. The European Commission's own forecast, published in July 2024, predicts that despite a significant expected increase in manufacturing capacity, the EU's overall share of the global value chain in a fast-growing market would increase only slightly, from 9.8% in 2022 to just 11.7% by 2030.

From promise to practice: how AI is shaping ECA audits

By Emanuele Fossati, ECA Directorate of the Audit Quality Control Committee



© European Union, source: ECA/Emanuele Fossati - AI generated picture

If we believe all the current hype, AI is everywhere, being used by many companies and organisations. AI seems to offer many possibilities but to what extent does it really change the work and processes of an organisation? Emanuele Fossati, Head of AI Strategy and Adoption in the ECA Data team, looks beyond the sales talk to see what really has changed at the ECA as a result of using AI tools for our institution's core work – external audit. Below he presents concrete changes in the way ECA staff employ AI tools.

Reconciling AI enthusiasm with professional scepticism

Enthusiasm about AI is real – but so is (a healthy) professional scepticism. Audit work proves its worth not through excitement or PowerPoint presentations, but through proper evidence handling, robust analysis, and insightful observations and recommendations. In

this article, I look at concrete audit tasks in which AI has already been used as an enabler – not as a 'magic button', but as a way to go deeper and cover more, while always remaining accountable.

AI as an enabler – not an autopilot

A useful way to position AI in audit work is as an 'analytical enabler'. A recurring theme in other articles in this edition of the ECA Journal is that while AI can support audit work, auditors remain fully responsible when using it – and that accountability requires traceability (which includes documenting prompts and outputs as part of audit evidence).

This framing matters, because it aligns with core audit principles: professional judgement cannot be outsourced to a tool, and quality control must remain rigorous. An error or hallucination could produce catastrophic results, which underlines the importance of constraining AI use by means of professional safeguards.

What does 'AI as an enabler' look like in practice? Below are concrete audit use cases where AI has supported audit teams in handling large-scale, complex and multilingual

evidence – enabling deeper analysis and more precise sampling.

Categorising large lists of EU financed projects

A recurring audit challenge is working with *very large project populations* that come with limited quantitative data but include detailed text descriptions, often in multiple languages. Auditors need to divide them into *relevant categories* and derive additional information that could support sampling based on specific criteria, and

better quantitative analysis. The ECA Data team has been working with multiple audit teams to help them design prompts based on their audit needs and criteria and then apply these to very long lists containing thousands of projects: a task impossible to perform manually with precision on that scale.

Figure 1 – Example of categorising large lists of EU-financed projects

Project description	AI Classification
[...] Within the framework of the My Current Programme, investments in the construction of new electricity and/or heat generation units using solar energy, consisting of the installation of photovoltaic cells on residential buildings or in the plot where the residential building is located for the Applicants' own needs, will be supported. [...]	deployment
The overall objective of the project is to increase the sustainable use of Craiova International Airport by expanding and modernising it in order to increase regional accessibility, in a safe manner and in line with the environmental protection objectives. [...]	unrelated
Roll-to-roll production of organic solar films with 1.2 m width (SolarRoll)	manufacturing
Solar thermal installation with concentration in Heineken Sevilla factory.	solar thermal
General objective – to increase capacity to produce electricity from renewable sources increasing energy efficiency, improving air quality and reducing greenhouse gas emissions, by using photovoltaic panels systems to produce electricity necessary for own consumption and delivery of surplus in the national energy system	deployment
EDF supports the energy transition in Martinique and puts its action in particular in the field of energy efficiency, the development and integration of renewable energies and its support for innovation. Through the CESI Martinique project, EDF aims to continue and accelerate the development of the individual solar water heater in Martinique by installing 17,500 units by 2020.	solar thermal
The overall objective of the project is the modernisation of RED and the increase of the distribution capacity of the 110/20kV Caracal West transformer station for the evacuation of electricity produced by CEFs (photovoltaic power plants). [...]	other solar / PV

In the above example, AI does not 'decide' categories, but rather *accelerates and scales the preparatory analytical work* required for categorisation, so auditors can spend

more time validating logic, testing assumptions and examining evidence.

Comparing different versions of the same legislative document

Audits frequently require an understanding of *how a legislative or policy text has evolved* across versions – changes in wording, obligations or scope following observations and remarks from other institutions. AI facilitates this process by highlighting differences more systematically, especially when documents are long and version histories are complex.

This does *not* replace legal analysis or policy judgement, but rather serves as a structured comparison aid: it helps auditors focus on relevant differences and document their review more transparently. That emphasis on transparency matters, because the ECA's approach to AI stresses accountability and repeatability – including documenting the chain from prompt to output as part of audit evidence.

Comparing previous ECA recommendations with new European Commission proposals

Another audit relevant application of AI is *comparing past ECA recommendations* with new legislative proposals from the European Commission. This supports an evidence based view of whether recommendations appear to be reflected in those proposals fully, partially or not at all – while recognising that such analysis represents input for judgement, not the judgement itself.

The interview with ECA Member Jorg Petrovič in this edition of the ECA Journal ([page 38](#)) provides a concrete illustration of this approach: it describes an analysis in which recommendations from published reports were compared with Commission proposals. The work of reading and assessing all drafts the Commission's legislative drafts manually would have taken weeks and

was done using AI for all the recommendations in less than 48 hours’ – with the auditors then carrying out the actual audit work to verify and assess evidence.

helps to provide breadth and speed in the evidence-gathering phase, while auditors remain responsible for substantiation, interpretation and conclusions.

This is a good example of the ‘AI as enabler’ logic: AI

Quantitative analysis of large surveys – especially with open, multilingual answers

Surveys remain a valuable audit tool, particularly when auditors need insight into implementation realities, stakeholder experience or operational constraints. Closed questions are relatively straightforward to analyse, open questions are not – especially when the answers are voluminous and multilingual.

thematic clustering of large bodies of free text responses, enabling auditors to carry out a more structured analysis of qualitative information than anecdotal citation alone. Again, the safeguard is the auditors’ professional judgement: themes suggested by AI must be reviewed, challenged and compared with other evidence before they can influence observations or findings.

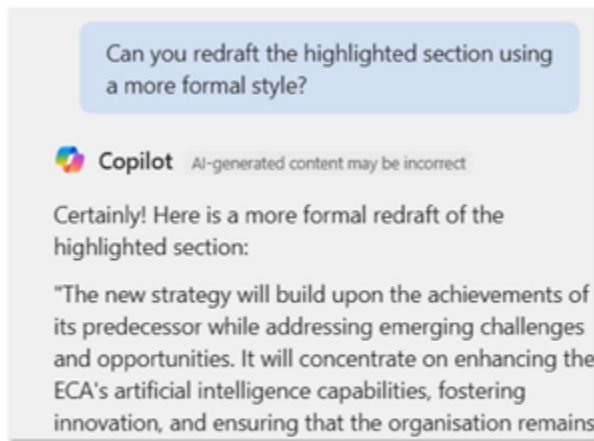
AI has been used to support the *systematic coding and*

Extracting and summarising key information from meetings, conferences and audit interviews

AI is increasingly used to generate structured summaries of meetings and conferences, enabling auditors to capture key points, decisions and follow-up actions with greater efficiency. This approach is particularly valuable for lengthy sessions with multiple topics, where manual note taking and synthesis can be time-consuming.

AI facilitates review by systematically organising discussion highlights, while auditors remain responsible for contextualising and extracting value from the AI summaries.

Figure 2 – Extracting and summarising key information



Improve written texts / draft new ones



Summarise meetings

© European Union, source: ECA/Emanuele Fossati - AI generated picture

Similarly, there is growing interest in using AI to extract insights from audit interviews, provided appropriate data protection guarantees and legal safeguards are observed.

The technology can help identify recurring themes, contradictions or actionable recommendations within interview transcripts.

AI agent supporting the completion of audit programmes in the area of public procurement

A more advanced (and still ongoing) pilot project concerns the development of an *AI-based agent to support auditors in completing audit programmes*, initially in the area of public procurement. Audit programmes are a core element of audit work: they translate audit objectives into concrete questions and sub-questions, which are then answered by analysing evidence data and information and consulting and referencing relevant legislation, all guided by audit standards and methodologies.

To create the context for the AI agent, and make it more precise and accurate, auditors initially need to define and prepare a relevant *knowledge base*. This includes applicable legal provisions, relevant audit manuals, and audit evidence in multiple formats.

The AI agent needs then to be properly instructed with all of the implicit and explicit professional knowledge that auditors possess but tend to assume to be 'obvious'. For an AI agent, nothing is obvious, and inaccurately defining the policy area context may lead to unintended results, or irrelevant observations.

A key part of the initial agent configuration work is *transforming traditional audit questions into effective AI prompts*. This is not a mechanical exercise: it requires

experimentation, reformulation and iterative testing. Initial outputs are systematically reviewed by auditors, leading to further refinement of prompts and repeated inputs until results meet expectations. This process highlights once again the recurring important lesson: meaningful use of AI in audits requires *active human calibration* between questions and results, not just an initial technical deployment.

Early results from this pilot are *impressive in their consistency and correctness*. At the same time, the project clearly remains a work in progress and further developments are planned this year to expand coverage, strengthen validation mechanisms, and ensure that outputs fully meet audit quality requirements before the AI agent is adopted more widely.

This use case illustrates especially clearly how AI can support — but not replace — professional judgement. The agent does not 'design the audit'. It assists auditors by proposing structured elements that auditors assess, adapt and ultimately take responsibility for. In that sense, it represents a controlled step towards more *interactive and agent based support*, firmly embedded within audit safeguards.

From analytical support to structured assistance — but no substitute for the auditor's judgement

From the above examples we can conclude that AI is *not* a magic button that produces an audit. It does not replace audit design, professional judgement or responsibility. What it changes is the feasible level of *coverage, granularity and speed* in specific analytical steps — which can enable auditors to work with *larger and more representative populations* and explore evidence in ways that would otherwise be constrained by time and manual effort.

Taken together, these use cases show that AI is already changing how auditors work, without changing *what an audit is*. Whether supporting large scale data categorisation, document comparison, recommendation

tracking, survey analysis, or the structured execution of audit programmes, AI acts as an *enabler* — expanding analytical reach, improving consistency and reducing certain mechanical constraints.

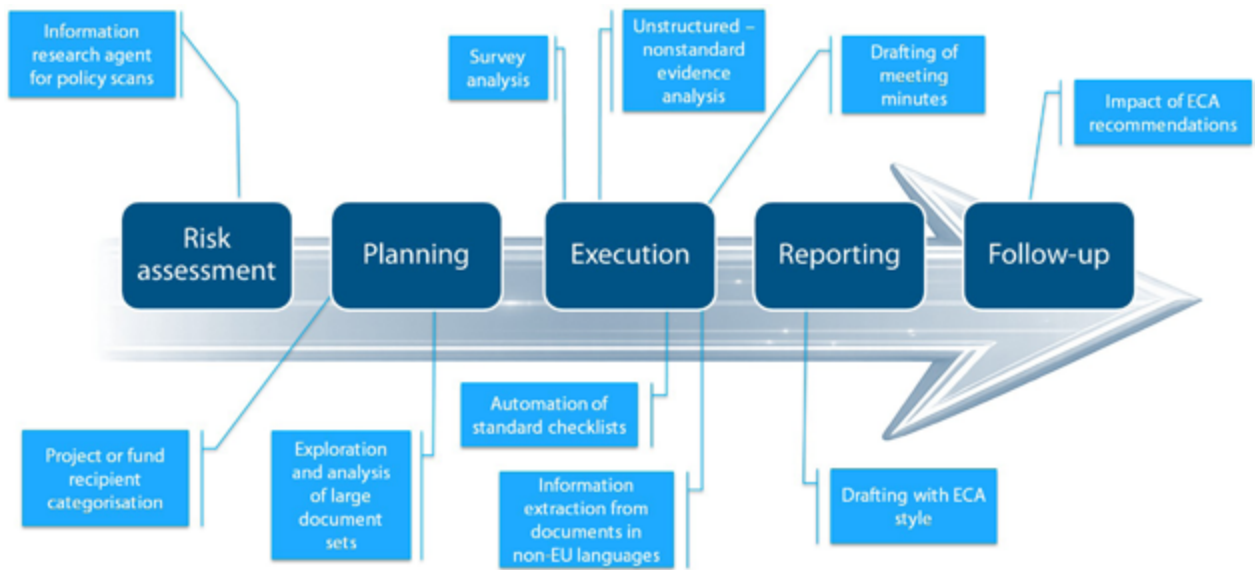
The emerging picture is therefore neither hype nor scepticism, but continuous progress. AI is not a shortcut to mass-produce audits. It is a set of tools — and increasingly, structured agents — that, when carefully designed and governed, allow auditors to focus their expertise where it matters most. In a profession built on evidence and judgement, that is where its real value lies.

The elephant in the room: measuring quantitative impact is harder than it sounds

AI is being adopted broadly and quickly across many organisations — including public audit institutions like the ECA. However, speed of adoption creates a paradox: *the faster and more widely a tool spreads, the harder it becomes to measure its precise impact cleanly*. In knowledge work,

outputs are rarely as standardised as those of a factory production line. Audit combines analysis, judgement, communication and quality control — activities that cannot readily be reduced to simple productivity indicators.

Figure 3 – AI applications during the audit life cycle



A second challenge is *heterogeneity of use*. Professionals do not interact with AI in a standardised way: people use it differently depending on their tasks, experience and personal working style. In practice, this often means using AI to complement one’s strengths and offset weak points

– which makes aggregated ‘AI productivity’ claims even less easy to define. The most reliable approach, therefore, is not to chase abstract key performance indicators, but to look at *observable changes in audit practice* and how these translate into better evidence and stronger audit outputs.

'Artificial intelligence as a thinking partner'

Interview with Ioanna Metaxopoulou and Marco Barros Lourenço, ECA Directors

By Gaston Moonen



© European Union, 2026. Source: ECA

The EU's digital transition also means a digital transition for its institutions, including the ECA. The real test of whether the ECA could go digital came in March 2020, when the institution had to go digital almost from one day to the next. And it did, forced by the COVID 19 pandemic. How has the ECA continued down this digital path, particularly considering the possibilities *and* constraints presented by rapidly changing artificial intelligence (AI) tools? Ioanna Metaxopoulou is the ECA Director responsible for the Directorate for Audit Quality Control (DQC), which includes the DATA (Data and Technology for Audit) team. Marco Barros Lourenço is the ECA Director of the Directorate for Information, Workplace and Innovation (DIWI), providing AI technology to the ECA. Both directors explain how the ECA is integrating AI into its work and how the synergy between the two directorates helps to put the ECA at the forefront of using digital technologies, including AI, to deliver its mission as the EU's external auditor.

Providing technology that caters to business needs

...AI adoption is mainly pushed by users, as they approach us regularly with the new use cases. ”

Within the ECA, DQC and DIWI are both largely responsible for introducing and integrating AI into the ECA's work, whether it concerns audit or administration tasks. Marco Barros Lourenço explains why: 'AI is a technology. In most organisations, AI is pushed by users and not the IT department. In the partnership between DQC and DIWI, DQC represents the business and DIWI the technology. We come together in what was established as the AI Steering Committee, chaired by our Secretary General and with representatives from DQC, DIWI and an audit directorate. The Committee provides a strategic overview of how AI is to be implemented at the ECA.' He adds that AI is not only applied to the ECA's core business but also relevant for its administration, ranging from HR issues, translation, legal, finance and procurement processes. 'What is happening now is that AI adoption is mainly pushed by users, as they approach us regularly with the new use cases. They see the added value of using this new technology in their daily work. To meet the demand, we are increasing the resources dedicated to this new technology so that we can respond to the high volume of requests to experiment with it. Compared with other institutions, we are not lagging behind, but we are still at the beginning, and we have plans to deliver more on AI still this year.' Ioanna Metaxopoulou adds that, in some areas, the ECA is actually at the forefront among supreme audit institutions (SAIs). She emphasizes that significant progress has been made, enabling the ECA to access and implement innovative technologies with rather limited resources. Some SAIs, such as the US Government Accountability Office, have many more resources. In her directorate, she has a dedicated team of 10 people. 'This team is responsible for data analytics, IT audits, AI and automation.'

AI is not just a new tool (...) it is collaboration between humans and AI, the machine. ”

For both directors, AI entails many things. For Ioanna, the first thing that comes to mind is change. 'Change in the way we have traditionally learned to work. Now, we are required to rethink and reshape our working methods.' She recalls something she heard recently: 'AI is not just a new tool, such as Excel or Word, but more than that – it is collaboration between humans and AI, the machine. The key question is: How can this collaboration be harnessed to achieve the best outcomes? There are tasks where AI excels beyond human capability, just as there are areas where humans outperform machines. We need to find how to work best together.' She recognises that perceptions of AI vary considerably among ECA staff and management. While some auditors champion AI, exploring its potential to transform and enhance workflows for greater productivity and efficiency, others remain sceptical.

It [AI] will be completely transformational over the next one to two decades. ”

For Marco, considering what is happening in the 'technology space', it is inevitable that AI is going to be more revolutionary than anything we have seen in the past. 'More than the computer, more than the internet. One of the reasons is that it is happening faster than any other technological revolution we have seen in the past, and it is more ubiquitous, covering more areas than ever before. This will have an enormous impact on the way we work, the way we live, the way we interact with each other and with technology. It will be completely transformational over the next one to two decades. And this transformation has already started.' Ioanna illustrates this with an example regarding generative AI: 'If you think about how it started in 2022 and how it works today, that is quite something.' Marco explains that AI will be transformative for everyone, not only people working with personal computers or mobile devices. 'Also, for other areas such as biology, aspects relating to agriculture, space, the weather – actually, for many human activities. It is so transformational that we don't know what its limits are.' He notes that, with other technologies, one could anticipate and simulate outcomes using mathematical algorithms. 'With AI, we don't know the capabilities of this technology, nor the impact.'

As for how to master a technology without knowing its limits, Marco observes that we are at the very beginning of the adoption cycle. 'At this moment it is difficult to evaluate the potential of this technology, it is beyond our understanding.' He also refers to the costs involved, such as energy costs. 'We don't produce enough energy to power all the AI data centres in the world for everyone to benefit from this technology and to implement all the uses cases we know today. It will require building many power plants.' He refers to an idea that recently surfaced: 'Releasing some of these data centres in space to capture solar power and be cooled at the same time. Our creativity, aided by technology and AI, will help to find solutions.' He believes that, by mid-century, the world will look very different from how it did at the start of the century. 'And not only because of the geopolitical situation.'

The key is knowing how to locate knowledge and apply it effectively... ”

Ioanna sees major changes coming not only because of the possibilities of AI, but also because of how the younger generation is embracing this technology. 'When they enter the job market, their approach is entirely different from ours.' She is not afraid AI will weaken certain human skills, such as drafting texts. 'In fact, AI has the potential to strengthen these abilities. When you use AI to draft a document, it often enhances the quality, and as a result, you further develop your own skills. This partnership between humans and AI creates a mutually beneficial, win win scenario.'

Marco offers a provocative example: 'Do you believe that in 10 years' time, people will read some of the lengthy reports we have?' For both directors, it is clear that communication will be profoundly affected not only by the use of AI but also by a new generation growing up with this technology. Ioanna sees this firsthand with her son. 'For him, memorizing facts is unnecessary, as information exists readily accessible at his fingertips. The key is knowing how to locate knowledge and apply it effectively, whilst also being able to discern that the information is accurate, unbiased and not misleading or false.'

From support tool to active AI agents talking with each other

As for the ECA's audit work regarding AI, Ioanna identifies two principal areas to focus: 'Currently, we conduct standards audits on the use of information technology by our auditees. But you will need to do audits also on the algorithms used. Together with our key stakeholders, we work on how to prepare our IT auditors to be able to audit these systems in the future. Learning what is happening in the private sector in this respect is also important.' The second area is AI agents and automation. 'The ECA is working actively in this area with the aim of developing and deploying intelligent agents that can streamline and enhance audit processes and administrative functions. This involves experimenting with AI assistants that can prefill audit programmes, generate findings from public reports, and even collaborate with other AI agents to produce documents and presentations in the ECA's style. I'm genuinely thrilled about the possibilities AI brings to our work.'

Everyone [in the ECA] currently has access to at least one generative AI solution. ”

For the two directors it was important that the need for digital transformation be prominently reflected in the [ECA's 2026-2030 strategy](#). Ioanna explains that this involved setting new objectives, such as upskilling and reskilling staff, equipping them with the necessary tools, and exploring ways to enhance specific processes through data through data analytics, automation and improved workflows. Her department is leading this effort, with valuable input from the audit side. 'We have invested significant effort into rolling out generative AI, data analytic capabilities and automation for ECA staff as quickly as possible. Everyone currently has access to at least one generative AI solution.' Furthermore, we have gone one step further and we have introduced a new [strategy for 2026-2030](#) focused specifically on AI for the institution.

...generative AI becomes more and more important to use in audit, responsibly and ethically. ”

Marco explains that he has recently created new teams in his directorate dedicated to the digital workplace, solutions, data and innovation, with the aim of responding to the business needs expressed by the users in various domains, including AI. 'Unfortunately, we lack the financial and human resources to provide support as quickly as we would like. Still, AI currently takes up about 5 % of our directorate's budget for IT.' He explains the three main pillars that he considers necessary for the successful implementation of an AI project: 'One is governance, which we have from the highest level at the ECA, the ECA College, supported by the Secretary General and with the involvement of directors. Another pillar is user engagement. Many colleagues already received training on the use of generative AI. And we have organised several workshops and constituted a network of AI champions.' Ioanna explains that they surveyed staff among others on what kind of AI related training they would like. 'We now have a clear roadmap, as generative AI becomes more and more important to use in audit, responsibly and ethically.' Marco confirms that the ECA has invested substantially in this and considers that it is ahead of many other institutions.

'However, as regards the third pillar, which is data, we are less advanced. Data is necessary for any successful implementation of AI. Our data is scarce, highly fragmented across different systems and databases. Ioanna and I have been working on a new project to modernise our various information systems so that we can have a new architecture for data, paving the way towards new use cases for AI.' Ioanna highlights the ECA's dependence on its auditees for data: 'The quality of the auditees' data is essential.' Data confidentiality is another aspect. Marco explains that the roll-out of Copilot at the ECA is cloud based. 'But it has a safeguard, provided to the EU institutions by Microsoft after the assessment conducted by the European Data Protection Supervisor and the approval by the Commission to Office 365. Still, we are working to have local models, replicating the approach of other EU institutions, including the Commission. These models are installed only locally and disconnected from the internet or air gapped, to ensure the protection of the ECA information. This is particularly important to meet the requirement of having results contextualised with ECA information and data that is stored locally.'

...it is essential that we proceed with caution. ”

Ioanna also highlights the introduction of additional safeguards to ensure responsible use of AI within the ECA. 'We have updated the guidelines for ECA staff on the use of AI, specifying what can be utilised and under which restrictions.' The same approach applies to the AI agents under development. 'Together with DIWI, we try to ensure that the development and deployment of these AI agents fully comply with all relevant regulations and are aligned with the provisions of the AI Act. As these technologies are still emerging, it is essential that we proceed with caution.'

...you can have a specialised team of AI agents doing different things... together. ”

Looking ahead, Ioanna envisions AI agents that will be able to actively carry out some of the ECA's tasks. Marco goes further. 'Imagine that you have agents that are absolutely specialised or have knowledge in topics that are relevant to our users. This is not science fiction, the technology is now available, and we can make it available at the ECA. For example, we can train an agent to produce PowerPoint presentations with the same style, look and feel, same colours, etc. and another agent specialised on drafting ECA style memos. The two agents can cooperate, and one can ask the other to prepare a PowerPoint presentation, without human intervention. So you can have a specialised team of AI agents doing different things... together.'

Ioanna provides an illustration of an AI agent at the ECA, capable of generating audit findings from ECA working papers and audit evidence. She notes 'we are actively pursuing several projects, and we hope to advance

...to protect the confidentiality of our own information, and our intellectual property (...) is relevant to our sovereignty...

”

towards their completion already by the end of 2026.' As for sensitive data, Ioanna clarifies that it should be handled with particular care.

In this respect, Marco refers to the significance of cybersecurity, an area he worked in before. 'All these things relating to AI make us think about our data being used by others for monetisation or how it can be used to improve the efficiency of cyberattacks. How important it will be for us to protect the confidentiality of our own information, and our intellectual property. This is currently being debated, as it is relevant to our sovereignty and ability to protect our organisation against external influence or manipulation. Also to protect the way we perform our work, our audits... In the end, protecting our mission and our role as the EU's external auditor. It is important to consider that today, we use mostly commercial and cloud-based AI models developed by companies from third countries. However, we have to ensure that we still maintain our own capacity to use local models with our information systems locally, where we can ensure a higher level of protection for our information.'

From staff productivity to corporate productivity

We need process automation systems that integrate with audit and document management systems, with AI models and with our information when needed.

”

Looking back, Marco sees that the ECA, as an institution that took important steps towards a digital transition, which will continue, but now through the integration of AI and process automation where applicable. 'In 2019, my predecessor continued the ECA's digital transformation with important initiatives, such as ensuring that everyone had a laptop and remote access to our systems. This enabled everyone to work remotely when the COVID 19 crisis came in March 2020. This initiative alone allowed the organisation to move entirely to fully digital audit, processed almost immediately, while respecting certain guard rails such as transparency, accountability and privacy.'

Marco explains that for the implementation of AI the ECA needs to consider various options, such as having Copilot and GTP@EC and local models for everyone. 'To reduce our dependency on one particular technology or provider.' He also refers to the need to replace the technical architecture of the ECA's audit data and management systems to ensure more integration, interoperability and usability with a better user experience for auditors. 'We need process automation systems that integrate with audit and document management systems, with AI models and with our information when needed.' Ioanna emphasises that these measures must be carried out in parallel to drive transformation, while at the same time upholding the ECA's core principles regarding ethics, accountability, and transparency. A key requirement is the delivery of reliable reports and audit conclusions. Ioanna is very clear on this. 'The trustworthiness of our products is paramount.'

It is essential that AI is employed responsibly and that human oversight remains integral throughout the process.

”

It is essential that AI is employed responsibly and that human oversight remains integral throughout the process. There have been instances in the private audit sector where certain products were AI generated and included erroneous or misleading information, commonly referred to as hallucinations. Such incidents could severely damage our reputation as an audit institution.'

Marco is optimistic on the possibilities in this area, seeing the potential for software coding to be replaced by machines. He also explains that he and Ioanna are planning meetings with colleagues to build ties with what he calls 'Champions.' 'Promote discussions on how they see the audit profession in 10, 15 years from now, how they see the impact of technology in the role of the ECA.'

Ioanna sees the effects of AI on the ECA's output. 'In my directorate, we use generative AI for improving our drafts and it has helped us tremendously to with the quality of our documents. We have also built a tool to check footnotes in line with our methodology. We are planning to develop an AI

assistant for carrying out an initial review of our draft reports. In addition, we are actively using AI looking for to enhance our work and streamline our processes.' She also sees ample possibilities for doing full population audits. 'We do that already when auditing agencies, in some areas of their activities.' Marco gives examples of generative AI used for text review. 'We implemented one of the first models here with our library, using AI algorithms to improve the ability to search our book catalogue. We use AI in translation, and in cybersecurity to identify attack patterns.'

...it [AI] could increase productivity up to 20%... ”

Both directors are convinced that AI has great potential to increase productivity within the ECA. Marco explains: 'From a survey held among ECA staff regarding the use of Copilot, we obtained the feedback that it could increase productivity up to 20%, which is similar to what we hear from other SAIs.' This was a major incentive to roll out Copilot to everyone at the ECA.

Information on AI's added value is essential for management when deciding on future investments. 'The ECA's AI Steering Committee decided to define specific KPIs, evaluating investments and their benefits.' For Ioanna, it is clear: 'During the last 18 months we have made huge progress, making leaps rather than taking small steps.' Marco agrees: 'We have improved a lot, particularly over the last six months, in AI literacy.' And according to both directors, staff are enthusiastic. Marco adds with a smile: 'To give you an idea, we have training sessions on Copilot with waiting lists.' Ioanna complements that our GPT coffees are always well attended and very popular.'

What I often do is use Copilot not to write but to give me ideas, giving me a different perspective. ”

Both directors underline how much they use AI for their daily activities. Ioanna: 'AI assists me in many things on a daily basis, for example, improving the quality of documents.' Marco concurs: 'It surprises me every day. We receive many reports to review, and we use AI to check consistency, but also to stimulate new ideas. To check things which would otherwise take a lot of time. What I often do is use Copilot not to write but to give me ideas, giving me a different perspective. I see AI as a thinking partner.' Ioanna adds: 'And to shape and enhance your ideas further.' For Marco, the impact of AI on work when it comes to productivity is a no-brainer. 'I consider it as low hanging fruit. The next thing we are working on is corporate productivity, to roll out AI in the organisation, improving the efficiency of how we operate.' He gives the example of the public procurement audit project which will be tested later this year, aiming to reduce the duration of the analysis and potentially increase the sample of procurement files. For the two directors, it is clear that AI will change the way audit work is carried out, will support processes within the organisation. Ioanna concludes: 'It can help us focus more on what really matters, on quality aspects. Let the routine parts be done by the machine.'


Ensuring quality will remain a challenge in the AI era

Data will remain key. And we are not at a point where we can ensure the quality of the data that we audit. ”

When asked to look ahead, Marco says that, over the past 10 years, digital audit has developed substantially, aiming to transforming data into information, using tools and databases, etc. 'Over the next 10 years, the focus will be on generating knowledge. For example, when you conduct a performance audit of cohesion funds and are reviewing the past ten years of all the data gathered and the information that is available for you, you will be able correlating the information and draw stronger conclusions, with different perspectives on the use of cohesion funds.'

He underlines that DIWI is there to provide the infrastructure and the tools to do this. 'And the DATA team is there to find the business and use cases and exploit it! The same way as DIWI runs the Dataware house and DQC provides the knowledge out of the data.' Ioanna is optimistic but also careful. 'We provide the tools for auditors to do their work. But be careful here – our auditees are also using AI. For example, the Commission may use AI to draft

policies and beneficiaries on the ground could use AI to develop their projects. So we have to be mindful about what we audit.' She warns that there is still some way to go: 'Data will remain key. And we are not at a point where we can ensure the quality of the data that we audit. This has been a long-standing issue. But now the fraudsters have more imagination and possibilities to use AI technologies to their advantage. We have to be ready to address new risks.' Marco concludes that this is why governance is so important. Ioanna concurs and adds: 'And to work together. If we work in silos, we will achieve very little. Introducing and leading innovation will require us to work together.'



Demonstrating alignment with the Digital Decade and its EU objectives

The years 2020 to 2025, also known as the Digital Decade, were intended to be crucial for the digitalisation of public administration services provided by central governmental bodies. Under the Czech Act on the Right to Digital Services, most registered administrative services were to be fully digitalised by February 2025. The law also embodied the commitment by the State that, from that date onwards, citizens would be able to manage most official administrative matters online – that is, without having to visit an office in person or resubmit documents already held by the State. The government has failed to honour this promise to its citizens. The reality? As at the end of January 2025, only 18 % of services were fully digital.

In our performance audits of the digitalisation of the Czech public administration, we systematically use the Digital Decade targets as a reference framework for assessing the success of the Digital Czech Republic strategy. This approach makes it possible to evaluate not only national compliance with the Act on the Right to Digital Services, but also whether the Czech Republic is on track to meet the EU objectives as regards fully online services, interoperability, and user-friendliness.

Our Comprehensive Report integrates three layers based on the Digital Decade:

- the normative layer – i.e. the Act on the Right to Digital Services, national strategies, and commitments arising from the Digital Decade (including targets for 100 % online delivery of key services and the availability of electronic health records);
- the performance layer – i.e. the efficiency with which funds are used and the extent to which procedures have genuinely been fully digitalised, including the use of central shared services;
- the user layer – i.e. the real ability of citizens and businesses to manage key life events entirely online, without needing to visit an office in person or submit duplicates.

Hallmarks of this period include the strengthening of the legal anchoring of digital services, progress in building key digital infrastructure, efforts to enhance the coordination of government digitalisation, as well as a number of unfulfilled ambitions in major e-government projects, mainly led by ministries.

Key barriers to successful digital transformation

Our audits have shown that the core issue is not a lack of funding, but rather structural barriers in the management of digitalisation. The most significant obstacles include the fragmentation of public administration information systems, weak coordination among departments, and insufficient use of central e-digital government services.

The SAO highlights that the State has failed to fulfil part of its statutory obligations, despite the adoption of key laws and strategies and the establishment of the Digital and Information Agency¹. A typical example is the e-justice strategy, where the Ministry of Justice achieved only two out of 20 specific objectives and several activities did not align with the central e-digital government solutions (shared services).

The Digital Czech Republic 2018+ strategy (including subsequently updated priorities) sets out goals similar to those of the current Digital Decade 2030 policy programme. In practice, however, the full digitalisation of the related services remains incomplete, and citizens are often forced to combine digital and paper-based channels. The inconsistency between political declarations and actual implementation is also reflected in the Czech Republic's only very gradual improvement

in EU comparisons of digital public services (Digital Economy and Society Index).

We saw a key pattern when auditing the Digital Decade programme: unless the internal digitalisation of public administration entities is part of projects that allow users to carry out procedures fully online through self-service client portals, digitalisation will not deliver efficiency commensurate with the money spent.

Across its five dozen audits, the SAO identified seven barriers (see **Figure 1**) which, by their nature, affected areas related to the results of the ICT/e-government focused audit activities. The most common barrier, identified in 87 audit findings, is the *non achievement of project objectives and low economic efficiency*. In second position, *incomplete data and complex data flows*. The SAO linked *legal, contractual and procedural constraints* to 59 audit findings. The next most relevant barriers were a *low level of project and quality management* and *strategic and design weaknesses*, with more than 30 audit findings for each of the two. Finally, *costly operational solutions* and *shortage of expert human resources* were linked to fewer audit findings, but this does not detract from their weight due to their cross-cutting nature.

¹ Act 12/2020 Coll., on the Right to Digital Services and on Amendments to Certain Acts; Act 365/2000 Coll., on Public Administration Information Systems and on Amendments to Certain Other Acts; the Digital Czech Republic 2018+ strategy – overarching objectives of three component strategies, ranging from the national contribution to the EU's digital agenda (1), through digital public administration (2), to the preparation of Czech society and the economy for engaging with and benefiting from digitalisation (3).

Figure 1 – Seven main barriers to the successful development of e-government



Source: Supreme Audit Office of the Czech Republic.

In their digital transformation, central administrative authorities have applied an undesirable ‘everybody in their own way’ approach. This has manifested, among other things, in high costs for ensuring interoperability between administrative registers in operation. Ensuring ‘value for money’ throughout the entire life cycle of the e government infrastructure and reducing the fragmentation of IT solutions across central administrative authorities and their services are prerequisites for achieving control over continuously increasing ICT expenditure while its effectiveness is deteriorating.

Some e-government projects were launched without clear responsibilities, accountability or enforceability of synergies between the parties involved. Legislative processes in the Czech Republic do not respect the principles of digital-friendly legislation, as is customary in other EU countries. The principles for creating digital-friendly legislation had already been presented to the government in 2017 but were discussed without a resolution. The principles mention the basic architectural principles of e-government (e.g. digital by default, only once, consolidation of information systems, and shared services), but do not provide for specific procedures for the preparation and evaluation of legislative proposals. The regulatory impact assessment (RIA) system – which builds on the general principles of RIA (ex-ante and ex post) and assesses what information is needed to administer solutions, how best to collect it, and whether

information and communication technologies can feasibly be used to obtain and handle it – does not effectively compensate for this shortcoming.

We identified an effect that goes against the idea of digitalisation itself: instead of analysing the procedures and optimising them to enable meaningful and efficient digitalisation, citizens are faced with more complex forms and civil servants are taking unnecessary steps.

Inspiration for the introduction of a digital readiness system in lawmaking can be found, for example, in the neighbouring Federal Republic of Germany, where proposals are screened using Digitalcheck as part of the legislative process. In Denmark, the verification of upcoming legislation follows the principles of digital-ready legislation, and this process is the responsibility of the Danish Agency for Digital Government.

The main body of the Summary Report can be summed up as follows. Findings from our audits confirm that the main barrier to successful digitalisation is a lack of understanding of client needs, i.e. those of citizens and, to some extent, also of officials. It is also the inability to use digitalisation as an exceptional opportunity to change established procedures and processes. Digitalisation in public administration is most often implemented at the level of forms, without a prior review of the necessity of each step in a given workflow. As a result, online services

are made available via public portals, while, behind the scenes, officials continue to work as before, carrying out most tasks manually. Digitalisation therefore delivers

neither the expected convenience and faster processing for users nor the desired savings in overall administrative costs.

Sharing perspectives for the years ahead

The Czech e-government will continue to face challenges in the coming years – from the limited effectiveness of fragmented solutions to the low usability of services, to problems in their implementation, showing that the readiness of digital friendly legislation has been underestimated. In 2025, research carried out on the Czech Republic's digitalisation as part of the Services. Digital project produced a number of positive findings that confirm the interest of citizens and businesses in online communication with the State².

However, a significant proportion of citizens and businesses have been dissatisfied with the State's digital services, and in this context the report of the Services. Digital project highlights, among other things, the lack of digital service quality management. A single framework must be set to measure efficiency, satisfaction and service quality, linked to the role of the product owner responsible for the e-government service. In particular, quality management of digital services requires continuous data collection on the performance of the administrative procedures.

The Czech Republic adopted its own National AI Strategy in 2019. At EU level, the Artificial Intelligence Act³ was adopted in 2024, and the Czech Republic followed up by adopting an updated version of its strategy, the [National Artificial Intelligence Strategy of the Czech Republic 2030](#) (NAIS).

The [Government AI Readiness Index](#), published annually

Trust is a key aspect

In conclusion, allow me to recall the words of Miloslav Kala, President of the SAO, who emphasised in our Comprehensive Report that trust is a key aspect of e-government. Czech and EU citizens alike expect that, for investments running into billions, they will

by the Oxford Insights consultancy, compares the preparedness of 188 countries for the implementation of AI. The index monitors 40 indicators across three pillars: government, the technology sector, and data and infrastructure. In the 2024 ranking, the Czech Republic was placed 31st worldwide and 18th among European countries. The assessment shows that the Czech Republic has high data availability and a clearly defined governmental strategy for artificial intelligence. However, it continues to lag behind in developing the technology sector and strengthening the digital capacity of its public administration, falling short of its strategic ambitions.

Consequently, one of the main challenges in the Czech Republic remains the coordination of AI implementation within projects and activities related to the digital transformation of the public administration. In addition to the existing responsibilities of the Digital and Information Agency and the Ministry of the Interior, the Ministry of Industry and Trade has also, under the NAIS, established an institutional framework for cooperation. Under this framework, the respective competent authorities are represented on the newly established Advisory Committee on AI. The resulting complex governance structure, with an unclear allocation of accountability for outcomes, continues to pose a risk. This underlines the need to create synergies in investments in the development of e government, including AI implementation.

gain a responsive, user friendly and efficient public administration, one that supports them instead of placing obstacles in their way. Trust in the digitalisation of public services is an integral part of trust in public institutions, and it is not something we can afford to gamble with.

² Česko.Digital, *Projekt Služby.Digital*, <https://en.cesko.digital/projekty/sluzby-digital/home>

³ Regulation (EU) 2024/1689 (Artificial Intelligence Act).

Digital transformation and the changing role of supreme audit institutions

By Rudi Turksema, Colin van Noordt and Burhan Gün (Netherlands Court of Audit)



© eabff / depositphotos.com

Digitalisation, particularly in the form of AI, provides a means of increasing the efficiency and effectiveness of governmental activities. However, it also comes with certain risks when deployed irresponsibly. Rudi Turksema, digitalisation programme manager, Colin van Noordt, an auditor, and Berhan Gün, a director, all at the Netherlands Court of Audit (NCA), reflect on the audits the NCA has undertaken on digital government over the last few years. They show what the findings reveal, and what the growing importance of and dependence on digital technology means for external auditors in their quest not only to provide oversight, but also to ensure that the government’s digital transformation serves the public good.

Digitalisation in the Dutch government: opportunities and risks

Many governments are increasingly deploying digital technology in their activities. There are good reasons for doing so: it can improve governmental processes, public service delivery, citizen engagement, and help to solve social challenges. In the Netherlands, most public services nowadays are only available digitally. A notable example is the pre-completed tax form. As the Tax Agency already has the information it needs, only a small effort is required

to proactively fill in the information for taxpayers. As a result, 99% of tax forms are completed digitally, taking up less time for citizens and business, and ensuring that fewer mistakes are made.

However, the benefits that digitalisation can bring to governments also entail several risks. If not addressed, these could lessen the benefits of digital technology,

infringe citizens' rights, and potentially cause harm. The advent of AI has made the situation even more acute. Increased dependence on digital technology certainly gives cause for concern in many areas, ranging from the use of cloud technology, the black box nature of algorithms and AI, and digital resilience, to excluding citizens with fewer digital skills.

The Netherlands is one of the countries that comes consistently high in digital government rankings, such as the e-Government Benchmark or the DESI index, being renowned for its high volume of digital public services and pre-filled forms. The Dutch government started its digitalisation drive relatively early. For instance, the (then) Tax and Customs Administration took its first major steps in this direction in the 1980s and 1990s, when it began to automate processes and introduce digital systems for tax returns.

The introduction in 2003 of DigiD, a digital identification tool that provides secure access to government services, was a major step towards digital government (*digitalisation*). DigiD enables secure interaction with, for example, the government or health insurers. We audited the functioning of DigiD, and found that it works well for the most part. DigiD has thus become an important building block for digital government in the Netherlands.

Over the years, more processes and services provided by

the Dutch government have been digitised. Nowadays, interaction with the government is mainly digital. When it comes to major social issues such as security, the economy, employment, and health, digitisation has become crucial for all activities.

Just as digitisation has grown in importance for government activities, so have the attendant risks and dangers of digital government. Issues such as overcoming the digital divide, placing too much trust in technology, a lack of transparency vis-à-vis the public, and information security vulnerabilities can quickly undermine public trust. This was intensified by a court ruling that SyRI (an algorithm-based fraud-detection instrument used by local government) was in breach of human rights. The system flagged potential welfare fraud, but its lack of transparency and accuracy resulted in wrongful accusations, disproportionately affecting vulnerable citizens.

The Dutch government's increasing dependence on technology and its suppliers is especially relevant in the current geopolitical context. The possible takeover of the provider of DigiD's underlying technical infrastructure by an American company has been causing considerable concern, and has raised questions about public control over vital digital infrastructure.

Digitalisation as a new audit object

At the Netherlands Court of Audit (NCA), we see digitalisation as a matter of strategic relevance. Digitalisation changes not only *what* is audited but also

how audits are conducted. Over the past few years, we have audited various aspects of digital technologies, including digital identity, algorithms, AI, cloud computing,

Box 1 - Recent publications on AI by the Netherlands Court of Audit

- Netherlands Court of Audit, [Understanding algorithms](#) (2021).
- Netherlands Court of Audit, [An audit of algorithms](#) (2022).
- Netherlands Court of Audit, [Digital identity demanding a lot from DigiD and eHerkenning](#) (2023).
- Netherlands Court of Audit, [Focus on AI in central government](#) (2024).
- Netherlands Court of Audit, [Dutch central government in the cloud](#) (2025).
- Netherlands Court of Audit, [Focus on quantum technology in central government](#) (2026).

Some other publications relating to AI and SAIs:

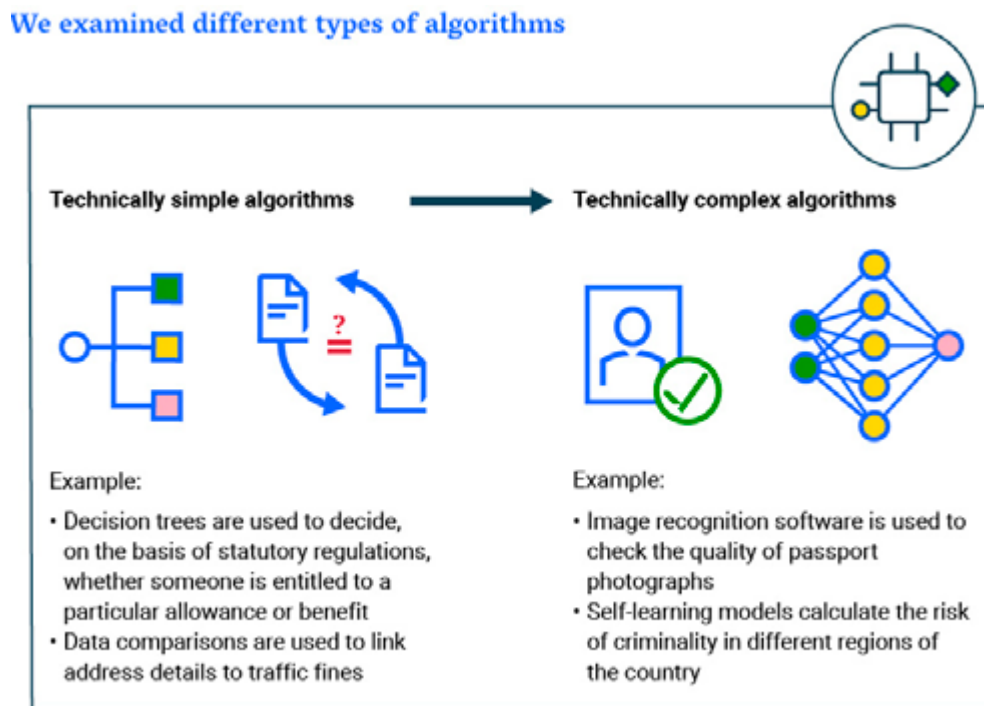
- European Court of Auditors, Special report 08/2024 [EU Artificial Intelligence ambition – stronger governance and increased, more focused investment essential in going forward](#).
- INTOSAI Journal, [Experiences on Auditing Algorithms and AI in the Dutch Government](#) (2024).

and even quantum technology. These audits have revealed that while the Dutch government is exploring opportunities, it is also struggling with technical debt, dependence on external suppliers, and limited in-house technical capabilities.

A particular focus of our audits has been AI and the algorithms being deployed by the Dutch government. This started with two separate performance audits on algorithms and the development of an audit framework

(2021 and 2022). The audit framework covers both simple rule-based systems and more complex systems based on machine learning. It is a multidisciplinary approach, including norms on governance, privacy, models and data, as well as general IT controls. See Box 1 for NCA reports on digitalisation, and a few other publications. As of 2023, audits of algorithms are part of our annual financial audit work (see **Figure 1**).

Figure 1 - Simple and complex algorithms



© Netherlands Court of Audit

Algorithms are used by many layers of government in the Netherlands. Governments deploy them in various ways, such as to detect tax fraud and errors, the assumption being that analysing data to identify irregularities is quicker and more accurate than manual checks. It might also improve compliance and eliminate the arbitrariness of human decision-making. These algorithms derive their explicit instructions from programmers. Artificial Intelligence (AI) involves specific algorithms whose instructions are based on data inputs. AI systems must meet additional requirements under the EU AI Act, if, for example, they are used in a high-risk area such as essential public services.

There are high expectations of the opportunities for society and government operations that AI could provide, both for the Dutch government and its counterparts in the European Union. Stimulating the use of AI in government is an important way to make

government more effective and efficient and to tackle social challenges. At European level, the public sector is also regarded as a key and strategic sector in which to strengthen AI capacity and use AI for the public good.

Although algorithms and AI both have considerable potential for improving government performance, they also entail risks when not deployed responsibly. They might contain biases that lead to discriminatory outcomes, or personal data may not be properly protected. A lack of transparency when using the technology might lead to governance challenges. Furthermore, a major question remains as to whether deploying AI and algorithms actually achieves the expected increase in performance and efficiency. In other words, are the costs of developing and deploying AI in government worth the gains?


Our audit results on AI and algorithms

These questions are central to our audits. However, answering them first requires insight into where, why and how AI and algorithms are being deployed by the Dutch government. A performance audit on the deployment of algorithms and the subsequent use of AI technologies by the Dutch government helped us to establish this broader overview.

Our audit showed that the deployment of AI is still in a preliminary phase despite several years of encouragement. While most of the 70 governmental organisations we audited have experience with AI, almost all of them deploy no more than three AI systems (see **Figure 2**).

Figure 2 - Number of AI systems in use per organisation

Government organisations use few AI systems



NB: UWV is the Dutch Employee Insurance Agency which helps people who are out of work or unable to work.

© Netherlands Court of Audit

In our audit, we highlighted that most of the AI systems that were trialled in the last few years have remained in the experimental phase, or have already been discontinued. Just over a quarter (28%) of the AI systems in question were actively being implemented. Surprisingly, it is often unclear if performance matched expectations. This raises doubts about whether AI systems meet the high expectations of increasing the effectiveness and efficiency of government.

We were also able to identify the key risks of deploying AI, such as limited transparency regarding where and why these systems were being deployed. Even though the Dutch government has an algorithm register, it lists only 5% of the 433 AI systems we identified in our audit. We also found that less than half of the systems currently under development or already in use had undergone a risk assessment. The potentially negative consequences of deploying AI thus remain hidden.

Most organisations classify their AI systems as minimal-risk AI under the AI Act (**Figure 3**), meaning that they are subject to no additional requirements under the AI Act. We note there is still considerable uncertainty surrounding risk classification under the AI Act, as


similar systems are classified in different ways. Most organisations describe such classification as a work in progress.

In our more in-depth audits of the deployment of AI and algorithms, we often conclude that algorithms do not always meet the basic requirements of governance. We see challenges in terms of the algorithms' performance, biased outcomes, data leaks, and unauthorised access. The good news is that looking at those algorithms that do meet requirements, having the proper controls in place is not mission impossible.

A recurring issue is the apparent disconnect between actual 'business' and algorithms. Algorithms that are not firmly linked to the performance of public tasks tend to contain flaws. The development of algorithms regularly takes place in a silo environment. As a result, it may be challenging to convey business aims to the development team. One example of this is the algorithm that is supposed to provide data-driven support to road inspectors. In practice, the inspectors hardly use it because they say that their knowledge and experience mean that they already know what to do.

Figure 3 – Organisations’ estimates of their AI systems’ risk classification

Organisations classify most of their AI systems as minimal risk



This figure presents information on only the 287 AI systems currently under development or in use.

© Netherlands Court of Audit

Another disconnect in this case was using an external partner to build the algorithm. This made it impossible for the relevant government department to monitor the algorithm’s performance. In our view, working with an external partner does not mean that public administrations are any less responsible for checking their algorithms.

Other frequently detected problems with algorithms relate to privacy, data and modelling, and general IT oversight. In the area of privacy, legal requirements often exceed the necessary capabilities, resulting in a backlog of under-documented algorithms and outdated

Data Protection Impact Assessments. Lastly, we often see that basic IT requirements have not been met, making algorithms vulnerable to outside threats.

In general, when it comes to deploying other digital technologies, digital government also struggles with persistent problems surrounding the provision of digital information and inflexible IT systems. And such government is equally unable to perform adequately due to a lack of sufficient IT knowledge. Government may very well have great ambitions, but it provides insufficient guidance and delivers insufficient results on this very important issue.

How digitalisation and AI can empower SAIs

For supreme audit institutions (SAIs), this presents a critical challenge: how to audit a government that is increasingly reliant on algorithms, while ensuring that digitalisation serves the public good without compromising fairness or accountability?

Our view is that SAIs must not only audit the way their auditees use these tools, but must also adapt their own methods to keep pace with the changing nature of government. Paraphrasing the Dutch painter Willem de Kooning: *‘We have to change, to stay the same.’* Without digital innovation at SAIs, there is a danger of their lapsing into irrelevance and becoming a ceremonial player in the audit system. Progress is necessary to continue to audit government performance.

We see five main areas where the work of SAIs can benefit:

- *Enhanced use of data (analytics):* advanced analytics enables SAIs to process and analyse large datasets more effectively than by using traditional methods. Algorithms can identify patterns, anomalies, or potential fraud in financial transactions, allowing auditors to focus on high-risk areas;

- *Automation of routine tasks:* many routine audit tasks, such as data collection, validation, and reporting, can be automated using AI and robotic process automation (RPA). This reduces administrative burdens, minimises human error, and accelerates the audit cycle, allowing SAIs to focus on more value-added activities;

- *Real-time and continuous auditing:* digital tools allow SAIs to move beyond periodic audits towards real-time or continuous monitoring. By integrating with government IT systems, SAIs can track public spending, programme performance, or compliance as events unfold, helping to prevent issues before they escalate and thus providing policymakers with timely insights; and

- *Strengthening transparency and public trust:* by leveraging digital tools, SAIs can make their findings more accessible and understandable to the public. Interactive dashboards, data visualisations, and open-data initiatives help to demystify complex audit reports, fostering greater transparency and trust in public institutions.

AI as a double-edged sword of opportunities and risks

To conclude: deploying AI and algorithms is a double-edged sword for the public sector. While it offers opportunities to improve efficiency, decision-making, and service delivery, it also introduces risks related to bias, transparency, privacy, and sovereignty. This highlights the need for robust oversight.

SAls have a unique opportunity to become digital guardians of public trust in the 21st century. The challenge lies in ensuring that AI is deployed responsibly, maximising benefits while mitigating risks. This requires robust audit frameworks, continuous monitoring, and a commitment to ethical AI principles. SAls can thus be a critical voice in a world that sometimes takes the benefits of digital technology for granted.

The question is not whether SAls will adapt, but how quickly and effectively they can harness the power of digitalisation to serve the public good. Audit institutions have to realise that this requires investment in technology, skills, and ethical frameworks. As SAls embrace these tools, they must also address challenges such as data privacy, algorithmic bias, and the need for continuous upskilling. Our work in the field of algorithms, for example, shows that such audits must be carried out by mixed teams. A combination of knowledge of governance, data, models, privacy rules, and IT is needed to provide a coherent assessment. If done rightly, digitalisation can transform SAls into more agile, insightful, and impactful watchdogs of public accountability.



From experimentation to integration: advancing AI in European public audit

By Alvar Nõuakas, National Audit Office of Estonia



© Andres Varustin

more deliberate action in this area, as he believes this will ultimately strengthen accountability and public trust.

AI is out there, and its potential is increasing every day. But how are public audit institutions using this potential? Among supreme audit institutions (SAIs), in meetings organised through their global meetings (e.g. in Egypt in October 2025) or regional discussions, there is clearly the realisation that AI offers potential, but also raises a lot of questions about how to use it. Alvar Nõuakas, the head of the EUROSAI IT Working Group Secretariat, which is currently led by the National Audit Office of Estonia, discusses the great capabilities that AI offers. Yet he also notes how reluctant many SAIs are to embed AI into their daily audit work. He recommends

Capability versus the utilisation paradox

We are very privileged to live in a transformative era, where artificial intelligence (AI) as a pervasive technological force is reshaping both the private and public sectors. Since 2020, we have seen an evolution of frequently articulated slogans in the public auditing community, ranging from 'AI is an emerging technology' to 'AI is no longer just an emerging technology'. Advances in generative AI have significantly expanded the scope of what is technically feasible – every new language model surprises us with added context and similarity to human reasoning. Technological maturity has reached a level where tasks such as document classification,

entity extraction, anomaly detection, and even drafting analytical summaries can be performed with a high level of accuracy.

As the Chair of the EUROSAI¹ IT Working Group (ITWG), representing a network of 46 supreme audit institutions (SAIs), the National Audit Office of Estonia (NAOE) has had the opportunity to observe this shift first-hand, while also observing developments in SAIs. AI capabilities are fundamentally changing the way people work in the private sector, where areas such as software development, efficient AI adoption, and augmenting human capacity

¹ EUROSAI is the European Organisation of Supreme Audit Institutions.

have become a matter of survival. By contrast, the public sector is taking longer to adjust and depends strongly on certain enthusiasts and supportive leadership. SAIs are no exception in this regard: the current state of AI in SAIs – even in Europe – can be described as a paradox of capability versus utilisation.

A recent survey among SAIs' chief information officers (CIOs), conducted by the EUROSAI ITWG Secretariat, shows that only 20% of audit institutions have actively integrated AI into their audit workflows. In these institutions, more than half of employees are using AI in their everyday tasks, while at the other end of the

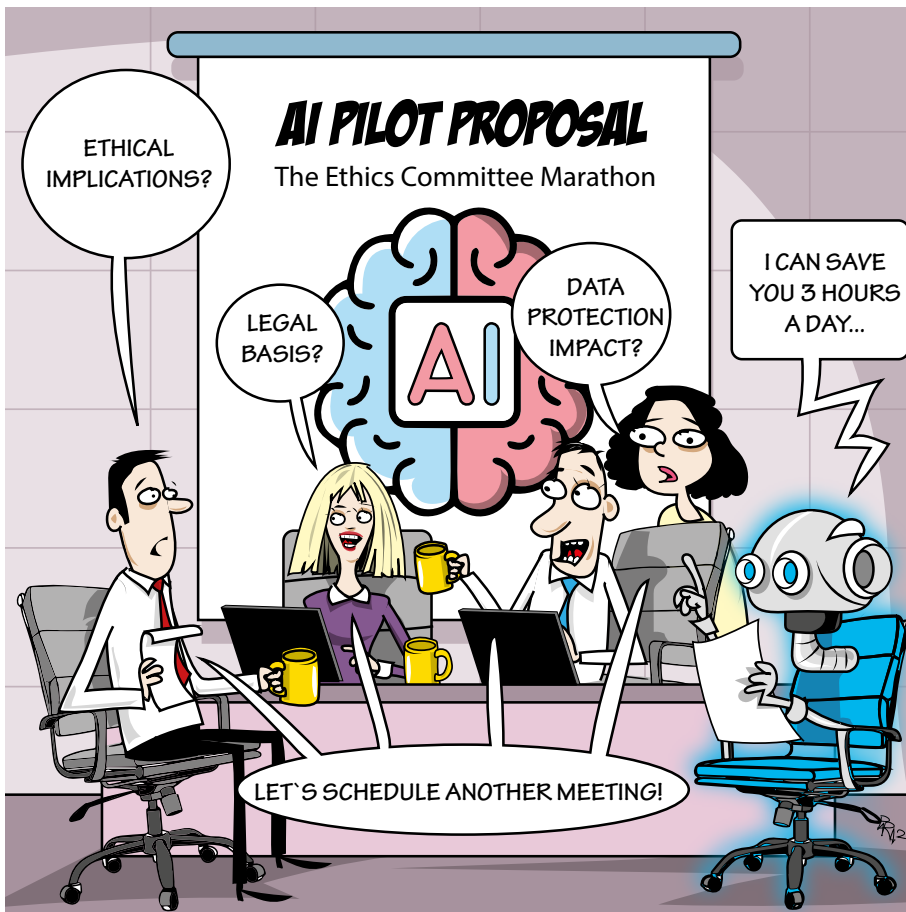
spectrum, over half of SAIs responded that AI is 'planned but not yet initiated'.

In this article, I argue that the time has come to move beyond experimentation and towards systematic integration of AI into the audit lifecycle. While informal exchanges and pilot projects remain valuable, they are insufficient to unlock the full potential of AI in public auditing. Instead, SAIs would benefit from a structured approach: identifying concrete use cases across the audit process, addressing governance and confidentiality risks, and investing in sustainable technical architectures without the risk of vendor lock-in.

Factors contributing to the slow roll-out of AI adoption in SAIs

Public audit institutions operate under strict legal and ethical frameworks, and the introduction of AI raises questions about accountability, transparency, and compliance, particularly in the light of emerging legislation such as the European Union's Artificial

Intelligence Act. At the same time, audit work frequently involves confidential or classified information, making the use of many cloud-based AI services problematic. These constraints are compounded by organisational inertia, as traditional audit methodologies are well-established and



© Andres Varustin

integrating new technologies requires both cultural and procedural change. Lastly, while interest in AI is high and practical expertise in deploying and maintaining such systems is in great demand, the SAIs are still developing their employees' skills.

The result is a landscape we are seeing at every international SAI-community meeting: AI is widely

discussed, but only selectively applied. Within the EUROSAI ITWG specifically, knowledge-sharing activities – such as e-seminars and annual meetings, and informal coffee-break exchanges – have played an important role in raising awareness. However, these activities sometimes remain at the level of 'what is possible' rather than 'what is operational'.

Further progress may depend on SAIs moving from ad hoc experimentation to a more systematic integration of AI. This requires a deliberate effort to map AI capabilities onto the existing audit lifecycle and SAI methodology, from monitoring and planning to follow-up. Each of the

audit stages contains tasks – many of them repetitive and data-intensive – that are suitable candidates for AI augmentation. Such comprehensive design would require full endorsement by management, as well as organisation-wide employee inclusion.

Identifying high-impact use cases

A next step after organisation-wide mapping would be prioritising use cases with high potential for efficiency gains and low implementation risk. I would like to point out some possible phases in auditing where we have experienced some success in adopting AI:

1. Document processing and classification

One of the most immediate opportunities lies in the automated processing of large volumes of documents. Audit work very often involves reviewing legislative texts, executive orders, contracts, and administrative records. In Estonia, we have tested AI-based classification of governmental executive orders. The results have been highly encouraging: models can categorise documents with near-perfect accuracy, requiring only human verification. This significantly reduces the time spent on initial sorting, and enables auditors to focus on higher-value analysis. Similar approaches can be applied across SAIs for classifying audit evidence, tagging documents by topic or risk area, and extracting key entities and relationships.

2. Data analysis and anomaly detection

AI techniques – particularly machine learning – can enhance traditional data analytics by identifying patterns and anomalies that may not be immediately apparent. Potential applications include detecting irregular spending patterns, identifying outliers in procurement data, and flagging inconsistencies across datasets. Financial auditing, with its highly standardised methodology and risk analysis, will certainly benefit from

AI being integrated into automated data processing. For example, we plan to move on with integrating AI into our data warehouse at the NAOE, enabling auditors to identify risks in new and more systematic ways. While these methods do not replace auditor judgment, existing and manually automated risk matrices are here to stay, and additional ones are being created. AI can serve as a powerful decision-supporting tool.

3. Natural language processing in audit work

Recent advances in large language models have opened up new possibilities for working with unstructured text. These include summarising lengthy documents, drafting preliminary audit findings, and comparing regulatory frameworks across jurisdictions. Most of the respondents to the ITWG survey among CIOs who stated that AI has been embedded into the audit process reported success especially in this segment. It is great to see professionals getting better at prompting, sharing their experience, and utilising natural language processing more efficiently. However, these applications must be carefully managed to ensure accuracy, traceability, and compliance with audit standards.

4. Supporting audit reporting

AI can also contribute to the reporting phase by assisting with structuring reports, ensuring consistency in terminology, and generating visualisations and summaries. This can improve both efficiency and clarity of communication with stakeholders.

Governance, ethics, and the European AI Act

Any discussion of AI in public audit must be grounded in a robust governance framework. The European Union's Artificial Intelligence Act provides a comprehensive regulatory baseline, emphasising principles such as transparency, accountability, and risk management.

For SAIs, several implications are particularly relevant:

- risk classification: certain AI applications may fall into high-risk categories, requiring stringent controls. For example, an AI system used to automatically assess compliance or flag potential legal or financial violations in audited entities could directly influence audit conclusions;

- transparency requirements: AI-assisted outputs must be explainable and auditable. Clear guidance should be

- established to reference materials generated by auditors with AI, both in-house and production documents; and

- human oversight: every aspect of audit output must remain under human control. There have been cases – especially in the private sector – of entities issuing misleading statements generated by AI which have affected the trustworthiness of the organisation. Fortunately, stringent controls have been established in SAIs to avoid such scenarios, but these need regular updates as AI possibilities evolve.

Incorporating these principles into SAI methodologies is essential to maintain public trust.

Ensuring secure and confidential data handling

Perhaps the most critical constraint for SAIs is the need to protect sensitive information. Many widely available AI tools rely on cloud-based infrastructure, where data may be processed outside an institution's control. For audit purposes, this presents unacceptable risks in many contexts.

A viable alternative, applied by at least five supreme audit institutions in the ITWG community, is to deploy open-weight models within secure, controlled environments. When combined with retrieval-augmented generation

(RAG), a method where the AI retrieves relevant internal documents before generating responses, these models can leverage internal know-how without exposing the sources externally. Key advantages of this approach include full control over data processing, compliance with confidentiality requirements, and the ability to customise models to domain-specific knowledge. However, implementing such architectures requires investment in infrastructure or agreements at governmental level (e.g. secure government cloud-based infrastructure solutions – IaaS) and technical expertise.

Avoiding misguided AI deployment

An additional risk that deserves explicit attention is the deployment of AI solutions without sufficient business analysis and feasibility testing. In several governmental contexts, there is a growing tendency to adopt AI functions that are readily available within existing office software ecosystems. While such integrations offer convenience and quick accessibility, they may not represent the most suitable or effective solutions for audit purposes.

This approach can lead to sub-optimal outcomes, particularly in smaller language environments where general-purpose tools may lack adequate support for local linguistic nuances. Without careful evaluation, SAIs risk locking themselves into ecosystems that prioritise vendor convenience over audit quality and

methodological rigour.

A more sustainable approach to AI adoption would be guided by clearly defined audit needs rather than by the availability of embedded features. This entails conducting thorough business analyses, assessing alternative tools, and rigorously testing feasibility prior to deployment. In practice, this means evaluating not only performance metrics but also linguistic accuracy, explainability, integration capability, and compliance with confidentiality requirements. By maintaining a technology-agnostic perspective and prioritising audit objectives, SAIs can avoid the trap of convenience-driven adoption, and instead implement solutions that genuinely enhance audit effectiveness.

Building capacity and collaboration

The successful adoption of AI in public audit is not only a technical challenge: it is also an organisational and cultural one. More than ever, it has become important to engage an SAI's full capacity – especially its people – to take the transformative leap. SAIs may consider developing internal skills in data science and AI engineering, and establishing multidisciplinary teams that bring together auditors, IT specialists, and legal experts. In this context, experienced auditors represent an especially valuable resource: their domain knowledge, professional judgment, and deep understanding of audit practice are essential for shaping, training, and validating AI applications in ways that are relevant and reliable. This suggests that the effective integration of AI should be approached as an inclusive institutional effort, in which expertise from across generations and professional backgrounds is brought into the development process. Admittedly, rapid digital transformation over the last decade has introduced 'change fatigue', and AI may be perceived as 'another initiative' without clear immediate payoff, leading to passive resistance. However, it is up to

leaders to promote engaging with adoption horizontally across their organization. By sharing experiences at a deeper technical level and exchanging best practices, SAIs can accelerate collective progress while avoiding duplication of effort. The current moment presents a unique opportunity. AI technologies have matured to a point where they can deliver tangible benefits in public audit. At the same time, regulatory frameworks such as the European AI Act provide guidance on responsible use. The challenge now is to move from discussion to implementation.

Leveraging international collaboration through networks such as EUROSAI is definitely important to support this transition (see also **Box 1**). Fostering structured dialogue between key digitalisation frontrunners in SAIs and promoting concrete initiatives (e.g. it has been suggested to share best practice in prompt development in auditing) can help to ensure that AI becomes a standard component of modern audit practice.

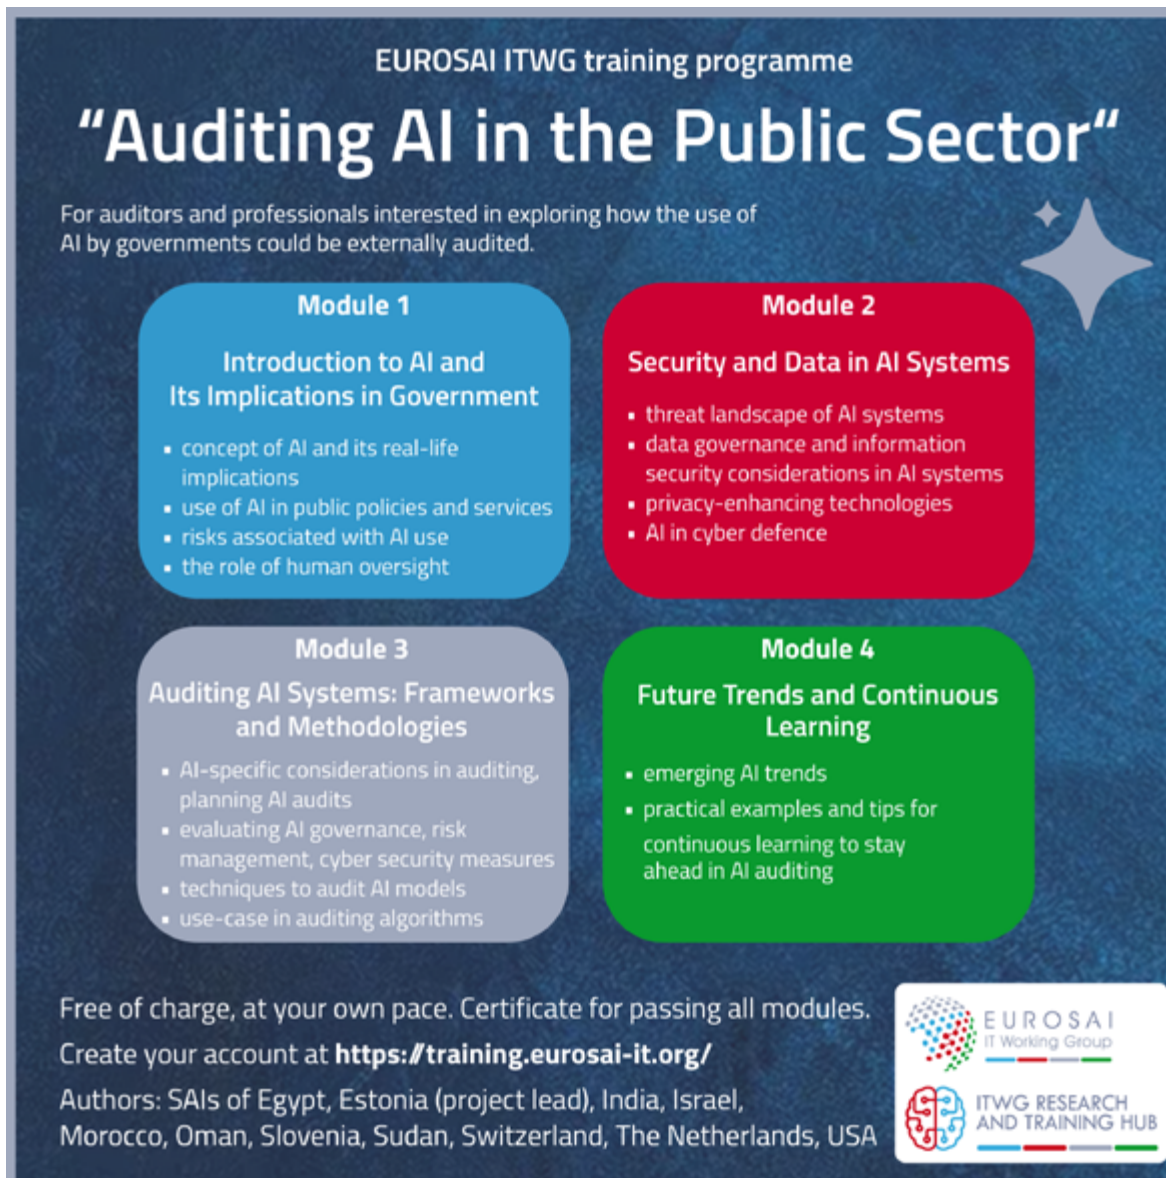
AI as tool for enhanced accountability and public trust

Artificial intelligence is not a distant prospect: it is an immediate tool with the potential to transform public auditing. However, realising this potential requires deliberate action. The paradox of high capability and low utilisation must be addressed through systematic integration, robust governance, and secure technological

solutions. By doing so, European SAIs can enhance both the efficiency and the effectiveness of their work – ultimately strengthening accountability and public trust. The path forward is clear and has been articulated well enough at dozens of meetings. The task now is to take concrete action.

Box 1 – EUROSAI ITGW e-learning course ‘Auditing AI in the Public Sector’

Besides promoting knowledge-sharing about the application of AI in supreme audit institutions, the EUROSAI ITGW training programme contains courses on auditing AI in government. At the end of March 2026, the EUROSAI ITGW launched the MOOC-based e-learning course ‘Auditing AI in the Public Sector’. Creating this has been an exceptional international cooperation project between 11 SAIs worldwide. See below for more details.



EUROSAI ITGW training programme

“Auditing AI in the Public Sector”

For auditors and professionals interested in exploring how the use of AI by governments could be externally audited.

Module 1

Introduction to AI and Its Implications in Government

- concept of AI and its real-life implications
- use of AI in public policies and services
- risks associated with AI use
- the role of human oversight

Module 2

Security and Data in AI Systems

- threat landscape of AI systems
- data governance and information security considerations in AI systems
- privacy-enhancing technologies
- AI in cyber defence

Module 3

Auditing AI Systems: Frameworks and Methodologies



- AI-specific considerations in auditing, planning AI audits
- evaluating AI governance, risk management, cyber security measures
- techniques to audit AI models
- use-case in auditing algorithms

Module 4

Future Trends and Continuous Learning


- emerging AI trends
- practical examples and tips for continuous learning to stay ahead in AI auditing

Free of charge, at your own pace. Certificate for passing all modules.
 Create your account at <https://training.eurosai-it.org/>
 Authors: SAIs of Egypt, Estonia (project lead), India, Israel, Morocco, Oman, Slovenia, Sudan, Switzerland, The Netherlands, USA

Governing AI while building it – accountability in the EU’s digital transformation

By Andreas Braun and Hazal Kantarci, PwC Luxembourg



© PwC Luxembourg

Professional services provided by private audit firms play a sensitive role in artificial intelligence (AI) adoption. Audit firms are among the more active adopters of AI solutions, but are also setting up and managing AI governance for their clients as part of the services they provide. In this article, Hazal Kantarci, Partner and leader of PwC Luxembourg’s AI Factory, and Andreas Braun, Managing Director and leader of AI Advisory for the public sector, discuss this dual role. They argue that this is already how things are operating in the private-sector audit, which is moving from disconnected experiments to production-grade use cases linked to real business pain points, while also setting up a suitable governance framework.

The builders are also the governors

What is holding back AI success? Recent evidence suggests that multiple organisations are still struggling to turn activity into measurable value. A Massachusetts Institute of Technology report entitled [The GenAI divide: State of AI in business 2025](#) highlighted that only a minority of generative AI initiatives by businesses were producing significant returns, despite substantial investment. PwC’s AI performance study, published in 2026, points in a similar direction. In a survey of 1 217 companies across

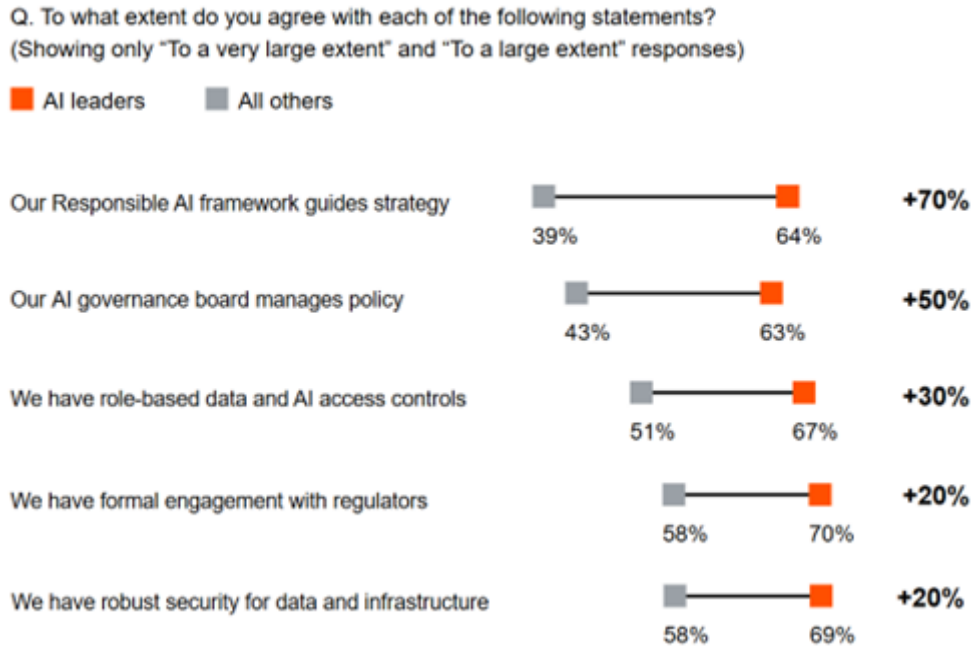
25 sectors, PwC found that 20% of companies captured 74% of AI-driven returns, with the most AI-fit companies performing 7.2 times better than the average, in financial terms (see **Figure 1**).

But why are they so far ahead? The organisations generating the greatest value from AI are those that have built the capabilities needed to translate experimentation into business outcomes. At PwC, we define AI fitness

as the ability to point the technology at what matters, build fit-for-purpose foundations and embed AI across an enterprise¹. AI leaders are not only doing more, but

also connecting technology, business and governance. They are building the operating model needed to convert activity into measurable outcomes.

Figure 1 – Governance and risk



Source: PwC's AI performance study.

This mirrors, in a sense, the findings of the European Court of Auditors in [special report 08/2024](#) on the EU's artificial intelligence ambition. The report does not suggest that Europe lacks talent or ideas. Rather, it points to the limited impact of several EU initiatives, insufficient coordination between the European Commission and member states, and governance and information exchange mechanisms that were not yet strong enough to support the EU's AI objectives.

Discussions on AI in Europe often focus on the lack of technology, investment and talent. We believe there is a need for broader discussions, encouraging also the capability to govern what we are building. For professional services firms, this means going beyond AI inventories and risk registers and looking more directly at the interaction between capability, cost and control.

Organisations like ours sit at the centre of this tension. We use the technology in a broad range of activities supporting clients, while at the same time helping clients assess their use cases and building governance around the technology. This dual role can create understandable concerns around independence, accountability and professional judgement – the core conditions of our audit and advisory work.

Nevertheless, we clearly see our practical experience with AI as an enabler for scrutinising and supporting our clients' use of that technology. A firm that builds and uses AI responsibly at production grade learns where risks emerge and how to set up the governance to address them. In practice we are both builders and governors, but how do we ensure trustworthy, disciplined use?

The reality of transformation and the moving goalposts

Understanding why both AI adoption and governance remain moving targets requires examining what AI deployment looks like in practice. At PwC Luxembourg, current solutions span generative AI assistants, knowledge retrieval tools, document review systems, workflow automation, and increasingly autonomous agents capable of executing multi-step tasks from pre-filling questionnaires and generating complex document-based output to navigating end-to-end multi-system

processes. Beyond eliminating repetitive work, these tools also increase quality and show lower error rates in our tests, compared to manual extraction.

Achieving the ability to measure this improvement and ensuring that results are explainable represent key aspects of responsible transformation. We have implemented a thorough testing benchmark that is applied to any tool before release and links any results to

¹ PwC, [Want ROI from AI? Go for growth: PwC's AI performance study](#), 2026, p.3.

their original sources, while also adopting measures to assess confidence in the outputs created.

Another core challenge of AI transformation is the technology's exponential development. Progress at this speed creates a moving goalpost. Increasingly, autonomous AI agents are becoming a reality, while analytical capabilities are also growing rapidly. Managing this requires revisiting your portfolio of ideas in light of what has become technologically feasible. But, more importantly, it requires having a governance framework that is agile enough to react to improved capabilities, by developing new measures of performance and new controls. This governance cannot be automated away: defining acceptable use, setting boundaries and exercising judgement remain inherently human responsibilities.

Two further critical aspects that we see in practice are cost and dependency. Deploying frontier AI requires access to external ecosystems, making pricing, resilience and autonomy central concerns. Not every use case justifies the most capable or expensive model; smaller or specialised options often provide a more efficient balance between performance, transparency and cost.

We, like many organisations, need to carefully balance

From principles to operational accountability

The language of responsible AI is becoming more familiar – from fairness and transparency to human oversight and accountability. But how are those translated into practice? How do you properly mark AI-generated content? How is human oversight effective if the operator needs to validate a thousand decisions taken by AI? Operationalising responsible AI is the real challenge. Below we focus on two key aspects: *visibility* and *measurement*.

Visibility is directly linked to accountability. We need to know at a minimum where and for what purpose we are using AI, to identify data concerns and the level of risk. Numerous organisations, ours included, grapple with the challenge of shadow AI. It is easy to blame staff not following guidance, but AI leaders recognise how increased workloads and having the right tools available are interlinked. This does not mean organisations should allow their staff to use any tool they like; they should allow AI use where it is effective and be able to successfully communicate a sound rationale when the use of certain systems is not allowed.

This leads us to *measurement*. How can we assess if AI is effective? The early indicators used tend towards

three dimensions: capability, cost and control. Capability expands what is possible, yet cost determines whether it scales, and control ensures accountability. These three dimensions are interdependent: scaling capability increases cost exposure, cost pressure constrains adoption, and focusing on cost optimisation alone can undermine both quality and oversight.

At PwC Luxembourg, these factors shape the way we approach AI deployment: not by optimising any single dimension, but by maintaining a workable balance among all three. Central to this is the recognition that maximum business value emerges at the intersection of AI capabilities and human judgement. This is consistent with the results of the PwC study referenced above, whereas the organisations benefiting most from AI are also those demonstrating the most mature governance, indicating that innovation and control evolve together rather than in opposition. Keeping professionals in the loop is not a limitation, but a deliberate design choice. Governance is not a constraint, but rather the mechanism that allows us to find the right trade-offs and responsibly increase the scope of AI use in our organisation, while maintaining trust and reliability – and thus make sure we are following the moving goalposts.

self-fulfilling prophecies. The target is often to use as many tokens as possible or to release as many tools as possible, even rewarding employees for high AI use. The predictable roll-back of this approach is already beginning². We argue that existing, even traditional, indicators remain appropriate. What is the quality of the output generated? What is the cost in time and resources? The way we have measured success over the last few decades remains appropriate in the age of AI, even though we need to consider the rising cost of AI use itself.

This has implications for AI governance and the way we perform assurance. Our current standards remain relevant but require adaptation, with more crucial processes being supported by AI. Visibility and measurement are essential in such decisions and need to be embedded in controls and evidence. Still, as any user of the technology can confirm, the creation of content can be overwhelming, and with agentic AI systems (see **Box 1**) it is becoming more challenging to trace. There is a risk that this may significantly increase the scope of assurance needed, if not managed appropriately within an AI governance framework that ensures auditability as a matter of principle.

² See for example : Gio Farley, 'Amazon Faces Tokenmaxxing Problem After AI Leaderboard Encourages Questionable AI Usage Metrics', in: [Tech Times, 30 May 2026](#), reporting on Business Insider coverage.

Box 1 - Example of an AI agent used at PwC Luxembourg

PwC Luxembourg has developed an AI agent within our client onboarding platform that automatically performs completeness checks and extracts information from documents uploaded by our clients. It then prepares a summary with the best next steps for our onboarding team, ensuring a managed human-in-the-loop process in which acceptance still requires human confirmation.

Is AI becoming the better referee?

Finally, we need to look at the human aspects of all this. The EU AI Act places strong emphasis on human oversight of AI systems but provides little detail on how to achieve this. We know from scientific research that this only works when we keep our cognitive limits in mind. Automation complacency and automation bias derive from the way attention is distributed under a heavy workload and are observed in experts as much as in novices³. This cannot be trained or practised away. Typical results include missing errors that the system does not flag or accepting incorrect recommendations despite evidence to the contrary. The effectiveness of oversight depends on design and operational conditions, with clear interfaces, management of alert fatigue, or even the introduction of friction in the process. Most importantly, however, it depends on giving the humans using AI systems enough time – a fact that is insufficiently appreciated by supporters of AI efficiency gains. Nevertheless, this is how

things stand ... unless AI becomes a better judge than humans.

This is a hotly debated topic and will remain so for the time being, given the exponential capability curve of AI. The effect of AI on audit quality is still considered mixed. A randomised field experiment with 100 professional auditors found that access to firm-specific AI improved documentation quality and reduced task time⁴. AI judges also deliver consistent results and traceable, reproducible rationales; they are also less exposed to certain human biases such as fatigue and being influenced by previous years' results. However, they introduce systematic biases of their own, including position bias (preference for information at the beginning or the end of input) and self-preference (rating their own outputs higher). Again, it seems that combining the strengths of both AI and humans can lead to improved outcomes, at least for now.

Looking forward: trust as European capability

The main question in our article is: *Can organisations govern a technology they are also building and using?* We argue that the answer is a clear yes, if AI adoption moves in lockstep with the disciplines required to control it. We should embrace the use of this technology, while ensuring responsible AI use through governance frameworks, recognised standards, traceable evidence, proportionate controls and human judgement.

The dual role of professional services firms will remain sensitive, and rightly so. However, withdrawing from AI is not the right answer; instead, capability building and discipline are required. Only those who understand the

technology can create the operational framework to deploy it responsibly, and only they will have the ability to monitor and audit it in the future

Likewise, Europe's AI ambition will be judged in multiple dimensions. Technology and innovation are an important aspect, while trust is another. The question will not only be how much the EU has enabled, but also whether it has done so in a way that is accountable and measurable. Only if we have the capability to use and monitor AI responsibly will it lead to positive societal change. This is where public auditors and professional service firms have a significant role to play.

³ Raja Parasuraman and Dietrich H. Manzey, 'Complacency and Bias in Human Use of Automation: An Attentional Integration', in: *Human Factors*, 2010, 52(3), 381–410, [DOI](#).

⁴ Markus Jezierski, Steffen Kaltenpoth, Oliver Müller and Barbara E. Weißenberger, *The Effect of Large Language Models on Audit Quality: Causal Evidence from a Randomized Field Experiment*, working paper, SSRN No 6095166 (2026)..

When the rulebook arrives before the road is built

A space engineer's perspective on AI regulation, sovereign infrastructure and the gap Europe must close

By Mustafa Elmoslhey, space engineer and entrepreneur, Luxembourg



© Mustafa Elmoslhey/AI generated

The general public is mostly aware of AI through its use of large language models such as ChatGPT and Mistral AI. However, the numerous AI developments for industrial use fly under most people's radar. What do industrial users in specific areas need in terms of AI tools, AI regulation and AI autonomy to ensure continued business operations? Mustafa Elmoslhey is a mechatronics engineer and senior space engineering executive based in Luxembourg, who served as R&D Engineering Manager at Amazon Robotics and subsequently as Head of Engineering and Managing Director of EnduroSat Luxembourg. He offers a practitioner's perspective on AI regulation and sovereign infrastructure from the vantage point of a space engineer working in Luxembourg, with a particular focus on the compliance asymmetry between EU-based deep-tech startups and their international competitors.

Navigating the EU's regulatory environment

There is a particular kind of cognitive dissonance that comes with building an AI system in Europe in 2026. On one side of your desk sits the most ambitious AI regulatory framework in the world the EU AI Act, fully in force since August 2024, with its tiered risk classifications, conformity assessments and transparency obligations. On

the other side sits your engineering notebook, full of the practical problems of making a machine-learning model work reliably on real-world satellite data, in real time, at a scale where a missed detection or a false positive has real operational consequences.

I have spent the last decade at this intersection – as R&D Engineering Manager at Amazon Robotics, where I began my journey into large-scale autonomous systems; as Head of Engineering and Managing Director at EnduroSat Luxembourg, where I led the building of engineering infrastructure from the ground up; and subsequently as a co-founder of several companies in the domain of AI and space technology. My observation from that vantage point is this: the EU has done something genuinely important by taking AI regulation seriously, and has simultaneously created a compliance architecture that is,

in its current form, better calibrated for regulating large technology platforms than for enabling the deep-tech startups that will determine whether Europe has strategic autonomy in AI in the first place.

This is not a complaint. I want to offer a practitioner's perspective that policy conversations about AI regulation rarely include – the view from someone who is simultaneously trying to build AI systems that work and navigate the regulatory environment within which they must operate.

What AI really means in the space and Earth-observation sector

Before discussing regulation, it is worth being precise about what artificial intelligence means in a domain like space-based Earth observation – because the word AI covers an enormous range of activities, not all of them equally well served by the current regulatory framework.

In the field of space-based Earth observation, AI does three distinct things. First, it processes enormous volumes of raw sensor data – satellite imagery, multispectral and hyperspectral measurements, synthetic aperture radar returns, atmospheric soundings – and extracts structured information from what would otherwise be an unmanageable noise floor. This is pattern recognition at scale, applied to physics. Second, it cross-validates multiple independent data sources against each other to detect anomalies and derive higher order intelligence. Consider a practical example from climate and environmental monitoring: an AI system ingesting thermal infrared satellite imagery alongside atmospheric composition measurements and ground-based sensor readings to detect and characterise industrial emissions in near-real time. The system must reconcile data from instruments with different resolutions, different revisit frequencies, and different calibration baselines – correlating a thermal anomaly in one dataset against a methane concentration spike in another and a wind direction model in a third, to generate a verified emission event with a confidence score. This is anomaly

detection applied to a global environmental dataset updated continuously from orbit. Third, and most consequentially, it issues a verified output – a machine-readable assessment of an observed phenomenon – that downstream users employ to make decisions with real operational stakes. Environmental regulators use it to identify industrial non-compliance. Climate scientists use it to calibrate emissions inventories. Agricultural agencies use it to assess crop stress and forecast yield at national scale. Disaster response coordinators use it to characterise the extent of flood or fire damage in the hours after an event, before ground teams can be deployed safely.

This third function is where the regulatory conversation becomes genuinely complex. The EU AI Act's risk-classification framework is largely designed around the question of direct harm to individuals – medical diagnosis, credit scoring, biometric identification of natural persons. Space-based Earth-observation AI sits in an uncomfortable space: the downstream decisions it informs can have significant consequences for communities, industries and governments, but the AI system itself is issuing a confidence-weighted observation about a physical phenomenon, not making a final regulatory or enforcement decision. Classifying this correctly under the AI Act's risk tiers requires legal analysis that no small engineering team should have to fund as a precondition to building a product.

The compliance asymmetry that keeps practitioners awake at night

Here is the competitive reality that European AI policymakers need to understand clearly: the companies operating in the global space AI market that European startups compete against are predominantly American or increasingly Chinese. They operate under no equivalent of the EU AI Act. They do not face conformity assessment requirements. They are not required to maintain the documentation that the Act demands. They are not required to disclose training data provenance. They are not subject to post-market monitoring obligations that require ongoing engineering resource to maintain.

This does not mean that non-EU competitors are irresponsible. Many operate to high technical standards.

But the compliance overhead – the legal review, the documentation, the conformity assessment, the post-market monitoring – falls entirely on the European side of the competitive equation. For a large technology company, this is a manageable cost of doing business. For a startup with a small engineering team and a finite runway, it is a material disadvantage that compounds at every development cycle.

I want to be precise about what I am not arguing: I am not arguing that AI should be unregulated. The risks that motivated the EU AI Act are real. Algorithmic bias, opaque decision systems operating in high-stakes environments, and the use of AI for mass surveillance are genuine

Navigating the EU's regulatory environment

There is a particular kind of cognitive dissonance that comes with building an AI system in Europe in 2026. On one side of your desk sits the most ambitious AI regulatory framework in the world the EU AI Act, fully in force since August 2024, with its tiered risk classifications, conformity assessments and transparency obligations. On the other side sits your engineering notebook, full of the practical problems of making a machine-learning model work reliably on real-world satellite data, in real time, at a scale where a missed detection or a false positive has real operational consequences.

I have spent the last decade at this intersection – as R&D Engineering Manager at Amazon Robotics, where I began my journey into large-scale autonomous systems; as Head of Engineering and Managing Director at EnduroSat Luxembourg, where I led the building of engineering

infrastructure from the ground up; and subsequently as a co-founder of several companies in the domain of AI and space technology. My observation from that vantage point is this: the EU has done something genuinely important by taking AI regulation seriously, and has simultaneously created a compliance architecture that is, in its current form, better calibrated for regulating large technology platforms than for enabling the deep-tech startups that will determine whether Europe has strategic autonomy in AI in the first place.

This is not a complaint. I want to offer a practitioner's perspective that policy conversations about AI regulation rarely include – the view from someone who is simultaneously trying to build AI systems that work and navigate the regulatory environment within which they must operate.

What AI really means in the space and Earth-observation sector

Before discussing regulation, it is worth being precise about what artificial intelligence means in a domain like space-based Earth observation – because the word AI covers an enormous range of activities, not all of them equally well served by the current regulatory framework.

In the field of space-based Earth observation, AI does three distinct things. First, it processes enormous volumes of raw sensor data – satellite imagery, multispectral and hyperspectral measurements, synthetic aperture radar returns, atmospheric soundings – and extracts structured information from what would otherwise be an unmanageable noise floor. This is pattern recognition at scale, applied to physics. Second, it cross-validates multiple independent data sources against each other to detect anomalies and derive higher order intelligence. Consider a practical example from climate and environmental monitoring: an AI system ingesting thermal infrared satellite imagery alongside atmospheric composition measurements and ground-based sensor readings to detect and characterise industrial emissions in near-real time. The system must reconcile data from instruments with different resolutions, different revisit frequencies, and different calibration baselines – correlating a thermal anomaly in one dataset against a methane concentration spike in another and a wind direction model in a third, to generate a verified emission event with a confidence score. This is anomaly

detection applied to a global environmental dataset updated continuously from orbit. Third, and most consequentially, it issues a verified output – a machine-readable assessment of an observed phenomenon – that downstream users employ to make decisions with real operational stakes. Environmental regulators use it to identify industrial non-compliance. Climate scientists use it to calibrate emissions inventories. Agricultural agencies use it to assess crop stress and forecast yield at national scale. Disaster response coordinators use it to characterise the extent of flood or fire damage in the hours after an event, before ground teams can be deployed safely.

This third function is where the regulatory conversation becomes genuinely complex. The EU AI Act's risk-classification framework is largely designed around the question of direct harm to individuals – medical diagnosis, credit scoring, biometric identification of natural persons. Space-based Earth-observation AI sits in an uncomfortable space: the downstream decisions it informs can have significant consequences for communities, industries and governments, but the AI system itself is issuing a confidence-weighted observation about a physical phenomenon, not making a final regulatory or enforcement decision. Classifying this correctly under the AI Act's risk tiers requires legal analysis that no small engineering team should have to fund as a precondition to building a product.

The compliance asymmetry that keeps practitioners awake at night

Here is the competitive reality that European AI policymakers need to understand clearly: the companies

operating in the global space AI market that European startups compete against are predominantly American or

¹ The US CLOUD Act primarily amends the Stored Communications Act (SCA) of 1986 to allow federal law enforcement to compel US-based technology companies via warrant or subpoena to provide requested data stored on servers regardless of whether the data are stored in the US or on foreign soil.

increasingly Chinese. They operate under no equivalent of the EU AI Act. They do not face conformity assessment requirements. They are not required to maintain the documentation that the Act demands. They are not required to disclose training data provenance. They are not subject to post-market monitoring obligations that require ongoing engineering resource to maintain.

This does not mean that non-EU competitors are irresponsible. Many operate to high technical standards. But the compliance overhead – the legal review, the documentation, the conformity assessment, the post-market monitoring – falls entirely on the European side of the competitive equation. For a large technology company, this is a manageable cost of doing business. For a startup with a small engineering team and a finite runway, it is a material disadvantage that compounds at every development cycle.

I want to be precise about what I am not arguing: I am not arguing that AI should be unregulated. The risks that motivated the EU AI Act are real. Algorithmic bias, opaque decision systems operating in high-stakes environments, and the use of AI for mass surveillance are genuine problems that deserve regulatory attention. The EU is

right to take them seriously, and right to distinguish itself from regulatory environments that have taken a more permissive approach.

What I am arguing is that the current implementation creates a compliance burden that is not well calibrated to the size, stage and risk profile of the entity being regulated. A European startup building Earth-observation AI tools and a hyperscale platform provider deploying analytical systems to government intelligence agencies are not in the same category of risk. They should not face the same compliance architecture.

The EU AI Act's SME provisions acknowledge this to some extent – reduced obligations for small and medium-sized enterprises, free access to regulatory sandboxes. In practice, however, sandbox access requires knowing where to apply, having the legal and administrative capacity to engage with regulatory bodies, and operating in a jurisdiction where the competent authority has the bandwidth to process applications. For a founder who is simultaneously hiring engineers, closing a funding round, managing supplier relationships and building an actual product, the sandbox is a theoretical benefit that is difficult to operationalise

Sovereign infrastructure: where the EU has done something genuinely right

I would like to balance the regulatory critique with an acknowledgement of what the EU has built, on which European deep tech genuinely relies – and which I believe represents a durable competitive advantage if the policy environment enables its full use.

Luxembourg's [MeluXina supercomputer](#) – 18 PetaFLOPS, GPU-accelerated, ranked in the EuroHPC Top 100 – is not a theoretical resource. It is where European AI models in space applications are trained at scale. The access conditions, pricing structure and institutional support make it genuinely useful for startups, not just large research institutions. The quantum-classical hybrid infrastructure coming online across the EuroHPC network in 2026 extends that capability into research that would otherwise require dependencies on US cloud providers operating under the US CLOUD Act¹.

SEAL-3-certified sovereign cloud infrastructure – exemplified by Post Telecom Luxembourg's DEEP by POST platform, recently awarded a major European Commission contract – means that when an AI system processes classified or sensitive Earth-observation data, operators can guarantee that data never touches infrastructure that can be compelled by a non-EU legal order. This is not a marketing claim. It is an architectural fact that determines which government and institutional contracts European

companies can credibly bid for and which they cannot. The equivalent sovereign guarantee is structurally unavailable to competitors using US-headquartered hyperscale cloud providers, regardless of the physical location of their data centres.

The EU Space Programme infrastructure – Copernicus Earth observation, Galileo navigation, IRIS2 secure connectivity – represents a genuine strategic commitment to sovereign space capability with no real parallel outside the US and China. For companies building space-based AI systems in Europe, this matters structurally. The EU has built an institutional environment for sovereign space intelligence on which the private sector can build.

The problem is that these infrastructure investments and the regulatory compliance architecture are not currently designed to work together coherently for the practitioner who needs both. The sovereign compute infrastructure sends one signal: we want you to build AI here. The AI Act compliance architecture sends another: before you build, navigate this. Bridging that gap – making the regulatory sandbox genuinely accessible to the startups that are actually using MeluXina and building on Copernicus data – is the most important near-term improvement the Commission could make.

What AI has changed in engineering practice

Coming from a background in hardware and robotics engineering, I observed a fundamental shift in how engineering decisions are made when AI is a core component of the system. In traditional engineering, you specify a system's behaviour in advance and verify that the implementation matches the specification. In AI-driven systems, you specify a desired outcome and train a system to achieve it – the intermediate behaviour is emergent, not specified. This is a profound difference, and it is one that current audit and oversight frameworks, including those being developed by European supreme audit institutions, are still learning to address.

The EU AI Act's requirement for technical documentation and post-market monitoring is a reasonable attempt to address this difference – to create an audit trail for systems whose behaviour cannot be fully specified in advance. But it assumes a level of internal documentation infrastructure that most startups do not have at the point

when they are building their first production system. The obligation arrives before the engineering organisation has matured to the point where maintaining it is natural practice rather than compliance overhead.

A practical recommendation that I would make, drawn from experience building AI systems in both the US and European regulatory environments, is: stage compliance obligations to match the development maturity of the system and the organisation. A startup in its first two years of AI product development should face lighter-touch obligations – disclosure of purpose, registration, basic technical documentation – with more rigorous conformity assessment triggered by commercial deployment at scale or deployment in genuinely high-risk contexts. This is how good engineering quality systems in fact work: you apply the level of rigour appropriate to the risk and the stage of development, and you scale it as you scale the system.

The network that Europe needs but does not yet have

One final observation that deserves direct attention: the gap between the EU's AI research ecosystem and its AI deployment ecosystem. Europe produces world-class AI research. The University of Luxembourg's Interdisciplinary Centre for Security, Reliability and Trust ([SnT](#)) produces algorithms and methodologies directly deployable in operational AI systems. [Imec](#) in Belgium, [Inria](#) in France and the [Fraunhofer Institutes](#) in Germany, for example, produce research output that is genuinely excellent. What is missing is the institutional connective tissue that moves research outputs into commercial products at pace.

In the US, the [Small Business Innovation Research \(SBIR\)](#) and [Small Business Technology Transfer \(STTR\)](#) programmes create direct financial pathways from federal agencies to startups commercialising research. In Israel,

the Innovation Authority creates co-investment structures that de-risk early-stage deep tech ventures. The EU's equivalents – Horizon Europe and the EIC Accelerator – are slower, more document-intensive and less calibrated to the pace at which AI development moves. By the time a Horizon Europe project completes its Phase 1 reporting obligations, a US or Israeli competitor has shipped three product iterations.

This is not an argument against accountability in public funding. It is an argument for designing accountability mechanisms that match the cadence of the technology being funded. AI does not develop on a three-year project cycle; nor should the funding mechanisms that support AI development in Europe.

The choice Europe is actually making

To conclude: Europe is making a genuine choice about what kind of AI it wants to build and what kind of rules it wants to apply to govern AI globally. The EU AI Act is an attempt to answer that question in a way that prioritises human rights, transparency and accountability over speed to market. That is a legitimate and important set of priorities.

The risk is that in pursuing those priorities, Europe creates an environment where the AI systems that matter most – those embedded in critical infrastructure, sovereign intelligence, environmental monitoring and space systems – are built elsewhere, by entities operating under fewer constraints, and then deployed in Europe anyway. Regulation without capability is not safety, but dependency.

The EU has the sovereign infrastructure – EuroHPC, Copernicus, IRIS2, world-class research networks, and committed national programmes in Luxembourg, France, Germany, the Netherlands and beyond. It has the political will to take AI seriously. What it needs now is a regulatory implementation that is as ambitious as its infrastructure investment – one that enables the startups and engineering teams building on that infrastructure to move at the pace that genuine strategic autonomy truly requires.

Europe will not win the AI race by regulating better than everyone else. It will win by building better than everyone else, in a regulatory environment that makes building here the natural choice rather than the compliant one.

AI for access to justice – an important opportunity

By Assistant Professor Hannes Westermann, Maastricht University



© denisismagilov/depositphotos.com

With the advent of generative AI systems such as ChatGPT, the reasoning ability of computers in relation to texts has exploded, opening up many new possibilities. One major positive impact this could have is helping billions of people to resolve legal issues, in other words, giving them *access to justice*. Below, Hannes Westermann, Assistant Professor at the Law and Tech Lab in Maastricht University’s Faculty of Law, gives insights into his research and actions exploring how to overcome these barriers with AI and discuss some implications of this research.

Access to justice and generative AI

Most people face legal issues at some point in their lives, such as problems with their housing, neighbourhoods or employment¹. When they do, they have various legal rights, such as entitlement to compensation or other remedies.

Unfortunately, several barriers often prevent individuals from asserting these rights. Even if people are aware of the rules, they may still struggle to understand how they

apply to their particular situation, or to fill out the forms or documents necessary to take legal action². People may also struggle to navigate court proceedings, which are often confusing, expensive and time-consuming³.

Overall, access to justice is a huge issue in society, often preventing individuals from obtaining a fair and speedy resolution to their legal problems. Estimates show that 1.5 billion people are unable to resolve their legal problems

¹ Ab Currie, The legal problems of everyday life, in Rebecca L Sandefur, ed, *Sociology of Crime, Law and Deviance*, Emerald Group Publishing Limited, 2009, p. 1; Laura Savage & Susan McDonald, Experiences of serious problems or disputes in the Canadian provinces, 2021, in: *Juristat: Canadian Centre for Justice Statistics* 1–28, 2022.

² Julie Macfarlane, *The National Self-Represented Litigants Project: Identifying and Meeting the Needs of Self-represented Litigants: Final Report*, National Self-Represented Litigants Project, 2013, at 9 Google-Books-ID: kynloAEACAAJ.

³ Noel Semple, *The cost of seeking civil justice in Canada*, 2015, 93 Can B Rev 639.

at any given time⁴. Increasing access to justice forms part of goal 16⁵ of the '2030 Agenda for Sustainable Development'.

Recently, generative AI has emerged as a technology with significant uses in the field of law. Unlike previous AI systems, 'large language models' (LLMs) can carry out many sophisticated text-based tasks across various domains, by generating texts based on a prompt⁶. At the

same time, they still face some important challenges, such as their propensity to sometimes 'hallucinate' (i.e. make up facts)⁷, limitations in addressing certain kinds of problems⁸, and concerns regarding the environment and authorship attribution⁹.

My research investigates whether and how such models can be used to address the important societal issue of access to justice.

Examples of AI for projects on access to justice

Let us explore some of the research projects I have worked on to explore the integration of generative AI and access to justice. To make the projects more concrete, let us imagine the situation of a tenant whose heating has stopped working, and who is not sure what to do.

Legal information

One way AI can support users is by providing them with legal information to help them understand their options. One project I worked on that does precisely this is the [JusticeBot](#), developed at the Cyberjustice Laboratory in Montreal in collaboration with the Quebec housing tribunal. Focusing on landlord-tenant disputes in Quebec, this system asks users a number of questions and then provides them with legal information¹⁰. In our example, the tenant would be asked questions to determine what had happened, then given a list of their rights in this situation (such as the right to compensation), as well as a list of steps they could take to seek to resolve their situation (such as writing a letter to the landlord or speaking to a lawyer). Since its launch in 2021, the system has been used over 50,000 times.

The JusticeBot is based on 'expert systems', a form of AI that is predictable and does not rely on generative AI. The advantage of this approach is that the answers will be consistent and can be verified by legal experts. However, generative AI can be used to enable the more efficient building of such tools for new legal areas¹¹.

Legal drafting

Sometimes, users may need to create documents or fill out forms to move their cases forward. For example, they may want to write a letter to the other party stating their claim or fill out a form with relevant details to bring a claim to court. However, understanding the relevant legal rules, selecting the relevant facts, and framing them in a way that corresponds to the rules can be challenging for laypeople¹². Generative AI has significant potential to assist with such tasks.

I explored this application as part of the 'Document Automation, Large Language Model Assisted' (Dallma) project¹³. This is a tool that can integrate expert knowledge and rules with LLM-drafted sections, allowing the automated creation of legal documents. The system will then ask the user questions and incorporate their answers into a document. To return to our example, let's imagine the tenant wanted to send a letter to their landlord. They would be asked questions such as when they first noticed the issue, how they communicated it to their landlord and how the landlord responded. Their answers would then be analysed, reformulated and incorporated into a legally accurate letter.

This system still requires users to describe their situation. In more recent research, however, we investigated whether LLMs could extract information directly from pictures of documents¹⁴ or even from videos taken by users¹⁵, with promising results. This would allow the

⁴ *Justice For All - Final Report*, New York: Center on International Cooperation: The Task Force on Justice, 2019. Online: <https://www.hiil.org/news/justice-for-all-a-report-by-the-task-force-on-justice/>.

⁵ *Transforming our world: the 2030 Agenda for Sustainable Development*, A/RES/70/1 (United Nations General Assembly) at 25 online: <https://documents-dds-ny.un.org/doc/UNDOC/GEN/N15/291/89/PDF/N1529189.pdf?OpenElement>.

⁶ Tom Brown et al, Language models are few-shot learners, in: *Advances in neural information processing systems* 33, *NeurIPS* 2020, 1877–1901.

⁷ Matthew Dahl et al, Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models, in: 16:1 *Journal of Legal Analysis* 64–93, 2024.

⁸ Parshin Shojaee et al, *The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity*, 7 June 2025, online: [arXiv.org](https://arxiv.org).

⁹ Emily M Bender et al, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, New York, NY, 2021, USA: Association for Computing Machinery.

¹⁰ Hannes Westermann & Karim Benyekhlef, *JusticeBot: A Methodology for Building Augmented Intelligence Tools for Laypeople to Increase Access to Justice*, Braga Portugal: ACM, 2023.

¹¹ Samyar Janatian et al, From Text to Structure: Using Large Language Models to Support the Development of Legal Expert Systems in: *Legal Knowledge and Information Systems Frontiers in Artificial Intelligence and Applications*, Volume 379, 2023, 167–176.

¹² Macfarlane, supra note 2; Karl Branting et al, Judges Are from Mars, Pro Se Litigants Are from Venus: Predicting Decisions from Lay Text, in: *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2020.

¹³ Hannes Westermann, *Dallma: Semi-Structured Legal Reasoning and Drafting with Large Language Models*, 2nd Workshop on Generative AI and Law, collocated with the International Conference on Machine Learning, 2024. Online: https://blog.genlaw.org/pdfs/genlaw_icml2024/58.pdf.

¹⁴ Hannes Westermann & Jaromir Savelka, *Analyzing Images of Legal Documents: Toward Multi-Modal LLMs for Access to Justice*, 16 December 2024. Online: <https://arxiv.org/abs/2412.15260>.

¹⁵ Lyra Hoeben-Kuil et al, Can LLMs Create Legally Relevant Summaries and Analyses of Videos?, in: *Legal Knowledge and Information Systems*, IOS Press, 2025, 98 DOI: 10.3233/FAIA251580.

tenant in our example to simply take a video of their home and be provided with a drafted legal letter.

Dispute resolution

If the letter is not sufficient to resolve the dispute, the tenant may want to enter into dispute resolution to resolve their problem. I also explored whether generative AI could help with this step, as part of the 'LLMediator' project¹⁶. Here, two parties are presented with an interface to discuss their dispute and try to find a solution. This interface is enhanced in various ways by generative AI.

Firstly, if the system detects that emotions are running high (e.g. users insulting each other), it can intervene and gently suggest a more positive way for users to convey

their message, thereby keeping the discussion on the right track.

Secondly, if the tenant and landlord are unable to come to a resolution themselves, a mediator may intervene in the discussion to help the parties in finding a solution. The LLM can be used here to provide a draft intervention message for the mediator. The mediator can then send or adjust this message.

Eventually, the system may support even more automated forms of mediation. While more research is required to demonstrate the safety of such an approach, initial results show that the LLM's ability to decide how to mediate disputes compares favourably even with some human-drafted messages¹⁷.

AI contributing to trust in the legal system

Above, I described a number of systems that could intervene in the dispute resolution process in various ways. As we can see, today's LLMs are getting to a point where they have many potential uses, such as providing legal information, drafting documents and helping to resolve disputes. Even now, laypeople are likely already using ChatGPT and other LLMs in practice to accomplish these legal tasks and more.

Such tools have tremendous potential to make the legal system more accessible. By making it easier to ask questions and get legal responses, they open the door for people who may not be able to afford a lawyer to interact with the legal system. Some people in this situation indicate that ChatGPT has helped them better understand and present their cases¹⁸.

At the same time, LLMs still have some limitations. For example, they sometimes hallucinate references to case law, which can lead to expensive fines if submitted to court¹⁹. Outside of the courtroom, hallucinated answers may induce users to do things that harm their case. Other potential issues include the lack of ability to explain

responses, issues with data privacy and environmental impact. Educating users of such models about their limitations and how to use them effectively is an important element of ensuring their safe and responsible use.

I am enthusiastic about uses of LLMs that support lawyers and legal aid clinics in accomplishing tasks more efficiently, as this would allow them to serve more of the population. Here, legal experts could verify the output of draft forms or other documents, lowering the risks of AI use. Information provided by LLMs could also help laypeople realise that their situation is a legal matter, and that they may want to speak to a lawyer.

Generative AI has the potential to create a legal system that is more approachable, inclusive and better able to serve people across various groups, which could increase trust in public institutions and social cohesion. I hope that my research exploring whether and how AI can and should be used for access to justice can contribute to this development, while ensuring that the focus remains on the needs of the humans using it.

¹⁶ Hannes Westermann, Jaromir Savelka & Karim Benyekhlef, *LLMediator: GPT-4 Assisted Online Dispute Resolution*, 2023, 3435 Proceedings of the ICAIL 2023 Workshop on Artificial Intelligence for Access to Justice (CEUR Workshop Proceedings). Online: <https://ceur-ws.org/Vol-3435/#paper1>.

¹⁷ Jinzhe Tan et al, *Robots in the Middle: Evaluating LLMs in Dispute Resolution*, in: *Legal Knowledge and Information Systems Legal Knowledge and Information Systems*, Volume 395, 2024, 168–179..

¹⁸ Jared Perlo & Angela Yang, *More people are using ChatGPT like a lawyer in court. Some are starting to win*, in: <https://www.nbcnews.com/tech/innovation/ai-chatgpt-court-law-legal-lawyer-self-represent-pro-se-attorney-rcna-230401>.

¹⁹ *Ibid*; Daniel Charlotin, *AI Hallucination Cases*, online: <https://www.damiencharlotin.com/hallucinations/>.

'Safeguarding technology is a shared responsibility that begins with those who develop and deploy it.'

Interview with Katarina Wallin Bureau,
Microsoft

By Gaston Moonen



© Katarina Wallin Bureau

The EU is one of the largest services markets in the world, with over 450 million potential customers. They mainly use digital – including AI – services provided by American Big Tech companies: over 70% of cloud services and around 80% of enterprise software come from these companies, which also raises various dependency concerns. With the EU taking on a trailblazer role globally when it comes to regulating digital services, we approached several American Big Tech companies to get their views on EU regulation in this policy area: what it means for them, how they deal with regional differences and what they see as the main challenges facing this EU governance framework. It turned out to be highly time-consuming to even make contact with an American Big Tech company, let alone obtain its views. However, Katarina Wallin Bureau, General Manager for Strategic Relations and European Government Affairs at Microsoft's Brussels office, was willing to answer some questions and provide insight into Microsoft's general views on matters including the European Commission's Digital Omnibus proposals and the corporation's ability to guarantee the confidentiality of its services vis-à-vis other jurisdictions.

Regulation as an enabler of innovation

What are the main advantages and disadvantages of EU regulations for Microsoft? What challenges could these regulations pose for Microsoft when deploying AI tools within the EU, and what benefits might they offer the company overall? In this context, what is your reaction to the Digital Omnibus Regulation proposal?

Katarina Wallin Bureau: The EU's regulatory framework plays an important role in building trust and creating a stable, rules-based environment for digital technologies. This is particularly valuable in areas such as artificial intelligence,

...regulation can be an important enabler of innovation, helping to provide clarity and set high standards globally.

”

where confidence from users, businesses and public institutions is essential for broader adoption. In that sense, regulation can be an important enabler of innovation, helping to provide clarity and set high standards globally.

At the same time, the European framework is ambitious and evolving, and its breadth can naturally bring a degree of complexity in implementation. Different legislative instruments, such as the AI Act, the Digital Services Act and broader data rules, address important policy objectives; ensuring alignment between them will be key as they are put into practice.

...we are supportive of the European Commission's Omnibus proposals to simplify EU rules in the digital domain...

”

This is one of the reasons why we are supportive of the European Commission's Omnibus proposals to simplify EU rules in the digital domain, part of which has recently been agreed and part of which continues to be under discussion among co-legislators. These initiatives provide an important opportunity to reduce unnecessary compliance burdens and legal uncertainty, while preserving the core objectives of the relevant regulatory frameworks. This is particularly significant in enabling businesses, including startups and SMEs, to confidently invest in and responsibly scale AI technologies across Europe.

At the same time, it will be crucial to maintain momentum in the legislative process. There is broad recognition at the political level of the need for simplification, as well as of the importance of translating this ambition into practice; operational outcomes will be key. Ensuring that the final framework is effective in delivering simplification, while maintaining robust safeguards, would represent a meaningful step forward.

Which safeguards has Microsoft put in place to block practices that can harm people or democratic processes – can you give a concrete example? Does Microsoft see this as its responsibility, or rather as the responsibility of others such as oversight authorities?

Beyond AI, Microsoft also invests extensively in cybersecurity (...) help protect democratic institutions from cyber-enabled interference and other malicious activities.

”

Katarina Wallin Bureau: Microsoft has put in place a range of safeguards to help prevent misuse of its technologies in ways that could harm individuals, organisations or democratic processes. These safeguards combine technical measures, internal governance and ongoing monitoring, and are integrated across the company's products and services.

One concrete example is Microsoft's approach to AI safety. We deploy content safety systems designed to detect and mitigate harmful outputs – including content that may facilitate illegal activity and promote violence – or designed to contribute to overall information integrity. These systems combine automated safeguards, continuous testing and human oversight to identify risks and respond to emerging threats. Beyond AI, Microsoft also invests extensively in cybersecurity and works with governments, civil society and industry partners to share threat intelligence and help protect democratic institutions from cyber-enabled interference and other malicious activities. Every day, we process over 100 trillion security signals and rely on more than 34,000 dedicated security professionals and 15,000 partners worldwide to detect, prevent and disrupt cyber threats on a global scale.

Microsoft has maintained a presence in Ukraine for more than 20 years and, since the start of the full-scale invasion in February 2022, has become a trusted cybersecurity partner. Our support spans three workstreams: we helped migrate Ukrainian government data from on-premises systems to secure cloud infrastructure in Europe; we supported efforts to detect and mitigate cyberattacks affecting Ukraine and the broader region; and we provided technical capabilities to document war crimes against civilians.

The private sector plays a critical role in embedding safeguards...

”

Microsoft believes that safeguarding technology is a shared responsibility that begins with those who develop and deploy it. The private sector plays a critical role in embedding safeguards early in the design and development process, continuously assessing and mitigating risks throughout the product lifecycle, and ensuring that innovation is pursued responsibly. Companies are often best positioned to anticipate how technologies may be used or misused

Effective governance therefore depends on both responsible corporate action and robust public institutions.

”

and to implement technical and operational measures that reduce potential harm.

At the same time, governments and oversight authorities have an essential role in establishing clear legal frameworks, setting societal expectations, providing independent oversight and ensuring accountability across the market. Effective governance therefore depends on both responsible corporate action and robust public institutions.

Deploying the same technology while adapting to different legal contexts

Could you provide a specific example illustrating the fundamental differences between Microsoft AI tools operating in the United States and those operating in the European Union?

Our AI tools are developed based on company-wide responsible AI principles and a global governance framework...

”

Katarina Wallin Bureau: Microsoft's AI tools are built on a common global technological foundation. Rather than considering them in a context of fundamental differences, a more meaningful distinction lies in the regulatory and governance context in which our AI tools are deployed. Our AI tools are developed based on company-wide responsible AI principles and a global governance framework, while also adapting implementation, documentation and compliance processes to meet local legal requirements, including in Europe. A concrete example is transparency and compliance under the EU AI Act, where certain AI systems are subject to specific legal obligations relating to transparency, risk management, documentation and, in some cases, restrictions on prohibited practices.

...the same technology is being deployed in different legal and societal contexts.

”

Microsoft has publicly explained that we are updating our contracts, internal policies and product governance processes to align with these requirements. In the United States, by contrast, there is currently no directly comparable single, economy-wide federal AI law with the same horizontal structure, so the compliance environment is less uniform and more sector- or state-specific. So, if one compares a Microsoft AI tool offered in both markets, the core model or service may be very similar, but the surrounding obligations can differ. This is not necessarily evidence of a fundamentally different technology; rather, it reflects the fact that the same technology is being deployed in different legal and societal contexts.

Microsoft has agreements with EU governments and institutions to provide services while maintaining independence and confidentiality. How can Microsoft guarantee it will not disclose data or sensitive information to authorities with jurisdiction over its core business?

...our services meet the high standards expected by European governments and institutions.

”

Katarina Wallin Bureau: Protecting the confidentiality of customer data is fundamental to the trust that governments, public institutions and businesses place in Microsoft. This is why we work closely with customers, policymakers and regulators to ensure that our services meet the high standards expected by European governments and institutions.

(...) we have made our [Digital Resilience Commitment](#) legally binding in contracts with European national governments and the European Commission...

In this context, we have made our Digital Resilience Commitment legally binding in contracts with European national governments and the European Commission, including a commitment to promptly and vigorously contest in court any order by any government to suspend or cease cloud operations in Europe. We are also committed to continuity measures, including expanded partnerships with European cloud partners that can support our customers' operational continuity in extreme scenarios. In order to reinforce this approach, we launched a European resiliency partnership with Delos Cloud to safeguard business continuity in Europe in times of crisis. This work also supports closer cooperation among Europe's sovereign cloud providers, including crisis response coordination and continuity options designed to help customers maintain operations even in the event of geopolitical disruptions.

These commitments are reinforced by technical safeguards, including customer-controlled encryption, in-region data processing and technologies such as confidential computing, which keep data protected while it is being processed. Our cloud solutions can operate across connected, hybrid and even fully disconnected environments, with consistent governance and local controls.

Going beyond sales – promoting AI literacy

In what ways does Microsoft engage with educational institutions, such as schools and universities, to equip the future workforce for an AI-driven workplace?

...AI capability is built not only through technical training, but also through practical use in teaching and learning.

”

Katarina Wallin Bureau: Microsoft engages with schools, universities, public institutions and nonprofit organisations across Europe to help prepare learners and workers for an AI-driven economy. This includes large-scale skilling commitments, partnerships with higher education institutions and targeted support for educators, so that AI capability is built not only through technical training, but also through practical use in teaching and learning.

For educators specifically, we launched Microsoft Elevate for Educators earlier this year. This is a global training programme that helps teachers learn the basics of AI and how to use it safely in the classroom by offering practical tools and a network to share what works. It comes with an official certificate, developed with education experts, that is aligned with the EU-OECD AI Literacy Framework, to help educators build the knowledge and practices needed to use AI in education.

On the ground, we are working directly with local education and training systems. For example, in the Netherlands we support a locally adapted AI curriculum that connects teacher training with student learning, while in Spain AI teacher training is built right into regional government school platforms. In Poland, a national programme managed to reach over 1.000 schools in just a few months, bringing AI skills to both general and vocational classrooms. We are also connecting these school programmes with the broader workforce, for example by running apprenticeship initiatives in Germany and teaming up with public employment services in France. Ultimately, it is about working with governments, schools and local partners to make AI literacy a normal, accessible part of how people learn and prepare for work.

...Microsoft's approach increasingly emphasises helping (...) how to use it [AI] critically, safely and responsibly.

”

Across these efforts, AI literacy is becoming just as important as technical proficiency. Microsoft's approach increasingly emphasises helping learners, educators and workers understand how AI works, where its limits and risks lie, and how to use it critically, safely and responsibly. This matters not only for employability, but also for ensuring that people can participate in an economy and society where AI is becoming part of everyday life.

Beyond direct skilling programmes, Microsoft also invests in research and policy insights through the AI Economy Institute, which helps analyse how AI is reshaping jobs and skills and can help inform how workforce policies need to adapt.

The curse of Big Tech: how digital giants are eroding the constitutional foundations of liberal democracy

By Prof. [Reijer Passchier](#), Professor of Digitalisation and Constitutional Democracy, Open University of the Netherlands, and Assistant Professor of Constitutional Law, Leiden University.¹



© Lumos Studio/stock.adobe.com

While acknowledging the benefits of Big Tech, European governments and citizens are increasingly aware of the risks this techno-economic model poses to the fundamental European values of privacy, electoral integrity, and sovereignty. How resilient must the EU be to protect its constitutional identity? And how can it restore what has already been lost? Professor Reijer Passchier, a specialist in digitalisation and constitutional democracy, examines the EU's current digital predicament and outlines actionable alternatives to reclaim independence from Big Tech companies and the governments that support them.

Impact of Big Tech on the EU's foundational pillars

The digitalisation of society, accelerated by the rise of artificial intelligence (AI), has spawned unprecedentedly

large, wealthy and powerful Big Tech conglomerates in the United States and China². Examples include Microsoft,

¹ This article goes back to a working paper, 'Big Tech en de weerbaarheid van de democratische rechtsstaat' ('Big Tech and the resilience of constitutional democracy'), prepared for a public hearing of the Dutch Senate on the resilience of democracy and the rule of law in the time of digital transformation on May 26, 2025. It is based on my book *De vloek van Big Tech: de juridisch-technologische wortels van constitutioneel verval en digitaal feodalisme* (Boom 2024), which I am currently translating and updating under the working title *The Curse of Big Tech: The Legal-Technological Origins of Constitutional Decay and Re-Feudalisation*.

² See e.g. <https://companiesmarketcap.com/>

Amazon, and ByteDance. While these firms have delivered convenience, increased efficiency in certain sectors, and offered vast amounts of affordable entertainment, they simultaneously pose a significant, perhaps even existential, threat to liberal democracy—the cornerstone normative ideal of Western states and alliances like the European Union (EU)³ and the Council of Europe⁴.

The stakes extend far beyond specific liberal-democratic values like privacy, copyright, or free markets. Increasingly, Big Tech also challenges the deeper and more general constitutional structures that safeguard specific values and make their existence possible in the first place. These foundational pillars include the sovereignty of the state (or the EU as a supranational entity), the distinction

between public and private spheres, the rule of law, political and economic citizenship, a vibrant civil society, and the very concept of constitutionalism (limited government) itself.

Call it the ‘Curse of Big Tech’, as I have referred to this problem elsewhere⁵, taking inspiration from the early 20th-century American Supreme Court Justice Louis D. Brandeis. As Brandeis wrote regarding the rise of large and powerful industrial conglomerates of his time: ‘Far more serious than the suppression of competition [or privacy, we may add] is the suppression of industrial freedom, indeed of manhood itself’. As digitalisation and the rise of Big Tech continue to progress, I would argue that this warning is now more urgent than ever.

The lobbying imbalance

To begin with, consider the constitutional implications of the massive, disproportionate lobbying exerted by Big Tech. In Brussels alone, these companies are estimated to spend over €150 million annually on lobbying—a figure that dwarfs the efforts of any other business group, let alone individual citizens or their collectives⁶. While attempts by large firms or other stakeholders to influence policy are not inherently incompatible with liberal democracy, the scale of Big Tech’s influence has reached a tipping point. Even more than the lobbying of previous generations of megafirms, Big Tech’s disproportionate sway seems to compromise the very ability of states and supranational policymakers to independently determine the public interest through balanced reflection.

Consequently, there is a genuine risk that state or EU law-making and decision-making processes could de facto degenerate into instruments used by a handful of massive private Big Tech firms and associated super-rich individuals to impose their own rules on society and safeguard their specific interests⁷. A stark illustration of this dynamic is the ‘Digital Omnibus Regulation Proposal’⁸ published by the European Commission in November 2025⁹. Arguably the most significant ‘success’ of the Big Tech lobby to date¹⁰, this proposal undermines substantial portions of the GDPR and the AI Act, effectively paving the way for the deregulation these companies have long sought.

The challenge of enforcement

Even when the EU and its member states succeed in enacting and maintaining stricter regulations, ensuring Big Tech’s compliance remains exceptionally difficult. Armed with unprecedented financial clout¹¹, these companies can engage in endless litigation or absorb most fines with relative ease. Furthermore, their unparalleled international and spatial mobility—partly facilitated by their own technologies—allows them to exploit the age-old competition among states for multinational favour, often accelerating the already devastating ‘race to the bottom’.

The particularly close ties between Big Tech and the current US administration do not make enforcing

European tech regulation any easier. The situation has reached a critical juncture where geopolitical leverage is weaponised. For instance, when a US Vice President threatens to dissolve NATO in response to EU fines against American tech giants¹², or when protecting American Big Tech becomes embedded in a US National Security Strategy¹³ that is openly hostile to European interests, the EU’s regulatory autonomy is severely compromised.

Furthermore, under the current asymmetrical power relations, extensive European frameworks like the General Data Protection Regulation (GDPR), the Digital Markets Act (DMA), and the Digital Services Act (DSA) risk producing counterproductive outcomes. Even

³ See, for example, the Preamble and Article 2 of the Treaty on the European Union (TEU).

⁴ See the European Convention on Human Rights.

⁵ R. Passchier, *De vloek van Big Tech: De juridisch-technologische wortels van constitutioneel verval en digitaal feodalisme* (‘The Curse of Big Tech: The legal and technological roots of constitutional decline and digital feudalism’), Boom 2024.

⁶ Corporate Europe Observatory, ‘Big Tech lobby budgets hit record levels’, October 14, 2025. <https://corporateeurope.org/en/2025/10/big-tech-lobby-budgets-hit-record-levels>

⁷ Cf. Y. Mounk, *The People vs. Democracy: Why Our Freedom is in Danger & How to Save it*, Cambridge, MA: Harvard University Press 2018, p. 243.

⁸ <https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-regulation-proposal>

⁹ Corporate Europe Observatory, ‘Article by article, how Big Tech shaped the EU’s roll-back of digital rights’, January 14, 2026.

<https://corporateeurope.org/en/2026/01/article-article-how-big-tech-shaped-eus-roll-back-digital-rights>

¹⁰ R. Mahieu, ‘The Ominous Omnibus: Dismantling the Right of Access to Personal Data’, in: *Verfassungsblog*, 5 December 2025.

¹¹ See, for example, <https://www.visualcapitalist.com/ranked-the-worlds-50-most-profitable-companies-in-2024/>

¹² G. Kilander, ‘JD Vance says US could drop support for NATO if Europe tries to regulate Elon Musk’s platforms’, in: *Independent*, 17 September, 2024.

¹³ <https://www.whitehouse.gov/wp-content/uploads/2025/12/2025-National-Security-Strategy.pdf>

when their content appears highly protective of liberal and democratic values, these laws often risk hindering potential competitors more severely than the entrenched

giants, thereby inadvertently reinforcing the very monopolies they aim to curb.

Extreme dependence and the erosion of sovereignty

Crucially, European governments and other key societal players have become so deeply dependent¹⁴ on Big Tech that even vital infrastructure—ranging from medicine distribution and financial transactions to national security operations—is now vulnerable to failure, manipulation, and blackmail. As this reliance intensifies, the locus of ultimate authority and enforcing power on critical social issues (classical sovereignty) is effectively shifting away from democratic centres like The Hague or Brussels, where such power constitutionally belongs, towards Washington, Silicon Valley, Seattle, Shenzhen, Beijing or some other major hub in China.

Increasingly, it is the US President or Chinese leadership, alongside a handful of unelected Big Tech CEOs and major shareholders, who dictate the operational reality of highly digitalised European institutions¹⁵. For example, as the coders of essential and widely used software - including leading 'platforms', as the (co-)authors of highly influential and consequential technological standards¹⁶, as the drafters of hard-to-evade general terms and conditions, and as mega-investors, they determine

whether ministries, hospitals, and schools function. These players together set the norms for political discourse and control the flow of information to European citizens. They influence electoral outcomes, define the practical value of European copyright, and decide which technologies are developed and for what purposes. In doing so, the fundamental prerogatives of the state and the EU as well as their citizens are being quietly outsourced to private, foreign entities.

The EU may regulate what it wants. But in the current situation, a few tech leaders, alongside the American or Chinese governments, ultimately decide to a large extent (if not almost entirely): how much privacy European citizens really have; to what extent AI companies may use them as guinea pigs for their latest products; how addictive those products are; whether children have access to them; and whether these companies will eventually get away with scraping the internet to 'train' their Large Language Models—a practice Naomi Klein describes as 'what may turn out to be the largest and most consequential theft in human history'¹⁷.

The pursuit of profit, growth, and power

When asked, representatives of Big Tech are likely to dismiss these concerns entirely, typically insisting that their actions serve technological progress and the world's best interests¹⁸. For years, Alphabet (Google's parent company) operated under the slogan 'Don't be evil'. Following every privacy scandal, Meta's CEO Mark Zuckerberg has reiterated the mantra that 'the future is privacy'. Similarly, Microsoft recently declared in a blog post that its 'support for Europe has always been—and will always be—steadfast', a statement clearly aimed at reassuring anxious European stakeholders¹⁹.

Yet, when tested, Big Tech and its leadership, even more so than earlier corporate elites²⁰, appear to pursue only three values, with near-pathological intensity:

profit, growth, and power²¹. Consequently, essential societal goods—such as a high quality of life, functional governance, meaningful innovation, livelihood security, and human dignity—are systematically subordinated to these imperatives²². This disregard was starkly illustrated when, despite warnings from Silicon Valley leaders (however exaggerated) about AI-driven mass extinction and calls from experts for a development pause, these companies simply accelerated their AI ambitions further.

Moreover, tech leaders are increasingly leveraging their wealth and influence to construct a 'post-democratic' West. This agenda, often described as an 'authoritarian stack'²³, includes the active support of 'patriotic'—that is, far-right—political parties across Europe²⁴. There

¹⁴ V. Gineikyte-Kanclere et al., 2025, *European Software and Cyber Dependencies*, publication for the Committee on Industry, Research and Energy, Policy Department for Transformation, Innovation and Health, European Parliament, Luxembourg.

[https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778576/ECTI_STU\(2025\)778576_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2025/778576/ECTI_STU(2025)778576_EN.pdf) See also: WRR, *Preparing for Digital Disruption*, The Hague 2019.

¹⁵ See also: M. Schaake, *The Tech Coup: How to Save Democracy From Silicon Valley*, Princeton: Princeton University Press 2024.

¹⁶ J. Bessen, *The New Goliaths: How Corporations Use Software to Dominate Industries Kill Innovation and Undermine Regulation*, Yale University Press 2022.

¹⁷ N. Klein, 'AI machines aren't hallucinating. But their makers are.' in: *The Guardian*, 8 May 2023.

¹⁸ See e.g. A. Becker, *More Everything, Forever: AI Overlords, Space Empires, and Silicon Valley's Crusade to Control the Fate of Humanity*, Basic Books, 2025.

¹⁹ <https://blogs.microsoft.com/on-the-issues/2025/04/30/european-digital-commitments/>

²⁰ J. Bakan, *The Corporation: The Pathological Pursuit of Profit and Power*, London: Constable 2004.

²¹ See, for example, R. Foroohar, *Don't Be Evil: The Case Against Big Tech*, Penguin 2019.

²² See, for example, K. Crawford, *Atlas of AI*, New Haven: Yale University Press 2021.

²³ <https://www.authoritarian-stack.info/>

²⁴ Corporate Europe Observatory, 'Article by article, how Big Tech shaped the EU's roll-back of digital rights', January 14, 2026. <https://corporateeurope.org/en/2026/01/article-article-how-big-tech-shaped-eus-roll-back-digital-rights>

is a growing convergence between certain Big Tech executives and the American government, openly

working in tandem to forge a techno-libertarian, post-democratic West²⁵.

The limits of corporate resistance

Even if Europe could trust Big Tech companies themselves, and even if these companies were to genuinely prioritise society's best interests, they would still be incapable of offering sufficient resistance to a US government hostile to Europe. While the US government does not (yet) exercise total control over these firms, it possesses, among other things, the authority²⁶ to compel them to grant access to the cloud environments to which European governments have increasingly migrated their data and virtual machines²⁷.

Consider the implications of a direct order: what would

Big Tech do if the US government mandated a complete blockade of their services in Europe as a sanction? We have seen precedents for such coercion, notably in the cases involving the Amsterdam Trade Bank and the International Criminal Court (ICC)²⁸. Faced with a choice between sacrificing their lucrative US operations—or risking the freedom and assets of their leaders and major shareholders—to defend Europe or the Netherlands, one cannot realistically expect them to choose the latter. Their loyalty, ultimately, lies with the jurisdiction that holds the keys to their existence.

Alternatives

Let us be candid: the American—and Chinese—Big Tech model is fundamentally incompatible with European liberal democracy, as is Europe's profound reliance upon this model. Regardless of what Big Tech's representatives and cheerleaders may claim, governments and critical infrastructure operators must, at the very least, develop a Big Tech exit strategy: a Plan B to swiftly switch to alternatives if circumstances demand it. Surely, it cannot be acceptable that institutions tasked with safeguarding our most fundamental values depend on the goodwill of a few profit-driven American multinationals and a geopolitical ally that is increasingly unpredictable, if not hostile, to democratic norms.

From the perspective of liberal democracy, European leaders must collaborate with European players to develop viable alternatives. Europe must seize this crisis as an opportunity to foster an economy where diverse new players can thrive.

Break the dominance of Big Tech, but beware of simply

replicating the American model: this is a pitfall that French President Macron's 'European Champion' policy and Rapporteur Mario Draghi's recent proposals²⁹ seem at risk of falling into. We must remember that the Big Tech model is also causing severe damage to American democracy and the constitutional structures underpinning it. Indeed, the Chinese government and its Big Tech have turned China in an almost 'perfect' police state³⁰.

Therefore, instead of creating European Big Tech, we should cultivate opportunities for smaller enterprises that do not pursue profit, growth, and power at all costs. We need models based on open source combined with steward ownership (e.g., Nextcloud, Fairphone or Proton), non-profit structures (e.g. Signal), or 'universal capitalism' and 'property-owning democracy' (a concept yet to be realised in the digital sphere³¹). By reducing size, increasing diversity and co-ownership, and prioritising public values over profit, digitalisation can finally live up to its promise—serving liberal democracy instead of undermining it.

The path forward

However, governments and important organisations must actively facilitate this transition. Smaller tech companies, lacking the billions in profits enjoyed by incumbents,

simply cannot deploy armies of lobbyists, sales teams, or representatives to every corner of the market. They are often less well-known and may not yet offer the

²⁵ See e.g. K. Hao, *Empire of AI: Dreams and Nightmares in Sam Altman's OpenAI*, New York: Penguin Press, 2025. P. Krugman, 'Is This The End of the Free World?', in: *Substack*, Dec 8, 2025.

²⁶ <https://edri.org/our-work/promises-unkept-the-eu-us-data-privacy-framework-under-fire/>

²⁷ V. Gineikyte-Kanclere et Al., 2025, *European Software and Cyber Dependencies*, publication for the Committee on Industry, Research and Energy, Policy Department for Transformation, Innovation and Health, European Parliament, Luxembourg.

²⁸ See e.g. A. Satariano and J. Smialek, 'Europe's Growing Fear: How Trump Might Use U.S. Tech Dominance Against It', in: *The New York Times*, June 20, 2025.

²⁹ https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en

³⁰ J. Bakan, *The Corporation: The Pathological Pursuit of Profit and Power*, London: Constable 2004.

³¹ G. Cain, *The Perfect Police State: An Undercover Odyssey into China's Terrifying Surveillance Dystopia of the Future*, Public Affairs 2021.

³¹ But note developments in the Energy Transition: J. Lowitzsch (ed.), *Energy Transition: Financing consumer co-ownership in renewables*, Palgrave Macmillan 2019.

seamless, integrated ecosystems that deliver ‘organic’ user experiences. Moreover, adopting a competitive, responsible tech model may necessitate reforms in tendering rules, procurement policies, and the retraining of staff accustomed to current platforms.

In order to save liberal democracy from Big Tech, leading figures, including elected representatives, must recognise and champion the fact that digital technologies are not merely a matter of business operations, economic growth or convenience, but a matter of utmost constitutional and strategic importance.

Admittedly, moving away from Big Tech may entail sacrificing certain conveniences, at least in the short term. Nevertheless, governments and important organisations—driven by their constitutional and treaty obligations to uphold liberal democracy—must be willing to strive for independence from Big Tech. In this regard, EU institutions play a pivotal role: they must encourage member states to resist the Faustian bargains offered by

Big Tech and urge them to seize the current crisis as an opportunity to realign economic and technological power with European norms.

In the end, the challenge posed by Big Tech is not merely a regulatory hurdle, but a fundamental test of our very constitutional identity. We face a choice that Louis Brandeis foresaw a century ago: ‘We can have democracy, or we can have great wealth concentrated in the hands of a few, but we cannot have both.’ Europe stands at this crossroads. Will we allow the ‘Curse of Big Tech’ to turn our constitutions into mere ink on paper, surrendering our sovereignty to private, foreign entities and ‘their’ authoritarian governments? Or will we seize this moment to build a digital economy where power is dispersed, and the average citizen remains the most important political player of all?

The path to constitutional resilience lies not in accommodating giants, existing or new, but in breaking their grip and reclaiming our democratic future.



“For democracies legitimacy comes from protecting the rights of their citizens, also on AI..”

Interview with Michael McNamara, Member of the European Parliament and European Parliament rapporteur on AI

By Gaston Moonen



© European Union, 2026. Source: EP.

In November 2025 the European Commission presented its Digital Omnibus proposals with technical amendments to a large corpus of digital legislation as part of its digital simplification agenda. The European Parliament followed with the appointment of Robert McNamara as AI rapporteur, representing the European Parliament’s viewpoints in the *trilogue* discussions between the European Parliament, the Council and the Commission during the last few months. While broadly satisfied with the outcome of the discussions so far, he underlines the need to implement the regulations and guidance as adopted, for the benefit of European companies, but, above all, for EU citizens.

Human rights as starting point

LIBE deals with AI from the point of civil liberties because obviously AI has an impact on fundamental rights...

Michael McNamara has been the European Parliament's rapporteur on AI since January 2026. However, this does not mean AI is his daily bread and butter. 'Consciously, I don't use a huge amount of AI. But I do use it occasionally, I suppose using a Google overview is essentially now an AI tool. I am a practicing lawyer, but since my election to the European Parliament, I haven't practiced very much. Obviously, AI is being used more and more by professionals as well as individuals.' He adds that unconsciously people use AI a lot, such as for finding holiday accommodation online.

The MEP explains that he became rapporteur being a Member the Parliament's Committee on Civil Liberties, Justice and Home Affairs (LIBE), responsible for protecting civil liberties and human rights, and also as co-chair of the joint Working Group of the Committee on the Internal Market and Consumer Protection (IMCO) and the LIBE Committee. 'LIBE deals with AI from the point of civil liberties because obviously AI has an impact on fundamental rights, rather than from an internal market perspective, which is the role of another committee.'

That focus is also related to his professional background. 'I have worked on human rights areas of law in the past, as a lawyer, as a human rights officer with the OECD. Obviously, AI has potential to be very relevant to that regard. One of the requirements for high-risk AI systems is that fundamental rights impact assessments are carried out whereas general-purpose AI (GPAI) model providers must follow a Code of Practice designed to protect fundamental rights by identifying systemic risks that could cause fundamental rights infringements. Depending on the use of particular data, you can end up with results that are biased and it is important to be vigilant against that.'

Outcome works but process questionable

..., I am not a particular fan of amending laws before they've even come into effect. (...) It does not reflect very well on either the legislators or the European Commission.

In May 2026 the European Parliament and the Council reached a provisional agreement, as part of the [Digital Omnibus Proposal](#), to streamline certain rules concerning AI. Michael McNamara is broadly satisfied with the outcome of the discussions but ventilates he is not all that happy with the process as such. 'Fundamentally, I am not a particular fan of amending laws before they've even come into effect. I mean, the AI Act was passed in 2024, we're now in 2026 and it's already been amended before it's even come into effect. I am not sure that that is a good way for Europe to proceed. It does not reflect very well on either the legislators or the European Commission.'

He has an understanding for Commission's position though. 'The harmonised standards envisaged to demonstrate compliance with the higher risk chapter were not developed. The Commission could develop common specifications but generally, it is reluctant to do that because it takes a lot of time and it does not have the expertise to do so.' Yet he thinks that it does not look good when the body responsible for ensuring compliance with the law, instead of ensuring compliance is requesting that the law be changed. 'I don't think that bodes particularly well for the rule of law in the European Union. However, this is beyond the Parliament's control.'

Where the European Parliament could have an influence, its AI rapporteur is fairly satisfied. 'Given the proposals that were being made, I think it was a good outcome. For example, there was no proposal around this banning non-consensual sexual imagery. That is something the Parliament introduced, and some member states had a similar requirement.' He also indicates that some Commission proposals were successfully turned down by the Parliament or changed. 'There was, for example, a proposal to postpone by six months the requirement that all AI-generated content be watermarked. The parliament brought that forward.' He also refers to the changes the Parliament introduced

[on the Digital Omnibus negotiations] Given our starting point, I am satisfied with the outcome. ”

regarding registration of companies operating in high-risk areas, such as biometrics, law enforcement, democratic processes. 'Given our starting point, I am satisfied with the outcome.'

Harmonised standards, are also needed to make it less difficult for smaller AI developers. ”

When asked whether the Digital Omnibus Proposal was triggered by pressure coming from Member States and Big Tech companies to slow down the AI regulation process, Michael McNamara observes that there is a lot of lobbying by Big Tech companies towards the European Parliament, Member States, and the Commission. Likewise civil society organisations also make their positions clear. Whether that lobbying is causative for the changes or merely correlates with the change in approach, is a question you have to ask the Commission.' In his view the main reason for the proposals was to stop the clock in the introduction of the high-risk chapter. 'Because the harmonised standards were not in place. It seems that the timeline was on the ambitious side when compared to the normal timelines for developing standards. Secondly, this is a novel area, different perhaps from developing standards in areas where standards already exist. Harmonised standards are also needed to make it less difficult for smaller AI developers.'

EU's AI governance – a work in progress

...70 percent of European data stored in tech companies and clouds is US controlled. ”

Making things easier for smaller AI developers, including those originating from Europe, is one of the policy objectives of this European Commission. But this may collide with the EU objective to regulate the AI architecture and its use more than in other global regions...while the EU's current dependency on US Big Tech is high. According to the European Parliament's AI rapporteur this will not change overnight. 'We are very dependent on American tech - 70% of European data stored in tech companies and clouds is US controlled. Only 15% of our data is stored in Europe and under European jurisdiction. That is clearly an issue. An issue that became very evident when the computer facilities of the International Criminal Court (ICC) prosecutor were stopped some time ago. It's clearly a very big problem over a longer period, addressing it quickly is not going to be easy nor painless.'

Regarding the AI Act, many parts haven't come into force at all yet in Europe so how can it be somehow responsible for AI lagging behind in Europe? ”

Michael McNamara refers to legislation the Commission is proposing around data sovereignty, the cloud and the AI development. 'It has been postponed several times, but it was recently published. Obviously, that's something that the Parliament will be very active on in terms of legislating in this area. However, I don't think it'll be particularly easy to define what is meant by 'sovereign', which types of data must be stored in Europe or why and how this should be done if Europe, if Europe simply doesn't have the storage capacity. It's going to be a difficult discussion, and Europe needs to prioritize the development of this infrastructure.' He also raises his concerns about environmental aspects that have to be looked at. 'Depending on where they're located, data centres required considerable amounts of water for cooling and above all, they require energy and there's an energy crisis right now. There are EU states that have perhaps a surplus of energy, but they're very few. And the lack of energy is creating competitiveness problems for their whole industrial sector already, without adding on the energy requirements of AI. It's going to be a very difficult conversation, even before you look at the difficulties in the electricity grid, something that has not been the subject of investment over a protracted period.' He adds that the latter is a major issue in Ireland, his motherland, and several others.

As to the dilemma of the EU choosing regulation possibly at the cost of innovation and marketing capacity and hence its autonomy regarding AI, the MEP is not convinced that it is the AI Act that puts Europe behind other regions. 'Regarding the AI Act, many parts haven't come into force at all yet in Europe so how can it be somehow responsible for AI lagging behind in Europe? I think that the causes are different.' He refers to the earlier mentioned energy constraints, but also to finance possibilities. 'There are many start-ups in Europe that move to the United States to scale up and get financing there. If you ask them whether they like regulation, of course,

they don't. But if you ask whether that was the primary reason they moved to the United States, the vast majority that I've spoken to say: "No, it's about financing." He underlines that bureaucracy, often related to regulation, needs to be reduced where possible and reporting has to be simplified, as is one of the aims of the current European Commission. 'But that doesn't mean that essential regulation to protect rights, to protect society from risks is not necessary.'

Subsequently the MEP searches for figures to support his view, providing financial investment data relating to securitisation of data centres¹. 'These figures are stark and just demonstrate how our financial markets work. From 2018 to 2024, data centre securitisation in the US was USD 36.8 billion. In 2025, it was USD 26.8 billion, which is a total of USD 63.6 billion in the United States. In Europe, over that period of time, it was USD 0.8 billion. Europe is hoping to invest €200 billion over the next five years to encourage AI development, of which €50 billion is to develop so-called gigafactories where there'd be computing power made available to European startups, researchers, etc. But I mean, OpenAI, which is one company, is planning to invest over USD 400 billion in the next couple of years. The figures are stark, the differences huge.'

Can democracy flourish on fake news?

I was in Bangladesh (...) deepfakes were also part of the election campaign and reported upon quite extensively by the election observation team... ”

One of the concerns regarding the influence of AI is how it can influence democratic processes. At EU level Michael McNamara has not seen any particular effects yet but his experience goes further. 'I would say two things. One relates to election observation I obtained because I work in a delegation oriented to South Asia. I was in Bangladesh for the recent elections and dealt with AI-generated postings on social media, primarily Facebook, but not just. And deepfakes were also part of the election campaign and reported upon quite extensively by the election observation team on the ground there. Part of what such teams report now regarding many elections relates to the use of deepfakes, the use of AI-generated content.' He explains that one could have somebody expressing a political view and then see a huge amount of either disparaging remarks or positive remarks building to a crescendo. 'But you don't actually know whether these people who are agreeing or disagreeing are real people or not, so you can easily create a perception of widespread support or rejection for a candidate or a view.'

...with AI, you can create an impression of considerable support or opposition, without it necessarily being mirrored by views held by real people. ”

Another concern relates to unprecedented amounts of funding by AI companies to candidates based on their view of the regulation of AI, for example for the midterm elections in the US in November 2026. 'This is a very interesting phenomenon, which I have not encountered yet in Europe, but it may not be far away.' He explains that in most European countries there are donations limits, including bans on anonymous donations over a certain amount, including campaign financing ceilings. 'Presumably, people adhere to them but just because we have those laws in place in Europe does not necessarily mean that large international actors are aware of those laws or are inclined to adhere to them.' He raises the role social media can now play in elections in every country in Europe. 'Again, the difficulty with bots is that, with AI, you can create an impression of considerable support or opposition, without it necessarily being mirrored by views held by real people.'

As to his own use of social media, he uses the most common tools to communicate with his voters. 'Every candidate in Ireland is directly elected, so I seek to communicate with my voters to the greatest extent, the most through US based media, and TikTok. I would use European-based social media, but my primary concern is to reach my voters, reaching out through media they use.'

¹Data centres are being financed in numerous ways to facilitate their rapid expansion. One of these options is data centre securitisation which has most commonly involved an asset-backed securities or commercial mortgage-backed securities structure.

Staying in the tech forefront...the European way

Industrial applications using AI are being used by various industries across Europe, giving us a competitive advantage.

”

While Europe may appear as lagging behind in AI in general, the Parliament's AI rapporteur sees also positive aspects regarding European AI developments, observing that much of the AI focus goes on large language models as developed by OpenAI, Anthropic, etc. 'Europe has been using machine learning for a very long time in sectoral uses where it has been embedded in industrial applications. I think Europe has an advantage in that, rather than in these very large language models that require huge amounts of data, huge amounts of energy, etc. AI has been used in various processes in Europe for some time, even going back to the code breaker that was used in World War II. Industrial applications using AI are being used by various industries across Europe, giving us a competitive advantage.' He underlines that it is important to recognise this, explaining that several European industries wanted this recognition also in the AI regulatory framework. 'And it remains to be seen if ultimately the huge investments in these large language models will bear fruit.'

...developing an alternative AI that is democratically funded and democratically motivated...

”

Michael McNamara does not necessarily share concerns that the EU's regulatory framework to ensure proper AI governance will stifle its competitive advantage, arguing that more people see the need for AI regulation to preserve democratic societies. 'Yoshua Bengio, a Canadian academic based in Montreal, talks about developing a democratic AI, one that is not driven by commercial interests and maybe one that is capable of behaving more responsibly. One of the problems with AI is that it's quite sycophantic and designed to please, aiming to keep its customers online, happy and engaged for the longest period of time possible whilst spending money and generating income for those large corporations. Bengio talks about the necessity of developing an alternative AI that is democratically funded and democratically motivated, also because ultimately, if we do, move to artificial general intelligence [AGI] or artificial superintelligence [ASI]. If somebody gets to that point, they will enjoy huge competitive advantages over those who don't have that. And I suppose that is part of what's motivating this race between China and the United States. I'm aware that as we speak, Trump is meeting Xi. I just don't know whether the governance structure for AGI or ASI will feature in their conversation. I rather doubt that.'

He observes that historically technological advantages have had decisive influence, giving the example of a Mughal army that came out of Central Asia and defeated the Sultan of Delhi. 'A much smaller army but they had a lot of technological superiority. I think we, as democratic societies, need to band together – with Canada, Australia perhaps, with the UK - and look at what we can create as an alternative to these large models that are attracting so much investment and preventing becoming vassals to countries having massive technological superiority.'

Of course they were complaining about regulation. That's what regulated entities do.

”

He emphasises that this doesn't mean that regulating AI isn't all so important. 'The EU is a very large and lucrative market for these large tech companies. If you look at the code of practice that was developed in August 2025 by independent experts – Joshua Bengio being one of them - for general purpose AI models there was a lot of discussion, with some of the bigger companies saying that they couldn't sign up if this or that was in it, etc. But ultimately almost all signed up. Meta and xAI didn't sign up, but only with regard to the copyright part, for the rest of it, the safety parts, they were happy to do so. Of course they were complaining about regulation. That's what regulated entities do.'

He reiterates that there are reasons for regulation. 'Our citizens expect that there will be safeguards in place, to protect for example their image or their data from being abused. We have very recent examples of states using information about their citizens to control them here in Europe. It's not a hypothetical fear, it's something that happened recently in Europe, and people want to avoid this, legitimately and understandably.' For the MEP, that's

ultimately where states in the 21st century derive their legitimacy from. 'For democracies legitimacy comes from protecting the rights of their citizens, also on AI, and the same is true of a superstate like the European Union. If it is just about regulating a market and ensuring a high GDP rather than looking after citizens' rights, then I think the EU starts to encounter legitimacy problems.'

Food for thought for auditors and politicians alike

When discussing what topic would be particularly useful to be audited by the ECA in this domain, the AI rapporteur has a few suggestions. 'One issue to look at is the correlation between funding and the development of AI, of startups, etc. In my view if start-ups initially grow thanks to EU funding or EU infrastructure, then that money would be well spent, even if they go on to be successful elsewhere.'

Another topic worth looking into is the argument raised above that regulation will stifle innovation. 'I suppose there are some studies around on how different data protection regimes have impacted the economy. The reasoning of large corporations may be that if such and such regulation was removed; we would all be much wealthier. But would we?' The parliamentarian does not take this for granted. 'These companies have a duty to their shareholders to maximise profits. Their duty is not to society at large or to foster democratic societies. He emphasises that for him this should not be the only determinant, but it would be good to know whether they have a point. 'During the course of developing the Digital Omnibus Proposals there were figures thrown around about the cost of AI compliance. When I asked this question to the Commission, I got very different figures than those presented by trade associations.'

Is it an existential must for European civilization to stay at the forefront of technological progress? (...) But if it is considered existential, how do we best target EU funding? ”

Finally, he raises two questions, while expressing doubts as to whether it would be an ECA task to focus on them. 'Is it an existential must for European civilization to stay at the forefront of technological progress? And if that is the case, how is money best spent?' He refers to the [ECA's special report 08/2024](#) on the EU's AI ambitions. 'The conclusions were not entirely flattering.' He underlines that money is finite and with so many different infrastructural requirements facing European governments and also the EU, this is just one of them. 'But if it is considered existential, how do we best target EU funding?' On the first question, he almost provides the answer himself, pointing to the substantial use of AI in defence industry. The second question remains to be answered, hopefully building on the lessons learned, reflected also in ECA reports.

Foresight and audit

The future of AI in audit – opportunities and challenges

By Anna Zygierewicz, Directorate of the Presidency



© BiancoBlue/Depositphotos.com

AI is a driver for change in audit work and in audit as a profession. But as with almost all changes, it may bring challenges as well as opportunities. Anna Zygierewicz, a foresight, risk and strategy officer at the ECA, outlines some ways AI could revolutionise audit, such as increasing quality and speed. However, she also presents the challenges auditors may when using AI, most notably the accuracy and hence trustworthiness of results, but also the various skillsets needed. Lack of transparency in the way AI algorithms work, as well as ethical and environmental aspects, are other issues that will also need to be addressed sooner or later.

AI use in audit

One definition of AI describes it as ‘the implementation of algorithms that can learn from data (...) and produce outputs (...) that typically would require human intelligence’, while stating that it ‘can also be characterised as the effective imitation of human (...) sensory abilities

(...)’¹. As well as AI, new technologies reshaping audit include big data and data analytics, blockchain, machine learning, deep learning, natural language processing, large language models, robotic process automation (RPA) and satellite imagery².

¹ Nicolas Epelbaum and Patricia Jackson Farrell, ‘Artificial Intelligence and Auditing: A Bibliometric Study’, in: *Accounting Perspectives*, 2026.

² European Parliament, *The future of digitalisation in budgetary control*, 2024; *The Application of Artificial Intelligence in Audit: State of the Art and Possible Future Developments*, in: *International Journal of Business and Management*, Vol. 20, No 5, 2025.

The use of new technologies in audit is progressing, though unevenly. Recent research reveals that, in large public accounting firms, 'simple AI' technologies (e.g. simple machine learning) are used widely in audits, whereas 'complex AI' tools (e.g. generative AI) are still being adopted experimentally and have not yet been fully deployed in production processes³. It also indicates that 'RPA is used to automate repetitive administrative processes while the use of RPA for audit tasks is not as common'⁴.

As of 2024, almost three quarters of businesses worldwide were already using AI in financial reporting, and this figure was expected to rise to 99% by 2027⁵. In terms of public audit, an Organisation for Economic Co-operation and Development (OECD) survey published in May 2026 shows that the use of AI is at an early stage, but that experimentation is expanding. AI forms part of the digital transformation process in many public audit institutions, and 80% of the institutions contacted by the OECD had established internal AI guidelines or policies⁶.

Opportunities arising from AI in audit

AI can help improve audit planning by enabling the analysis of larger volumes of data and the identification of potential risks and trends. This allows auditors to conduct scenario modelling and implement mitigation strategies.

During the audit itself, AI can help auditors collect and analyse data, as well as detect abnormalities and irregularities and generate lists of findings based on its analysis. Experts suggest that 'many artificial intelligence tools also involve machine learning. AI algorithms can learn from historical data and identify patterns that might indicate fraud, errors, or control weaknesses'⁷.

Using AI may also enable auditors to collect and analyse significantly larger volumes of data than they could ever process manually. As a result, AI-driven audits may replace the traditional sampling approach with full-population testing, while also increasing productivity and cost-efficiency in the audit process⁸.

The involvement of auditors themselves remains indispensable to the audit process, however. This is mainly because AI output is still not fully trustworthy. Reviewing AI-generated output remains, at least for the time being, essential for ensuring reliable audit results.

Using AI can make audit work more engaging. Some researchers point out that AI changes both the way auditors approach audit and their role in the process, as AI 'allows auditors to focus on high-value tasks, fostering a more strategic and analytical approach to the audit process'. Other benefits cited include:

a) Automation of routine audit tasks: AI can replace auditors in repetitive and time-consuming tasks, such as data extraction, classification and reconciliation, enabling them to focus on higher-value analysis.

b) Enhanced data analysis and anomaly detection: using sophisticated algorithms to detect anomalies, AI can help identify irregularities, discrepancies and potential instances of fraud.

c) Risk assessment and management: AI-driven tools can help identify potential risks, allowing auditors to conduct scenario modelling and to implement mitigation strategies; and

d) Real-time auditing and dynamic insights: AI can allow real-time analysis of financial records, transactions and other critical information, the results of which can be used in decision-making processes⁹.

The use of more advanced AI tools, such as agentic AI, is also being observed in audit work and may have even bigger impact, as 'agentic AI can perform advanced audit tasks easier and with less human input than generative AI' (see **Box 1**).¹⁰ Such systems can not only achieve specific goals, but also proactively contribute to problem solving.¹¹

The foresight element of AI use in audit should not be overlooked. Governments and businesses – like individuals – try to anticipate and prepare for future events and possible change in a structured and systematic way. The discipline of strategic foresight – which uses different methodologies to explore, anticipate and shape the future – is one way of achieving this. AI can feed into strategic foresight in various ways that are useful for public auditors, including the ECA, such as developing multiannual strategies or designing and carrying out relevant audit tasks.

³ Julia Kokina, Shay Blanchette, Thomas H. Davenport, and Dessislava Pachamanova, '[Challenges and opportunities for artificial intelligence in auditing: Evidence from the field](#)', in: *International Journal of Accounting Information Systems*, Volume 56, December 2025.

⁴ Ibid.

⁵ KMPG, *AI in financial reporting and audit: Navigating the new era*, 2024.

⁶ Maria Eugenia Heyaca and Andrea Pallotta, *The State of Artificial Intelligence in Public Audit. Evidence from Selected Countries and the European Union*. OECD Artificial Intelligence Papers No. 58, May 2026.

⁷ Wolters Kluwer, *Artificial intelligence in auditing: Enhancing the audit lifecycle*, 17 April 2024.

⁸ Olubusola Odeyemi, Kehinde Feranmi Awonuga, Noluthando Zamanjomane Mhlongo, Ndubuisi Leonard Ndubuisi, Funmilola Olatundun Olatoye, and Andrew Ifesinachi Daraojimba, '[The role of AI in transforming auditing practices: A global perspective review](#)', in: *World Journal of Advanced Research and Reviews*, 2024.

⁹ Ibid.

¹⁰ Andrew, Kenney, '[How AI is transforming the audit — and what it means for CPAs](#)', in: *Journal of Accountancy*, 2026; and Hamish Bowen and Michael Chainey, *From firefighting to futureproofing: How AI is reshaping internal audit functions across New Zealand*, Grant Thornton New Zealand, 2026.

¹¹ '[Adoption of AI and Agentic Systems: Value, Challenges, and Pathways](#)', in: *California Management Review*, 2025.

Box 1 - Agentic AI and its potential for audit

Most of the AI tools used today – ranging from Microsoft Copilot to Mistral – are reactive: they respond to prompts instead of taking initiatives. Rather than waiting to be asked, agentic AI systems can take sequences of actions, engage with external data sources and adjust, based on what they find. The change would go towards a model where the AI goes out, gathers leads, makes decisions, and return with a synthesised picture. For audit, be it internal or external, agentic AI could:

- monitor transactions flows and date in real time and do risk scores as conditions change. This would be a shift from point-in-time assurance to ongoing oversight;
- rather than testing a sample, it could test 100% of transactions against control requirements, identifying exceptions and flag anomalies;
- identify patterns that human review would miss, such as third-party behaviour that deviates from historical norms, with implications for fraud detection; and
- provide reporting and follow-up advice, autonomously draft findings, accelerating the audit cycle.

Agentic AI can automate the compliance work, understand goal context and goal setting, freeing auditors to focus on review and judgement-intensive work.

Challenges arising from AI in audit

KPMG points out that new technologies automate manual work in audit, making the audit profession more enjoyable, but also raise questions ‘about the nature of the audit process, the reliability of AI systems, and the changing role of the auditor’¹².

One of the challenges of using AI in audit concerns the quality of data and metadata used. Supreme audit institutions (SAIs) stress that ‘in audit practice, AI is only as good as the data that feeds it. And AI is inseparable from the metadata that describes, qualifies, and governs that data.’¹³ Using AI ‘entails a set of risks related, among other areas, to privacy, confidentiality, security, transparency or bias and discrimination. (...) SAIs, when auditing public services and policies, must adapt their human and technological resources to this new scenario’¹⁴.

These challenges may be addressed as follows¹⁵:

- Privacy, confidentiality and security*: the large volumes of data collected need to be protected from unauthorised access and misuse. Citizens need to be informed about their rights in relation to data collected on them.
- Transparency*: there needs to be algorithmic transparency, so that ‘citizens can know how autonomous decision systems make decisions that impact their lives’.

c. Algorithmic discrimination and equity: AI algorithms must not be influenced by the bias or prejudices of the people collecting and processing the data.

The quality of AI-generated output poses another challenge. Currently, AI output may not be fully reliable or trustworthy and therefore needs to be reviewed by auditors. They may nevertheless rely excessively on this output and fail to apply professional scepticism, which could lead to inadequately founded conclusions and recommendations. Another risk is that auditors may become very dependent on AI tools, which poses a systemic risk to audit quality if AI tools fail or are unavailable.

Cybersecurity poses another challenge, with AI ‘transforming cyber on both sides of the fight – strengthening defence while enabling more sophisticated attacks’¹⁶. Cyberattacks may therefore pose a threat to data collected in the course of audits.

The ethical challenges of using AI should not be overlooked either. To mention only three:

- AI may pose a threat to human rights*: the European Union Agency for Fundamental Rights has stressed that ‘the use of AI may reveal private information about people and can put vulnerable groups at further disadvantage.

¹² KPMG Netherlands, *The far-reaching impact of Artificial Intelligence on the audit profession*, 2024.

¹³ *Artificial Intelligence and Ontologies*, in: *EUROSAI Innovations*, No12, Winter 2025.

¹⁴ Dolores Genaro-Moya, Antonio Manuel López-Hernández, and Mariia Godz, *Artificial Intelligence and Public Sector Auditing: Challenges and Opportunities for Supreme Audit Institutions*, 2025.

¹⁵ Ibid.

¹⁶ World Economic Forum, *Global Cybersecurity Outlook 2026*, 2026.

It may also be used without fully understanding its risks, which in turn presents challenges for remedying any adverse impacts¹⁷.

b. *Autonomous AI may act in an unethical and unsafe way:* according to a recent study, autonomous AI agents have exhibited behaviours such as unauthorised obedience of non-owners, disclosure of sensitive information, uncontrolled resource consumption, cross-agent propagation of unsafe practices, and partial system takeover. This raises questions ‘regarding accountability, delegated authority, and responsibility for downstream harms’¹⁸.

c. *AI may increase unethical behaviour by users:* researchers discovered that AI ‘may facilitate the delegation of unethical behaviour’ such as ‘lying and cheating for profit’, as ‘machine delegation may reduce the moral cost of cheating’¹⁹.

AI is less resistant to bias than one might hope, and our ability to evaluate algorithms is still limited^{20,21}. Tests on algorithms have revealed bias that raises concerns about the trustworthiness of AI output. Then there are the environmental challenges related to AI use, such as its high energy consumption and significant carbon footprint²².

AI is likely to have significant implications for auditors, affecting not only their methods but also their role, skills and job structure²³.

The American AI company Anthropic estimates that ‘AI could wipe out half of all entry-level white-collar jobs’²⁴,

while Ford predicts that AI will ‘cut in half the number of white-collar jobs in the U.S.’²⁵. In the case of auditors, however, AI is more likely to transform their role than replace them, shifting their focus towards higher-value activities such as applying their professional judgement, assessing risk, interpreting findings, undertaking strategic analysis, communicating with clients and performing ethical oversight. New skillsets, particularly AI skills, may also play a growing role in the audit profession²⁶. In 2025, almost 7% of jobs posted in English-speaking countries by the Big Four accounting firms (Deloitte, EY, KPMG and PwC) required AI skills²⁷.

The growing use of AI by private and public organisations will also result in growing pressure to audit AI. Audit frameworks providing standards and guidance were in many cases written before AI use became widespread, and standard-setters may still need to play catch-up when it comes to addressing AI-assisted procedures. Some of the challenges auditors may struggle with in their work include difficulty understanding algorithms and AI decision-making processes; the data used in these processes; detecting in-built bias possibly lead to unfair or discriminatory output; or the ownership, when AI systems operate autonomously, etc.²⁸.

The EU AI Act imposes requirements on AI users in Europe, including with regard to the technical auditability of AI systems pursuant to the General Data Protection Regulation²⁹. This is intended to enhance the detection of data misuse by the audited AI systems.

AI in audit – getting to grips with the future

The introduction of AI is a driver for change in audit work and in audit as a profession. The process will accelerate, with new and advanced AI tools being used and auditors becoming increasingly skilled in AI. However, for the process to be a success, auditors will need to adapt their work methods and tools to the AI era and to upgrade their AI skills to be able to perform their work in the AI environment in efficient and reliable manner. Using AI in audit work should be viewed as an opportunity to enhance its:

- *quality:* more testing, consistent execution of audit procedures, enhanced detection of anomalies and errors;

- *efficiency:* more audit procedures executed automatically and in less time; and

- *effectiveness:* deeper data analysis, more accurate risk identification, and enhanced fraud detection.

Audit work will increasingly be performed with the support of AI tools, while continuing to be monitored, reviewed and ultimately assessed by human professionals. However, the implementation of AI in audit needs to be carefully planned and managed in order to address key challenges and mitigate potential risks.

¹⁷ European Union Agency for Fundamental Rights, *Assessing High-risk Artificial Intelligence: Fundamental Rights Risks*, 2025.

¹⁸ Natalie Shapira, Chris Wendler, Avery Yen et al, *Agents of Chaos*, Northeastern University, Harvard, MIT, Stanford, and Carnegie Mellon & Co., 2026.

¹⁹ Nils Köbis, Zoe Rahwan, Raluca Rilla et al, ‘Delegation to artificial intelligence can increase dishonest behaviour’, in: *Nature* 646, 126–134 (2025).

²⁰ Test results of the New Zealand Institute of Skills and Technology showed that AI models tend to favour women for jobs over men. See: Jesus Mesa, ‘AI Hiring Favors Women Over Equally Qualified Men’, Study Finds, in: *Newsweek*, 15 July 2025.

²¹ Tests showed significant racial, gender and intersectional bias in how large language models ranked resumes. See: *AI tools show biases in ranking job applicants’ names according to perceived race and gender*, University of Washington, 2024.

²² James O’Donnell and Casey Crownhart, ‘We did the math on AI’s energy footprint. Here’s the story you haven’t heard’, in: *MIT Technology Review*, 2025.

²³ European Training Foundation, *The AI impact on labour markets. What we know so far*, 2025.

²⁴ Shivaune Field, ‘Entry-level jobs are under threat: Inside the AI shift reshaping the workforce’, in: *Forbes*, 2025.

²⁵ Jake Angelo, ‘Microsoft AI chief gives it 18 months—for all white-collar work to be automated by AI’, in: *Fortune*, 2026.

²⁶ ‘How AI skills and experience are transforming the workplace’, idem.

²⁷ Clara Murray, Ellesheva Kissin, ‘Big Four post more job ads for AI specialists than auditors’, in: *Luxembourg Times*, 19 May 2026.

²⁸ KPMG, *6 key challenges of auditing AI and how to approach them*, 2025.

²⁹ [Regulation \(EU\) 2016/679](#), OJ L 119, 4.5.2016.

Foresight and audit

What does AI say about its impact on the future of audit?

AI-assisted article, reviewed by Daniela Hristova, Directorate of the Presidency



© Andrey Popov/stock.adobe.com

When discussing potential contributions to our ‘Foresight and audit’ section, which explores what AI might bring to the future of audit, we decided to ask our foresight, risk and strategy officers for their perspectives. But what does AI itself think it could bring to public audit? Writing a *prompt* to find out is easy, but getting AI to produce an article that passes the test of being based on real sources proved much more difficult. Daniela Hristova, assistant to the director in the Directorate of the Presidency, explains how she went about the task, and presents the outcome of a prompt that seems rather straightforward and delivers an interesting read. However, she explains, the process of generating and reviewing the article leaves a feeling of uncertainty and raise questions – precisely the opposite of what audit should do. And as for AI’s views on its own potential impact on the future of public audit: tellingly, they shift from enthusiasm to caution, in light of AI’s numerous risks and limitations.

Getting an AI-generated perspective – with some guardrails in mind

When preparing this edition of the ECA Journal, we decided – guided by both curiosity and professional scepticism – to conduct a small and rather straightforward experiment. Our objective was to quickly

assess the reliability of AI-generated academic content by asking AI systems to provide an article based on the following prompt:

How do you think AI will change external audit of the

public sector in the next ten years and which major change for public audit will stand out? Write an article, length between 1800 and 2000 words, including source references.

For the test, we selected four (fairly well-known) large language modules (LLMs), including a European one. For the purposes of this article, we will keep the names of these LLMs anonymous.

The results raised concerns. At first glance, the articles generated appeared convincing: they were well structured, coherent, and contained several examples supported by numerical data referenced to specific publications and authors – and in some cases even to specific pages or hyperlinks. All four AI platforms generated broadly similar arguments, observations and conclusions. However, upon closer examination, only one of the four results presented numerical data, and we were unable to verify the figures provided. More concerningly, the extent to which the existence of the sources cited

could be confirmed varied widely, ranging from all sources being verifiable in one case to none in another, with the remaining cases falling somewhere in between. In one instance, the system referenced a genuine ECA special report, but with an entirely fictitious title.

For one LLM, we also repeated the experiment using the same prompt and the same tool. The second time, the AI tool produced a different article, notably without the numerical data found in the first version. However, verification of the sources cited revealed similar problems.

The output of the LLM whose sources were all verifiable is reproduced in full below, changed only to apply ECA Journal formatting. It exemplifies the current capabilities and limitations of AI-assisted writing and source attribution. Sometimes the information provided is too extensive and repetitive. At the same time, it provides an interesting perspective on AI's potential implications for the future of public audit. We will let you, the reader, be the judge of whether this article is useful.

Artificial Intelligence and the future of external audit in the public sector: what will change most in the next decade?

Introduction

Artificial intelligence (AI) is rapidly transforming governments around the world. Public administrations increasingly use machine learning, predictive analytics, natural language processing, and generative AI to manage public services, detect fraud, allocate resources, process applications, and support decision-making. As governments become more dependent on AI, external auditors must adapt accordingly.

The implications for public sector auditing are profound. Supreme audit institutions (SAIs), national audit offices, courts of audit, and other public audit bodies have traditionally focused on assessing legality, regularity, economy, efficiency, and effectiveness. Over the next ten years, however, they will increasingly be required to evaluate not only financial transactions and administrative processes but also complex algorithmic systems that influence public decisions.

AI will change external audit in two distinct ways. First, auditors themselves will use AI to perform audits more effectively. Second, and perhaps more importantly, auditors will increasingly be required to audit AI systems used by governments. These two developments will reshape audit methodologies, skills, evidence gathering, and accountability mechanisms.

The most significant change likely to emerge over the next decade is the transition from periodic, sample-based auditing to continuous, data-driven assurance supported by AI. This transformation could fundamentally alter how public accountability is exercised.

AI as a tool for public auditors

Historically, public auditors have relied heavily on sampling techniques because reviewing every transaction was impractical. AI significantly changes this constraint. Machine learning systems can analyse entire populations of transactions rather than samples. They can identify anomalies, unusual patterns, high-risk expenditures, procurement irregularities, duplicate payments, and indicators of fraud at a scale previously impossible. AI can also automate time-consuming tasks such as document classification, contract review, invoice matching, and extraction of information from unstructured records.

According to a recent OECD study examining public audit institutions in fourteen countries and the European Union, emerging AI applications in public audit already include anomaly detection, document processing, knowledge management, and predictive risk assessment. The report concludes that although adoption remains at an early stage, experimentation is accelerating across many audit institutions¹.

These capabilities could dramatically improve audit efficiency. Instead of spending substantial resources on routine testing, auditors may increasingly focus on interpreting results, investigating anomalies, and exercising professional judgement. AI will not eliminate auditors, but it will alter the balance between mechanical review and analytical assessment.

Furthermore, generative AI systems are likely to become valuable assistants for audit teams. They can summarize large document collections, identify inconsistencies

¹ Maria Eugenia Heyaca and Andrea Pallotta, *The State of Artificial Intelligence in Public Audit: Evidence from Selected Countries and the European Union*, OECD Artificial Intelligence Papers No. 58, OECD Publishing, 2026

across reports, draft working papers, and support research on regulatory requirements. Such tools may significantly reduce administrative burdens and allow auditors to concentrate on higher-value activities.

The emergence of AI auditing as a core audit function

While AI will help auditors conduct audits, the more profound challenge is that governments themselves are increasingly using AI to make decisions. Many public authorities already employ algorithmic systems in areas such as tax administration, social benefits, immigration, healthcare prioritisation, policing, environmental regulation, and procurement. As these systems become more sophisticated, external auditors must determine whether they operate lawfully, fairly, effectively, and transparently.

Traditional audit approaches are not sufficient for this task. Auditors must understand issues such as training data quality, algorithmic bias, explainability, model governance, cybersecurity, and human oversight.

Researchers examining AI and public sector auditing argue that SAs must adapt their human and technological resources to address risks related to privacy, confidentiality, security, transparency, discrimination, and bias associated with AI systems².

Consequently, a new audit domain is emerging: algorithm auditing. Future public auditors may routinely assess questions such as:

Does the AI system comply with applicable laws and regulations?

- are decisions explainable and traceable?
- is the underlying data reliable and representative?
- does the system create unintended discrimination?
- are accountability mechanisms clearly defined?
- is adequate human oversight maintained?

These questions extend beyond traditional financial auditing and move public auditors into the broader field of digital governance assurance.

From financial accountability to algorithmic accountability

One of the most important shifts over the next decade will be the expansion of accountability itself. Traditionally, external audit has focused on whether public money

was spent appropriately and whether government programs achieved their objectives. AI introduces a new accountability challenge because decisions may increasingly be influenced by complex algorithms rather than directly by public officials.

Citizens may find it difficult to understand why an AI system denied a benefit application, prioritised one patient over another, or flagged an individual for investigation. When such systems affect rights and opportunities, public trust depends on effective oversight. Public auditors are uniquely positioned to provide this oversight. Their institutional independence, investigative powers, and experience in assessing public administration make them natural guardians of algorithmic accountability.

The INTOSAI Journal has emphasised that AI should be integrated into auditing within a framework grounded in professional judgement, ethical standards, institutional oversight, and public trust³. These principles are equally relevant when auditors examine AI deployed by government agencies. Over time, citizens may increasingly look to public auditors not merely as guardians of financial integrity but also as independent evaluators of digital governance systems.

Continuous auditing and real-time assurance

The most transformative development likely to occur during the next decade is the emergence of continuous auditing. Traditional public audits are retrospective. Auditors typically review activities months or years after they occur. Reports often describe problems that have already happened and may be difficult to correct.

AI makes a different model possible. Through automated monitoring of transactions, procurement systems, grant programs, and financial controls, auditors can receive near-real-time indicators of emerging risks. Rather than waiting for annual reviews, audit institutions may continuously analyse data streams from government systems.

This shift resembles the difference between periodic medical check-ups and continuous health monitoring. Instead of identifying problems after substantial damage has occurred, auditors can detect issues while corrective action remains possible. Predictive analytics may further enhance this capability. AI systems can identify patterns associated with fraud, waste, cost overruns, or project failure before these outcomes materialise. Audit institutions could therefore move from a reactive role toward a preventive one.

The result would be a fundamental redefinition of assurance. Public audit would increasingly become a continuous process rather than a periodic event.

² Dolores Genaro-Moya, Antonio Manuel López-Hernández, and Marija Godz, 'Artificial Intelligence and Public Sector Auditing: Challenges and Opportunities for Supreme Audit Institutions', in: *World*, Vol. 6, No. 2, 2025.

³ Jessica Du, *The Use of Artificial Intelligence Techniques in Auditing*, in: *INTOSAI Journal*, Q1 2026.

New skills for public auditors

The future audit workforce will require capabilities that differ substantially from those of today.

Traditional audit competencies—accounting, law, economics, public administration, and performance evaluation—will remain essential. However, they will need to be complemented by expertise in data science, information technology, cybersecurity, statistics, and AI governance.

Audit teams are likely to become more multidisciplinary. A future audit engagement examining an AI-driven welfare allocation system might include:

- financial auditors;
- data scientists;
- machine learning specialists;
- legal experts;
- ethics specialists; and
- public policy analysts.

The OECD has identified limited technical expertise as one of the major obstacles to scaling AI within public audit institutions⁴. Addressing this capability gap will become a strategic priority for SAIs worldwide. Universities and professional bodies may also revise audit education programs to include algorithmic accountability, AI governance, data analytics, and digital ethics.

Challenges and risks

Despite its potential, AI presents significant risks for public auditing.

First, AI systems can produce inaccurate or misleading outputs. Generative AI models may hallucinate facts, create fictitious references, or misinterpret context. Auditors cannot simply accept AI-generated conclusions without verification.

Second, transparency remains a major concern. Some machine learning models operate as ‘black boxes’, making it difficult to explain how decisions were reached. This creates challenges both for government agencies and for auditors attempting to evaluate them.

Third, excessive dependence on AI could weaken professional scepticism. Auditing has always relied on independent judgement. If auditors become overly reliant on automated tools, critical thinking may decline.

Fourth, data quality problems remain widespread in government. AI systems are only as reliable as the information they process. Poor-quality data can generate inaccurate risk assessments and misleading audit findings.

Finally, governance frameworks are still evolving. International discussions increasingly emphasise the need for robust institutional structures, standards, and oversight mechanisms to ensure trustworthy AI auditing⁵.

Therefore, successful adoption will require careful balancing of innovation and professional responsibility.

International standards and the evolution of audit methodology

Over the next decade, international audit standards are likely to evolve significantly. INTOSAI's International Standards of Supreme Audit Institutions (ISSAIs) already provide principles for public sector auditing. However, growing AI adoption will likely generate new guidance covering algorithm audits, AI-assisted audit procedures, model validation, and governance assessment⁶.

Audit methodologies may increasingly include requirements to evaluate:

- data lineage and provenance;
- model governance frameworks;
- human oversight arrangements;
- bias testing procedures;
- transparency and explainability mechanisms;
- cybersecurity protections; and
- AI risk management systems.

The development of internationally recognized methodologies will be crucial. Without common standards, audit approaches may vary significantly across jurisdictions, reducing comparability and potentially weakening public confidence.

The major change that will stand out

Among all anticipated developments, one change stands out above the rest: the transition from sample-based retrospective auditing to continuous, AI-enabled assurance.

Historically, auditors have accepted a fundamental limitation: they could only examine a fraction of available

⁴ Maria Eugenia Heyaca and Andrea Pallotta, *The State of Artificial Intelligence in Public Audit: Evidence from Selected Countries and the European Union*, OECD Artificial Intelligence Papers No. 58, OECD Publishing, 2026.

⁵ Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel E. Ho, *Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance*, 2022.

⁶ International Organization of Supreme Audit Institutions (INTOSAI), *Professional Standards and ISSAI Framework*.

information. AI removes much of this constraint. As AI systems gain the capacity to analyse entire datasets continuously, audit institutions will increasingly monitor public operations in near real time. Auditors will spend less effort identifying problems and more effort interpreting risks, assessing governance, and recommending corrective action.

This transformation will have implications far beyond efficiency. It could fundamentally alter the relationship between government, auditors, legislators, and citizens. Instead of learning about failures long after they occur, stakeholders may receive timely assurance and earlier warnings. Public accountability could become more proactive, preventive, and dynamic.

In many respects, this shift resembles the transition from paper records to digital systems several decades ago. The underlying purpose of auditing—providing independent assurance—will remain unchanged. However, the methods, timing, and scope of that assurance will be radically different.

Conclusion

Artificial intelligence will reshape external audit in the public sector during the next decade. Auditors will use AI to analyse vast quantities of information, improve risk assessment, and increase operational efficiency. At the

same time, they will assume a new responsibility: auditing the AI systems increasingly used throughout government.

The future public auditor will need stronger technological capabilities, deeper understanding of digital governance, and greater ability to evaluate complex algorithmic systems. New standards, methodologies, and multidisciplinary teams will emerge to support these responsibilities.

Yet the most important transformation will not simply be technological. It will be conceptual. AI will enable public auditing to move beyond periodic, sample-based reviews toward continuous assurance based on comprehensive data analysis.

If implemented responsibly, this change could strengthen transparency, accountability, and public trust. In an era when governments increasingly rely on algorithms to make decisions, independent audit institutions may become more important than ever—not only as guardians of public money but also as guardians of democratic accountability in the digital age.



‘Supreme audit institutions are pillars of democracy!’

Interview with Pierre Moscovici, ECA Member since 1 January 2026

By Gaston Moonen



On 1 January 2026, Pierre Moscovici took over from François-Roger Cazala as the ECA Member from France, bringing with him a wealth of experience in both national and EU government branches. Looking forward to the learning curve ahead of him as a new ECA Member, he gladly shares his experience and his vision of the European Union, as well as his thoughts on the role of audit within the European constellation and potential ambitions for the ECA to consider.

A new ECA Member, but a seasoned EU public servant

Pierre Moscovici expects to conclude his professional career in the branch in which he started: as an external auditor of the public sector – albeit in a different position, having served the public cause for 42 years. For him, the choice to go into public service was clear: ‘When I started my career and when I studied in the early 80s, the vocation of public service was probably stronger than it is today. We considered, especially in France, that public service was a noble vocation, working for the general interest of the people. As a student at Sciences Po, I chose to be in the section dedicated to public service.’ This was followed by studies at the Ecole National d’Administration (ENA). He points out that the first career choice he made was to work at the French Court of

...independence and freedom of judgement. That's what I found at the Court of Audit.

”

Audit. 'Because of my interest in the subject matter and secondly, because of its independence and freedom of judgement. That's what I found at the Court of Audit. And that's what I'm also looking for here at the ECA.'

Pierre Moscovici has served as a Member of Parliament, Member of the European Parliament, minister, EU Commissioner, and most recently First President of the French Court of Audit, from 2020 to 2025. Many of these roles also concerned European policy, the ECA being the fourth EU institution he has worked for. 'I am a European. My European years represent half of my career.'

Europe has power, a different power, of course, than some other global players such as China or the US, but power. (...) Europe was made for reconciliation and peace.

”

He observes that when he started working as a Member of the European Parliament in 1994, Europe was very different from what it is today. 'First, Europe has not only expanded, it has also unified. He explains his involvement, including as MEP rapporteur for the accession of Romania, where his father was born. 'This was absolutely historical, as a student I could not even have imagined that it would happen. Secondly, we have a single currency, an embryonic European defence set-up, a foreign security policy. We have much stronger EU institutions than before, including a parliament with the right to be a co-legislator. We have an economic policy linked to the Eurozone. And we have responded correctly to all the crises we have been through – the financial crisis, the migration crisis, the COVID pandemic, the Ukraine crisis. This explains why this Europe has become, as Jean-Claude Juncker said, political, and as Ursula von der Leyen says, geopolitical.' He refers to Jean Monnet, who said that Europe was a response to crises. 'We are Europeans and the EU is a good machine for solving crises. It's less brilliant sometimes when it needs to have a vision. We need to combine both.'

For him, all these developments mean that the EU can now think about strategic autonomy and European sovereignty. 'Europe has power, a different power, of course, than some other global players such as China or the US, but power. Different because we are made up of various countries and we must respect our nations. Different also because our values are probably different. We are peaceful. Europe was made for reconciliation and peace. It's impossible, inconceivable, to see Europe as an aggressor.' He observes that as the EU, we must become stronger than we are today. 'Whether it is economically, including on the defence side, but also concerning climate change. We must be stronger to tackle artificial intelligence. We must be stronger, together, to address the various transformations that our countries, that our world is now enduring, whether in terms of demography, economy, greening, or geopolitics.'

We must be stronger, together, to address the various transformations that our countries, that our world is now enduring...

”

Pierre Moscovici served for six years as a Commissioner in the Juncker Commission. 'As Commissioner, I had a significant portfolio, overseeing economy, finance, customs, and tax, as well as two Directorates-General.' He specifies that this was at the end of the financial crisis and involved trying to find solutions for the crisis in Greece .

'I had three targets. The first one was to ensure that our deficit would be reduced, that our debt would be controlled as the treaties stipulated. But the idea we had with Jean-Claude Juncker was to introduce flexibility in the mechanism. That's what we did during this mandate. All the countries, except France, were out of Excessive Deficit Procedure - EDP, but France did that some two years after.' He is proud that the Commission never had to sanction a member state. 'Trying to incentivise, to convince, is much better than sanctioning. I was not a dove, but nor was I a hawk.'

Trying to incentivise, to convince, is much better than sanctioning.

”

The second target he mentions was increasing tax transparency and fighting tax fraud and tax evasion. 'As the G20 and the OECD incited us to do. In that role, I pushed forward some 20 directives, anti-tax fraud and tax evasion, which increased transparency, and also established a blacklist of tax havens, convincing a lot of external countries to improve their own rulings. We signed agreements of exchange of information with countries such as Switzerland,

Now we don't have any banking secrecy on the European continent, which represents major progress.

”

Monaco, Liechtenstein, San Marino, which was not done before. Now we don't have any banking secrecy on the European continent, which represents major progress.'

A third target was to be the negotiator for the European Commission regarding Greece. 'Jean-Claude Juncker and I, we tried to soften the way the Commission intervened with Greece, which was highly unpopular. Strong dialogues with some frugal member states. Just to mention a telling figure: my predecessor, who did a great job, went once to Greece during his mandate. I went 17 times because I thought that the debate or the dialogue between Greece and the Commission also had to be driven at a political level, dealing with ministers instead of sending heads of units to do so.' Looking back, he is proud Greece was able to get out of its programme in 2018. 'With success. I'm happy to see that this country is reducing its public debt, that it has regained traction on growth, and that it's now clearly recovering. A success story for Greece is a success story for Europe.' What also strikes him is that in Greece, the European Commission, and hence the EU, is seen as a friend. 'A friend asking for efforts to be made, but still a friend.'

Placing the citizen at the heart of it all

...external audit is key for any institution, whether it is national or international.

”

As Europe has evolved, so has, in his opinion, external audit between when he started in 1984 and now. 'My feeling and conviction is that external audit is key for any institution, whether it is national or international. As First President of the French Court of Audit, I've been also the external auditor of the UN, and of the World Trade Organisation. I could see from the inside how important it was, and is, to have an independent look from the outside. That is precisely our role. Not to decide - we are not decision makers - but to give an independent opinion that must be taken into account by the institutional organisations in order to improve decision-making and to improve also their capacity to act.' He then corrects himself: 'To be frank, I say "should" because I want to be frank on that: I am still a new Member, still in my observation phase, looking into the functioning and real processes in this house.'

...the citizens' knowledge, even about the mere existence of the ECA, is way too little.

”

However, Pierre Moscovici does not hold back. 'Yet one thing I know is that the knowledge of the citizens, even about the mere existence of the ECA is way too little. And there must be decisive improvements there. As far as I see, here we've got fantastic staff, high quality reports. However, the impact of what is done here is insufficient and we must absolutely improve that. That is key.' He recalls one of his focal points in his former function. 'For me it was always about impact. My strategy was to improve the impact and multiply it by three as far as we can measure it through social networks or the press. For the ECA its impact on the public opinion is not strong enough and must be reinforced.'

We are working for the citizens, all of us in Europe.

”

For Pierre Moscovici the added value of policy assessment is clear. 'I very much promoted this as First President. 'When I started in 2020, this activity represented some 3% of all the activities. When I left, we were at 12%. I incentivized that because I think having a broad-based approach, a horizontal approach, also shared with the academics is something very important for the public debate. We are working for the citizens, all of us in Europe.'

Pierre Moscovici refers to the transformation process he launched at the French Court of Audit to increase its visibility. 'I would mention two aspects here. The first one, and it is not a problem for the ECA because we already do that, was to publish all the reports of the French Court. Because until then, only one fifth of our reports were published.' He refers to the 180 reports delivered each year. 'Only 35 of them were published - more than the one report published per year when I started in 1984. As First President I decided that 100% of our reports were to be published, except those which were confidential, related to for example defence. This is a huge transformation and the role of the French Court of Audit in the French public debate is much higher now.'

”

[on his role as ECA rapporteur] In the end it is not my report but the report of this institution.

His second point relates to citizens. 'I established a citizen platform in which citizens were authorized to propose things of control, and it represents now some 5% of the French Court's programming, so deriving from citizens initiatives. I think this is something that could be strengthened here. I can tell you this works, with proposals on themes the French Court itself would not have picked up.' He gives examples relating to fighting tax fraud and tax evasion, the role of private audit offices inside the public administration, or equality between men and women.

Balancing vertical and horizontal responsibilities

...the ECA College is more a board than what we have at the French Court of Audit. (...) The point is to find the right balance with the staff: neither let staff do whatever they want nor redo everything that the staff have done.

”

While both the French Court and the ECA are external audit institutions, Pierre Moscovici is aware of the differences. 'I am still in the observing mode. But the French Court is a jurisdiction institute, with members themselves writing reports, supervising the audits, discussing them. Hence their investment is different than here. As ECA Members we have a different function, which is providing an impetus to the staff, trusting the staff, accompanying the staff, and finally deciding on what the staff have presented, sometimes correcting what the staff have done. Obviously, the ECA College is more a board than what we have at the French Court of Audit.' He realises that he needs to get used to a different approach. 'The point is to find the right balance with the staff: neither let staff do whatever they want nor redo everything that the staff have done.'

As an ECA Member, Pierre Moscovici serves in the 'Financing and administering the Union' audit chamber, which he thinks is an excellent chamber to start with. 'It spans the whole machine of the Union's finances from revenue to budgetary and financial-instrument management, the regularity of the recovery and resilience facility, the protection of the EU's financial interest, administrative and efficiency performance audits. I hope that with my background, which is also one of a politician and external auditor, I can help towards a better understanding of how public money is raised, allocated, and accounted for.' He explains that his preferred topic is revenue and budgetary management. 'But I'm equally drawn to performance frameworks in public policies. I hope that in the years to come I will gradually enter that.'

The only way to survive is to transform, if not you die. (...) we, as the ECA, must make and can make further progress on raising awareness of our publications.

”

Specifically, he is now reporting Member for Chapter 10 of the ECA's annual report (regularity of spending of the EU public administration). 'And I have started my first visits to the European Parliament for the audit of the management of buildings of EU institutions. Very interesting, also after the COVID crisis.' He adds that in the future he will also lead an audit on reporting on fraud in essential areas for the protection of the Union's financial interests.

Clearly Pierre Moscovici's interests also extend to the horizontal aspects of ECA management issues. He believes that institutions need to be open to change and innovation to provide added value. 'The only way to survive is to transform, if not you die.' Almost apologetically, he adds: 'You know, I'm not a diplomat. Perhaps something I say may be too blunt and I will learn how to be less blunt in the future. But I think we, as the ECA, must make and can make further progress on raising awareness of our publications. And particularly in France, where the role of the ECA - and that's one of my own roles - is perhaps not unknown but not very visible.'

Prioritising transparency to inspire the citizens' trust

When discussing key elements for supreme audit institutions (SAs), Pierre Moscovici is clear. 'The most important factors that underpin the credibility of a supreme audit institution are the clarity and practical relevance of the recommendations, as well as the ability to monitor this implementation. When I was First President, I was often told that the Court produces useful

In a period when illiberal democracy is a threat to all of us, it is even more important that these institutions are perceived by the European citizens as trusted third parties. (...) we are the anti-fake news office.

”

work. But it is really important to be listened to. This is why tracking the implementation of the recommendations is another important aspect, which I also worked on to improve as First President of the French Court. So that our work has an impact on the citizens.'

For Pierre Moscovici this clarity also includes transparency in the work of all the institutions, not only saying what you have done, but also how you have done it. 'Having started as a young junior auditor in 1984 and recently working as First President in 2025, I gradually became convinced, partly through my executive responsibilities, that supreme audit institutions are pillars of democracy! In a period when illiberal democracy is a threat to all of us, it is even more important that these institutions are perceived by the European citizens as trusted third parties. That's what I tried to do at the French Court of Audit. The ECA should play the same role as a trusted third-party vis-a-vis European public opinion. Because many people tell our citizens lies about the way money is collected, the way money is spent, about corruption. At the French Court I said that we are the anti-fake news office. In my opinion, the ECA must also play this role, increasing the impact on citizens. Because we live in a period of doubt, uncertainty, and fake news. And in that kind of period, you need to have SAIs as a reference.'

Proudly, he shares that the French Court of Audit is perceived today as the national public institution that people trust the most. '90% of the French people know the French Court, 70% have a good opinion of it; 45% trust the French Court, while only 20% trust the French Parliament or the President of France. It takes time but it is important to get there. It is something we must have on our mind.'

...I started as an auditor. I'm finishing as an auditor. I've come full circle. And I always was a European.

Ultimately, for Pierre Moscovici this focus on the citizens, showing them impact, independence, relevance and trustworthiness, will also be the focal point of his six-year mandate. 'It is always the citizen I want to reach. Being an ECA Member now I will try to bring something from my experience.' I want to dedicate the remaining time of my professional life to what I've done. And basically, I started as an auditor. I'm finishing as an auditor. I've come full circle. And I always was a European. And this institution, I think, is useful to Europe and first and foremost, I hope, also to Europe's citizens.'

© European Union, 2026. Source: ECA.

Cour des comptes
européenne



Cour des comptes
européenne



‘Professionalism is the key trademark of the ECA’

Interview with Daniel Caspary, ECA Member since 1 March 2026

By Gaston Moonen



On 1 March 2026 Daniel Caspary succeeded Klaus-Heiner Lehne, who had served as ECA Member, of which six years as ECA President, since 2014. The new ECA Member from Germany, who originates from Karlsruhe, has a background as economist and was active in politics before joining the ECA, having served as MEP since 2004 until starting as Member at the ECA. Besides sharing his first impressions, he underlines how important the sense of impact is to keep people motivated to make a difference in a democracy through their work, be it as ECA staff member or ECA Member.

Committed to make a difference

Most of the people I met are engaged and committed to delivering their part of the work for a free and democratic EU.

”

Engagement and commitment are words that regularly come up when speaking with the new ECA Member coming from Germany. Daniel Caspary is clearly eager to *engage* in his new tasks. ‘My first impression is that the ECA is a very lively institution, as I actually expected. Already in the first weeks I met many engaged and knowledgeable staff members, which makes me optimistic. Whether in the audit teams, or at management level, they have a clear dedication, showing great interest in the topics they work on. Through their work they clearly aim for impact on what is going on in the EU!’ He finds this very positive. ‘Most of the people I met are engaged and committed to delivering their part of the work for a free and democratic EU.’

It was not my plan to become a full-time politician. But I always wanted to make a difference.

”

Such engagement is well suited to Daniel Caspary who got involved in public policy making at a young age. 'I realised at an early stage that if you get engaged in politics in general, you can make a difference and improve situations.' He recalls his first success when he was serving as pupils' president during secondary school level. 'Then we fought for public transport during night hours. When I got in contact with the responsible people in the mayor's office and the city council... we made a difference. It was not my plan to become a full-time politician. But I always wanted to make a difference.'

...the contacts with the citizens always were and also in my new role as ECA Member they play a key role.

”

In his late twenties Daniel Caspary was elected Member of the European Parliament. 'When I first got elected there the plan was to stay – provided that voters and the party would elect me again - for two or three terms.' He explains that in his third term, while thinking about a change after that term, he got elected chair of the German Christian Democratic delegation in the EP. 'With this new responsibility there was a totally different possibility to make an impact. It increased my opportunities to spark change.' Not that it was always easy sailing as an MEP. 'If you want to be a dedicated and influential MEP, then you have to pay the price on another side.' He subsequently refers to what this implies for his family, being his wife, who also works, and their five children. 'In every job you have to be committed, you must not be demotivated if you see a problem, you must overcome it. And you have to find allies and partners in other people who know how to do it. This is exactly the same in the political life. You try to stay engaged. As in the European Parliament, it is the same here at the ECA. The one who is drafting reports also sometimes has to overcome hurdles, has to stay tough in getting the data and information.'

Daniel Caspary explains that as MEP he always appreciated the contacts with voters. 'If you want to serve people you have to know what citizens want you to do. So, the contacts with the citizens always were and also in my new role as ECA Member they play a key role. Because if things are done at the EU level and we don't tell the citizens, they don't know. The second thing is, having success in legislation is seeing that you can change things; perhaps in a few cases you can trigger *the* big change, but we can influence and bring change in many small things, in whatever role and situation.'

On average, in these 22 years as MEP, every week I saw at least one group from a school or university.

”

As MEP Daniel Caspary has been particularly active regarding trade related issues and he was happy that he could get them into the European Commission's working programmes. 'It is important to bring people together, talking about for example doubling the budget of the Erasmus programme or the introduction of Interrail tickets, now there for young people. So, some concrete changes.' Another aspect is that one tries to help citizens when they come with problems, suggesting ways to resolve them. 'Including reaching out to citizens, such as to the younger generation. On average, in these 22 years as MEP, every week I saw at least one group from a school or university. I met in total about 60 000 people in Brussels and Strasbourg in those 22 years – not mentioning all the citizens I met at all the different occasions in my home region. I think it is important to reach out, also because for many European citizens it is difficult to understand how the EU actually works and how helpful and useful the EU is in our daily life.'

A continuous effort to reach out

The European Parliament should be granted the right to introduce fully fledged inquiry committees...

”

While working as an MEP the new ECA Member had often stressed the importance of parliamentary oversight. When discussing whether he sees elements missing for the European Parliament to properly exercise such an oversight, Daniel Caspary is rather determined. 'What is definitely missing in the Parliament, as a democratic right in a democracy, is the right of initiative - because controlling the European institutions also could mean the possibility to remove or withdraw legislation. At the moment the European Parliament does not have such a possibility.' He sees another issue to consider for the next Treaty revision, whenever that may come. 'The European Parliament should be granted the right to introduce fully fledged inquiry committees,

Too many of the key players in the institutions and member states are not yet aware of the high-quality reports of the ECA ... ”

such as national parliaments have. This is also a very good tool to force the Commission to deliver more data and information to the Parliament.'

He thinks the ECA has a major role to play in supporting the Parliament in its oversight role. 'Cooperation between the Parliament and the ECA has improved a lot, in many areas, during the past few years. But I think this is a kind of Sisyphus work, and one of the main tasks ECA Members should focus on.' He underlines this should not only be a priority when new MEPs or Commissioners come in. 'It is also up to the 27 ECA Members to reach out to the decision makers in whatever institution and make them aware - at the right time - what findings and recommendations the ECA has published in concrete policy areas.' He underlines the great expertise, with a rich diversity in background and experience, of both the College of ECA Members and ECA staff members. 'My first impression, after three months in the house and after 22 years in the Parliament, is that the ECA officials have a great expertise in many different fields.' He believes that this diverse knowledge - on standards, audit techniques, economy, the political network, what is ongoing and which are the needs of the legislator - offers great opportunities for the ECA as an institution to deliver relevant work. 'Too many of the key players in the institutions and member states are not yet aware of the high-quality reports of the ECA and the opportunities that arise from this. Therefore, I think there is still a good possibility to increase the influence of this institution.'

He gives a concrete example on such an opportunity. 'We, as ECA Members, have as a task to make members of other institutions aware of our output at the right time. As ECA we try to set up our timetables so that our reports are published when they may be relevant for the legislators or other key players in the institutions. But MEPs get such a flood of information that if they get an ECA report, one or two years later - when a revision of the legislation is discussed and which we audited two years before - they are not aware that this evaluation is available and contains relevant information. Nevertheless, most likely key findings are still relevant.' He observes that he cannot change MEPs behaviour of not being aware of ECA reports. 'But I can change my action as an ECA Member to deliver relevant information at the right moment of decision-making.' With this he does not want to dismiss the ECA's efforts to plan the publication of its reports when they may have most relevance. 'But many things in politics you cannot predict. It is in fact an ongoing attempt to reach others.'

As to giving a value judgement, the ECA should stay on the safe side, be impartial and not adopt an unfounded position for which it can be attacked. ”

As ECA Member he wants to find the right balance between the overall responsibilities one has as a member of the ECA College and as reporting Member for a specific task. 'The 27 ECA Members are carrying together this overall responsibility. The other one is that we have to focus on specific areas. With my background, my idea is *not* to be engaged in exactly the same area as I used to be.' This means a focus on a relatively new area for him. 'Just recently I was on a mission in Romania with a team auditing coastal protection in relation to the human part of climate change. It was my first audit mission. Reading files is one thing and seeing what happens on the spot is something else. This is what no artificial intelligence will ever replace.' Participating in this audit mission further added to the picture he got of the ECA in recent years: 'That is the huge knowledge and dedication of the audit team. They know which questions to raise, where to dig deeper. I saw again the professionalism of this institution.'

Daniel Caspary is quite explicit on what distinguishes the ECA. 'Professionalism is the key trademark of the ECA. On the one hand, we need to know what the legislators wanted to achieve. On the other hand, we need to find the facts and the effects.' He considers both aspect essential to remain factual when looking at impact and avoid becoming political. 'As to giving a value judgement, the ECA should stay on the safe side, be impartial and not adopt an unfounded position for which it can be attacked. The more we stick to this - and I think this institution did a great job over the last few years - the higher is the trust that other institutions will have in our output.' He considers proper standards and audit evidence as the appropriate safeguards to keep such trust.

”

[on his role as ECA rapporteur] In the end it is not my report but the report of this institution.

Daniel Caspary, as Member of the Sustainable use of natural resources Audit Chamber, is responsible for an audit task relating to coastal erosion, which he took over from his predecessor Klaus-Heiner Lehne, is also reporting Member for an audit task relating to land protected areas and one relating to fisheries. 'I have a strong interest in various issues. One part is having the leading role as rapporteur, in close cooperation with the audit teams. The other is keeping the whole picture and the institution in mind.' For him this means keeping a thorough eye on what else goes on at the ECA. As rapporteur, he sees his role as promoter of the audit teams. 'They have the knowledge and the background. I can give them support with my background. But I also see my role as supporting them whenever they need to get information.' He does not want to overemphasize his role as rapporteur for a report. 'In the end it is not my report but the report of this institution.' He explains that no matter who is the rapporteur, he will try to reach out to people in his network to make them aware of an ECA report where needed. 'For example, I did so recently for a report on the Just Transition Fund, for which my colleague João Leão is the rapporteur.'

Aiming for impact that keeps people motivated

”

AI has the potential to help us speed up processes without losing quality.

For the future of public audit, the new ECA Member expects further substantial impact of the digitalisation of audit work, something he also highlighted in his communication to the European Parliament when nominated as ECA Member. 'I see huge chances of making life easier for all of those working in this institution. It can offer strong support in administrative tasks, great support for your email traffic. It can also substantially support the inclusion of even more evidence and more knowledge in the work of this institution. I am quite sure that Artificial Intelligence will never replace the human part of the tasks. You have to get a personal impression; you have to be on the spot. But it can be a great support for all of us in digging deeper, in gaining more time to think while more simple tasks are taken over by software.'

He thinks that at least for the next years AI should not and must not replace human judgement. 'It must not replace the critical review by ECA members, nor the critical review and gut feeling of experienced auditors.' He sees AI for audit as a possible game changer. 'AI has the potential to help us speed up processes without losing quality.' He underlines that providing judgement should remain with humans. 'At least for the next years every output of AI will have to be checked and double checked. All of us must not make the mistake to trust AI output and stop thinking. But of course, in collecting data, AI data processing offers great potential.'

Concerning risks related to EU's autonomy when it comes to AI in relation to Large Language Models, Daniel Caspary responds with his current function in mind. 'If you had asked me this at the beginning of this year, I would have answered you. But in my new role as ECA Member, I would only give an answer based on the evidence found.'


When discussing whether he has set, as new ECA Member, specific goals for his mandate of six years and for which he can be held accountable for, Daniel Caspary explains he has two key goals. 'First, that the external stakeholders have the conviction that, together with the 26 colleagues, I have made a difference. And the second one is that, after those six years, the officials working in this institution have the impression that I was part of moving things forward and helping and supporting the almost 1 000 people working for this institution, fulfilling their duties and performing their job. If I achieve both, then I think it is a success.'

As to where he would like to see the ECA in five years from now, he responds with prudence. 'For me it is too early to express myself on this. I am always

Now I am in the process of realising what I still have to understand, learn and be curious about.

”

sceptical if people arrive somewhere and think that they know everything after two weeks. Now I am in the process of realising what I still have to understand, learn and be curious about.' However, he is rather explicit about a task he sees for ECA Members: to reach out and use their network to further increase the impact of the ECA. 'I think this will also help keep our staff engaged and motivated, when they see that their work also has relevance and impact in decision-making of the EU. With my background, that is one of the assets I can give to this institution. And I will!'



Maintaining forward momentum on gender equality

By Annemie Turtelboom, ECA Member and Dean of Chamber II, the ECA audit directorate covering investment for cohesion, growth and inclusion



© titima157 / Depositphotos.com

In January 2026, the heads of 11 supreme audit institutions (SAIs) in Europe issued the Madrid Declaration – a commitment to continue to promote female leadership within their institutions. That declaration provided the inspiration for the ECA’s International Women’s Day event this year. Following on from that event, which she chaired, ECA Member and Equality Ambassador Annemie Turtelboom takes a look at the progress the institution is making on its own internal gender equality targets, as well as at developments and the work being done at SAIs around the world.

The Madrid Declaration and ECA actions

In January of this year, the heads of 11 supreme audit institutions gathered in Madrid. The occasion? With the accession of Lithuania’s Irena Segalovičienė and Romania’s Mirela Călugăreanu to their respective offices¹, the proportion of female heads of SAIs in the EU, together with Ukraine, had reached 40% (11 out of 28). To commemorate this moment, the Spanish Court of Audit and the Portuguese Court of Auditors arranged a meeting of the female heads of SAIs. They followed this up by issuing the Madrid Declaration² – a commitment to continue to promote female leadership within their institutions.

The Madrid Declaration provided the inspiration for ECA’s own event to mark International Women’s Day a few months later, on 8 March. The ECA held an online event, which I moderated in my capacity as ECA Equality Ambassador, to discuss the Madrid Declaration with two heads of SAIs – Irena Segalovičienė and Belgium’s Hilde François. ECA President Tony Murphy and Diversity and Inclusion Officer Olga Ioannidou gave introductory remarks. The event was a useful exercise – not only for informing our staff about the work being done at the ECA and by our colleagues at national SAIs, but also for gathering the views of our staff themselves. During

¹ EUROSAI Highlights, [Newsletter 01/2025](#), October 2025.

² [Madrid Declaration](#), High-Level Meeting of Women Heads of Supreme Audit Institutions of the European Union and Ukraine, January 2026.

our Q&A session, ECA staff gave their opinions on the issues they found most important: translating rhetoric into concrete action; gender quotas; diversity in audit evaluations; and much more.

The Madrid Declaration reminds us that, as per Article 8 of the Treaty on the Functioning of the European Union, the EU is committed to eliminating inequalities, and promoting equality, between men and women – including within its institutions. In 2018, the ECA set itself three targets (see **Box 1**).

By the end of 2025, all three had either been met or were close to being met (40%, 39% and 37%, respectively). Other actions taken at the ECA include issuing notices to our directors and principal managers to ask them to encourage women to become heads of task, and organising informal sessions for experienced female heads of task to share that experience with other women. But as the mathematically inclined among you will have noticed, 40% is not even close to parity, especially when you consider that 54% of the ECA's staff are female³. While the ECA is meeting its targets – and indeed the national

SAIs are meeting theirs – is a positive development, it is only one step forward. We still have a long way to go.

Box 1 - ECA targets for the share of women in head of task and management positions

At a meeting of its Members on 15 November 2018, the ECA set three targets to be achieved by the end of 2027, each of 40%, for the share of women:

- acting as heads of task;
- in middle and senior management positions in audit; and
- in middle and senior management positions in non-audit functions.

International developments

The picture at international level is changing. The INTOSAI Development Initiative notes, in the most recent edition of its [Global SAI Stocktaking Report](#)⁴, that 60% of SAIs have some kind of action on gender equality at an institutional level. A total of 40% state that their strategic plans promote gender equality at an institutional level, while 37% say they address it in their HR strategies. These proportions have generally increased since the last time the survey was carried out in 2020. In addition, 31% say they have conducted at least one gender audit, and 21% say they have mainstreamed gender into their audits – one of the principles of the Madrid Declaration. Those two terms are thrown around a lot, but it's worth taking a moment to examine what they mean (see **Table 1**).

Institutionalisation of gender equality into organisations, including in their policies, programmes, projects and/or provision of services, structures, proceedings and budgets.

The ECA has a working group on gender-responsive and inclusive reports, which will continue to monitor activities in this area. At the time of writing, a follow-up to the ECA's International Women's Day event was planned for May 2026, when a delegation from the Spanish Court of Audit will visit the ECA to discuss their methods of promoting gender equality.

Maintaining the momentum behind gender equality action

The Madrid Declaration, and the ECA's event based around it, has given us a good opportunity to look at the current state of play regarding gender equality in audit, and at developments and trends in this area, both at the ECA and beyond.

Gender equality at an institutional level is not, perhaps, the most directly personal way we encounter it in our everyday lives. But it is one area in which we, as

institutions, can – and must – take action to improve. For an audit institution like the ECA, that means recognising the importance of equality, not just in our organisational structure, but also in our reports, through techniques such as gender audits and gender mainstreaming. The Madrid Declaration can give us fresh impetus to make sure we do not rest on our laurels, but instead continue to work restlessly for improvement.


³ ECA social balance sheet for 2025.

⁴ INTOSAI Development Initiative, [Global SAI Stocktaking Report 2023](#).

Table 1 – Key terms around gender actions and audit

<p>Gender audit</p>	<p>A gender audit is a tool to assess and check the institutionalisation of gender equality into organisations, including in their policies, programmes, projects and/or provision of services, structures, proceedings and budgets.</p>
<p>Gender mainstreaming</p>	<p>Gender mainstreaming means incorporating gender responsiveness and inclusion into audit reports. An internal ECA analysis completed in 2024 found that different audit institutions had different practices. Spain’s SAI, for example, stated that it systematically checked all audits from a gender perspective, as required by national law 3/2007. In its 2022 annual activity report, the Spanish SAI reported that 25 of its reports that year (51 %) had included a gender equality analysis. For some European SAIs, it depended on the topic, while others did not include any such analysis at all. Other means used to varying degrees by European SAIs in their reports or communications included: using gender-responsive and inclusive language, covering gender and inclusion topics, and providing sex-disaggregated data.</p>

Source: European Institute for Gender Equality, Gender Audit, 2019, and internal ECA analysis completed in 2024.



Our outreach work on the EU's next seven-year budget – 12 ECA opinions published this winter

By Damijan Fišer, Directorate of the Presidency



Photo: © butenkow / stockadobe.com

This winter, the ECA published a dozen opinions on the new ways proposed by the European Commission for financing EU policies and programmes from 2028 to 2034. In these opinions, the auditors give their views on a wide range of areas, ranging from competitiveness, research and culture to cohesion, agriculture, and international support. However, publication is not the end of the process. This spring, the ECA has presented a number of the opinions to the EU's legislators, the European Parliament and the Council of the European Union, as they embark on negotiating the outcome of these proposals. This is a golden opportunity, a once-in-a-decade chance, to ensure that the EU learns lessons from past budgets and ensure that the next one starts off on a sound financial footing. As the EU's financial watchdog, the ECA has issued several warnings to EU policymakers, highlighting risks and challenges. How and where? Damijan Fišer, an ECA senior communications officer and deputy spokesperson, explains the intense publication schedule of our opinions between January and March 2026, and the ECA's work in the following months.

MFF – what's in a name?

The EU's long-term budget, technically known as the multiannual financial framework (MFF), provides financing for programmes in all EU policy areas, including

agriculture, regional policy, new technologies, migration and defence. Traditionally, long-term budgets have provided a stable framework with overall spending limits;

this approach aligns spending with the EU's political priorities, increases the predictability of EU finances, and ensures budgetary discipline. Ultimately, the long-term

budget also makes it easier to agree on the annual EU budget. The EU's long-term budgets typically cover seven-year periods; the next one will run from 2028 to 2034.

The Commission's proposal

In July and September 2025, the European Commission made several legislative proposals (see **Figure 1**). To start with, it proposed a financial allocation of almost €2 trillion, an increase of 59 % compared to the current 2021-2027 budget of €1.2 trillion. As a result, annual national contributions to the budget would increase by 81 %, to €235 billion.

To finance the bloc's policies, the EU's executive proposed that own resources (the main sources of revenue for the EU budget should be increased from four to nine. These include new resources based on non-collected e-waste, tobacco excise duties, and a corporate resource

for Europe. At the same time, it proposed a marked decline – of 20 % – in the proportion of EU funding to be implemented together with the member states. It also suggested a sizeable new European Fund of €865 billion for cohesion and agriculture, centred around a single national and regional partnership plan, and a substantial increase in funding to strengthen the EU's defence-industrial base and enhance its defence capabilities. In addition, there would be a major shift towards financing not linked to costs, and an option for member states to finance their plans through repayable EU loans of up to €150 billion, which is a significant novelty at this scale.

The ECA's views – in a nutshell

At the end of April 2026, we published a [document](#) that summarised our concerns. In it, we outlined our position on the legislative package as a whole, bringing together the main messages from our individual opinions that centre around eight key thematic areas such as accountability, simplification and EU added value. At that time, ECA President Tony Murphy stressed that the legislative proposals for the EU's next multi-year budget showed that this was not business as usual, but a major overhaul. Many of the changes proposed were no guarantee of better spending in the future, he added. Indeed, as certain parts of the proposed arrangements would fundamentally change the way EU spending is planned, managed and scrutinised, we reiterated our warnings of risks to sound financial management and called for stronger safeguards.

In particular, we warned that if the new revenue streams were not approved, there would be a significant budget shortfall, meaning that member states' contributions based on their relative wealth (gross national income) would have to be increased, or the budget's ambition

scaled back. In addition, we noted the large increase in EU debt that would result from the proposed borrowing. In terms of spending, merging different policies could compromise the achievement of their objectives and require trade-offs between priorities. For large parts of the budget, spending priorities would be in the hands of member states with diverging interests.

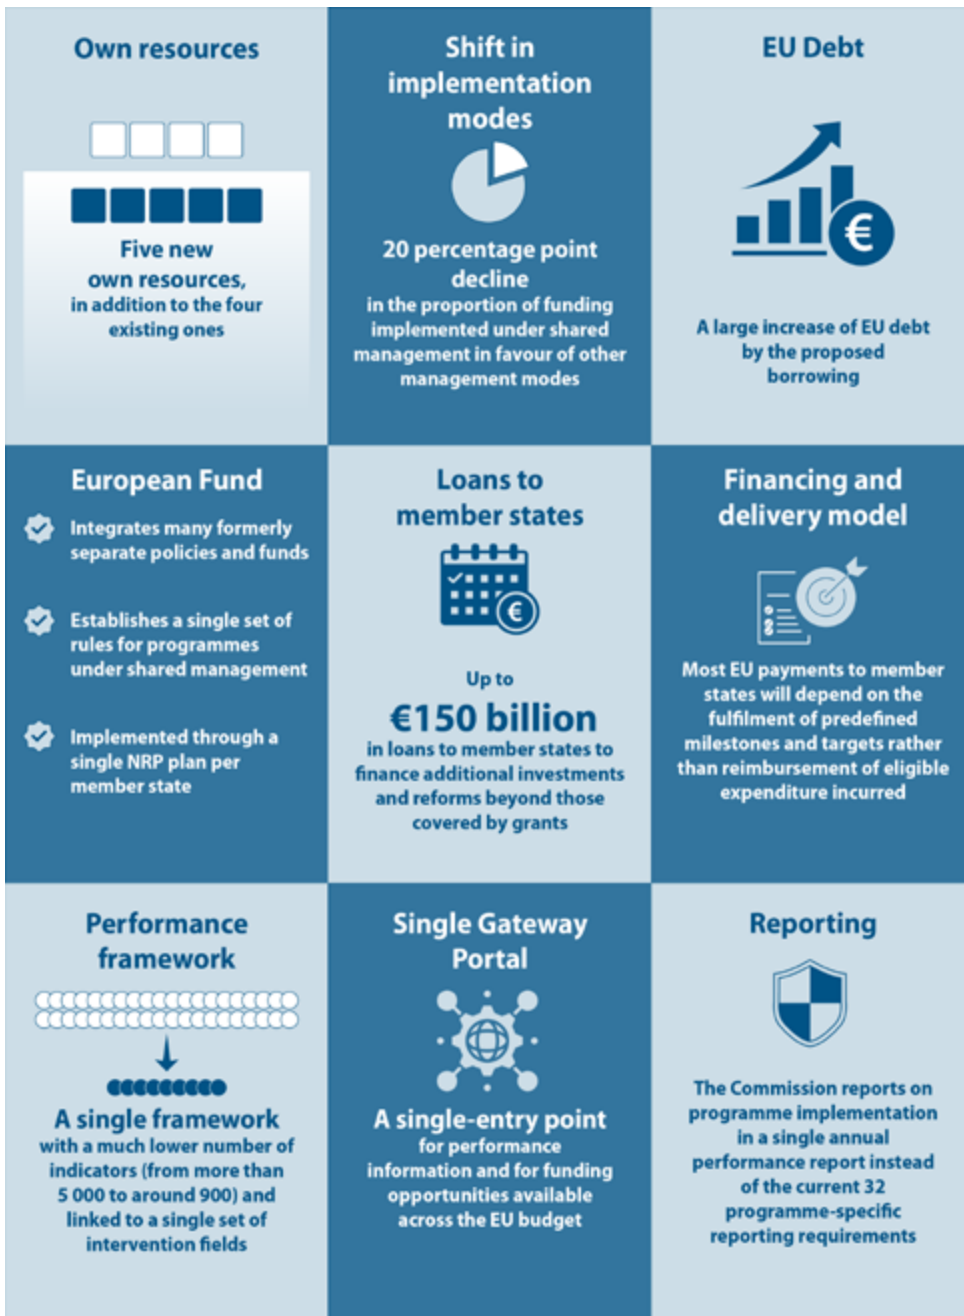
Furthermore, greater flexibility should not mean spending more money without securing more effective outcomes. The proposed performance framework suffers from weak design, which does not allow one to measure what results EU spending has achieved and what EU citizens are ultimately getting for their money. At the same time, arrangements for providing assurance that EU money is being spent soundly are, for significant parts of the budget, too reliant on what is often weak oversight by member states. Lastly, we stressed that the proposals were not clear enough in giving the ECA's auditors the unrestricted right to access the information they need to discharge their mandate.

Informing the media

Throughout the year, we publish audit reports on a wide range of EU policies, programmes and projects. We regularly present these reports – which are the core of our work – to the media through press briefings. Although they are more technical in nature than our usual reports, our opinions on the Commission's future budget proposals also received considerable interest from the press. In hindsight, this was perhaps unsurprising: the negotiations of the EU's long-term budget are one of the key points in the news agenda for the EU-focused media

this year. We held a couple of press briefings with our reporting Members: Iliana Ivanova hosted the briefing on the Common Agricultural Programme (CAP) opinion, and Annemie Turtelboom and Alejandro Blanco Fernandez hosted a briefing on the opinion on the European Fund. Our Members also held a number of interviews on other opinions. The media coverage varied from one opinion to another: the opinion on the CAP attracted the largest number of media articles across the EU.

Figure 1 – Main changes proposed by the Commission (2028-2034 MFF)



© European Union source: ECA. Based on the 2028-2034 MFF legislative proposals.

Discussing with the experts and stakeholders

Some opinions served as a basis for our discussions with experts and stakeholders. For example, the European Fund largely replicates the delivery model of the EU's pandemic recovery fund (the Recovery and Resilience Facility), and our audit work has repeatedly highlighted lessons to be learned from the RRF to enhance the performance orientation, accountability and transparency of similar instruments in future.

This spring, the ECA held two online discussions about lessons learned from the RRF for the future national and regional partnership plans that underlie the European Fund, one in late April about whether reforms should be linked to EU budget funding, and another in early June on the future of cohesion and agricultural policies. Both were organised by ECA Member Ivana Maletić in cooperation with ECA Members Alejandro Blanco Fernandez and Iliana

Ivanova, who – together with guest speakers from think tanks, academia, the Commission and other EU bodies – interacted with several hundred participants from EU institutions and member state authorities.

Our Members also brought our messages to the attention of politicians and authorities in member states by organising high-level discussions and stakeholder events, for example in Spain and Croatia.

Presentations to EU policy makers

Under the Treaty, the ECA is required to give opinions on these proposals (**Box 1**). The opinions were also requested directly by the European Parliament and the Council of the European Union. These institutions also invited us to present some of the opinions to their committees and working parties. We made 9 presentations of individual opinions at the Council of the European Union, and 17 at the European Parliament.

At the end of March, ECA President Tony Murphy presented an overview of our opinions to the joint meeting of the European Parliament's Committee on Budgetary Control and its Committee on Budgets, and ECA Member Hans Lindblad presented our opinion on the Single Market and Customs Programme to the Parliament's Committee on the Internal Market and Consumer Protection. Our Members presented further opinions to the national ministers and MEPs, while our auditors presented a number of opinions to working parties at the Council, including the opinions on Horizon Europe, the European Competitiveness Fund, the European Fund, expenditure tracking, the CAP, Global Europe, the Union Civil Protection Mechanism and the single market.

In May 2026, having been involved in drafting the overview document and in communicating many of our opinions to the press, I had the pleasure of presenting our overview to the Council's ad-hoc working party on the MFF in May, following an invitation from the Cyprus Presidency of the Council of the EU. If we have one overarching message, I stressed, it would be that making sweeping changes to the EU budget does not necessarily

make it better.

While our opinions are not legally binding, they do help promote sound financial management, accountability, and simplification in future programmes. As the European Parliament and Council have already begun to consider the proposals, we hope our input into the process will be useful.

And this may not be the auditors' last word on this topic. In autumn this year, we are planning a conference on the MFF, based on our review of lessons for the next EU budget. To find out the date when it is confirmed, please follow us on social media (@EUauditors).

Box 1 – ECA opinions on the 2028-2034 MFF

The ECA's 12 opinions on the proposed 2028-2034 EU budget, with links to the complete documents and the related press releases, both in 24 EU languages, are available in a **[dedicated section on the ECA's website](#)**.



ECA publications from January to June 2026

Audit reports, reviews and opinions

Opinions 01 - 12/2026

Published between
12/01 - 12/03/2026



Opinions on 2028–2034 MFF Proposals

On 16 July 2025, the European Commission put forward its first set of legislative proposals for the 2028-2034 Multiannual Financial Framework (MFF), followed by a second set of sectoral proposals on 3 September 2025. From mid-January to mid-March 2026, we issued twelve individual opinions in response to requests from the European Parliament and/or the Council of the European Union. We provided views on a wide array of areas of EU action, ranging from competitiveness, research and culture, to cohesion, agriculture, and international support.

[See all opinions on 2028-2034 MFF proposals](#)

Opinion 01/2026 proposal for a regulation on establishing the **European Competitiveness Fund**

Opinion 02/2026 proposal for a regulation on establishing **Horizon Europe**

Opinion 03/2026 Proposal for a regulation laying down the **Multiannual Financial Framework**

Opinion 04/2026 on Proposal for a decision on the system of **own resources of the European Union**

Opinion 05/2026 on proposal for a regulation establishing the conditions for the implementation of the Union support to the **Common Agriculture Policy** and a proposal for a regulation as regards the school fruit, vegetables and milk scheme ('**EU school scheme**'), sectoral interventions, [...], rules on availability of supplies in time of emergencies and severe crisis and securities

Opinion 06/2026 on proposal for a regulation on the **Union Civil Protection Mechanism and Union support for health emergency preparedness and response**

Opinion 07/2026 on proposal for a regulation establishing **Global Europe**

Opinion 08/2026 on proposal for a regulation establishing the **Single Market and Customs Programme 2028-2034**

Opinion 09/2026 on proposal for a regulation establishing the **European Fund for economic, social and territorial cohesion, agriculture and rural, fisheries and maritime, prosperity and security**

Opinion 10/2026 on proposal for a regulation establishing a **budget expenditure tracking and performance framework and other horizontal rules**

Opinion 11/2026 on proposal for a regulation of the European Parliament and of the Council establishing the **AgoraEU** programme for the period 2028-2034

Opinion 12/2026 on Proposal for a regulation of the European Parliament and of the Council establishing the **Erasmus+** programme for the period 2028-2034



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 01/2026

Published on 14/01/2026



Control systems for olive oil in the EU

The EU is the world's leading olive oil producer, consumer and exporter. The Commission and member states have put in place control systems to ensure that olive oil sold in the EU is genuine (i.e. it corresponds to the declared olive oil category), safe to consume, and can be traced back to its origin. We examined control systems' effectiveness, and whether the Commission oversees these mechanisms in member states and provides support. We found that there is a comprehensive EU legal framework for checks on olive oil, but member states apply it unevenly. We recommend that the Commission strengthen its oversight, clarify and improve certain rules and requirements, and support member states to improve the traceability of olive oil.

[See our special report](#)

Special report 02/2026

Published on 19/01/2026



EU transport infrastructure

Megaprojects are key to the completion of the EU trans-European transport network. In 2020, we published a special report showing major delays, cost increases, weak coordination between member states, and weaknesses in the Commission's oversight. This report provides an update, taking into account developments since then. We observed a further increase in the combined cost of the megaprojects, mainly driven by two of them, and additional delays which imply that the EU core network will not be completed by the 2030 deadline. In 2024, new legal provisions were introduced with the potential to improve the Commission's oversight of the implementation of the network, although the changes will mostly be relevant for projects that started later than the megaprojects we audited.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 03/2026

Published on 26/01/2026



Specific measures to support agriculture in the EU outermost regions

Since 1990, the EU has had a specific fund to support agriculture in its outermost regions, which also aims to guarantee the supply of essential agricultural products. We assessed whether this fund addressed the needs and constraints of the outermost regions. We found that it contributed to maintaining the competitiveness of some traditional sectors and supported the diversification of agriculture, though with modest results, while neglecting sustainability considerations. Finally, we found that it partially compensated for additional import costs, but it was difficult to assess whether these benefits were passed on to end users. We recommend re-examining support for traditional agricultural activities, enhancing crop and livestock diversification, and improving the assessment of whether the benefits of support for imports are reaching end users.

[See our special report](#)

Special report 04/2026

Published on 02/02/2026



Critical raw materials for the energy transition

For a successful energy transition, the EU requires increasing amounts of critical raw materials. We assessed measures to secure their supply, such as diversifying imports, increasing domestic production and improving resource management. We found that the EU faces an array of challenges. While the legislation sets a strategic course, its targets lack justification. Import diversification has not produced tangible results and bottlenecks hinder production and recycling. Despite faster permitting, many strategic projects will struggle to secure supply by 2030. We recommend that the Commission strengthen the foundations of the EU's raw materials policy, ensure that diversification efforts lead to more secure supply, address financing bottlenecks, make better use of sustainable resource management and increase the added value of strategic projects.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 05/2026

Published on 04/02/2026



Sustainable commuting around urban areas

Three quarters of the EU's population live in urban areas, which attract jobs and economic activity, generating heavy commuting traffic. The EU's urban mobility policy promotes sustainable transport, with legislation, guidance and funding. The EU's legal framework has been strengthened recently, requiring 431 cities to adopt sustainable urban mobility plans. However, we found shortcomings in the plans we audited, including gaps in their coverage of commuter flows and limited ambition to get commuters out of their cars. Future improvements in urban mobility will depend largely on local action. We recommend enhancing guidance and monitoring, promoting comprehensive coverage of commuter flows in plans, and establishing a methodology to measure changes in greenhouse gas emissions resulting from implemented projects.

[See our special report](#)

Special report 06/2026

Published on 11/02/2026



Tackling fraud in the RRF

Due to the size and delivery model of the RRF it is important that the risk of fraud is well managed. We therefore examined the RRF's anti-fraud systems, at both Commission and member state level. We found that while the Commission has taken steps to improve the RRF's anti-fraud framework, there are still weaknesses. These include lack of detailed anti-fraud requirements, insufficient anti-fraud measures in some member states, and incomplete data on suspected fraud cases. These weaknesses increase the risk that EU funds may be lost to fraud. We recommend that the Commission provides more detailed guidance on tackling RRF fraud, and strengthens its own audit and verification processes. We also recommend that, for future similar instruments, detailed requirements are included in the regulations from the outset.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 07/2026

Published on 19/02/2026



Europe's Beating Cancer Plan

Each year, cancer causes nearly 1.1 million deaths and generates over €100 billion in costs in the EU. The Commission's latest response is Europe's Beating Cancer Plan (EBCP), a €4 billion initiative launched in 2021 and covering prevention, early detection, cancer care and quality of life. In this audit, we examined whether the EBCP is an effective EU-wide response. We looked at its design, implementation, sustainability and monitoring arrangements, with a particular focus on cancer inequalities. We found that it supports coordination among member states, but its effectiveness could be compromised by overlapping initiatives, sustainability concerns, monitoring weaknesses and an unclear outlook after 2027. We recommend that the Commission adjust key EBCP initiatives and strengthen its monitoring and evaluation.

[See our special report](#)

Special report 09/2026

Published on 26/02/2026



European innovation partnership in the common agricultural policy

The European innovation partnership for agricultural productivity and sustainability is a key policy tool intended to foster bottom-up, collaborative innovation. We assessed whether the innovation partnership in the common agricultural policy has effectively contributed to innovation in EU agriculture. We found that it is falling short of its full potential due to an insufficient focus on farmers' needs, a selection process that failed to prioritise projects with innovation potential, and weak dissemination of project results. To address these shortcomings, we recommend improving the partnership's focus on farmers' innovation needs, strengthening project selection procedures and enhancing the dissemination of results.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 08/2026

Published on 04/03/2026



International nuclear safety cooperation

We examined whether the Commission, together with the European External Action Service, has been effective in enhancing nuclear safety in non-EU countries. Overall, we conclude that the Commission remains an important player in international nuclear safety cooperation, having helped deliver a wide range of actions, some of which are large and complex. However, these achievements often came later and sometimes at a higher cost than initially planned. For some, long-term sustainability remains a significant challenge. The Commission's effectiveness is also hampered by the lack of a comprehensive and up-to-date strategy, robust prioritisation of proposals, and shortcomings in the monitoring of the financed actions. We put forward recommendations to address these areas.

[See our special report](#)

Special report 10/2026

Published on 09/03/2026



Energy communities

Citizens are projected to generate half of the renewable energy necessary to achieve the EU's goal of climate neutrality. Energy communities are legal entities that empower citizens, small businesses and local authorities to produce, manage, share and consume their own energy. Communities can contribute to the energy transition, and to increasing affordability and citizen involvement. The EU has only reached 27 % of its objective of having at least one energy community per municipality with more than 10 000 inhabitants by 2025. This objective lacks relevance, support and monitoring. EU definitions remain unclear about the participation of apartment owners' associations. Moreover, governments have not created the conditions needed to support energy communities, such as incentives for electricity storage to ease grid connections. Our recommendations address these shortcomings.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 11/2026

Published on 19/03/2026

Innovation Fund



The Innovation Fund supports targeted investments in innovative technologies with the potential to reduce greenhouse gas emissions. We carried out this audit to assess whether the Innovation Fund is helping to scale up demonstration projects which contribute effectively to the EU's decarbonisation objectives. We found that the Innovation Fund promotes large-scale innovative technologies, but it is not delivering the expected level of reductions in greenhouse gas emissions. We also found that the level of financial resources is uncertain, there is no structured analysis underlying their allocation, and projects face delays and cancellations. We recommend that the Commission should establish a structured analysis to guide resource allocation, examine additional measures capable of accelerating the deployment of funds, and improve project assessment.

[See our special report](#)

Special report 13/2026

Published on 25/03/2026

Single market for services



Businesses still face considerable barriers when providing services in another EU member state. Focusing on the EU Services Directive, this audit assessed whether the Commission was effective in removing barriers to cross-border services hindering economic growth and development in the EU. We found that Commission action in this field remains insufficient. We recommend that the Commission develop a clearer and more ambitious strategy for the single market for services, use the European Semester more actively and better incentivise member states to carry out necessary reforms, it should clarify legislation, focus enforcement of infringements of rules on cases with considerable impact, reinforce tools facilitating cross-border services, and monitor and evaluate progress in completing the single market for services.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 12/2026

Published on 26/03/2026



The EU's advisory bodies

Within the EU institutional framework, the European Economic and Social Committee and the European Committee of the Regions provide opinions on proposed EU legislation. The former represents employers, workers and civil society organisations, while the latter represents elected local and regional authorities, thereby both supporting participatory governance. We examined whether these Committees produce relevant, timely and high-quality opinions and effectively assess their impact. We found that both Committees have quality management procedures that support the production of their opinions, but there are gaps for the selection of experts. Delays in producing the opinions may sometimes reduce their influence on decision-making, and monitoring systems emphasise outputs and visibility rather than systematically assessing whether opinions influence draft legislation.

[See our special report](#)

Opinion 13/2026

Published on 14/04/2026



Opinion concerning the proposal for a regulation of the European Parliament and of the Council establishing the Temporary Decarbonisation Fund

This opinion was issued pursuant to Article 322(1) TFEU which requires the European Court of Auditors to be consulted before the adoption of legislative acts relating to the EU's financial rules.

The purpose of this opinion is to comment on the design, governance and funding provisions of the proposed Temporary Decarbonisation Fund. It is intended to help ensure that the future fund promotes sound financial management and contributes to the objectives of the EU's climate and industrial policies.

[See our opinion](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 14/2026

Published on 06/05/2026



RRF traceability and transparency

The €577 billion Recovery and Resilience Facility (RRF) was established in February 2021 in response to the COVID-19 pandemic. We audited whether the Commission and member states ensured sufficient traceability and transparency of RRF funds. We found that while most member states can trace RRF funds, some data is not always collected in a systematic way, resulting to delays and sometimes incomplete information. Actual cost data is not consistently used by member states to update estimates or requested by the Commission to manage the RRF. Regarding transparency, published information on recipients, actual costs and results achieved is insufficient. We recommend that future instruments provide for systematic collection, use, and publication of all relevant data on the use of EU funds.

[See our special report](#)

Special report 15/2026

Published on 04/06/2026



Cohesion policy support for youth employment

Bringing young people into sustainable employment has been one of the biggest employment challenges for many EU member states during the past decade. To address this challenge, the EU established a strategic framework, and allocated around €25 billion through cohesion policy since 2014. We found a wide range of EU-funded measures in member states supporting youth employment, but these lacked a focus on long-term labour market integration, and little information exists on their lasting effects. Inactive young people face particular barriers to entering the labour market and remain difficult to reach. We recommend enhancing the monitoring and evaluation of cohesion policy funding for youth employment, improving the effectiveness of hiring incentives, and providing more targeted support to inactive young people.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 16/2026

Published on 09/06/2026



Western Balkans Investment Framework

The Western Balkan countries committed to completing the core EU trans-European network by 2030. We examined the effectiveness of EU support through the Western Balkans Investment Framework for transport infrastructure in the Western Balkans, focusing on connecting the region to the EU's core transport network by 2030. We found that while the audited projects match the stated connectivity priorities, progress is slowed by immature projects being selected and shortcomings in the supervision by financial institutions. Monitoring, reporting and EU visibility were also insufficient, with delays and sustainability issues persisting. Given the implementation delays and operational challenges, we concluded that the 2030 deadline is unlikely to be met. We recommend that the Commission should improve the selection, monitoring, sustainability and visibility of projects.

[See our special report](#)

Special report 18/2026

Published on 22/06/2026



The Commission's anti-fraud strategy

Fraud risks affecting the EU budget continue to evolve, requiring a robust and coordinated anti-fraud framework across the Commission. This report assessed whether the Commission's anti-fraud strategy is appropriately designed to manage fraud risks and support consistent implementation. While the strategy broadly follows recognised fraud risk management principles and covers the full anti-fraud cycle, we found weaknesses in risk assessment, action planning, monitoring and reporting. In particular, risk assessments do not systematically use all relevant information sources or sufficiently consider emerging risks, actions are often not sufficiently ambitious or lack measurable outcomes and clear timelines, and monitoring focuses more on implementation than on results or impact. We therefore make recommendations to strengthen oversight, risk assessments, action plans, and reporting.

[See our special report](#)



ECA publications from January to June 2026

Audit reports, reviews and opinions

Special report 17/2026

Published on 25/06/2026

The Union Civil Protection Mechanism's rescEU reserve



We examined the Commission's planning and implementation of projects aimed at developing and stockpiling rescEU reserves, as well as their deployment. Such rescEU reserves represent three quarters (€2.9 billion) of the Union Civil Protection Mechanism's 2021-2027 funding.

We conclude that the rescEU reserves provided a valuable contribution to European civil protection through the acquisition and leasing of different rescEU capacities and through the deployment of those capacities when a crisis so required. However, we found that weaknesses in planning the calls for proposals hampered project implementation and sustainability. In some instances, the procedures associated with claiming reimbursement for response actions were cumbersome. We put forward recommendations to address these areas.

[See our special report](#)

Next edition

EU efforts to combat fraud and corruption

Weeding out fraud and corruption from the complex public finance landscape is a challenging task – and a potentially hazardous one. It is a key responsibility of policy makers and public institutions to safeguard public funds.

But arguably, the responsibility of preserving trust in public institutions and public policymaking is even more important. There is thus a strong element of self-interest for public-sector bodies in preventing and combating corruption.

Public auditors in the EU, from EU to local level, are highly attentive to the risks of fraud and corruption. Citizens often expect supreme audit institutions (SAIs) to function as a sentinel in the fight against crimes against public integrity. Since their work is primarily focused on sound financial management, external auditors collaborate closely with anti-fraud bodies, investigative authorities and policymakers. In the course of this collaboration, they report irregularities, uncover systemic weaknesses, and help create strong safeguards against abuse.

Seven years have passed since we last covered this important topic. In our last exploration of this topic, we also highlighted the principles that are essential for reducing fraud and corruption: ethics and integrity. Since 2019, these two principals have grown even further in importance, also in the light of global developments, as policy initiatives and responsibilities have evolved. Both the EU's financial mechanisms and the ECA's collaboration with partner organisations, such as the EU's anti-fraud office (OLAF) and the European Public Prosecutor's Office (EPPO), have further developed. The Recovery and Resilience Facility (RRF) introduced a new financial-management model; EU and national responsibilities in the fight against fraud and corruption have become more clear-cut. In the meantime, the ECA has published

several reports on this theme, covering, for example, fraud in the RRF ([special report 06/26](#)), the EU transparency register ([special report 05/24](#)), the rule of law in the EU ([special report 03/24](#)) and blacklisting ([special report 11/22](#)). In the coming years, we plan to publish several related audits.

In our next edition, we will explore the evolving world of fraud and corruption in depth. We will examine how the EU and its member states are addressing the risks affecting this subject area, how audit institutions are adapting their approaches, and which tools and practices are proving most effective. From prevention strategies to real world cases, from institutional cooperation to emerging threats, we will explore various perspectives on action against fraud and corruption, as well as initiatives to strengthen integrity and the rule of law in an increasingly complex environment. strengthen integrity and the rule of law in an increasingly complex environment.



© European Union, source: ECA.

Highlights

- 7** The Architecture of Intelligence
- 17** Artificial intelligence and protecting knowledge systems in the European Union
- 24** 'I want this continent to become an AI Continent'
- 34** Advancing EU public administration readiness for AI Act compliance
- 38** 'I have no doubt: AI's advantages far outweigh all the negatives!'
- 48** The ECA's audit of the EU's microchips ambitions – examining an industrial ecosystem
- 75** From experimentation to integration: advancing AI in European public audit
- 84** When the rulebook arrives before the road is built
- 96** The curse of Big Tech: how digital giants are eroding the constitutional foundations of liberal democracy
- 101** 'For democracies legitimacy comes from protecting the rights of their citizens, also on AI...'

The contents of the interviews and the articles are the sole responsibility of the interviewees and authors and do not reflect the opinion of the European Court of Auditors.

Editor in chief: Gaston Moonen
Tel.: +352 4398 - 45716
E-mail: antonius.moonen@eca.europa.eu

Desktop publishing: Lucie Peterkova
Distribution: Directorate of the Presidency

© European Union, 2026



The European Court of Auditors' reuse policy is set out in Decision No 6-2019 on the reuse of documents.

Unless otherwise indicated (e.g. in individual copyright notices), ECA content owned by the EU is licensed under the Creative Commons Attribution 4.0 International (CC BY 4.0) licence.

As a general rule, therefore, reuse is authorised provided appropriate credit is given and any changes are indicated. Those reusing ECA content must not distort the original meaning or message. The ECA shall not be liable for any consequences of reuse. Additional permission must be obtained if specific content depicts identifiable private individuals, e.g. in pictures of ECA staff, or includes third-party works. To use or reproduce content that is not owned by the EU, it may be necessary to seek permission directly from the copyright holders. The EU does not own the following:

Cover image: © European Union, source ECA. Created with Copilot.



The contents of the interviews and the articles are the sole responsibility of the interviewees and authors and do not reflect the opinion of the European Court of Auditors.

For more information:
European Court of Auditors
12, rue Alcide De Gasperi
1615 Luxembourg,
LUXEMBOURG

eca-journal@eca.europa.eu
Published in June 2026

Past editions of the Journal
can be found on ECA's website:
eca.europa.eu/en/journal



@EUauditors
eca.europa.eu

