

AI COALITION 4 NL

POSITION PAPER

EUROPEAN AI INDUSTRY

**A EUROPEAN AI INDUSTRY CAN ONLY
BE BUILT ON COHERENT POLICY**

August 2026



AI is one of the most widely discussed topics of our time, and the debate is sharply polarised. On one side of the debate, it's a miracle technology that will tackle the greatest challenges of our age with superintelligent systems. On the other, a modern Frankenstein's monster that will inevitably turn against its own creators. Leaders of tech companies, heads of state and even the Pope are weighing in on the issue. What almost everyone shares is the assumption that AI is a general-purpose technology that, like the steam engine and the internet, will fundamentally change the economy and society. Moreover, the adoption and transformation driven by AI appear to be progressing significantly faster than those associated with previous general-purpose technologies.

The US and China regard technological superiority in AI as an important means of exercising influence and power on a global scale. In both countries, technology, capital and talent are readily available to support large-scale AI investment. The outcome could have a major impact on financial markets, where there is currently widespread speculation about which business models will work for AI.

Europe is attempting to chart its own course. While Europe possesses technology, capital, and talent, these assets are fragmented, making it harder for large-scale AI initiatives to get off the ground. Moreover, neither Europe nor the Netherlands has many large, well-capitalised digital companies capable of, and willing to, shoulder such investments themselves. The influential report¹ by Mario Draghi released in the summer of 2024 was intended to serve as a 'reset' for Europe. A growing awareness of dependence on other global powers, coupled with concerns about future competitiveness, spurred European policymakers into action. They accelerated efforts in an area where Europe has traditionally excelled: regulation. European and Dutch regulations were already rarely viewed as conducive to starting and scaling up businesses, and for AI companies, the phased implementation of the complex AI Act adds another layer of difficulty.

Europe is particularly finding a sound balance between the public and private sectors, carefully weighing economic and societal interests. Europe is explicitly unwilling to abandon its core values, emphasising fundamental human rights and protecting human dignity and autonomy. In the short term, this may result in falling behind, as is currently the case with AI development. But in the long run, it can lead to greater inclusive prosperity. However, a key condition is that finding the right balance does not devolve into a mere status quo and that sufficient attention is paid to innovation and change. Draghi's message was clear: we are prosperous now, but how can we remain so a decade from now? Especially against the backdrop of a disruptive technological transformation driven by AI.

The fact that a distinct European approach is desirable, and even necessary, is evident from the growing resistance to AI Big Tech. Even the leaders of these major AI companies acknowledge that unchecked growth and concentration of power will lead to problems.² Voluntary self-regulation is an inadequate response. This is where government has a role to play: creating a balanced and fair AI playing field without market dominance, while ensuring the responsible development and deployment of AI. This technology is too important and too far-reaching to be left to the leaders of a handful of major AI companies. The question, then, is: what should Europe do, and how can a relatively small country such as the Netherlands hold its own amid these enormous forces?

¹M. Draghi, 'The future of European competitiveness', september 2024.
https://commission.europa.eu/topics/competitiveness/draghi-report_en

²<https://www.economist.com/business/2026/07/14/demis-hassabis-has-a-plan-to-harness-ai-safely>

Exponential Acceleration: AI and Economic Power

Since ChatGPT first emerged in late 2022, increasingly advanced language models have followed one another at a breakneck pace; the most recent include Claude Mythos 5 and Fable 5, as well as the Chinese Kimi K3. This increase in intelligence requires an exponential increase in scale. GPT-1 had just over one hundred million parameters in 2018; by July 2026, Kimi K3 had reached 2.8 trillion.³ The volume of training data is also increasing: from 34 billion tokens for Google T5 in 2019 to more than 32 trillion for DeepSeek V4 in 2026. Training and deploying these models therefore requires enormous investment in data centres. McKinsey estimates global capital expenditure on AI data centres through to 2030 at around \$5.2 trillion. Data centres, in turn, require large amounts of energy. The International Energy Agency expects the electricity consumption of all data centres to roughly double between now and 2030, reaching around 950 TWh, almost three per cent of global consumption, with AI the main driver.

The largest investments in AI computing infrastructure and the development of the most advanced models are taking place in the US and China. In response, Europe is increasing its own investment: at the end of July 2026, the EU opened a tender for up to seven AI gigafactories, backed by up to €10 billion in public funding intended to unlock at least €20 billion in private investment, totalling more than €30 billion. This is on top of the existing network of nineteen AI factories.⁴ With Mistral Large 3 (675 billion parameters), Mistral is Europe's largest model developer. In the Netherlands, work is under way on GPT-NL, which is modest in scale at 26 billion parameters.

Alongside an understanding of the computing power and scale of AI models, a deeper understanding of AI applications is needed to determine how Europe and the Netherlands should position themselves. These applications are broad, ranging from software development and the automation of cognitive work through Agentic AI, to the design of proteins⁵, new materials⁶, complete industrial designs⁷ and humanoid robotics⁸. European companies are particularly well represented at the application layer, while companies in the capital- and compute-intensive layers beneath it are predominantly American and Chinese. AI applications are the most visible part, but underneath them lies a stack consisting of data centres and networks, AI chip technology, data storage and cloud infrastructure, critical raw materials, and foundation models.

³<https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers>

⁴Europese Commissie / EuroHPC JU, 'EU launches AI Gigafactories call', 30 juli 2026.
<https://digital-strategy.ec.europa.eu/en/news/eu-launches-ai-gigafactories-call-boost-europes-computing-capacity-and-unlock-more-eu30-billion>

⁵<https://www.cradle.bio/>

⁶<https://cusp.ai/>

⁷<https://www.prometheus.ai/>

⁸<https://www.1x.tech/>

Regulation alone is not industrial policy

On August 2, 2026, the AI Act entered a new phase. The transparency obligations under Article 50 have applied since then. Anyone providing an AI system that interacts directly with people must ensure that users know they are dealing with AI, unless this is already evident to a reasonably informed and attentive user. Providers of systems that generate synthetic text, images, audio or video must label that output with a machine-readable watermark and, in the case of deepfakes, with a visible logo. For systems already on the market before August 2, 2026, the transition period runs until December 2, 2026. On the same date, the European Commission⁹ was given enforcement authority over providers of general-purpose AI models. This gives the AI Office significant supervisory powers over frontier AI models, in addition to the ability to request documentation, require access to AI models and carry out inspections. The AI Office can also impose measures requiring providers to take risk-mitigation or corrective measures, and market restrictions may be imposed, meaning certain frontier AI systems could be denied access to the European market. All of this is backed by fines of up to three per cent of global annual turnover or €15 million for non-compliance.

At the same time, criticism of this legislation is mounting. This is of course coming from non-European Big Tech, but also from European businesses, where overly stringent regulation is seen as a risk to competitiveness. Six days earlier, on July 27, 2026, the Digital Omnibus on AI came into force.¹⁰ This pushes back the requirements for high-risk AI under Annex III to December 2, 2027, and for AI used in regulated products such as medical devices, even to August 2, 2028.

So in the same week, Europe began enforcing the rules while also postponing them. And although the Digital Omnibus is intended to simplify regulation, its practical effect is that businesses experience even greater regulatory complexity.

Such an erratic course is plainly wrong. It rewards inaction: companies that have invested in compliance are put at a disadvantage compared with those that have sat on their hands. This wait-and-see approach undermines legal certainty in European digital legislation and, as a result, discourages investment while moving the objectives of the regulation further out of reach. Moreover, the notion that regulation is inherently bad for innovation is also misleading.

Regulation was never the problem

The Porter Hypothesis¹¹ argues that well-designed regulation can actually stimulate innovation by encouraging companies to develop new solutions. Europe has already proven this hypothesis. In the 1980s, Brussels mandated a single mobile standard, GSM, and reserved the frequency range for it. That decision created a single market, within which Nokia and Ericsson grew into global leaders. This is regulation designed to drive innovation: a rule that creates a market.

The AI Act doesn't do this. Although the AI Act does pay attention to innovation with AI, its core purpose is to prevent societal harm and protect fundamental rights. It's a legitimate and valuable objective. But it's not industrial policy, nor was it ever intended to be.

⁹<https://digital-strategy.ec.europa.eu/en/policies/enforcement-ai-act>

¹⁰Verordening (EU) 2026/1744

¹¹https://en.wikipedia.org/wiki/Porter_hypothesis

The legislation cites innovation, adoption and the internal market as objectives in its recitals, and since August 2, 2026 has required Member States to have at least one operational regulatory sandbox. These will be managed by the national supervisory authorities responsible for enforcing the AI Act (in the Netherlands, the Dutch Data Protection Authority and the Digital Infrastructure Inspectorate). And although the AI Act explicitly gives regulators an innovation mandate as well, it's unreasonable to expect regulators to take responsibility for developing a European AI industry. Regulators will approach innovation primarily from the perspective of mitigating harm; they can remove barriers, provide clarity and make innovation easier, but driving European AI innovation cannot come from regulators. This creates a risk of mission drift.

The more fundamental problem is that Europe is not doing enough to build a strong domestic AI industry. The scope for innovation offered by the regulatory sandboxes established under the AI Act is limited. After decades of purchasing digital products from American suppliers, Europe has come to conflate regulation with industrial policy. The hope is that regulation can impose global standards: the Brussels Effect.¹² This worked with the GDPR, because personal data is subject to the law of the market in which it is processed. Computing power and frontier AI are different: those who own the chips, data centres and models cannot simply be disciplined through market access. The Brussels Effect may hold up for high-risk applications, but that is far from guaranteed for the fundamental layers underneath them. Meanwhile, new guidelines and amended legislation are increasing complexity and the regulatory burden on businesses.

Nobel Prize winner Philippe Aghion recently summed it up succinctly: the EU is a regulatory giant and a fiscal dwarf. While the United States and China have managed to combine industrial policy with competition policy, Europe has instead constrained industrial policy in the name of competition.

The European approach: AI Fair Tech

AI is a general-purpose technology with its own dynamics that will lead to fundamental changes. This requires adaptation on the part of citizens, businesses, governments and institutions, as well as changes to the research and innovation process itself. 'Value alignment' refers to the idea that AI models should conform to the values of the context in which they are used. But a technology that intervenes so rapidly and so deeply also requires the reverse: citizens, businesses, governments and institutions must adapt to what AI makes possible. We call this inverse alignment.¹³ If Europe is to benefit from greater inclusive prosperity, it must ensure both the effective integration of AI (value alignment) and the effective adaptation to AI (inverse alignment).

The AI industry Europe needs to achieve and absorb both forms of alignment effectively is what we call AI Fair Tech. It rests on three elements. It is lawful and ethical, with fundamental rights and human dignity as its baseline. It is sovereign, by reducing excessive dependencies, either by developing capabilities ourselves or by working with other power blocs. And it is resilient, by preparing citizens, businesses, governments and institutions for an innovation that will inevitably bring disruption.

The lawful and ethical dimension of AI Fair Tech is rooted in the European approach, beginning with a high-level AI expert group¹⁴ focused on fundamental rights and human dignity and resulting in the AI Act, which is intended to ensure the responsible development and deployment of AI. Its success will depend in part on the operational strength of the institutions entrusted with supervisory responsibilities, such as the AI Office and the national supervisory authorities.

¹²<https://scholarship.law.columbia.edu/books/232/>

¹³'What Humanity Needs To Flourish In The Next Decade', Noema Magazine. <https://www.noemamag.com/what-humanity-needs-to-flourish-in-the-next-decade/>

¹⁴<https://digital-strategy.ec.europa.eu/en/policies/expert-group-ai>

The second aspect of AI Fair Tech concerns European digital sovereignty. While the Draghi report already warned that Europe risked becoming a digital colony, the publication of the report *EuroStack: A European Alternative for Digital Sovereignty*¹⁵ in early 2025 made our dependencies tangible and laid out how they could be reduced. The report contains proposals for building sovereign European infrastructure and puts forward an extensive investment plan. It calls first for buying European, then for developing interoperable standards as a basis for global standards, and finally for pursuing targeted industrial policy.

A year and a half later, precious little has come of the latter point regarding the AI layer. Substantial public investment has failed to materialise, there is barely any European frontier AI infrastructure, and on top of the existing dependence on cloud and software markets, a new lock-in around OpenAI, Google and Anthropic is now emerging. Sovereignty requires a European position of its own that can help curb this dominance. Sovereignty does not mean that Europe must control every layer of the AI value chain itself; that is neither feasible nor desirable. It does mean, however, that Europe needs its own position in strategically important areas. AI infrastructure and frontier technology certainly belong in that category. Given the speed of developments, there is an urgent need to act.

The third aspect of AI Fair Tech is AI resilience, among citizens, governments, institutions and businesses. Resilience requires safe and trustworthy AI, protection of institutions against subversion, and reducing economic dependencies. Legislation alone cannot deliver this. Resilience requires a clear vision of inverse alignment and of how the AI transformation can be absorbed into our democratic rule of law and its institutions. At its core, this is about power and countervailing power. A Europe under strong influence from a small group of extraordinarily wealthy technology idealists is not a future that Europe should want. The presence of strong and effective institutions is also essential for sustainable prosperity.¹⁶

A distinctive European AI Fair Tech approach, based on the pillars of responsible, sovereign and resilient AI, is rooted in Europe's public-private tradition and consists of a coherent approach to regulation, innovation and investment. Regulation requires inverse alignment: alongside the current focus on prohibiting undesirable AI, there must also be an effort to promote AI that is desirable. Innovation is a shared responsibility of governments and businesses, which must jointly establish and fund priorities, including AI infrastructure, AI technology and AI talent. Investment can take many forms: grants, investment funds, direct corporate investment and, above all, procurement.

From vision to policy: inverse alignment in practice

The importance of inverse alignment in creating a European position in a general-purpose technology can hardly be overstated. Writing, the wheel, the steam engine and the internet have fundamentally transformed economies and societies. The emergence of new technologies is often accompanied by inflated expectations, disillusionment and, ultimately, stable adoption.¹⁷ But with general-purpose technologies such as AI, it goes further: their impact on society is so fundamental that a reassessment of norms and values, ultimately resulting in new regulation, is necessary.¹⁸ This requires a careful and continuous process in which a societal balance must be found around the adoption and embedding of the technology. It is inevitable that regulation will lag behind technological development, because general-purpose technologies affect so many aspects of the economy and society in such fundamental and diverse ways that it is impossible to foresee all their implications in advance. The current European AI Act is primarily aimed at curbing undesirable excesses of AI, rather than being driven by a vision of how regulation can promote the responsible adoption of AI.

¹⁵<https://www.bertelsmann-stiftung.de/en/publications/publication/did/eurostack-a-european-alternative-for-digital-sovereignty/>

¹⁶Acemoglu et al, *The Role of Institutions in Growth and Development*, 2008.

¹⁷https://en.wikipedia.org/wiki/Gartner_hype_cycle

¹⁸<https://www.theguardian.com/technology/article/2024/aug/24/yuval-noah-harari-ai-book-extract-nexus>

Regulation and innovation cannot be viewed in isolation. Innovations such as AI inevitably lead to regulation, while regulation in turn has a significant influence on innovation. Where government has the mandate for regulation, innovation is a public-private partnership, in which knowledge institutions and industry jointly drive innovation through research, development and experimentation, and subsequently scale it up; government establishes priorities with them and contributes funding. Mobilising talent and capital is essential in this process. Innovation is crucial to delivering future inclusive prosperity. And precisely because innovation, particularly when it concerns a general-purpose technology such as AI, plays such an important role in broad societal prosperity, regulating AI carefully is crucial.

Alongside regulation and innovation, inverse alignment above all requires investment, and given the fundamental and far-reaching impact of AI, the sums required are enormous. The fact that current investment in AI infrastructure may appear excessive is a matter of perspective: canals, roads and railways required similarly vast investments in their time, and later the same was true of fixed-line and mobile telephony, the internet and data centres. The difference lies in the pace: because the technology is developing more rapidly, the investments have to be made within an increasingly compressed timeframe. Europe is lagging behind in this respect, both in the private and public sectors.

Businesses are taking a wait-and-see approach, partly because of regulatory uncertainty, while governments are investing too much in isolation, even though they have purchasing power that, if pooled, could significantly accelerate the scaling of innovations. As a result, sovereign AI infrastructure in Europe and the Netherlands is proving difficult to establish, even though government can take the lead here, whether independently or together with strategic sectors such as energy and finance.

Inverse alignment therefore requires coordinated action by government and business, based on an integrated AI industrial policy in which regulation, innovation and investment reinforce one another. Given AI's impact on the economy and society, government has an important role to play. It has three core instruments at its disposal: legislation, funding for research and innovation, and public procurement. These instruments will need to be aligned and deployed in a coordinated manner. AI is here to stay. The fragmented and sometimes contradictory approach of resisting AI while providing ad hoc incentives must therefore give way to a coherent adoption strategy for AI Fair Tech: responsible AI with built-in safeguards for sovereignty and resilience.

What Europe needs to do

In the area of AI infrastructure, Europe is working on an EU-wide network of AI computing centres, with the announced AI gigafactories as the latest step: together representing a public-private investment of €30 billion over several years. This is a step in the right direction, but without a coherent AI industrial policy, it will be the equivalent of building cathedrals in the desert.¹⁹ Government is taking the lead, but European industry should not be left behind, since the bulk of the investment will have to come from businesses. The difference with the US is substantial: the four largest American technology companies alone invested \$415 billion in AI infrastructure in 2025,²⁰ with estimates for 2026 rising to between \$660 billion and \$760 billion. Whether the €30 billion from Brussels will be enough to establish sovereign European AI infrastructure remains highly questionable. Scarcity can be seen as an advantage, because it pushes us towards smaller, specialised foundation models and towards the joint development of open-source and open-weight models. But even along this more efficient route, it remains important for European companies to invest in sovereign AI infrastructure.

¹⁹<https://www.ceps.eu/ceps-publications/eu-plans-for-ai-gigafactories-sanctuaries-of-innovation-or-cathedrals-in-the-desert/>

²⁰<https://www.intotheminds.com>

The development of European AI foundation models - including training, fine-tuning, monitoring and further development - also requires closer cooperation between government, business and knowledge institutions in the form of AI industrial policy, as well as mobilising the right talent. In this area, Europe is making good progress, particularly in the Netherlands.²¹ At the same time, Europe, including the Netherlands, lacks a strong concentration of that talent in powerful innovation centres focused on pushing the boundaries of AI. In China and the US, much frontier AI technology is emerging from such centres. Europe should learn from this.

AI regulation is and will remain important, but so is proportionality, as are well-functioning regulatory sandboxes. The AI Act should evolve towards a regulatory framework that promotes desirable AI and thereby drives innovation and investment rather than holding them back. Consistent and proportionate development, implementation and enforcement are essential to achieving this. This applies to national supervisory authorities, which oversee prohibited applications, high-risk applications and AI systems subject to transparency requirements; and equally to the AI Office, which controls access to the European market for frontier AI, can remove models from the market on those grounds, and can impose fines when AI companies fail to behave in ways that we consider acceptable in Europe.

Simplifying the AI Act is essential in this regard. More text and more paperwork lead to greater complexity, and the Draghi report shows that it is not the well-formulated foundation of a law that holds back innovation, but the additional guidance and explanatory material layered on top of it, which can cause innovation projects to grind to a halt. The AI Act itself runs to around 140 pages and is of a high standard, but it is increasingly accompanied by hundreds of additional pages of guidance, guidelines and technical standards. On top of that will come interpretations from national supervisory authorities, with Europe expected to have around 2,000 national AI Act regulators in total.²² Harmonising the supervision of those 2,000 independent authorities appears to be an almost impossible task. Together, these elements create enormous institutional complexity, precisely at a time when Draghi is calling for simpler policies and regulation.

We call on the AI Office to reflect on the practical feasibility of supervising and enforcing the AI Act. We advocate a strong foundational law, the AI Act itself, while allowing sufficient discretion for national supervisory authorities and ensuring that their powers are exercised proportionately so as to promote innovation. Regulatory sandboxes are a welcome initiative, provided they are properly positioned within the supervisory authorities. Academic research shows that such sandboxes work well only when regulators have both a harm-mitigation mandate and an innovation mandate.²³ National supervisory authorities must have the people and resources needed to fulfil this innovation dimension of their mandate in practice.

A consistent approach to the development, deployment and enforcement of the AI Act is essential to all of this. So the European AI Office should above all continue with the implementation and enforcement of the AI Act. Since this month, it has had exclusive authority over AI models and AI on very large platforms, which is a major responsibility for an organisation that is still being built up. Leading scientists, policymakers and civil society organisations have called on the AI Office to take responsibility for the supervision and enforcement of frontier AI models.²⁴ It deserves the people and resources it needs to do so.

Europe recognises that more needs to be done than AI regulation alone and is working on rolling out the AI Continent Plan, with coordinated investment in AI infrastructure, regulatory simplification, and ambitions for AI talent and technological innovation. It is in the latter area that most needs to be done. Europe's current approach to innovation is based on research carried out by universities and national RTD organisations. Digital innovation, and AI innovation in particular, builds on this foundation but has a different dynamic, which is why it requires a different approach.

²¹<https://www.interface-eu.org/publications/talent-in-talent-out>

²²<https://cms.law/en/nld/legal-updates/there-are-around-2-000-ai-market-surveillance-authorities-in-the-eu>

²³Fahy, L. (2022). Fostering regulator-innovator collaboration at the frontline: A case study of the UK's regulatory sandbox for fintech. *Law & Policy*, 44(2), 162-184. <https://doi.org/10.1111/lapo.12184>.

²⁴<https://www.euractiv.com/news/civil-society-urges-eu-to-make-full-use-of-ai-enforcement-powers/>

This different approach rests on two pillars: challenge-based innovation and AI frontier centres. Under current R&D&I funding, competition takes place mainly at the outset, when applicants write the best proposal and secure project funding. The result is that there is little competition during implementation and little scope to adjust course. Challenge-based innovation differs in that it operates as a continuous competition in which funding follows the achievement of results. Teams compete to develop solutions to a predefined challenge. Along the way, there are stages with milestones, with only the best-performing teams receiving funding for the next phase.

Challenges are well known from the US DARPA programme, but European countries are also experimenting with them, for example through SPRIND in Germany, while the EU is taking its first steps in this area. We have been advocating a challenge-based approach at both EU and Dutch level for some time, and welcome the first initiatives, including NADI. The scale, however, particularly when it comes to AI innovation, remains far too limited and needs to be expanded substantially, with a very significant AI track within NADI.

Technology research and technological innovation are closely connected, while at the same time being fundamentally different. Technology research is about discovery and invention; technological innovation is about scaling and bringing products and solutions to market. This distinction is now recognised, and alongside R&D programmes, innovation programmes are emerging, for example those aimed at start-ups. But the approach remains too traditional and too heavily influenced by classical R&D practice, with complex projects at both European and national level being carried out by consortia. These projects are large and bureaucratic, and often fail to turn a shared interest into a shared outcome. As a result, their immediate impact falls short and becomes fragmented across a multitude of largely individual workstreams.

It is high time for Europe to rethink the way it innovates. Frontier AI technology emerges in particular from relatively small innovation centres in the US and China that are directly or indirectly part of, and funded by, large technology companies. These centres have a strong focus and combine research and development with innovation and product development. Examples include the AI frontier centres of Google, Meta, Microsoft, Amazon, OpenAI, Anthropic, DeepSeek, Tencent and Alibaba. Within Europe, but outside the EU, such centres exist in the United Kingdom and Switzerland.

The strength of industry-driven innovation centres is historically proven, from corporate R&D centres such as the famous Philips NatLab and today's ASML. The public-private Imec is an example of a highly successful, largely privately funded innovation centre at the frontier of semiconductor technology. In AI, however, European industry is largely failing to step up, with a few exceptions such as ASML; even the French flagship Mistral is funded predominantly by non-European companies.

The conclusion is that Europe must pursue a renewed approach to AI innovation, combining challenges with AI frontier centres. We therefore propose, by analogy with the EU network of AI factories, establishing a European network of relatively small AI frontier centres (of the order of 100 FTEs) where top AI talent works on challenges at the frontier of AI innovation. These centres should be primarily privately funded, with targeted public co-financing, and serve as incubators for frontier AI technology, products, start-ups and technology transfer. They should be given special access to the infrastructure of the AI (giga) factories and operate with a highly dynamic workforce, in which talent moves through rapidly, ensuring a continuous inflow and outflow of new ideas and innovations.

In addition to strengthening the approach to AI innovation, the coordination of key government instruments (legislation and enforcement, subsidies, investments, and procurement) must be significantly improved. A digital leader like Estonia²⁵ demonstrates that strong coordination in these areas pays off. The government should therefore consolidate its AI policy expertise into AI policy centres, which would be assigned a formally established coordinating role regarding the integrated approach to AI regulation, AI oversight, AI research and innovation, and government procurement policy.

²⁵<https://e-estonia.com/estonia-is-building-an-ai-powered-state/>

We therefore advocate establishing public AI policy centres at both European and Member State level, complementary to the AI frontier centres, to coordinate AI regulation and supervision, innovation policy, and policy relating to public procurement. National centres should have a legally established role in relation to government bodies and should complement the AI frontier centres: they should draw on the necessary AI technology expertise from AI frontier centres, avoiding competition for top technical talent.

A public-private AI policy centre should be considered against the backdrop of the existing institutional landscape. With the European AI Office and eight national AI Act supervisory authorities in the Netherlands alone, the institutional governance of AI is already complex. To prevent that complexity from increasing further, the centre's statutory mandate should explicitly include the objective of reducing complexity, whether by reducing the number of rules or by clarifying, translating and simplifying them. This applies both to regulation and supervision and to the organisation of public procurement of AI. Coordinated procurement can centralise buying power and reduce costs, while enabling governments to strategically select suppliers that advance AI Fair Tech. Academic research shows that public-private collaboration is particularly well suited to delivering innovation.^{26 27}

There is momentum already. In June, the European Commission appointed Jim Hagemann Snabe as Special Envoy for Industrial AI; his mandate runs until March 2027 and concludes with a report. But a report is not industrial policy. In September, von der Leyen will deliver her State of the Union address, while negotiations on the EU's multiannual budget will begin. This is the moment to turn the envoy's mandate into an institution with a dedicated budget.

The AI race has only just begun, and Europe is more than capable of establishing a position of its own on the global stage, pursuing its own AI Fair Tech course, driven by innovation and investment, alongside enabling regulation for the responsible development and application of AI.

²⁶Lena Brogaard, 'Innovative outcomes in public-private innovation partnerships: a systematic review of empirical evidence and current challenges', *Public Management Review* 23(1) (2021), pp. 135–157, DOI: 10.1080/14719037.2019.1668473.

²⁷Nunzia Carbonara & Roberta Pellegrino, 'The role of public private partnerships in fostering innovation', *Construction Management and Economics* 38(2) (2020), pp. 140–156, DOI: 10.1080/01446193.2019.1610184.

