POLAND / AI POLICY / INSTITUTIONAL RECORD

The Exascale Gap

An analysis of Poland's infrastructure shortfalls in the face of the AI Act and frontier AI

SIMON MIOTK

Submitted to Poland's Ministry of Digital Affairs

Official paper of

Studenckie Koło Naukowe Analiz Rynkowych

Market Analysis Student Research Society | SGH Warsaw School of Economics

English archival edition | 28 October 2026

simonmiotk.com/record/exascale-gap

READING THE RECORD

The paper at a glance

35 PFLOPS	The capacity of Poland's fastest supercomputer, Helios. Europe's exascale JUPITER system has approximately 28 times as much (1,000 PFLOPS). [6][8]
5.9%	The share of Polish enterprises using AI in 2024. The EU average was 13.5%; Denmark was at 27.6%. Poland ranked second from last in the EU. [5]
EUR 200bn / ~EUR 0.4bn	The EU's InvestAI package compared with the sum of Poland's direct instruments: a difference of several hundred times. [12][4]
~4.4 times per year	The growth rate of the computing power required to train AI models. The policy provides for a review once a year. [9][4]
EUR 35m / 7%	The maximum AI Act fine, or 7% of global turnover. Obligations for high-risk systems apply from 2 August 2026. [1][2]
415 to 945 TWh	Global electricity consumption by data centres by 2030, a 2.3-fold increase. This creates tension with EU climate policy. [15]

ABOUT THIS EDITION

This is a complete English translation and visual re-edition of the supplied Polish paper. Its arguments, figures, estimates and policy timelines are preserved as an archival record and have not been updated for this edition.

STRUCTURE

Introduction

I. Systemic gaps in Poland's AI policy

II. Political and regulatory risks

III. Recommendations

Conclusion

References

PURPOSE AND CONTEXT

Introduction

The pace of artificial intelligence development over the past two years has changed the global balance of economic and political power. The AI Act entered into force on 1 August 2024, and European Union member states are moving from declarations to implementation: the first prohibitions have applied since February 2025, while obligations for high-risk systems will apply from 2 August 2026 [1]. Against this background, Poland must present a coherent, feasible and evidence-based strategy for AI development. The scale of the challenge is exponential: according to Epoch AI, the computing power required to train leading models grows by approximately 4.4 times a year. Even a one-year delay therefore translates into a loss of advantage measured in orders of magnitude [9].

The coincidence of dates is symbolic. On 17 November 2025, the day on which the Ministry of Digital Affairs published its Policy for the Development of Artificial Intelligence in Poland to 2030 [4], EuroHPC announced that Germany's JUPITER had become Europe's first exascale supercomputer: 1.0 EFLOPS and fourth in the world [7]. Poland's fastest machine, Helios, offers 35 PFLOPS and ranks 155th on the TOP500 list [6][8]. The government document contains ambitious declarations and allocates responsibilities, but reveals several systemic gaps: no operational timetable; insufficient and imprecisely defined computing capacity; fragmented coordination mechanisms; and a monitoring system that does not match the pace of change in AI.

This paper analyses the key gaps in the government document and proposes realistic measures that can be implemented immediately. Each argument is accompanied by numerical data from public sources: Eurostat, EuroHPC/TOP500, the IEA, Epoch AI, Stanford HAI and the European Commission. The proposals rest on continuous monitoring and policy evaluation, a measurable increase in computing capacity, the establishment of a team of young technology experts, measures to counter talent emigration, and precise budgets and deadlines. Poland has real potential to benefit from the current transformation, but this requires a strategy grounded in data, speed and responsibility towards future generations.

PART 01

Systemic gaps in Poland's AI policy

1.1 / The strategic gap: no operational timetable

The AI Development Policy to 2030, published in November 2025, contains a crucial statement: an action timetable "will only be developed within three months of the document's adoption" [4]. At the time of publication, there is therefore no operational roadmap defining:

• delivery dates and milestones

• responsible bodies and progress indicators

• dependencies between projects

• investment priorities

• cost estimates for individual stages

From the perspective of public-project management, the document is thus declaratory rather than operational. In practice, this leads to delays, a lack of coherence and a risk of losing EU funding, which requires precise plans and accountability. The risk is twofold: delays across the ecosystem, as downstream activities - AI Factories, talent programmes and HPC investments - slip by further quarters; and a lack of accountability for results, because progress cannot be monitored without a public calendar. At a growth rate of approximately 4.4 times a year [9], each six-month delay more than doubles the gap that must be closed relative to countries already implementing their roadmaps, including France, Germany, Estonia and the United Kingdom.

1.2 / The delivery gap: annual monitoring versus monthly breakthroughs

The government document provides for annual evaluation and updates to AI policy [4]. Yet the training compute of leading language models grows approximately fivefold a year, doubling about every 5.2 months; for significant models overall, the rate is around 4.4 times a year [9]. An annual review cycle works for road infrastructure or heritage conservation, but not for a technology whose capabilities change within weeks. During a single government review cycle, global demand for computing capacity grows severalfold. A system that looks only once a year is structurally unable to keep up.

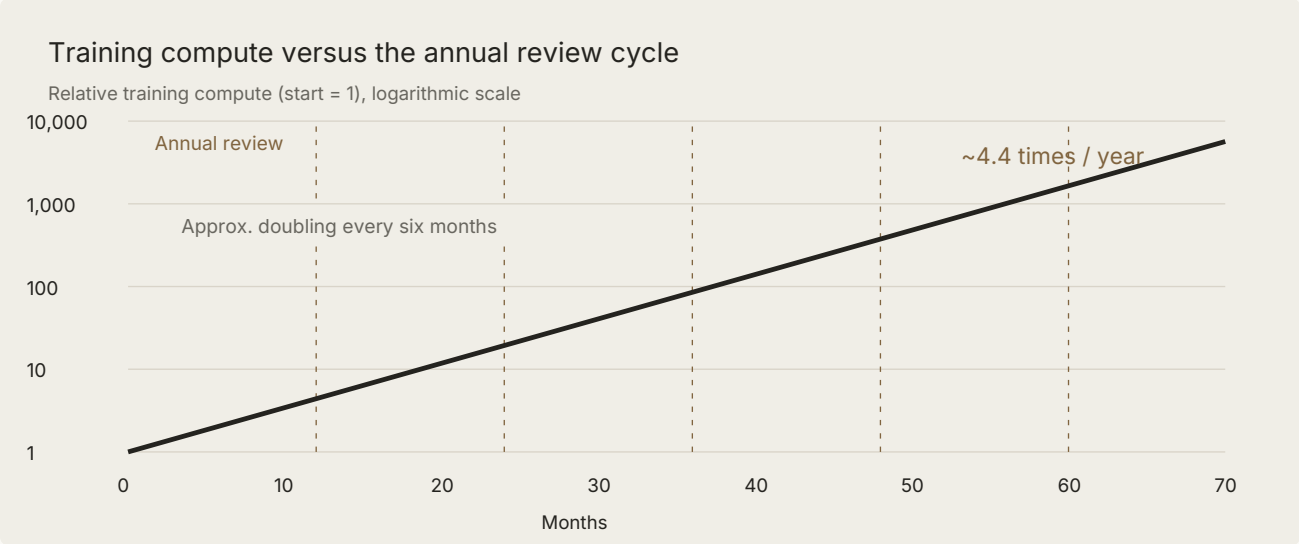


Figure 1. Exponential growth in training compute versus an annual policy review cycle. Source: Epoch AI.

The source paper uses approximately 4.4-fold annual growth, doubling around every six months, and around fivefold growth for frontier models.

Consequently, the government's monitoring system reacts too slowly, misses emerging risks and cannot adapt policy to market dynamics. The solution is continuous assessment based on dashboards and monthly evaluation of indicators, a team of young technology analysts, and Rapid-Response Units: small decision-making groups that respond immediately to breakthroughs, including new frontier models, optimisation techniques and changes in compute allocation. This model is used by countries with high levels of digital readiness [17].

1.3 / The infrastructure gap: computing capacity without specifications

Computing capacity is a key strategic resource of the twenty-first century. AI Policy 2030 does declare an increase in the capacity available to businesses from 3 to 40 exaflops through the PLGrid network and a voucher system [4]. This is progress compared with earlier versions. However, it leaves the most important parameters unspecified: numerical precision, whether FP64 or AI-oriented FP16/FP8; architecture; chip generation, such as H100/H200, B200 or Gaudi; energy requirements; and a GPU procurement plan. Without these parameters, '40 exaflops' remains a slogan rather than an engineering project. For comparison, today's flagship, Helios, provides 35 PFLOPS in FP64 and 1.8 EFLOPS in AI mode [8].



Figure 2. European flagship supercomputers and Helios: FP64 performance in PFLOPS. Source: TOP500 / EuroHPC, November 2025.

Values and world rankings reproduced from the Polish paper. Helios is also stated to provide 1.8 EFLOPS in AI mode.

The scale of the problem is best illustrated by a set of reference points, from the regulatory threshold for systemic risk and the compute required to train a single model to one European AI gigafactory:

Reference point	Scale	Significance
AI Act systemic-risk threshold	10^{25} FLOP for training	Regulatory boundary for frontier models; approximately 11 providers worldwide exceed it. [3]
Training a GPT-4o-class model	~3.8×10^{25} FLOP	Approximately four times the systemic-risk threshold; Epoch AI estimate. [9]
Helios, AI mode (Poland)	1.8 EFLOPS	Training a GPT-4o-class model: approximately eight months at theoretical peak, and over a year in practice, occupying the national flagship. [8]
National 2030 target (PLGrid)	3 to 40 EFLOPS	Declared increase in compute for businesses; precision and architecture unspecified. [4]
One EU AI gigafactory	~100,000 latest-generation GPUs	Hundreds to thousands of EFLOPS in AI mode, many times current national capacity. [12][14]

Table 1. The compute ladder: from the regulatory threshold to an AI gigafactory.

Without HPC infrastructure at an appropriate scale, Poland cannot train its own language models, develop deep tech, compete with the global private sector or fully benefit from EU programmes such as AI Factories and AI Continent. Notably, the three European systems in the world's top ten - JUPITER, at 1,000 PFLOPS in Germany; LUMI, at 380 PFLOPS in Finland; and Leonardo, at 241 PFLOPS in Italy - are hosted by Germany, Finland and Italy. Poland hosts none of them [6][7]. An AI strategy without a precise account of computing capacity is like a military strategy without information about the army.

1.4 / The skills gap: low AI adoption and talent emigration

The document commits to developing skills but underestimates two interconnected trends: very low AI adoption among Polish businesses and the departure of the most capable engineers. According to Eurostat, only 5.9% of Polish enterprises with ten or more employees used AI in 2024, compared with an EU average of 13.5% and 27.6% in Denmark. In 2025, the Polish and EU figures rose to 8.4% and 20.0% respectively, but the gap remained [5].

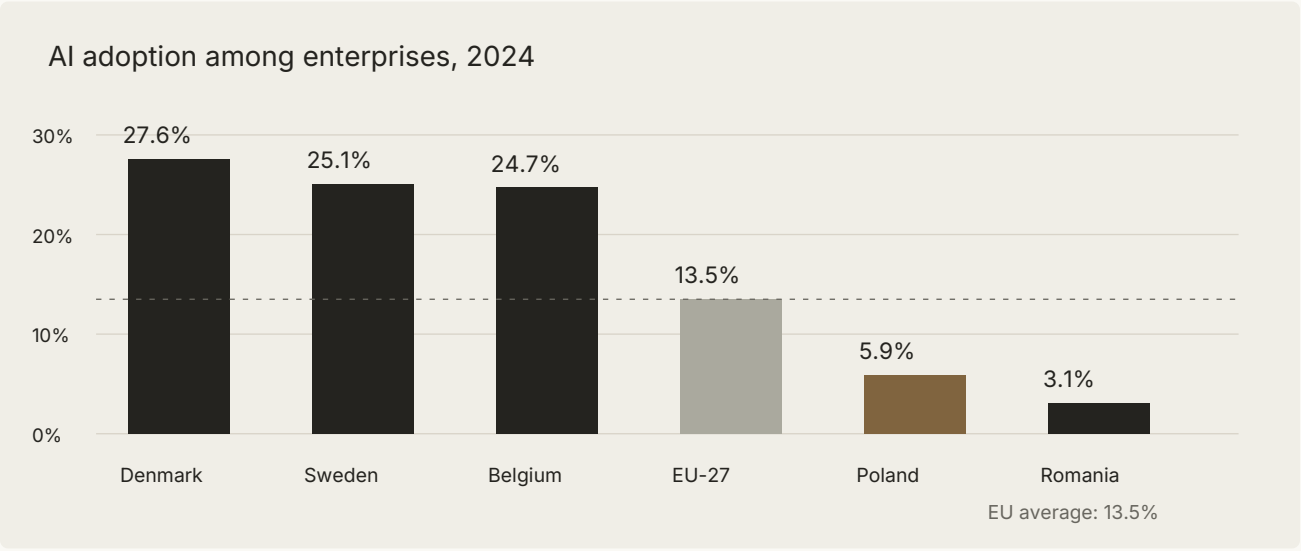


Figure 3. Share of enterprises using AI in 2024. Source: Eurostat.

Enterprises with ten or more employees. Poland: 5.9%; EU average: 13.5%. Values reproduced from the source paper.

An important claim in earlier versions of this analysis therefore needs correction: Poland is not 'behind Romania'. Romania ranks last, at 3.1% in 2024, while Poland is second from last. The distinction is largely symbolic; both are at the bottom of the EU ranking. On the workforce side, Poland's ICT sector has around 250,000 specialists, with an estimated shortfall of 40,000-50,000 people [18]. Emigration deepens the gap: pay differences with Western Europe lead the most highly qualified specialists to work in the United States, the United Kingdom, the Netherlands or Germany [18]. Without an ecosystem that retains talent, commitments to skills development remain a fiction.

1.5 / The technology gap: from gigafactory aspirations to hard KPIs

The policy identifies specific direct allocations: PLN 60 million in 2025, PLN 162 million in 2026, PLN 300 million in PFR Ventures, matched by at least the same amount of private capital, and funding for language models [4]. The flagship infrastructure project is Baltic AI GigaFactory. Contrary to an earlier assessment that it lacked specifics, some parameters are already known. It is an application submitted to the European Commission under InvestAI in 2025 by a Poland-led consortium with Lithuania, Latvia and Estonia. Its estimated value is approximately EUR 3 billion, with 65% private capital and 35% EU and national funding. It envisages up to 100,000 latest-generation GPUs and training models with more than one trillion parameters, including Poland's PLLuM and Bielik.AI [14].

The problem is therefore not a lack of ambition, but a lack of delivery precision and funding certainty. For a project of this scale, the following have yet to be publicly specified:

- target FLOPS and chip generation, rather than only an indicative GPU count
- a binding timetable and stages, including milestones, risks and total cost
- power demand in MW and the cooling model
- the access model, including APIs and priority sectors: research, SMEs and industry
- integration with the AI Act Service Desk
- funding status: the project remains an application to the Commission, and 65% of its budget depends on private capital [14]

An investment in a gigafactory without a binding technical specification and secured funding remains, in analytical terms, a political declaration rather than a technology-infrastructure project [11].

PART 02

Political and regulatory risks

2.1 / EU green policy and the energy requirements of HPC

Developing generative and frontier models requires enormous computing resources, which in turn require energy. According to the International Energy Agency, data centres consumed approximately 415 TWh in 2024, around 1.5% of global electricity consumption. By 2030, demand is expected to more than double to approximately 945 TWh, as much as Japan consumes today. Consumption by AI-optimised data centres is expected to increase more than fourfold, while AI's share of data-centre power consumption is projected to rise from 5-15% to 35-50%. In Europe, the increase is expected to be approximately 45 TWh, or 70% [15].

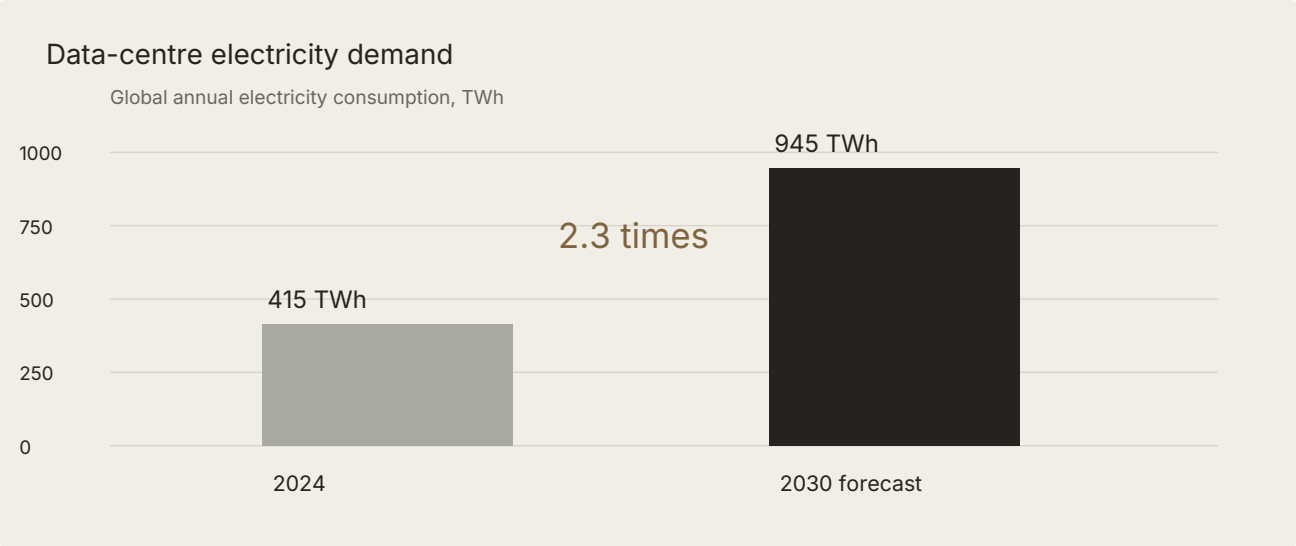


Figure 4. Global electricity consumption by data centres, 2024 and 2030. Source: IEA, Energy and AI (2025).

2030 is a forecast. The source paper states that AI's share of data-centre consumption rises from 5-15% to 35-50%.

A structural conflict of interests emerges: AI requires extreme computing capacity, while the EU requires emissions reductions. The government document refers to 'green AI infrastructure' but does

not specify how to provide capacity competitive with the United States and China while complying with climate regulations. This lack of precision is a real strategic risk. Access to green energy and liquid cooling could be an advantage for Poland and the location of Baltic AI GigaFactory, provided it is planned in numbers rather than slogans.

2.2 / Fragmented funding against the scale of competition

AI funding in Poland is highly fragmented across NCBR, NCN, PARP, PFR, BGK, regional programmes, EU funds, the National Recovery Plan and PFR Deep Tech. This means a lack of central coordination and a shared project pipeline, duplicated activities and difficulties in setting priorities. The problem becomes more significant when set against the scale of the competition. At the AI Action Summit in Paris in February 2025, the European Commission launched the EUR 200 billion InvestAI initiative, including EUR 20 billion for four or five AI gigafactories with approximately 100,000 chips each. France announced EUR 109 billion [12]. Poland's direct instruments total hundreds of millions of zlotys: a difference measured in hundreds of times.

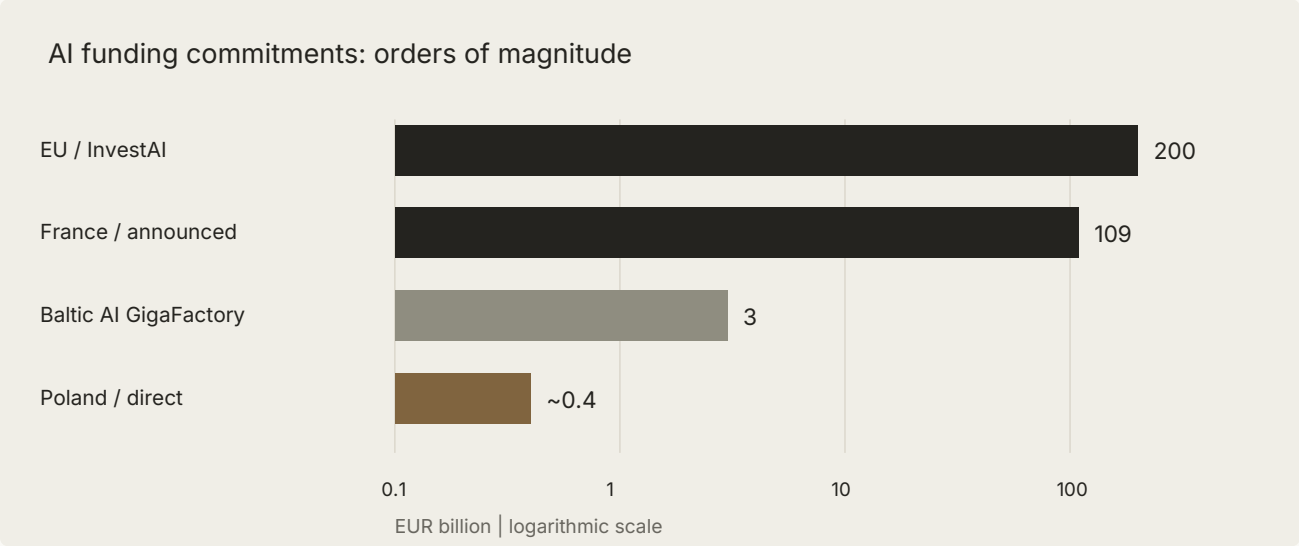


Figure 5. Scale of AI funding commitments, EUR billion, logarithmic scale. Sources: European Commission InvestAI, AI Action Summit, Baltic AI GigaFactory consortium and AI Policy 2030.

InvestAI is funding mobilisation; France is a multi-year announcement; Baltic AI GigaFactory is a proposal with 65% private capital. The Polish figure is an approximate EUR conversion of direct instruments. The amounts have different scopes and time bases.

For a complete picture, public funding should be compared with the private capital flows that ultimately determine where talent and compute are located:

Entity / programme	Amount	Basis
United States, private AI investment	USD 285.9bn (2025); USD 109.1bn (2024)	Annual private investment; approximately 23 times China. [10]
China, private investment	USD 12.4bn (2025)	Annual private investment; excludes state funds. [10]
United Kingdom, private investment	USD 4.5bn (2024)	Annual private investment. [10]
EU, InvestAI	EUR 200bn, including EUR 20bn for gigafactories	Mobilisation of funding, 2025. [12]
France, announced commitments	EUR 109bn	Multi-year public and private commitments, 2025. [12]
Baltic AI GigaFactory, led by Poland	~EUR 3bn; 65% private	Application to the European Commission, 2025. [14]
Poland, direct instruments	PLN 60m (2025); PLN 162m (2026); PFR PLN 300m	AI Policy 2030. [4]

Table 2. Public and private AI funding: orders of magnitude. USD and EUR follow the original sources.

The recommended direction is to consolidate funding flows within an 'AI Fund of Funds', drawing on PFR Deep Tech's capabilities, with a unified system for submitting and assessing projects and a public map of priorities: compute, medical technology, industry and security. This model reduces duplication, increases transparency and shortens the investment cycle [16].

2.3 / Implementing the AI Act without a prepared domestic ecosystem

The AI Act imposes structural obligations on member states: registration systems, audit processes, reporting, supervision of high-risk applications, compliance support and implementation of an AI Act Service Desk. The timetable is unforgiving, and the penalties are substantial:

Obligation / infringement	Maximum penalty	Application date
Prohibited practices (Article 5)	EUR 35m or 7% of turnover	From 2 February 2025
General-purpose AI models (GPAI)	EUR 15m or 3% of turnover	From 2 August 2025
High-risk systems	EUR 15m or 3% of turnover	From 2 August 2026
Incorrect or incomplete information	EUR 7.5m or 1.5% of turnover	Not specified
Full application of the Regulation	Not applicable	2 August 2026; exceptions until 2027

Table 3. AI Act penalties and timetable, reproduced from the Polish paper. Source: Regulation (EU) 2024/1689, Article 99.

Meanwhile, Poland does not yet have the full institutional capacity and regulatory workforce to perform these tasks. An implementation timetable has yet to be developed; the technical computing infrastructure needed for compliant AI deployment is insufficient; and a rapid risk-evaluation system is not functioning. A serious gap emerges: regulation will outpace the state's ability to implement it. The consequences may include delays in implementing EU policies, legal problems for businesses, business migration to countries with simpler compliance, and reduced investment attractiveness for Poland, against a background of fines for non-compliance reaching 7% of global turnover.

PART 03

Recommendations

3.1 / Establish a Young AI Monitoring Team

At a growth rate of approximately 4.4 times a year, Poland needs a structure capable of rapid analysis and response. A Young AI Monitoring Team is proposed, bringing together highly capable young experts in technology, economics and law. Its responsibilities would include continuous monitoring of the AI market and assessment of the impact of new technologies; quarterly rather than annual reports that reflect the real pace of change; oversight of spending and information flows between institutions; and rapid recommendations when new breakthroughs occur. Such structures operate in countries with high digital resilience and shorten strategic response times from months to weeks [17].

3.2 / Increase computing capacity measurably: HPC and GPUs

Competitiveness in AI requires faster investment in HPC. The priority is to specify a capacity target in FLOPS together with numerical precision - FP64, FP16 or FP8 - and chip generation, rather than only budget allocations or an indicative number of GPUs. An expansion plan should cover 24-36 months, with quarterly milestones. Investments should also be integrated with EU-compliant energy models, combining renewable energy and liquid cooling, with a specified MW requirement. The benchmark should be the scale of a single European gigafactory, not today's Helios capacity [12][14].

3.3 / Consolidate funding through an AI Fund of Funds

Consolidate funding flows from NCN, NCBR, PARP, BGK, PFR and regional programmes within a single structure based on PFR Deep Tech, with a unified system for submitting and assessing projects and a map of funding priorities. The objective is to reduce duplication, increase transparency and shorten the investment cycle, so that domestic resources, although orders of magnitude smaller than EU funding, achieve the greatest possible impact [16].

3.4 / Retain talent in Poland

Given the shortage of 40,000-50,000 ICT specialists and wage pressure from the West [18], the necessary measures include grants and scholarships for AI students and doctoral researchers; a fast route into government projects, modelled on British and Dutch approaches [19]; cooperation with universities and research centres; and the development of public-private AI laboratories. Retaining talent is a necessary condition for effective implementation of the AI Act and for building a domestic ecosystem.

3.5 / Set hard KPIs for Baltic AI GigaFactory

The gigafactory is a flagship project, which is precisely why it requires the most demanding indicators. The recommendations are to specify target capacity parameters - FLOPS, GPU number and generation, and throughput; project KPIs covering the timetable, stages, risks, total cost and financing structure, including the conditions for securing 65% private capital; sectors with priority access to compute; and public quarterly implementation reports. Without clear parameters, the project will remain a declaration rather than infrastructure.

Conclusion

The rapid development of AI confirms its exponential nature: changes that took a decade in traditional sectors now occur within months, while the computing power required to train leading models grows approximately 4.4 times a year [9]. In this environment, effective public policy cannot rely on static documents or annual reviews. It requires structures capable of continuous monitoring and rapid response.

The analysis of the AI Development Policy to 2030 identifies a document that represents an important step forward, including the 3-to-40 EFLOPS target, PFR Ventures and Baltic AI GigaFactory, but remains burdened by critical gaps: no operational timetable; unspecified numerical precision and compute architecture; fragmented funding at a scale orders of magnitude below the EU's; low AI adoption, at 5.9% of businesses compared with 13.5% across the EU [5]; and talent emigration. If these gaps remain unaddressed, Poland risks lasting technological lag and losing the greatest development opportunity in its recent history, against a regulatory background of AI Act fines reaching 7% of global turnover.

The conclusions are nevertheless optimistic on one condition. Poland can not only catch up but move ahead of Central and Eastern Europe if it adopts a systemic approach: strengthen HPC infrastructure with measurable targets, implement continuous monitoring, establish an interdisciplinary team of experts, consolidate funding and create conditions that retain talent. Only this model can turn strategic declarations into action and make Poland a significant participant in Europe's AI ecosystem.

SOURCE RECORD

References

Reference numbering and source descriptions are retained from the Polish paper. Where the source provides only an organisation or a domain, this edition does not invent a more specific publication link.

[1] Regulation (EU) 2024/1689 of the European Parliament and of the Council (AI Act). Official Journal of the European Union, 2024. Entry into force: 1 August 2024; full application: 2 August 2026, with exceptions.

[2] AI Act, Article 99: administrative fines, up to EUR 35m / 7% of turnover; EUR 15m / 3%; EUR 7.5m / 1.5%. artificialintelligenceact.eu/article/99/

[3] AI Act, Article 51, and European Commission guidance on GPAI: systemic-risk threshold of 10^{25} FLOP; notification within two weeks.

[4] Ministry of Digital Affairs (2025). Policy for the Development of Artificial Intelligence in Poland to 2030, version dated 17 November 2025. ai.gov.pl

[5] Eurostat (2025). Use of artificial intelligence in enterprises. Data for 2024: Poland 5.9%, EU 13.5%; and 2025: Poland 8.4%, EU 20.0%. ec.europa.eu/eurostat

[6] TOP500 (November 2025). The TOP500 List. top500.org/lists/top500/2025/11/

[7] EuroHPC JU (2025). JUPITER Officially Propels Europe into the Exascale Era: 1.0 EFLOPS, fourth worldwide. eurohpc-ju.europa.eu

[8] ACK Cyfronet AGH / HPE (2024). Helios supercomputer: 35 PFLOPS (FP64), 1.8 EFLOPS (AI), TOP500 rank 155. cyfronet.pl

[9] Epoch AI (2024-2025). Trends in AI / Training compute: compute doubles approximately every six months, around 4.4 times a year; GPT-4o at approximately 3.8×10^{25} FLOP. epoch.ai/trends

[10] Stanford HAI (2025; 2026). AI Index Report, Economy: US private investment of USD 109.1bn in 2024 and USD 285.9bn in 2025; China USD 12.4bn; UK USD 4.5bn. hai.stanford.edu

[11] McKinsey & Company (2023). The Economic Potential of Generative AI: USD 2.6-4.4 trillion annually across 63 use cases.

[12] European Commission (2025). InvestAI: mobilising EUR 200bn, including EUR 20bn for AI gigafactories. France: EUR 109bn, AI Action Summit, Paris. digital-strategy.ec.europa.eu

[13] EuroHPC JU / European Commission (2025). AI Factories: Poland's PIAST-AI in Poznań and GAIA in Kraków within the PLGrid network. eurohpc-ju.europa.eu

[14] Baltic AI GigaFactory consortium / European Commission InvestAI (2025). Poland-Lithuania-Latvia-Estonia application: approximately EUR 3bn, 65% private capital, up to 100,000 GPUs, models exceeding one trillion parameters, including PLLuM and Bielik.AI.

[15] International Energy Agency (2025). Energy and AI: data centres at 415 TWh in 2024 and 945 TWh in 2030; AI share of 5-15% rising to 35-50%. iea.org/reports/energy-and-ai

[16] PFR (2024-2025). PFR Deep Tech / PFR Ventures: structure and financing instruments.

[17] OECD (2023). GovTech Maturity Index: models of continuous monitoring and rapid response.

[18] Eurostat / CEDEFOP (2024-2025). ICT specialists in employment; Poland: skills mismatch. ICT sector of around 250,000 people and a shortage of 40,000-50,000.

[19] The Alan Turing Institute (2023). Fellowship Structure and AI Talent Programmes: talent-retention models in the United Kingdom and the Netherlands.