



**Prof. Włodzisław
Duch**

Co-founder of the Polish Neural Networks Society and the Polish Artificial Intelligence Society. He served two terms (2006–2011) as President of the European Neural Network Society. In 2014–2015, he was Undersecretary of State at the Ministry of Science and Higher Education in Poland.
is.umk.pl/~duch

AI – HIGH TIME FOR A WAKE-UP CALL!

The use of artificial intelligence for scientific purposes requires deep expertise and systems tailored to specific applications. As technology accelerates, countries must keep pace or risk falling behind.

Włodzisław Duch

Department of Informatics
Faculty of Physics, Astronomy and Informatics
Nicolaus Copernicus University in Toruń

a minister-level official, with deputy ministers representing the ministries of finance, economy, health, energy, education, science, and technology. As a result of such measures, thousands of scientists are now involved in implementing AI technologies across a wide range of fields.

The world has gone mad about artificial intelligence. Investment plans for AI centers in the United States, Europe, and Asia already exceed one trillion dollars. Companies such as OpenAI have incurred losses of hundreds of millions of dollars, spending approximately \$3 billion on training new models and around \$2 billion on services that, by mid-2025, were being used weekly by 700 million users. With 15 million paid subscriptions, ranging from \$20 to \$200 per month, such companies can afford to offer a wide array of free tools – from automatic translation and support for application development to weather forecasting. Google DeepMind, in turn, provides access powerful tools for building applications, creating virtual worlds, retrieving and analyzing information, and processing molecular biology and medical data.

Science also benefits from this frenzy. Following a statement issued by the European Commission in April 2018, most European countries introduced numerous initiatives and established new AI research centers: 13 in Belgium, 11 in Finland, 9 in France, 12 in Spain, 11 in the Netherlands, 9 in Germany, 14 in Russia, 11 in Sweden, and 14 in Italy. In the United States, 29 National AI Research Institutes have been created. China and several Arab countries plan to allocate \$100 billion each to AI development.

The United Arab Emirates have gone further by establishing a government council for artificial intelligence and blockchain development, chaired by

Waking Up – Without Hurry – in Poland

In 2018, we established the Polish Alliance for the Development of Artificial Intelligence (PP-RAI), intended to coordinate the operations of five organizations: the Polish Neural Networks Society (PTSN), the Polish Artificial Intelligence Society (PSSI), the Polish Machine Learning Group, the Polish Chapter of the IEEE Systems, Man, and Cybernetics Society, and the Polish Chapter of the IEEE Computational Intelligence Society. Unfortunately, the initiative did not receive meaningful institutional support, and PP-RAI's activities were therefore largely limited to organizing international conferences.

At the same time, Poland's research evaluation system does not incentivize participation in conferences that do not generate publication points. As a result, even researchers interested in engaging with such initiatives often struggle to obtain funding to attend them.

Poland has spent the last seven years largely dozing and is only now beginning to wake up. At the end of 2020, the *Policy for the Development of Artificial Intelligence in Poland* (a 71-page document) was adopted, yet no new scientific centers were established. Ideas NCBR – which initially focused on AI research and science–business cooperation as a limited liability company – was only transformed into the Ideas Research Institute in 2025.



The Red Queen
Hypothesis: It Takes All
the Running You Can
Do to Stay in the Same
Place

GENERATED BY AI

Let us compare Poland's situation with that of other EU countries and their investments in new research units focused on artificial intelligence in recent years. Companies based in Denmark and Sweden use AI technologies up to four times more frequently than those operating in Poland. According to the Global Innovation Index (2024), which is based on publication output and patent activity, the Warsaw science-and-technology cluster ranks 90th worldwide, a position driven primarily by patents filed by Samsung R&D Poland. Eurostat data on AI adoption in enterprises further indicates that Poland remains at the bottom of the EU rankings, ahead only of Romania.

At the beginning of 2025, the Ministry of Digital Affairs prepared a 64-page document entitled *Policy for the Development of Artificial Intelligence in Poland until 2030*. The document draws on several prior studies, including the Polish Economic Institute's report *Gdzie jest polska nisza AI i jak ją wykorzystać* (*Where Is the Polish AI Niche and How to Use*), the *Substantive*

Recommendations of the Working Group on Artificial Intelligence (GRAI), and a number other analytical papers. In parallel, the Ministry of National Defense has developed its own departmental artificial intelligence strategy extending to 2039.

Against this backdrop, the question arises: is there any realistic prospect for Poland to become one of the leaders of the "AI continent"? Although the declared objectives are ambitious, achieving them will be difficult. Poland has not made sufficient investments in human capital, nor does it possess AI development funds comparable to those available in leading countries. Moreover, the pace of progress in artificial intelligence over the next two to three years remains highly uncertain. OpenAI is expected to drive profound restructuring across entire industries, yet even its strategic planning currently extends no further than six months ahead. In recent months, numerous new approaches to building more capable systems have emerged, enabling AI to perform increasingly complex tasks. Two years ago, GPT-4 was able to complete



tasks that required five to six minutes of human effort; today, GPT-5 successfully solves a substantial share of tasks that would take humans approximately up to 5 hours!

Under these conditions, a key challenge remains: how can Poland create an environment in which “effective cooperation between science, the public sector, and business will enable the development and implementation of the latest digital technologies and their appropriate use,” as envisioned by the new AI development policy for the coming five years?

Science Requires Specialized AI Systems

The role of the basic sciences that underpin new technologies should not be overlooked. Implementing technologies that perform well in benchmark tests but are not yet mature enough for real-world applications, as well as deploying systems such as ChatGPT that appear to “know everything about anything” yet often provide only superficial knowledge, frequently leads to disappointment. These are not omniscient systems capable of independently solving complex problems. Meaningful success depends on close collaboration between domain experts and AI systems tailored to specific tasks.

Effective AI applications require the continued development of new algorithms rather than the mere adoption of existing, general-purpose models. Research, education, and the adaptation of organizational procedures within both businesses and research institutions are essential to the productive use of artificial intelligence. The success of AlphaFold – developed by DeepMind, for which Demis Hassabis and John Jumper were awarded the Nobel Prize

in Chemistry in 2024 – was made possible by a deep understanding of molecular biology processes. The latest version, AlphaFold 3, is capable of predicting the structures of protein complexes and their interactions with RNA, DNA and other molecules.

Such tools are built on advanced machine learning approaches, including deep learning architectures, transformer models, and diffusion-based models. Over the past few years, the application of artificial intelligence in drug discovery has accelerated nearly every stage of the process: from the analysis of genomic, transcriptomic, metabolic, proteomic and epigenetic data, through drug repurposing, to the design of novel compounds and the prediction of their properties prior to synthesis, followed by the optimization of clinical trials. As a result, success rates in identifying effective drugs have approximately doubled, development timelines have been reduced by more than half, and costs have fallen by up to 70%. Numerous studies published in *Nature Medicine* document breakthroughs made possible by such specialized AI systems.

The Lancet recently published an article about groups of AI agents in medical applications. Unleashing the full potential of AI requires treating such systems as partners in healthcare. Large language models alone are insufficient; multi-agent systems capable of using appropriate tools are required. Analyses of difficult cases requiring diagnosis and deep reasoning, conducted by 25 experts from U.S. clinical hospitals, have demonstrated that large AI models have reached a level that clearly exceeds human capability. In the case of highly complicated real-life cases published in *The New England Journal of Medicine*, Microsoft’s AI Diagnostic Orchestrator correctly diagnosed 85% of cases – more than four times as many as a group of experienced physicians. The procedures and tests recommended by AI to confirm diagnoses were significantly less expensive than those recommended by medical doctors.

TxAgent performs multi-stage reasoning using 211 ToolUniverse tools to analyze drug interactions, contraindications, and treatment strategies. It evaluates drug interactions at the molecular, pharmacokinetic, and clinical levels, identifies contraindications based on the patient comorbidities and medications, and adapts treatment strategies to individual patient characteristics, thereby improving treatment recommendations. A randomized trial of a generative AI chatbot known as Therabot, used in the treatment of mental disorders, has demonstrated its effectiveness in adults suffering from severe depression, anxiety, or eating disorders.

In July 2025, for the first time, the Grok 4 model succeeded in operating a vending machine in a real-world setting. A long-term business scenario requires tracking inventory, placing orders, setting prices,

transferring fees, and maintaining consistency of decisions over time. Until now, AI systems deployed in such applications have begun to make mistakes relatively quickly. Managing even a small business of this kind requires extensive knowledge and situational understanding. Within a month, all large systems are expected to cope with such tasks, and within two months, small, specialized local systems are likely to perform just as well. In a year's time, all owners of such machines will entrust their management to AI systems. Whoever moves too quickly will run into trouble. The fact that not everything works smoothly right now does not mean that this will be the case next year.

These are not simple applications. They have become possible thanks to new AI systems that use large language models to identify associations, as well as systems capable of reasoning. Their development and deployment would be much easier if IT students were able to choose additional coursework focused on a discipline of particular interest to them. I proposed such a solution 30 years ago. Although our students may select numerous elective courses required to obtain 30 ECTS credits, a computer science student who chooses foundational courses in law or medicine does not receive formal certification of the acquired skills. In the United States, the system of awarding diplomas of both major and minor fields of study is widely used. As a result, graduates are better prepared for professional work.

Two Years of Acceleration Right Ahead of Us

Although many people find this difficult to believe, no fundamental barriers to the development of artificial intelligence are currently visible. Large AI models are not merely upgraded calculators. Human brains are not calculators that simply count neuronal impulses. Electronics and biology represent two distinct domains; however, at the functional level, the processes occurring in the brain and in artificial neural networks exhibit a substantial degree of similarity. Neural networks learn and form associations, yet they cannot be reduced to programs that merely execute a sequence of instructions.

We construct large networks and supply them with knowledge derived from text, images, and video, but it remains impossible to predict their responses with precision. If such predictability were achievable, traditional computer programs capable of engaging in dialogue or generating images would have existed long ago. The next two years are likely to bring even more rapid progress. We have learned to design algorithms that are far more effective than those available just a year ago. Artificial systems face fewer inherent constraints than biological brains, and as a result, we

can expect to observe superhuman capabilities across a growing range of applications, not limited to games. The HRM (hierarchical reasoning model), published in August 2025, generates abstract plans that are subsequently executed by specialized modules. Sudoku puzzles and challenging ARC (Abstraction and Reasoning Corpus) tasks – which had resisted attempts by large-scale systems with trillions of parameters – were successfully solved by HRM with only 27 million parameters and just 1,000 training examples. This result challenges the prevailing belief that progress in artificial intelligence necessarily requires extremely large models trained for weeks on costly hardware and massive datasets.

Poland Has a Lot to Do

There is a great deal of work to be done, and the government is well aware of this. The Ministry of Digital Affairs has made significant progress over the past year. However, it is not enough to purchase AI laboratories for schools or build AI Factories. Follow-

Although many people find it hard to believe, barriers to the development of artificial intelligence are nowhere to be seen. Large AI models are not upgraded calculators.

ing the example of other European Union countries, we should invest in people: in doctoral programs, interdisciplinary studies, and new research centers preparing for large-scale AI applications, as well as in the development of senior academic staff, whose aging we have witnessed all too clearly. We lack data repositories and a program for systematic data accumulation that could be used to train AI models. We have the National Data Warehouse (KMD) and the dane.gov.pl platform, but there is little data of practical use there.

We need to significantly expand the PLLuM and Bielik models to make them suitable for social science analysis and to develop AI-based exercises that spark curiosity, teach how to formulate good questions, and create challenges that encourage the creative use of AI. That would be a far better solution than purchasing off-the-shelf software from Google or Microsoft.

Fidji Simo, the new General Manager for Applications at OpenAI, wrote: “I believe AI will create more opportunities for more people than any other technology in history. If we do our job well, AI can empower all people to a greater degree than has been possible so far.” If – and only if – we do it well. ■

Further reading:

Brennan K., Kak A., Myers West S., *Artificial Power: AI Now 2025 Landscape*, <https://ainowinstitute.org/2025-landscape>

Policy for the Development of Artificial Intelligence in Poland until 2030, <https://ai.gov.pl/media/2025/06/Polityka-rozwoju-sztucznej-inteligencji-w-Polsce-do-2030-r.pdf>

The 2025 AI Index Report, Stanford University Human-Centered AI, <https://hai.stanford.edu/ai-index/2025-ai-index-report>