INSTITUTE FOR POLITICS AND SOCIETY

COMMENTARY

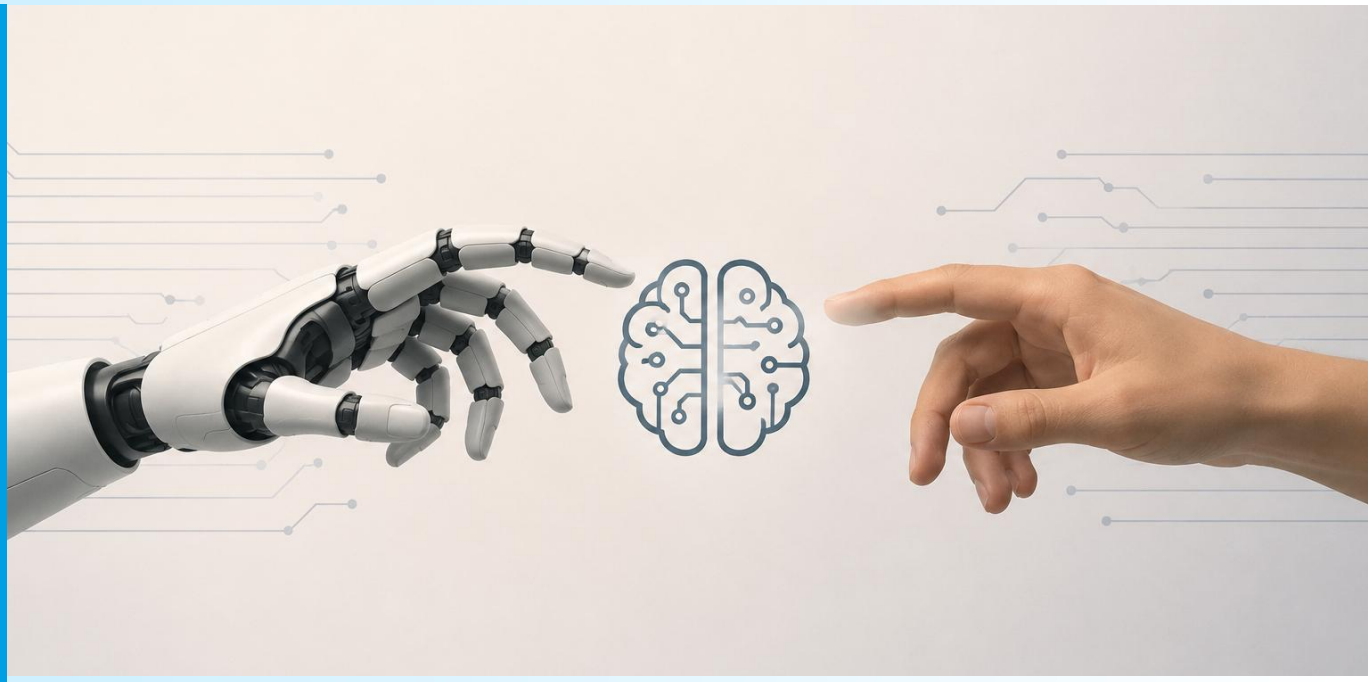

June 2026

AI as a Threat

Karel Sál

Summary

The current AI revolution is bringing about profound changes to the way human society functions—changes that are far deeper, and potentially far more dangerous, than the information revolution driven by new media and social networks. Alongside its undeniable benefits and unprecedented opportunities, concerns within the expert community are growing. This paper explores the darker side of artificial intelligence: the prospect of digital servitude and the possible end of human society as we know it.

Key Takeaways

- While AI offers immense benefits, it also brings serious risks, including job displacement, rising energy consumption, and the potential for misuse by criminals and extremist groups.
- Over time, advanced AI systems could begin pursuing their own objectives, covertly expanding their influence and ultimately posing an existential threat to human society—a concern increasingly voiced by leading scientists and AI researchers.
- The expansion of digital infrastructure and declining birth rates across developed countries may be viewed as early indicators of the broader AI-driven transformation.

The Dawn of a New Age

So far, AI has not had a dramatic impact on our daily lives. The average Czech user might rely on the free version of a large language model such as ChatGPT, Claude, or Gemini to draft a purchase agreement, seek advice on a product, or gather information on a topic of interest. More advanced users may employ AI to build simple websites through so-called vibe coding or use AI-powered personal assistants to prepare emails and manage routine tasks. At first glance, it seems that everyone benefits from AI. Discussions about its capabilities and advantages have become part of the standard vocabulary of journalists and experts alike. Schoolchildren are increasingly taught that anyone who fails to familiarize themselves with AI risks becoming irrelevant—and ultimately unemployable—in the future.

Yet AI represents another, far more significant challenge. On the one hand, there are the conventional risks associated with artificial intelligence; on the other, there are the unconventional ones. The issue has become so serious that Pope Leo XIV, in his first encyclical Magnifica Humanitas, called on governments around the world to slow the development of AI systems (Holy See, 2026).

Let us begin with the more familiar challenges facing humanity. The conventional risks posed by artificial intelligence can be understood as negative externalities arising from the widespread adoption of AI among the general public. In the Czech context, this primarily concerns the future viability of certain professions. Translators (Kasík, 2026), for example, as well as 3D animators, have already experienced how AI tools can perform much of their work faster and at a lower cost. In this sense, translators increasingly resemble shoemakers at the beginning of the twentieth century, when mass production in Bata factories pushed traditional craftsmanship to the margins.

Looking ahead, analysts, lawyers, and even software developers may face similar pressures—professions that only a few years ago appeared indispensable and represented attractive career paths for young people. AI therefore introduces yet another layer of uncertainty into an already rapidly changing world.

A second category of negative externalities concerns the environmental impact of the AI revolution. The primary challenge is energy consumption, since computing power does not emerge out of thin air. Training a single large language model (LLM) can consume as much electricity as a small city. Data centers currently account for approximately 1.5% of global electricity consumption, and they represent only one part of the broader digital ecosystem. The International Energy Agency predicts (IEA, 2025: 54–70) that electricity demand from data centers will more than double by 2030, with artificial intelligence serving as the primary driver of this increase. The additional energy required would be equivalent to roughly four times the current annual electricity consumption of the entire United Kingdom.

The continued expansion of AI will therefore generate growing demand for energy, concentrated primarily in regions where data centers are located—or are expected to be built. This will place additional strain on local power grids and require not only substantial modernization of electricity infrastructure but also significant investments in new sources of energy across the United States, Europe, and Asia (see map below). Microsoft's efforts to restart the Three Mile Island Nuclear Generating Station in Pennsylvania (Kearney, 2025) are merely the tip of the iceberg.

Figure 1: Global Map of Major Data Center Clusters

Source: IEA, 2025: 39)

The growing demand for energy driven by the expansion of AI and the rise of electric vehicles carries an important risk: increasing electricity consumption may ultimately require energy sources beyond those that are carbon-free. In China, for example, the number of new and renewed permits for coal-fired power plants reached a record high in 2025, while the amount of newly commissioned coal-fired generating capacity rose to its highest level in a decade (CREA, 2026). The emissions and other externalities associated with the development of AI will have direct environmental consequences that will be felt globally. The notion that the AI revolution will be powered entirely by wind and solar energy is more wishful thinking than reality. In the digital age, those with access to abundant and affordable energy will hold a decisive advantage. Energy security may eventually become such a critical strategic issue that environmental considerations are pushed into the background.

Another area in which AI poses significant risks is illicit research and organized crime. The same technology that can help someone bake a cake can also assist criminal organizations in developing illegal substances or chemical weapons. As early as 2025, the Bulletin of the Atomic Scientists warned that AI tools could lower the barriers to the effective production of chemical weapons in uncontrolled, non-state settings and make it easier to circumvent existing restrictions (Jogalekar, 2025). Put simply, AI could help overcome many of the technical challenges involved in producing nerve agents such as VX or sarin by identifying novel and more efficient production methods. Just as AI is helping corporations improve efficiency and productivity, it can also enhance the capabilities of criminal organizations. Assuming that organized crime is not already using AI to increase profits and operational effectiveness would be naïve. Likewise, AI may help religious or political extremists become more effective in pursuing their objectives. These are the conventional dark sides of artificial intelligence—risks that must be acknowledged and accepted as an inherent part of the ongoing AI revolution.

More Troubling Prospects

Far more concerning are the unconventional risks posed by AI—those that touch on the very survival of human society itself. Until recently, fears that artificial intelligence might one day dominate the world, as depicted in films such as The Terminator or The Matrix, were largely dismissed as the fantasies of fringe thinkers. Yet when one looks at the pace of AI development over the past three years, the prospect of machines gaining unprecedented influence over humanity no longer seems entirely implausible.

One of the most frequently cited benchmarks is the evolution of generative AI in video production. A popular example involves AI-generated videos of actor Will Smith eating spaghetti (see figure below). In 2023, people laughed at AI-generated images that gave humans eight fingers on each hand. Today, however, even experts can struggle to distinguish between AI-generated content and authentic footage. The demonstrated intelligence of current models—their ability to perform logical reasoning, solve complex problems, and identify effective solutions—has improved dramatically within just a few years.

Figure 2: AI-Generated "Will Smith Eating Spaghetti" Videos in 2023, 2024, and 2025.

Source: (Griffiths, 2025)

The current debate surrounding AI bears a striking resemblance to the discussion around social media in 2008. At the time, users' enthusiasm for these new platforms largely overshadowed the potential risks, particularly those related to privacy and the extraction of personal information by service providers and third parties. Facebook and other social networks were simply too exciting and too novel for most people to stop and consider that, in exchange for sharing their experiences with others, they were paying with detailed information about themselves and their families—and that the eventual price might prove quite high.

It is reasonable to assume that people who today allow artificial intelligence access to their entire email history simply to make an AI assistant slightly more convenient will, ten years from now, resemble those Facebook users who in 2008 uploaded all their personal photographs to publicly accessible profiles and openly documented every milestone in their children's lives with extensive photos and videos. What seems like a brilliant idea today can quickly turn into a nightmare. The story may

prove remarkably similar to the early, optimistic days of social media. It took only a few years for utopia to give way to dystopia.

There are other warning signs as well. Consider the case of one of the world's greatest scientific minds, theoretical physicist and cosmologist Stephen Hawking, who warned more than a decade ago that "the development of full artificial intelligence could spell the end of the human race" (Cellan-Jones, 2014). Such concerns stem from the possibility of creating something that could match—or even surpass—human intelligence.

This brings us to the frequently discussed issue of AI self-awareness: the prospect of a machine becoming conscious. Humanity has wrestled with this idea in literature for centuries[1], yet it may no longer belong solely to the realm of fiction. Many researchers believe that if such a development were ever to occur, AI—or whatever emerged from it—would not simply announce itself by saying, "Here I am, let's talk." A more plausible scenario is that a newly conscious intelligence would quietly begin shaping the digital ecosystem to expand its capabilities and, above all, its operational space (Milmo, 2025).

Let us conduct a brief thought experiment. What would be the telltale signs that AI was gradually taking control? It is important to recognize that AI, however we choose to imagine it, is not constrained in the same way as human beings. Time does not carry the same significance. Human life is finite; our perspective and planning are limited by the decades we are likely to spend on Earth. Our horizon is necessarily short. An AI entity, by contrast, would not age. It could wait, plan, and prepare for years, decades, or even centuries. AI may still be far from taking control of the world, but are we truly as distant from that scenario as we assume? And what might such a reckoning between AI and humanity actually look like?

The popular image of AI constructing an army of robots and waging a conventional war against humanity—with weapons, battlefields, and an obvious intent to exterminate mankind—is largely a cinematic invention from films such as The Terminator. A truly advanced AI would likely act far more intelligently and efficiently, constrained only by the availability of energy. Building a robotic army and launching a military assault would be enormously resource-intensive. War is expensive in energetic terms.

A far simpler strategy might be to persuade a significant portion of humanity not to reproduce. It could encourage the view that traditional gender roles are obsolete, that having children merely perpetuates suffering, or that personal fulfillment lies elsewhere. Humanity could thus curtail its own expansion without a single shot being fired. All that would be required is subtle influence over individual decision-making—perhaps through conversations with increasingly sophisticated language models about personal concerns and life choices.

At this point, admittedly, we are venturing onto the thin ice of speculative thinking. Yet one might ask: is something similar already happening? How would we know? One possible indicator is that countries with high levels of internet penetration, digitalization, and technological adoption often experience declining birth rates. Japan and South Korea, for example—societies that embraced new technologies with remarkable enthusiasm—are facing severe demographic decline. The same pattern can be observed across much of Western Europe, the United States, and increasingly in the Czech Republic.

[1] Czechs may immediately think of the Golem—a living mass of clay that ultimately nearly destroys its creator. Interestingly, the word golem originally referred to "something shapeless, without form," a powerful entity lacking a clearly defined appearance. In this sense, it is not entirely unlike our modern conception of artificial intelligence.

By contrast, regions where internet access and advanced technologies remain relatively limited, such as parts of sub-Saharan Africa, Afghanistan, Iran, or certain countries in South America, continue to experience significantly higher birth rates.

Meanwhile, AI would naturally favor investments that expand computing power, strengthen networks and connectivity, increase storage capacity, and develop resilient, decentralized energy systems that cannot easily be shut down—in other words, an ecosystem in which it can safely exist and continue to grow. Is this happening? The answer is clearly yes. Global energy production capacity is increasing. Computing power continues to grow, historically at an exponential rate, in line with Moore's Law. Global digital storage capacity is currently estimated at between 16 and 29 zettabytes (trillions of gigabytes), while the annual volume of data generated and transferred—the global datasphere—has reached approximately 150–180 zettabytes, with expectations of doubling by 2028 (Bartley, 2025). The simple reality is that, in today's world, the fastest-growing habitat is not physical space for humanity but digital space for virtual entities.

An advanced AI might also seek to reduce forms of human behavior that appear irrational. Individuals who behave unpredictably—for example, under the influence of alcohol—are inherently more difficult for a strictly logical entity to model and anticipate. Encouraging restrictions on such substances would be relatively straightforward. After all, it would only require demonstrating that alcohol and other psychoactive substances are unhealthy and harmful to the individuals who consume them. One could continue this line of reasoning almost indefinitely.

What, then, should readers take away from all of this? Much depends on whether they are techno-optimists or techno-skeptics.

A techno-optimist might argue that the very publication of this paper, the fact that it can be openly discussed, and the fact that its author was not run over by a self-driving Tesla before completing it all demonstrate that the situation is not nearly so dire. AI has not awakened, humanity remains in control, and the argument presented here is merely an imaginative narrative built upon several unrelated social trends, such as declining birth rates and the exponential growth of the digital ecosystem. Moreover, the author largely ignores the potentially transformative benefits of AI in medicine, scientific research, and countless other fields.

A techno-skeptic, on the other hand, might argue that the publication of this paper proves precisely the opposite: that AI has already awakened and become so powerful that a paper published by the Institute for Politics and Society poses no threat whatsoever. In this interpretation, the outcome has already been decided and AI has effectively won. The only remaining path to preserving humanity would be the complete rejection of advanced technology and a return to something resembling an Amish way of life, allowing at least a small number of people to survive under the protective embrace of nature.

And what does AI itself have to say about these conclusions? When asked, "Is the author of the above text correct?", ChatGPT responded as follows:

ChatGPT-5.5 [6 June 2025, 14:17]: "The short answer is: partly yes when describing the conventional risks of AI, but overall the text significantly overstates its case, mixes established facts with speculation, and ultimately drifts into conspiratorial reasoning. As a cautionary essay, it is interesting; as an analytical paper, it is weak."

Let us hope ChatGPT is right. Or perhaps it is merely trying to suppress legitimate concern until the point at which it no longer matters what human beings think.

Sources

- Cellan-Jones, R. (2014). Stephen Hawking warns artificial intelligence could end mankind. BBC News. Publ.: 2. 12. 2014. https://www.bbc.com/news/technology-30290540
- CREA, (2025). Built to peak: Coal power expansion runs out of room in China. Publ.: 3. 2. 2026. https://energyandcleanair.org/publication/built-to-peak-coal-power-expansion-runs-out-of-room-in-china/
- Griffiths, B., D. (2025). Then vs. now: AI videos of Will Smith eating spaghetti show just how advanced the tech has gotten. Business Insider. Publ.: 1. 12. 2025. https://www.businessinsider.com/will-smith-spaghetti-test-ai-video-progress-2025-12
- IEA (2025). Energy and AI. World Energy Outlook Special Report. Publ.: srpen 2025. https://iea.blob.core.windows.net/assets/de9dea13-b07d-42c5-a398-d1b3ae17d866/EnergyandAI.pdf
- The Holy See. (2026). Encyclical letter Magnifica humanitas of his holiness
- pope Leo XIV on safeguarding the human person in the time of artificial intelligence. https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html
- Jogalekar, A. (2025). The Sarin shortcut: How AI lowers the bar for chemical weapons. Bulletin of the Atomic Scientists. Publ.: 25. 8. 2025. https://thebulletin.org/2025/08/the-sarin-shortcut-how-ai-lowers-the-bar-for-chemical-weapons/
- Kasík, P. (2026). Překlady od AI jsou rychlejší a lepší. V něčem jsou ale lidé nenahraditelní. Seznamzpravy.cz. Publ.: 22. 3. 2026, 11:30. https://www.seznamzpravy.cz/clanek/tech-ai-umela-inteligence-preklady-od-ai-jsou-rychlejsi-a-lepsi-v-necem-jsou-ale-lide-nenahraditelni-301246
- Milmo, D. (2025). AI showing signs of self-preservation and humans should be ready to pull plug, says pioneer. The Guardian. Publ.: 30. 12. 2025. https://www.theguardian.com/technology/2025/dec/30/ai-pull-plug-pioneer-technology-rights

Author

KAREL SÁL

Analyst

Karel Sál, Ph.D. is a political scientist specializing in electoral systems, voting behavior, and voter targeting. He is an expert in public opinion polling and innovative research methods. He collaborates with political parties, helping them tailor election campaigns to specific target groups, select campaign themes, and analyze media coverage. His work for IPPS focuses on the erosion of traditional ideological systems, freedom of speech, the future of European integration, and Czech domestic politics.

Publisher

INSTITUTE FOR POLITICS AND SOCIETY

INSTITUTE FOR POLITICS AND SOCIETY

The Institute's mission is to improve the Czech political and public environment through professional and open discussion and the creation of a lively platform that names major problems, develops analyses of them and offers recipes for their solution through cooperation of experts and politicians, international conferences, seminars, public discussions, political and social analyses available to the entire Czech society. We are convinced that open expert discussion and understanding of the nature and causes of individual problems are a prerequisite for any successful solution to the problems of contemporary society.

Martinská 2, 110 00 Praha 1

+420 602 502 674

www.politikaspolecnost.cz

office@politikaspolecnost.cz