Article

Regulating artificial intelligence: A technology-independent approach

Žiga Turk

University of Ljubljana, Faculty of Civil and Geodetic Engineering, Ljubljana, Slovenia

Abstract

Successful applications of artificial intelligence (AI), such as ChatGPT, have been prompting regulators to speed up the related regulation processes. China and the EU have been particularly ambitious in this regard. The EU AI Act has been swiftly progressing through the institutions and is expected to be officially adopted in spring 2024. This article argues that its overall approach is wrong, that it extends EU regulation into policy areas which come under national competences and that it will hurt European AI innovation in particular and society in general. Instead of regulating AI per se, the EU or the member states should regulate the use of AI in specific sectors or, better still, regulate it in technologically independent ways—by specifying what is allowed or prohibited, regardless of the technology used.

Keywords

Artificial intelligence, AI Act, Regulation, Subsidiarity, Internet policy, Innovation policy

Introduction

In 2023 artificial intelligence (AI) made it into the consciousness of the general public. First ChatGPT and then a few other tools gave the impression that there was an intelligent being behind the screen. The answers to factual questions, summaries of articles, draft essays, colourful paintings and so on looked like the results of an intelligent human author. Even the AI scientific community was caught by surprise by how well these systems work (Hassenfeld 2023).

Corresponding author:
Professor Dr Žiga Turk, University of Ljubljana, Faculty of Civil and Geodetic Engineering, Jamova 2, 1000 Ljubljana, Slovenia.
Email: ziga.turk@gmail.com

This hype and success also gave rise to all kinds of speculation and fears. Stephen Hawking, Elon Musk and Yuval Harari, to name just a few, have long been warning about the existential threat that AI could pose to the human race (Taylor 2023). In the public perception, this evoked memories of Hollywood science-fiction films in which computers and robots take over the world and humans fight them for survival.

Meanwhile, the EU has been observing how China is using digital technology to monitor its citizens, how advertisers are using God knows what to serve up well-targeted ads on the Internet and how public opinion can be swayed by social media channels such as X, Facebook, Instagram and Tik Tok.

EU lawmakers have also been remembering how the Internet developed below their radar, and how it dramatically changed the media landscape, the consumption of news and, some believe, the results of referenda and elections (Rose 2017). The disinformation panic of the late 2010s was seen by some lawmakers as a reminder that one should regulate new technology before it is too late, before it outgrows politics. This mistake was not to be repeated with AI.

It would seem, then, that European politicians were responding to an urgent public need when, on 9 December, the Parliament and the Council agreed on a 'historic deal' on 'the first rules for AI in the world' (Council of the EU 2023). The EU might not have any competitors in the global AI race—all the main players are American or Chinese—but the EU's global relevance in this field is in regulating it.

These were not, however, the first rules for AI in the world. In 2022, 37 countries passed regulations related to AI, and more followed in 2023 (Lynch 2023). In August 2023 China adopted a comprehensive generative AI regulation. It imposes restrictions on the training data used and on the outputs produced by tools that offer generative AI services to consumers (Roberts and Hine 2023).

Given all this innovation activity, on the one hand, and regulatory ambition, on the other, consideration should be given as to whether the EU's approach to AI regulation is appropriate. And if not, what is the alternative?

Artificial intelligence

There is a lot of confusion about AI. To put it simply, AI is a computer program that, like any other computer program, given some input, provides some output. The particular branch of AI that has been making waves recently is very good at two kinds of tasks. It can find patterns in the data given as input, and it can generate patterns similar to those it has been trained with, given a prompt. The first feature makes it very efficient, for example, at diagnosing illnesses, finding potential criminals or sifting through job applicants. The second feature allows it to write summaries, reports and essays; generate pictures; and compose music.

Patterns are everywhere. Intelligent human behaviour is about recognising and creating patterns. When someone says 'hi', and we respond with 'hello', this is a pattern. When a glass breaks when it hits the floor, this is a pattern. That day follows night is a pattern. That this article has an abstract, introduction, a couple of sections, conclusions and references is a pattern. That a human face has two eyes, two ears, a nose and a mouth is a pattern. Humans know some patterns from experience—for example, how to recognise faces or animals. And they have reduced other patterns to abstractions and mathematical formulae. The velocity of that glass hitting the floor is governed by a formula that takes the square root of twice the height from the floor multiplied by the acceleration of gravity. Current AI is very good at learning from experiences but not good at all at reducing specific problems to general theories and creating abstractions (Wolfram 2023).

Because patterns are everywhere and because language and pictures are the means by which we communicate, the AI under discussion here is generic technology. It can be applied to any issue that can be explored through language or pictures. The generation of content is protected by freedom of expression, and humans should not be limited in terms of how AI helps them to be creative. Recognising patterns should not be prohibited either, because it is the essence of rational, scientific work. It is therefore questionable whether regulating AI—as a generic technology—makes sense at all, because this would mean regulating what humans are allowed to discover and what they are allowed to create.

The EU's approach to AI regulation

The areas in which AI has the potential to be used are huge, and the need to somehow partition these into manageable pieces is understandable. The EU approach is to classify AI applications into four categories according to the risk they pose (Whyman 2023):

- *Unacceptable risk*. Here belong AI systems that are prohibited due to clear threats to safety, livelihoods and rights. Examples include social scoring by governments and real-time biometric identification for law enforcement in public spaces.

- *High risk*. AI applications falling under this category are subject to strict pre- and post-market regulations. These include applications used in critical infrastructure, education, employment, law enforcement and judicial systems.

- *Limited risk*. This concerns AI systems that require transparency to ensure that users know they are interacting with AI. An example is AI chatbots.

- *Minimal or no risk*. Most AI systems can be freely used without legal obligations but its providers are encouraged to follow codes of conduct. Examples include AI for video games and spam filters.

The EU AI Act expects that all AI 'services' would be registered and stipulates prohibitive fines for businesses that use AI in breach of the Regulation. While at first glance this

seems reasonable, one wonders what a similar list would have looked like for two just-as-revolutionary technologies from the past—the printing press and the Internet.

Earlier approaches to regulating generic technologies

Early printing-press regulations included ad hoc censorship of some materials, as exemplified by the *Index Librorum Prohibitorum*, first published by the Catholic Church in 1559. The Licensing of the Press Act 1662 in England required printers to obtain a licence, effectively enabling the government to control what could be published. The late eighteenth century saw a liberalisation of the rules and the introduction of the first 'freedom of expression' acts, most notably the American First Amendment (1791) and the French Declaration of the Rights of Man and of the Citizen (1789). Subsequent legal instruments, such as the Universal Declaration of Human Rights (1948), further internationalised the idea of freedom of expression (Barendt 2005).

Thus, press regulation matured into limitations on how the state could use its power to interfere with the press and, more generally, with communication. The state was allowed to intervene and limit speech in just a short list of situations including matters of national security, the incitement of violence, Holocaust denial and child pornography. Otherwise, it was told to stay out of the way. The founding fathers of modern democracies did not elaborate on or categorise the ways in which the technology of the printing press could hurt citizens.

Another general-purpose technology is the Internet. In the early 1990s it was at a stage of development and public awareness similar to that of AI today. There were also ideas to regulate it in a general way, modelled on telecoms regulations. However, in the end, only a few targeted regulations emerged, and these actually protected the Internet from being governed by laws more suited to the era of the printing press. Most notable was the Telecommunications Act of 1996, Section 230 of which provided immunity to Internet service providers and platforms from liability for content posted by their users.

Neither with the printing press nor with the Internet did lawmakers try to come up with a list of possible uses and then classify them as ranging them from 'unacceptable risk' to 'minimal or no risk'. The exceptions to this were the Church in the sixteenth century and some twentieth-century dictatorships, which classified quite a few books as being of 'unacceptable risk' to their potential readers.

Regulation to do good

Two patterns of AI regulation are emerging globally. The *ex ante* kind is trying to foresee all possible evils and potential for wrongdoing. China and the EU are in this lawmaking camp. Not only are they trying to protect their citizens from all kinds of real and perceived dangers, but they are also explicitly stating the noble goals that AI should be expected to pursue.

The OECD (2022) recommends that AI should contribute to inclusive growth and sustainable development, that it should enhance the well-being of people, and that the

design and deployment of AI systems should respect human-centred values and fairness, including privacy, dignity and diversity. Furthermore, it notes that AI systems should operate transparently and that their workings should be understandable and explainable.

These are all nice-sounding goals, but they quickly become restrictive. What if an AI discovers a pattern that, for example, shows that the presence of the Y chromosome is by far the best predictor of someone's gender? How will that work with people who believe that gender is a choice? That results should be explainable is a particularly high-order expectation because the mathematics behind some branches of AI are simply too complex for us to understand why one answer is given rather than another.

Technology is values-neutral. It is the use of technology, from a stone knife to nuclear energy, that can be good or bad, moral or immoral, not technology per se.

Other approaches to AI regulation

Those countries which have their legal basis in common law are counting on case law to set precedents once specific disputes come to court. The US, Canada, India and Switzerland are all taking a more laissez-faire approach. Rather than creating a separate law for AI regulation, Switzerland plans to adapt existing laws to regulate the use of AI systems in areas including data protection, equal treatment, competition, product liability and civil law (Kohn and Pieper 2023).

This approach seems reasonable. AI is a general-purpose technology. Patterns can be discovered anywhere. Patterns can be created about anything. Regulating technology as such limits what the tool can actually do. It limits the kind of innovation that takes place. This is why sweeping AI regulation is bound to slow down AI research and development.

Even French President Macron is critical of the EU's approach to AI regulation, stating that it is 'not a good idea' and that it is a risk to innovation, potentially putting EU technology companies at a disadvantage compared to their counterparts in the US and the UK (Caddle 2023). Macron has also emphasised that regulation should be controlling and not punitive. His focus here is on ensuring that the new laws facilitate innovation rather than penalising or overly constraining AI development.

The general European regulation of AI also hides the danger of a further reduction of subsidiarity within the Union. It regulates the use of a technology that is applicable in many different policy areas, including internal security, justice, education, healthcare and culture. These are not European but member state competences. Thus, Brussels has found yet another way to extend its powers and create an even closer Union.

The alternative to regulating AI as such is to regulate where, how and by whom AI can be used. For old-fashioned liberals, it would make sense to introduce regulation that limits the use of AI systems by governments, thus preventing the creation of a Big

Brother–style dystopia. Alternatively, regulation could be introduced, along the lines of that introduced for the printed press, that limits governments' ability to restrict businesses and individuals from using AI.

As for where and how AI can and cannot be used, two approaches can be taken: a technology-specific approach and a technology-independent approach. A technology-specific approach would prescribe how a specific technology—that is, AI—can or cannot be used in this or that area. It would, for example, amend health care regulations to include provisions on how AI can or cannot be used. In the same way, regulations on public safety, taxation and equality before the law would be amended. The law on policing would state, for example, that AI cannot be used for face recognition in public spaces. A social security law would prohibit (or not prohibit) social scoring. Work and employment legislation would prohibit (or not prohibit) the use of AI in job-application screening.

In contrast, a technology-independent approach would specify which behaviours are prohibited, regardless of which technology is used. For example, it would prohibit the impersonation of people, be it by a human actor or by an AI-generated avatar. It would prohibit the discovery of patterns in crime and predictive policing, regardless of whether the patterns were discovered by AI, Excel or just by an old cop's experience.

Conclusion

The author believes that the most reasonable approach is a technology-independent one that regulates the specifics of particular issues. In addition, case law should be derived from real-world disputes. By taking the alternative approach, the EU is likely to end up even further behind the US, the UK and China, for two reasons: (1) its AI development would not be as ambitious as that of the competition and (2) AI users in all areas of life and work would be deprived of the best tools available globally.

AI is just a tool—a complex and powerful tool, but in the end, just a tool. Regulation should not be concerned with what kinds of tools are built, but how and where such tools may or may not be used. Or to be more precise, regulation should continue to define what governments are allowed to do and what citizens and businesses are not allowed to do—whether with AI tools, other tools or manually, it does not matter.

'Thou shalt not kill' is not concerned with the tool used for killing. It applies to stones, knives, guns, AI or any other tool or weapon.

References

Barendt, E. (2005). *Freedom of speech*. Oxford: Oxford University Press.
Caddle, P. (2023). EU's AI regulation framework 'not a good idea', Macron says. *BrusselsSignal.eu*, 12 December. https://brusselssignal.eu/2023/12/eus-ai-regulation-framework-not-a-good-idea-macron-says/. Accessed 13 February 2024.

Council of the EU. (2023). Artificial Intelligence Act: Council and Parliament strike a deal on the first rules for AI in the world. Press release, 9 December. https://www.consilium.europa.eu/en/press/press-releases/2023/12/09/artificial-intelligence-act-council-and-parliament-strike-a-deal-on-the-first-worldwide-rules-for-ai/. Accessed 13 February 2024.

Hassenfeld, N. (2023). Even the scientists who build AI can't tell you how it works. *Vox*, 15 July. https://www.vox.com/unexplainable/2023/7/15/23793840/chat-gpt-ai-science-mystery-unexplainable-podcast. Accessed 13 February 2024.

Kohn, B., & Pieper, F.-U. (2023). AI regulation around the world. *TaylorWessing.com*, 9 May. https://www.taylorwessing.com/en/interface/2023/ai—are-we-getting-the-balance-between-regulation-and-innovation-right/ai-regulation-around-the-world. Accessed 1 January 2024.

Lynch, S. (2023). 2023 state of AI in 14 charts. *Stanford University, Human-Centered Artificial Intelligence*, 3 April. https://hai.stanford.edu/news/2023-state-ai-14-charts. Accessed 13 February 2024.

OECD. (n. d.). OECD AI principles overview. https://oecd.ai/en/ai-principles. Accessed 13 February 2024.

Roberts, H., & Hine, E. (2023). The future of AI policy in China. *EastAsiaForum.org*, 27 September. https://www.eastasiaforum.org/2023/09/27/the-future-of-ai-policy-in-china/. Accessed 13 February 2024.

Rose, J. (2017). Brexit, Trump, and post-truth politics. *Public Integrity*, *19*(6), 555–8.

Taylor, C. (2023). 'Sapiens' author says AI is an alien threat that could wipe us out: 'Instead of coming from outer space, it's coming from California'. *Fortune.com*, 12 September. https://fortune.com/2023/09/12/sapiens-author-yuval-noah-harari-ai-alien-threat-wipe-out-humanity-elon-musk-steve-wozniak-risk-cogx-festival/. Accessed 13 February 2024.

Whyman, B. (2024). AI regulation is coming—what is the likely outcome? *Center for Strategic and International Studies*, 10 October. https://www.csis.org/blogs/strategic-technologies-blog/ai-regulation-coming-what-likely-outcome. Accessed 13 February 2024.

Wolfram, S. (2023). What is ChatGPT doing . . . and why does it work? *StephenWolfram.com*, 14 February. https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/. Accessed 13 February 2024.

Author biography
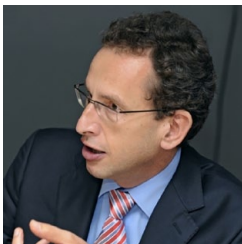
Professor Dr Žiga Turk *is a professor at the University of Ljubljana, Slovenia. He was both minister for growth and minister of education, science, culture and sport in the government of Slovenia, and was secretary general of Felipe Gonzalez's Reflection Group on the Future of Europe. He is a member of the Academic Council of the Wilfried Martens Centre for European Studies.*