

AI - to legislate, or not to legislate, that is the question

Brian Williamson¹, February 2023



The European Union² propose legislation encompassing Artificial Intelligence (AI), the US National Institute of Standards and Technology (NIST) has unveiled a voluntary risk management framework³, whilst the proposed UK approach is context specific supported by a set of cross-sectoral non-statutory principles.⁴ Other countries are also contemplating legislation or other measures in relation to AI.

The question this note addresses is whether it makes sense to legislate in relation to an emerging general-purpose technology akin to fire, steam, electricity, computing or AI, or to adapt as required as applications and challenges emerge in specific contexts.

Is legislation for a general-purpose technology the answer?

We should take a step back and consider whether legislative rulemaking in relation to general purpose technologies is sensible. We did not legislate in relation to other general-purpose technologies including steam, electricity, and computing.

Nevertheless, their applications were subject to existing legislation and regulation, and as applications developed existing rules were modified and new rules developed. For example, we did not have a law of steam, but we did regulate in relation to rail safety.

The key to benefiting from AI is entrepreneurial experimentation in a diverse range of applications, and successful use cases and policy challenges are unlikely to be easy to predict. For example, 'creative' endeavours such as art and writing were not seen as ripe for early disruption by AI, yet generative AI looks set to do so and is raising questions over, for example, how to interpret Copyright.⁵

A law of AI cannot hope to correctly identify and address the specific challenges that will arise, whilst it is almost guaranteed to reduce innovation and market entry; and perpetuate less efficient, and potentially more risky, existing ways of doing things.

What about ethics for AI?

The application of ethics to AI is viewed by some as distinct and specific. Yet there is a risk that,

¹ Partner, Communications Chambers (brian@commcham.com). This paper represents the view of the author only, and does not represent a corporate view of Communications Chambers.

² The Artificial Intelligence Act. <https://artificialintelligenceact.eu>

³ NIST, AI Risk Management Framework, January 2023. <https://www.nist.gov/itl/ai-risk-management-framework>

⁴ Department for Culture, Media and Sport, National AI Strategy, December 2022. <https://www.gov.uk/government/publications/national-ai-strategy/national-ai-strategy-html-version#pillar-3-governing-ai-effectively>

⁵ The Verge, The scary truth about AI copyright is nobody knows what will happen next, November 2022. <https://www.theverge.com/23444685/generative-ai-copyright-infringement-legal-fair-use-training-data>

without a comparison to existing ways of doing things and without consideration of the ethical underpinnings of law and economics, that the pursuit of ethics for AI could prove counterproductive.

Both an efficient market and public policy appraisal based on cost-benefit analysis embody the Pareto principle, namely that progress involves outcomes where some are better off and could in principle compensate anyone made worse off. Further, the State plays a role in correcting for market inefficiencies and in redistributing the gains from progress. There are ethical underpinnings to the operation of markets and the State, irrespective of the technology in question.

Placing specific requirements on the application of AI would mean not only foregoing more efficient ways of doing things, but also potentially perpetuate higher risk ways of doing things. Paradoxically placing even greater requirements on the application of AI in areas judged high-risk, such as medicine, may perpetuate risk by limiting or slowing scope for AI to reduce overall risk by complementing or replacing human decision making.

Should the 'promise' of AI prompt new limits?

There will be instances where AI is so much more powerful than existing ways of doing things that we may wish to impose limits on what can be done, either by machines or humans or both.

An example is autonomous vehicles, which may substantially reduce accident rates. In the first instance this may be a reason to apply the human driver standard in allowing autonomous vehicles, to facilitate transition. Ultimately, however, it is likely to prompt the question of whether standards for autonomous and human-driving should be raised, potentially to the point where humans find it difficult or impossible to pass the requirement. In other words, should we sometimes keep humans out of the loop, a question that will likely prove politically fraught.

Another example is state surveillance (surveillance socialism?). AI offers the prospect of an expansion of surveillance which some might argue would be justified on grounds of crime reduction or pursuit of other public goods. However, this potential can be expected to prompt debate about limits on surveillance, a debate that a rights-based analysis may not settle given the likely existence of competing rights. However, if the potential of AI does prompt us to set new limits on surveillance, they should arguably not be AI specific. After all the Ministry of State Security (Stasi) in the former East Germany was able to conduct extensive surveillance without the help of AI.

Finally, we may in some instances choose to keep humans in the loop even where their performance is poor, at least judged against some criteria. For example, we may prefer to keep trial by jury even were AI more efficient, less noisy (variable) and perhaps even less prone to bias. Trial by a jury may nevertheless be regarded as more legitimate and less prone to non-transparent manipulation.

The above examples illustrate that limits are likely to prove context dependent, and that we may in some cases choose to limit the scope for human decision-making rather than that of machines.

Conclusion

A law of AI or any other general-purpose technology involves a category error, introducing a barrier to the use of an input rather than focussing on outputs and outcomes.

A law of AI would not only forego potential productivity and income benefits, but also perpetuate more risky ways of doing things in areas such as health care and transport, and that would hardly be ethical. Instead of crafting a law of AI we should facilitate the use of AI and reinterpret the application OF existing rules, and adopt new ones where require, prompted by the development of new applications in specific contexts.