

# How AI Can Be Regulated Like Nuclear Energy

AI labs compare AI risk on par with the existential and safety risks of nuclear, yet they dismiss regulations.

---

By Heidi Khlaaf

5 min. read · [View original](#)

---

Prominent AI researchers and figures have consistently dominated headlines by invoking comparisons that AI risk is on par with the existential and safety risks that were posed with the coming of the nuclear age. From statements that AI should be subject to [regulation akin to nuclear energy](#), to declarations paralleling the risk of [human extinction](#) to that of [nuclear war](#), the analogies drawn between AI and nuclear have been consistent. The argument for such extinction risk has hinged on the hypothetical and unproven risk of an [Artificial General Intelligence \(AGI\)](#), imminently arising from current Large Language Models (e.g., ChatGPT), necessitating increased caution with their creation and deployment.

Sam Altman, the CEO of OpenAI, has even referenced to the well established nuclear practice of “licensing”, deemed [anti-competitive](#) by some. He has called on the creation of a federal agency that can grant licenses to create AI models [above a certain threshold of capabilities](#). This is akin to how operators of nuclear facilities are required to be licensed by a nuclear regulator, such as the US Nuclear Regulatory Commission and the UK Office for Nuclear Regulation. As part of their licence conditions, nuclear operators must comply with regulatory safety guidelines and a legal duty to reduce risks so far as is reasonably practicable.

The incessant references to nuclear safety should naturally lend urgency in applying nuclear-level regulations and assessments to the very AI systems being deliberated. Yet, there is only resistance to be met at the mention of the EU AI Act, a whittled down document compared to the intricacies and rigour of the regulations required for nuclear systems. Indeed Altman has previously noted that the EU AI Act would be “over-regulating”, specifically noting OpenAI's systems as they exist are not high-risk, and that OpenAI will cease to operate if they [cannot comply with the upcoming EU AI Act](#).

Two conflicting messages are thus being presented. One being that the future of AI

capabilities present such a significant risk, that an exclusive nuclear-level regulatory licensing scheme is required, but also that existing systems do not warrant any regulation, despite [well-documented harms](#), and deemed insignificant. And herein where the analogies to nuclear safety crumble.

***[Read More: The Most Influential People in AI 2023](#)***

Safety Engineering, the discipline which assures the safety of nuclear systems, is built on the foundation that system failures or misbehaviours, no matter how small, can compound and lead to catastrophic events. Consider that all components of a nuclear system, that being uranium or plutonium, hardware, and software, down to literal concrete, are individually and rigorously regulated. These elements separately do not pose a threat of nuclear catastrophe, but in orchestration lead to the emergence of a controlled nuclear fission reaction. Conversely, failures in said orchestration can cascade into catastrophic or high-risk events. For example, a minuscule software bug known as a race condition led to the Northeast blackout of 2003.

Our inability to prevent today's AI harms, such as algorithmic discrimination and reducing the cost of disinformation or cybersecurity attacks, only entails that we are ill-prepared to trace and

grasp any cascading implications and control of AI risks. And if we lack the technical and governing facilities to control or measure AI mechanisms which output these harms, then we would be mistaken to believe we have the technical foundation to resolve larger-scale avalanching risks that AI may hypothetically enable.

Then why have AI labs shown resistance to following through with the nuclear-level rigour required to keep the predecessors (e.g., ChatGPT) of “extinction level” technology safe? In a [blog post](#) by the OpenAI founders, they note that “it would be important that such an [regulatory] agency focus on reducing existential risk and not issues that should be left to individual countries.” This statement attempts to not only detach the existing harms AI systems pose from hypothetical existential risks, but fails to recognise the fundamental engineering principle that harms compound. If current AI systems are indeed the foundations of a hypothetical AGI, as is often implied, then as they stand they too must be regulated to address current harms in preparation for further emergent behaviour, whether it be AGI or otherwise.

A pattern emerges that invoking the analogy to nuclear war often serves to inflame specific narratives, whether it be exclusive licensing, or an attempt to divert the regulatory conversation

to focus on speculative future threats instead. Both having potentially harmful impacts. For the former, AI-based systems do not possess any unique software components that warrant a generalized licensing scheme that would not heavily impede the use of software and hardware as a whole. Indeed, any implementation of such a scheme would likely result in significant overreach due to the broad definition and software components of AI systems, lending basic blocks for technological advancements to be available to a privileged few. After all, current AI systems are built using traditional hardware and software components. Even Generative AI technologies utilise Deep Neural Networks, a technique that dates back to the 1970's.

For the latter (abiding by the foundations of nuclear safety aside), current AI systems have demonstrated enough harms to [warrant regulation](#) even without an extinction threat. Nuclear-level regulation methodologies already exist for the very same hardware, software, and data components, ready to be adopted any day for AI labs to utilise for high-risk applications. Yet we have observed that the very pushback and lobbying by AI labs against the EU AI Act has proven to be fruitful, having already distracted from regulatory efforts relative to the harms posed by these AI systems today. A [post](#) by the European Commission reflected the exact wording of the [extinction letter](#),

"Mitigating the risk of extinction from AI should be a global priority", a far cry from the risks outlined in the initial drafts of the act.

If AI labs are to consistently invoke exaggerated fears through comparisons to nuclear hazards, then they must be willing to take the analogy to completion. Ironically, if they were to explore the readily available safety mechanisms for nuclear components, they would be pressed to find inflammatory language supporting the use of terminology such as "extinction" and "existential risks". Indeed, nuclear catastrophe is supported by known scientific and geo-political capabilities that have been theoretically and meticulously studied even during the Manhattan Project. The risks identified are tethered to tractable and compounding risks. Meanwhile, there exists no scientific basis or evidence for how or when AGI will emerge (if ever), leaving us with only a hypothetical risk that has capitulated many of our regulatory efforts today.