Balancing Innovation and Regulation: A Critical Look at the AI Act's Research Exemptions

Pietro Giovanni Bizzaro
Digital Commons Lab Unit of Digital Society Center of Bruno Kessler Foundation. Mail: pbizzaro@fbk.eu

Maurizio Napolitano
Digital Commons Lab Unit of Digital Society Center of Bruno Kessler Foundation. Mail: napolitano@fbk.eu

1. Introduction

The European Union's Artificial Intelligence Act represents a milestone in global technology regulation. Its ambition is twofold: to safeguard fundamental rights while fostering innovation in artificial intelligence. Yet, this dual mandate contains an inherent tension. On the one hand, the Act explicitly recognizes the need to "promote scientific and technological progress".[1] On the other, the framework it establishes for research exemptions risk becoming so complex and ambiguous that it may stifle precisely the research activities it seeks to protect. [2]

The challenge reflects a broader dilemma in contemporary technology governance: how to design rules that protect society without discouraging innovation. The AI Act attempts to resolve this tension through two distinct research exemptions: a general provision in Article 2(8), which excludes AI systems or models developed solely for research and development provided they are not placed on the market or put into operation, and a more restrictive scientific research exemption in Article 2(6),[3] which frames scientific research within narrower conditions. Yet, as this analysis will demonstrate, these provisions suffer from conceptual ambiguities and interpretive uncertainties that render them inadequate tools for their intended purpose.[4]

2. AI Act and research

AI systems and models developed solely for research and development are excluded from the AI Act's scope until they are placed on the market or put into operation. While this is meant to give freedom to innovate, the unclear boundaries of exemption create uncertainty and may even push research activities outside the EU's regulatory sphere. The effectiveness of the Act in promoting innovation depends on legal certainty, something the current exemption fails to provide.

The research privilege applies only while an AI system or model is used exclusively for research, testing or development. Once it is used beyond those purposes, the full provisions of the AI Act apply. This raises difficult questions: must a trained model be checked for compliance all the way back to its origin? Self-evidently, such a requirement would be practically impossible.

[1] https://european-union.europa.eu/principles-countries-history/principles-and-values/aims-and-values_en#:~:text=promote%20scientific%20and%20technological%20progress,and%20solidarity%20among%20EU%20countries.

[2] L. Colonna, *The AI Act's Research Exemption: A Mechanism for Regulatory Arbitrage*?, In E. Gill-Pedro, A. Moberg, (directed by) *YSEC Yearbook of Socio-Economic Constitutions*, Cham, 2023. https://doi.org/10.1007/16495_2023_59.

[3] H. Hermanns, A. Lauber-Rönsberg, P. Meinel, S. Sterz, H. Zhang, *AI Act for the Working Programmer* in B. Steffen (directed by) *Bridging the Gap Between AI and Reality 2024*. *Lecture Notes in Computer Science,* 2025, Cham, 15217. https://doi.org/10.1007/978-3-031-75434-0_6.

[4] João Pedro Quintais, *What Is a 'Research Organisation' and Why It Matters: From Text and Data Mining to AI Research*, in *GRUR International*, May 2025, 74, 5, 397–398. https://doi.org/10.1093/grurint/ikaf030.

An AI system or model is considered "placed on the market" when it is first made available commercially (see also the definition in Art. 3 para. 9 EU AI Act). Products or services sold, licensed or passed to external users fall under the Act's requirements and research institutions that release an AI model as a finished product must ensure compliance.

"Putting into operation" means actual use beyond research and testing (see also the definition in Art. 3 para. 10 AI Act). Once a system is used in a real environment with real data, it is then no longer covered by the research privilege.

Open-source AI systems are excluded from the Act's obligations until they are placed on the market or put into operation as part of a high-risk, prohibited, or transparency-bound system. AI systems released under open-source licences may be exempt from transparency duties if their parameters, weights, architecture and usage are public. Nevertheless, exemptions are lost for high-risk, prohibited, or systemic-risk applications, even if open source, and do not waive all transparency duties, particularly for GPAI models with systemic risk.

This position paper argues that the AI Act's research exemptions, solely from the document reading, represent a regulatory hiccup that demands immediate remedial action. The fundamental problem lies not in the boundaries of the exemptions, but in their conceptual incoherence and practical inapplicability to modern research practices. The exemptions' reliance on outdated distinctions between commercial and non-commercial research, their failure to account for the collaborative nature of contemporary AI development, and their ambiguous relationship to core AI Act requirements create a regulatory environment that is both unpredictable and potentially counterproductive.[5]

While initial assessments show a clear regulatory conclusion, it remains ambiguous regarding downstream modifications of general-purpose AI foundation models and systems. Current interpretations suggest that fine-tuning might transform the developer entity into a downstream model provider, though precise thresholds remain undefined.[6]

Much like the AI Act, previously the GDPR has introduced special rules for scientific research through Article 5(1)(b) and Article 89, seeking to create safe harbors for researchers by allowing further data processing for scientific purposes that would otherwise be restricted.[7] However, as with the AI Act, these exemptions have proven to be seldom functional, being more formalistic than effective.[8] In practice, GDPR's safeguards (such as strict requirements for anonymization, proportionality assessments, and technical-organizational measures) were intended to

[5] P. Hacker, A. Engel, M. Mauer, *Regulating ChatGPT and other Large Generative AI Models* in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '23), Association for Computing Machinery, 2025, New York, NY, USA, 1112–1123. https://doi.org/10.1145/3593013.3594067.

[6] European Commission, *Guidelines on the scope of obligations for providers of general-purpose AI models under the AI Act*, 18 July 2025 https://digital-strategy.ec.europa.eu/en/library/guidelines-scope-obligations-providers-general-purpose-ai-models-under-ai-act.

[7] European Parliament, Directorate-General for Parliamentary Research Services, *How the general data protection regulation changes the rules for scientific research*, Publications Office, 2019, https://data.europa.eu/doi/10.2861/17421.

[8] C. Staunton, S. Slokenberga, D. Mascalzoni, *The GDPR and the research exemption: considerations on the necessary safeguards for research biobanks*, Eur J Hum Genet, 2019 Aug, 27(8), 159-1167. doi: 10.1038/s41431-019-0386-5. Epub 2019 Apr 17. PMID: 30996335; PMCID: PMC6777499.

provide flexibility but have often erected bureaucratic hurdles, especially for smaller research institutions or in highly sensitive fields such as health and bioengineering.[9]

The AI Act seems to repeat similar mistakes by anchoring its research exemptions in rigid and outdated distinctions between commercial and non-commercial research, while failing to reflect the deeply collaborative and cross-organizational dynamics of modern AI research. Just as the GDPR exemptions fell short of providing a meaningful operational pathway despite their protective intentions, the AI Act's exemptions risk creating uncertainty rather than clarity. In both frameworks, the tension arises from a conceptual gap: the law acknowledges the importance of fostering research but provides exemptions in overly narrow or impractical terms, leading to environments where researchers face disproportionate compliance burdens.

Drawing this connection shows that the issue at stake is not only one of regulatory alignment but also of regulatory design. Both the GDPR and the AI Act recognize research as a public good worthy of special treatment, yet both risk undermining this goal by creating exemptions that are ambiguous in scope and onerous in practice. This suggests that without corrective action, the AI Act research exemptions may follow the same trajectory as those under GDPR, well-intentioned provisions whose bureaucratic constraints hinder, rather than support, scientific progress.

3. Opportunities

Despite its challenges, the AI Act offers significant opportunities to foster a more ethical, globally influential, and responsibly innovative AI ecosystem in Europe. It raises ethical awareness by requiring transparency, oversight, and risk mitigation, particularly for high-risk systems[10]. This pressure pushes organizations to consider not only technical performance but also societal impacts, bias mitigation, and fundamental rights at every stage of AI development.

By introducing some of the world's strictest AI rules, the Act positions the EU as a global standard-setter creating the so-called "Brussels effect"[11]. This both harmonizes standards within the EU and encourages global firms to follow EU norms, fostering a race towards higher regulatory benchmarks.

The AI Act also sets demanding rules for open-source AI. While many developers have claimed openness, few meet these strict definitions. This helps prevent "open washing"[12], where systems labeled as open without real transparency or accessibility. Such practices risk undermining genuine innovation and eroding the public.

[9] *Ibid*. See also, C. STAUNTON, S. SLOKENBERGA, A. PARZIALE, D MASCALZONI, *Appropriate Safeguards and Article 89 of the GDPR: Considerations for Biobank, Databank and Genetic Research*, Front Genet, 2022 Feb 18;13, 719317. doi: 10.3389/fgene.2022.719317. PMID: 35251121; PMCID: PMC8896881.

[10] V. BOLGOURAS, A. ZARRAS, C. LEKA, et al, *Eu regulatory ecosystem for ethical AI. AI Ethics*, https://doi.org/10.1007/s43681-025-00749-x.

[11] A. BRADFORD, *The Brussels Effect: How the European Union Rules the World*, 2020, Faculty Books. 232. https://scholarship.law.columbia.edu/books/232.

[12] A. LIESENFELD, M. DINGEMANSE, *Rethinking open-source generative AI: open-washing and the EU AI Act*, in *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '24), June 03–06, 2024, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3630106.3659005.

4. Drawbacks

On the other hand, the AI Act's implementation creates major operational and strategic challenges for research in Europe. Approval processes and documentation duties add bureaucratic overhead that can delay time-sensitive projects, sometimes leaving questions outdated before results emerge. These demands shift resources away from innovation and toward compliance. These administrative demands disproportionately impact SMEs and small research centers, which often lack compliance teams and the financial capacity to manage complex, evolving regulatory requirements, exacerbating inequalities in access to innovation pathways[13].
This environment risks fostering "compliance-driven research," where regulatory safety becomes more important than bold experimentation, potentially steering projects away from high-impact, but complex or sensitive areas. For critical sectors such as medicine, neurotechnology, and bioinformatics, the layering of AI Act obligations atop existing frameworks makes compliance particularly burdensome, threatening to stifle progress in fields where rapid, flexible research is most urgently needed[14].

5. Conclusions

In sum, the AI Act's research exemptions, while rooted in a praiseworthy ambition to balance scientific progress with rights protection, may suffer from a lack of conceptual clarity and operational feasibility that threatens to undermine innovation, especially in critical fields like health and bioengineering[15]. Without prompt action from policymakers to issue detailed guidelines and establish clear, standardized models and predictable safe harbors for research, Europe risks stifling the very scientific vitality it aims to foster, placing undue burdens on small research institutions and sensitive sectors while paving the way for a compliance-first, rather than discovery-driven, research environment. The principle that research is a public good must be matched by regulatory frameworks that nurture, and not hinder, the advancement of knowledge and societal benefit.

[13] A. Wernick, K. Meding, *Beware! The AI Act Can Also Apply to Your AI Research Practices*, 2024, https://arxiv.org/pdf/2506.03218, https://doi.org/10.48550/arXiv.2506.03218.
[14] M. Aboy, T. Minssen, E. Vayena, *Navigating the EU AI Act: implications for regulated digital medical products*, in *NPJ Digit Med*, 2024 Sep 6;7(1), 237. doi: 10.1038/s41746-024-01232-3. PMID: 39242831; PMCID: PMC11379845.
[15] M. Finck, *In Search of the Lost Research Exemption: Reflections on the AI Act*, in *GRUR International*, 2025; ikaf100. https://doi.org/10.1093/grurint/ikaf100.