

Introduction

Christoph Winter

Artificial intelligence is being deployed at scale across almost all domains of economic and social life. Until recently, no major jurisdiction had enacted a comprehensive legal framework designed specifically to govern AI. Existing law, such as product liability, data protection, or sector-specific regulation, reached AI incidentally rather than by design. The EU AI Act (Regulation (EU) 2024/1689) changes that, subjecting AI to a dedicated legal regime — with all the promise and risk that entails.

1. Why GPAI law

Within the AI Act, the provisions on general-purpose AI (GPAI) models occupy a uniquely consequential position. GPAI models, such as the large language models that underpin systems like ChatGPT, Claude, or Gemini, give rise to risks and regulatory challenges that are distinct from those of more specific AI applications. The AI Act imposes baseline obligations on all providers of such models placed on the EU market. A subset of those models — those classified under the AI Act as posing systemic risk — face a more demanding set of requirements. It is this subset, and the legal questions it raises, that form the primary focus of the Cambridge Commentary.

The provisions governing GPAI models with systemic risk constitute, in substance, the first body of *frontier AI law*: targeting the small group of models at the leading edge of known capabilities, whose potential for large-scale societal impact distinguishes them from AI models and even GPAI models more broadly. The AI Act defines this category through the concept of ‘high-impact capabilities’ — ‘capabilities that match or exceed the capabilities recorded in the most advanced general-purpose AI models’ (Article 3(64)). Critically, the systemic risks that the AI Act targets are defined as risks ‘specific to’ those high-impact capabilities; they are not risks that could arise from AI generally, but those that are specific to models at the frontier.

What makes frontier AI — or GPAI models with systemic risk — distinctive, and what justifies treating it as a separate regulatory category, is the particular problem of emerging capabilities. As these models improve, they acquire capabilities that were not explicitly trained for, that often cannot be predicted in advance, and that may not become apparent until after a model has been deployed. This makes GPAI models with systemic risk fundamentally different from most other AI models — including models that have been on the market for years and may themselves pose severe risks — whose risks, albeit serious, are at least knowable and assessable in advance. Frontier AI law accordingly serves a particular function: it imposes obligations, such as conducting model evaluations, adversarial testing, and incident reporting, that serve as epistemic tools for understanding these models’ capabilities and risks. The 2026 International AI Safety Report underscores the depth of the challenge that gaining this understanding

presents: how and why GPAI models with systemic risk acquire new capabilities is often difficult to predict even for their developers.¹

The EU is not the only jurisdiction to have recognised this regulatory challenge. California's SB 53 and New York's RAISE Act represent parallel efforts in the United States to impose obligations on developers of the most capable AI systems. Each of these regimes reflects a similar underlying recognition — that the most capable AI systems warrant specific regulatory attention — while differing substantially in scope, legal character, and enforcement. The EU's approach is, in some sense, the most legally developed: it is binding, enforceable, and backed by a newly-created and dedicated institutional infrastructure in the form of the European AI Office. But whether the EU's approach proves effective depends significantly on whether that infrastructure is adequately resourced. The European AI Office will need substantially greater technical, legal, and policy capacity if it is to rigorously evaluate companies' disclosures under the AI Act and the EU GPAI Code of Practice, and to enforce the obligations relating to GPAI models with systemic risk effectively.

2. Why this commentary

The GPAI provisions were not part of the AI Act's original architecture. The Commission's 2021 draft proposal did not address GPAI models at all; it was only through successive additions by the Council in December 2022 and the Parliament in June 2023 — as ChatGPT's release thrust foundation models into public and political consciousness — that Chapter V took shape, with its final form settled under considerable time pressure in the December 2023 trilogue. The result is a set of provisions that are, in places, notably underspecified. This is part of what makes both the EU GPAI Code of Practice and a dedicated, scholarly commentary necessary. The Cambridge Commentary is the first to focus exclusively on the GPAI model provisions of the AI Act. Prior commentaries — including Michèle Finck's comprehensive article-by-article treatment of the full Act (Oxford University Press, 2026), the volume edited by Martini and Wendehorst (Beck, 2nd edn, 2026), and the volume edited by Pehlivan, Forgó, and Valcke (Kluwer Law International, 2024) — address the AI Act in its entirety, which necessarily limits the depth any single chapter or provision can receive. By concentrating entirely on the GPAI provisions, the Cambridge Commentary offers a depth of analysis that a general commentary cannot provide.

The Cambridge Commentary is also, to our knowledge, the first on the AI Act to engage systematically with the technical AI safety and governance literature that the GPAI provisions presuppose. Interpreting several provisions requires more than doctrinal legal analysis. Determining what counts as a 'general-purpose AI model' under Article 3(63) — and in particular what 'significant generality' and 'competently performing a wide range of distinct tasks' mean in practice — requires engagement with the technical literature on model capabilities and benchmarking. The distinction between a model and a system, which determines the scope of the entire Chapter V AI Act regime, cannot be resolved by legal analysis alone. Understanding what Article 51's compute threshold was designed to capture, and why it

¹ Yoshua Bengio and others, 'International AI Safety Report 2026' (DSIT 2026/001, 2026) <<https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026>> accessed 15 April 2026.

may require updating, requires engagement with the empirical literature on the impact of training compute on capabilities. Interpreting what Article 55's requirements for adversarial testing and systemic risk assessment demand in practice requires familiarity with the state of the art in AI evaluation and safety science. We have tried throughout to make this interdisciplinary engagement transparent and accessible to lawyers charged with interpreting and applying the Act, without assuming technical background.

The Cambridge Commentary is also the first to analyse the GPAI provisions in light of the EU GPAI Code of Practice as adopted; no prior commentary engages substantively with it. It does so at a moment when the European AI Office is making foundational interpretive decisions that will shape enforcement for years. A commentary available during this formative period can contribute to those decisions rather than merely record them. That is one reason for publishing online first.

3. Method and approach

The GPAI provisions of the AI Act are marked by a degree of legal uncertainty that goes beyond what newly enacted legislation typically presents. As noted above, the provisions were negotiated at speed, in novel territory, and several of their most important concepts — what counts as a ‘systemic risk’, what ‘adequate’ mitigation requires, how the compute threshold should adapt to algorithmic efficiency gains — admit of more than one reasonable legal interpretation. The contributors have taken that uncertainty seriously, and it has shaped the approach in two related ways.

The first concerns what we call *interpretive optionality*. Rather than advocating throughout for a single preferred reading, the Cambridge Commentary on each provision aims to map the interpretive landscape: to identify where the text is clear, where it is ambiguous, and what the strongest arguments on each side look like. Where contributors favour a particular interpretation, they say so explicitly. The goal is to equip readers — practitioners, regulators, and scholars alike — to reason well about these questions and to adjust their views as the law and technology develop, rather than to hand them conclusions that may not endure over time.

The second concerns the relationship between legal and policy analysis. The two are easily — and sometimes deliberately — conflated in a field as novel and contested as this one, and no commentary can claim to separate them with perfect precision: the Court of Justice of the European Union's strongly teleological approach to EU legislation means the line between interpretation and evaluation is not always clear. But there is a meaningful difference between asking what a provision requires and asking what it *should* require, and contributors were instructed to keep that distinction in view — to lead with the legal question, to flag explicitly when analysis shades into normative territory, and to resist the temptation to resolve legal ambiguity by substituting a policy preference. Where the law is silent or unclear, the Cambridge Commentary says so rather than papering over the gap. The ambition is to foster legal clarity during the formative period of the AI Act's application, and in doing so contribute to a more constructive dialogue among regulators, providers, and the wider community of lawyers and scholars working in this field.

Each chapter of the Cambridge Commentary has undergone multiple rounds of expert review before publication. The views expressed are those of their authors; as editor, I have not sought to align them with my own interpretive preferences, and in some instances I disagree with the positions taken. The review process was designed to ensure that the interpretations advanced are legally reasonable and properly supported — not to produce uniformity of view.

4. Publication and updates

The Cambridge Commentary launches with Chapter V AI Act, which forms the core of the GPAI model provisions. Further articles and chapters will be added on a rolling basis: the definitions in Article 3 most directly relevant to GPAI models, the scope provisions of Article 2, the enforcement powers in Chapters IX and XII insofar as they bear on GPAI model providers, and the transitional provisions of Articles 111 and 113 to the extent relevant. The online-first format makes this incremental approach possible, appropriate, and preferable, given the pace of regulatory development, to waiting for comprehensive coverage before any chapter is published. The GPAI provisions are, at the time of writing, actively evolving: the EU's Omnibus Simplification Package is considering amendments that could alter the scope of several provisions analysed here. We are committed to revising individual chapters as material developments warrant; dates of submission, publication, and subsequent updates are noted for each.

Within each chapter, the commentary follows a broadly consistent structure: general remarks covering purpose, legislative history, and systematic positioning; a structural overview of the provision; and substantive analysis proceeding paragraph by paragraph. Overarching questions — including the definition of a GPAI model, the model/system distinction, and the status of AI agents — are addressed under general remarks at the outset.

The provisions of the AI Act related to GPAI models with systemic risk will govern what may prove to be the most powerful technology ever built. Getting the analysis right matters for the regulators charged with enforcing it, for the lawyers and compliance teams advising the companies that build and deploy these models, and for academics working to understand this emerging field of law. The contributors have approached that task with both legal rigour and a genuine sense of its importance.

Above all, I wish to thank the section editors Emily Gillett and Maarten Herbosch, the assistant editors Madalina Nicolai, Gregor Gindlin, and Zlatko Grigorov, and all contributing authors, without whom this project would not exist. Thanks are also due to Andrew Leeke, whose work on the website has made it significantly more accessible to readers. I hope that the Cambridge Commentary will serve practitioners, the European AI Office, and the wider scholarly community as a reliable guide to a body of law that is still, in the most important sense, being written.

Christoph Winter
University of Cambridge, 2026