

Regulating General Purpose AI in the EU: An Analysis of the AI Code of Practice

Giulia Torchio
Samuel Goodgerⁱ

INTRODUCTION

The EU stands at a pivotal moment for the future of digital governance. The implementation of the AI Act (AIA), one of the world's first comprehensive and democratically legitimised pieces of Artificial Intelligence (AI) legislation, has started, and is due to be completed by 2027. Aiming to support companies with this staggered process, the European Commission has recently published a Code of Practice¹ (CoP) for providers and developers of General Purpose AI (GPAI) models. The tool is meant to facilitate compliance with the AIA while external bodies develop the technical standards. Despite the seemingly technical nature and limited scope of the CoP, ensuring that companies adhere to it is of the utmost importance in order to boost the safety of AI in the EU. This policy brief outlines the development of the Code and explains its political relevance. This research aims to kickstart a process of reflection on necessary amendments to the text ahead of future revisions of the CoP.

Despite the seemingly technical nature of the Code, ensuring that companies adhere to it is of the utmost importance in order to boost the safety of AI in the EU.

I. BACKGROUND: THE NEED FOR A CODE OF PRACTICE

On 1 August 2024, just over three years after the initial proposal was presented, the EU Artificial Intelligence (AI) Act came into force. A landmark piece of legislation, the AI Act established the world's first comprehensive regulatory framework for AI. It follows the model of the 'New Legislative Framework' for goods, outlining a risk-based approach that categorises risks into four levels: unacceptable, high, limited and minimal; different requirements are mandated, based on the classification. Rather than being enshrined in the Act itself, these risk-based, mandated measures are being developed by external standardisation bodies.

For this specific task, the European Commission has entrusted the European Committee for Standardisation (CEN) and the European Committee for Electrotechnical Standardisation² (CENELEC), which are private entities acting in accordance with the EU Standardisation Regulation.³ CEN and CENELEC are expected to take up to three years to develop the technical standards. The Commission decided to create external criteria to establish a technologically neutral and future-proof legislative framework as well as to leverage technical expertise, given the complexity of the technology.

Frontier, or General Purpose AI (GPAI) models, defined as models capable of performing a wide range of tasks, such as generating text, images, predictions and other forms of content, as well as making decisions on someone's behalf, fall outside this risk hierarchy.⁴

ⁱ ACKNOWLEDGEMENTS

The authors would like to thank experts from Pour Demain and the Future of Life Institute for their valuable contributions to this piece.

In the legislative text, a distinction is made between these models in the form of two classes of risk: systemic and non-systemic. Providers of GPAI must meet requirements based on their class, with stricter rules for models posing systemic risk.

To facilitate the compliance process while the standardisation authorities decide on the technical benchmarks, the AI Act allows for the development of interim measures. These are outlined in the Code of Practice (CoP) for GPAI providers and deployers, which was developed through a multi-stakeholder process spearheaded by the Commission and involving nearly 1,000 actors.

The CoP applies to providers of GPAI models and GPAI models with systemic risk and entered into force on 2 August 2025. It details three categories of provisions: (1) transparency; (2) copyright-related rules; (3) and safety and security. For further guidance, the Commission has also released GPAI Guidelines,⁵ which clarify definitions as well as relevant AIA articles.

Far from being a temporary, inconsequential and highly technical text, the Code is pivotal for the AI Act's full implementation. It is a dynamic testing ground for the AI Act's standards, addressing safety, transparency, copyright and fair use concerns ahead of the full enforcement of formal legal obligations expected for 2027. It also supports Europe's efforts to become a leading AI continent⁶ by reducing the administrative burden for companies and providing more legal certainty through easier paths towards compliance.

WHAT IS A CODE OF PRACTICE?

CoPs are steering instruments that supplement EU law with new rules that are not legally binding. Contrary to other community soft law tools, such as standards – whose practice is defined in a 2012 Regulation⁷ – the development of CoPs is not regulated in any EU legislation. In practice, this means that much is left to the Commission's discretion.

Why was the Code so contentious?

The provisions on transparency, copyright, and safety and security detailed in the Code are applied proportionately. Providers of GPAI models will have to abide by the measures outlined in the first two chapters. Providers of GPAI models characterised by *systemic risks*ⁱⁱ will also have to comply with safety and security requirements.

Out of the three areas, provisions related to copyright, and safety and security proved to be some of the most contentious aspects of the CoP, having undergone

significant changes throughout the drafting process. This is because they are both connected to fundamental policy debates, with far-reaching implications for the future of AI development and its societal impact.

Data is the basis for the development of AI models, especially GPAI, of which generative models are a typical example. The CoP's provisions on copyright will directly affect both AI developers, who might see their rights to data collection severely limited, and rightsholders, who might be the recipients of repeated infringements. In the first case, cumbersome measures could potentially hinder innovation by creating data scarcity. In the second instance, creative industries will lose out to generative models that unlawfully appropriate content.

Similarly, the chapter on safety and security is also crucial in striking a balance between innovation and protection. Lax measures regarding systemic risks will lead to dangerous scenarios for EU citizens. Malicious actors could acquire enhanced cyber offensive capabilities, use AI to develop chemical, biological, radiological or nuclear (CBRN) weapons, or easily spread disinformation. Conversely, an overly broad definition of systemic risk could lead to an excessive regulatory burden on developers, hindering the deployment of even low-risk GPAI.

Four evolutions of the Code of Practice

Analysing the Code's four iterative drafts is crucial in order to understand its development and increased political relevance. Examining the changes made at each round of consultation provides insights into evolving policy priorities and the overall fairness of the process.

The first draft of the CoP⁸ attempted to reconcile safety and innovation concerns. The text outlined important copyright-related rules in Section 2 and on risk definition and taxonomy in Section 3. Regarding copyright, it outlined sub-measures for both upstream and downstream compliance. For the taxonomy, it listed six instances of systemic risks. These proposed provisions quickly became some of the most contentious aspects of the CoP.

Departing from the approach of the first draft, the second one⁹ introduced different tiers of systemic risks into the taxonomy. It distinguished between "systemic risks" (such as nuclear risks and loss of oversight, for example) and "additional risks for consideration" (such as infrastructure, fundamental rights, public security, etc.). This decision was strongly criticised for creating a hierarchy that contradicted the purpose of AI Act Recital 110 and Art. 3(65) on systemic risk¹⁰ and labelled serious concerns, such as large-scale discrimination, as *optional*.

Upon publication, the third draft¹¹ of the CoP elicited backlash on multiple fronts. On the contentious issue

ⁱⁱ According to the AI Act, GPAI models are considered to have systemic risk if they have high-impact capabilities which could be evaluated based on appropriate benchmarks, such as the threshold for cumulative amount of computation used for the training.

of the risk taxonomy, the third iteration of the CoP changed the name of the second tier of risks from “additional risks for consideration” to “other types of risks for *potential* consideration in the selection of systemic risks”. The new wording was strongly criticised¹² for prioritising corporate interests over human rights. On copyright, the draft made significant concessions to model providers.

The fourth and final draft of the CoP was published on 10 July 2025, after weeks of intense speculation.¹³ The most notable changes in the final text were the removal of the whistleblower protection, of the contentious risk taxonomy previously outlined in Annex 1, and of adequacy assessments for providers. At the same time, requirements for public transparency mechanisms and external model assessments were preserved in the final Code.

II. STATE OF PLAY: AN ANALYSIS OF THE CODE

Upon closer examination of the CoP and its provisions, it is apparent that significant concessions were made during the multistakeholder process, and particularly in the last few months of it. While it is impossible to exhaustively judge the validity of the Code before it is enforced and its interplay with other digital rules is analysed, a preliminary assessment can be drawn based on changes made to crucial provisions in the final text.

Significant concessions were made during the multistakeholder process, and particularly in the last few months of it.

Compared to the three previous versions of the AI Code of Practice, which introduced and expanded upon the concept of a tiered risk taxonomy, the final version does not distinguish between systemic risks. This is a particularly welcome change, as the hierarchy contradicted the broader spirit of the AIA and placed excessive trust in companies with a questionable track record of adherence to EU law.¹⁴

Also on systemic risks, the final CoP establishes that signatories will have to systematically analyse if other risks, such as those associated with major accidents, child sexual abuse material (CSAM) or mental health, are systemic in scope. This is a particularly relevant provision to enforce, as concerns arise over users’ conversations with GPAI models described as sycophantic,¹⁵ which in

some cases have led to delusions and even self-harm.¹⁶ Currently, the process must be documented in the Model Reports and Frameworks and providers will be required to make summaries of their findings publicly available.

However, while these measures on systemic risks and transparency raise the bar for scrutiny, they regrettably only require documents to be provided after the release of the models, which increases the cost of post-deployment action or enforcement. Additionally, the accountability and regulatory oversight of GPAI providers are weakened by the removal of whistleblower protections or any safeguards against systemic risks throughout model development.

In terms of copyright, the final CoP builds on the changes introduced in the second version to reverse the traditional burden of proofⁱⁱⁱ in relation to copyright law. It establishes measures to ensure that web crawlers do not use copyright-protected content to train models, even in the case of third-party datasets. Despite the novel approach, the measures are still unlikely to substantially change unfair data collection practices that are at the heart of the business models of big tech companies.

This is because the language chosen in the CoP is both too vague and too restrictive to oblige GPAI providers to respect EU copyright law. Specifically, the Code’s measures essentially restrict the scope of the prohibitions to data obtained through web crawling. This could undermine the effectiveness of copyright measures, as crawling is not the only data acquisition method for training generative AI (others include, e.g., real-time data streaming, data acquired through the Internet of Things (IoT) and user-generated data on platforms).

Finally, in a welcome change, the CoP establishes that newly released GPAI models with systemic risks^{iv} will be required to undergo third-party systemic risk assessments and external evaluations, with evaluators given free access to the most capable versions.

Challenges along the way

Ever since it was decided that interim rules would be developed through a soft law instrument, the AI governance community has been split into several camps.

For some experts,¹⁷ the CoP’s provisions are overstressing the principles of the AI Act. On the contrary, for rightsholders¹⁸ and civil society organisations (CSOs), the final Code still contains significant compromises that undermine principles of transparency, fail to protect citizens’ rights and give providers excessive leeway in shaping their compliance.¹⁹ Industry, and particularly foreign companies, strongly opposed any strict measures on

ⁱⁱⁱ It is responsibility of providers to demonstrate that they comply with EU copyright laws by having the necessary rights to access material. Traditionally, it was the complainant’s responsibility to prove that a breach of copyright had occurred.

^{iv} Except those “similarly safe” to already compliant models. While this exemption may cause concern, the incentive for AI companies to differentiate new releases may make external assessment the de facto norm for new models placed on the market.

both copyright and risks from making the final version of the text, repeatedly expressing their intention not to sign it.²⁰ Backing up the private sector, representatives of the US administration also tried to influence²¹ discussions surrounding the CoP.

However, issues of interest and disagreements over the final wording of the provisions are not the only ones that emerged during the process. It is also essential to reflect on challenges associated with the process's methodology and its transparency.

A report from the NGO Corporate Europe Observatory (CEO) revealed²² that industry benefited from privileged access to meetings with the Commission and dedicated workshops. While not surprising, given the intense lobbying efforts²³ surrounding the AI Act's negotiations, it remains concerning. Additionally, experts directly involved in the working groups spoke of a distinct lack of transparency from the Commission's side during the final months of the process. In pointing this out, they emphasised that this was not the case at the beginning of the drafting process and that the situation only changed around the time of the third draft's publication.

For some experts, the CoP's provisions are overstressing the principles of the AI Act.

It is partly due to this lack of transparency that, in late May 2025, a report²⁴ alleging that the Commission was considering pausing the implementation of the AI Act sparked public conjecture. According to the article, this would have been possible by mobilising the 'stop-the-clock' Directive, as occurred²⁵ earlier this year with the Corporate Sustainability Reporting Directive (CSRD) and the Corporate Sustainability Due Diligence Directive (CSDDD). However, it was later revealed that the Commission had not entertained such a possibility and that a pause had only been contemplated²⁶ if issues with technical standards had arisen.

III. PROSPECTS: WHAT HAPPENS NEXT?

Despite delays and pressure from interest groups until the final moment,²⁷ the final version of the Code was presented on 10 July 2025, followed by the release of the complementary GPAI guidelines on 18 July. After the final CoP was published, the Commission's AI Office announced that it would implement a one-year grace period for signatories releasing models after 2 August 2025, as well as adopt a more lenient approach when considering future proceedings and fines against signatories.

Although there was a decline in the level of transparency late in the multi-stakeholder process, the Commission, for the most part, skilfully managed the CoP's drafting through an inclusive multi-stakeholder process. Nevertheless, this does not mean that the final Code is without flaws. As explained in this policy brief, there is a need for review from the AI Office. This should be carried out as follows: a technical review of the Code and a procedural review of the *modus operandi* to develop such soft law instruments.

Although there was a decline in the level of transparency late in the multi-stakeholder process, the Commission, for the most part, skilfully managed the CoP's drafting through an inclusive multi-stakeholder process.

On the **technical** level, the next review of the GPAI CoP – tabled for 2027 at the latest – should focus on the following issues:

- ▶ **The lack of systemic risk mitigation measures targeting internal or pre-deployment models.** While the Code outlines serious incident reporting timelines (Measure 9.3) and requires GPAI model developers to share information with the AI Office promptly, full reports must only be shared with regulators at the time of the model's release (Measure 7.7).²⁸ This approach is dangerous as regulators may discover certain risks posed by novel capabilities only when it is too late. A future revision of the CoP should include new measures to mandate pre-deployment sharing of interim Model Reports with the AI Office, as well as reinstate the whistleblower protections.
- ▶ **The list of systemic risks in Appendix 1 might not be broad enough.** While the inclusion of fundamental rights as a potential systemic risk (Appendix 1.1) and a model's ability to impact them (Appendix 1.3.1) are a step forward, more will be needed. Through future revisions, the AI Office will have to ensure that the list is updated to include new types of fundamental rights risk, such as discrimination, in the Appendix.
- ▶ **The language on copyright measures is too restrictive.** Measures I.2 and I.3 on the extraction, reproduction and rights reservations of copyright-protected material obtained from the World Wide Web are limited to web crawling practices. In future revisions, the CoP should expand the scope of these measures to regulate other data acquisition practices beyond web crawling, such as real-time data streaming and IoT devices.

On a **procedural** level, the Commission should:

- ▶ **Develop a clear set of guidelines informing the drafting of Codes of Practice.** Policymakers should reflect on the successes and shortcomings of the AI CoP multi-stakeholder process and establish official processes to regulate the future creation of these soft law instruments. As part of this initiative, the Commission should specifically focus on measures to ensure greater transparency, fairer representation of interests and better coordination.
- ▶ **Increase the staffing capacity of the AI Office.** A limited staff is not conducive to the effective enforcement of the AI Act and the CoP. Following the example of foreign regulators or agencies with comparable mandates, such as that of the UK AI Security Institute,²⁹ which can boast 250 employees, the AIO should intensify recruitment, especially of technical experts. Failure to recognise the need for more enforcement capacity will only spell greater risks for EU citizens.

CONCLUSION

The GPAI Code of Practice’s multi-stakeholder drafting process revealed inherent tensions between codifying AI safety measures, protecting rightsholders and addressing concerns of the AI industry. Despite a demonstrated commitment to balanced representation of interests at the start, our assessment is that the Commission’s declining transparency and industry concessions — leading to accusations of regulatory capture — posed a risk in terms of undermining the CoP’s objectives.

Nevertheless, the final Code represents an acceptable compromise and first step towards enhancing the safety of EU-based GPAI users and promoting a culture of compliance. Following the release of the final document, a majority of Western GPAI model providers publicly supported⁵⁰ the CoP or announced their intention to do so. Among the largest companies, only one has so far refused to sign the Code of Practice.⁵¹

The final Code represents an acceptable compromise and first step towards enhancing the safety of EU-based GPAI users and promoting a culture of compliance.

However, a mark of true success goes beyond a compromise text and first iteration. As part of future revisions, the AI Office should focus on better prioritising users’ protection, increasing pre-deployment monitoring and expanding the scope of the copyright section. At the same time, it will be essential to increase the Office’s capacity to allow for proper enforcement of AI rules in the EU as well as to draw up official procedures guiding the development of soft law instruments. Lastly, the Commission should focus on ensuring legal coherence with other documents, such as the upcoming CoP on Transparent AI Systems,⁵² similar in format to the subject of this paper, which will also provide the Commission with an opportunity to showcase improvements in procedure.

The support the European Policy Centre receives for its ongoing operations, or specifically for its publications, does not constitute an endorsement of their contents, which reflect the views of the authors only. Supporters and partners cannot be held responsible for any use that may be made of the information contained therein.

- ¹ European Commission (n.d.), *General-Purpose AI Code of Practice*, <https://digital-strategy.ec.europa.eu/en/policies/ai-code-practice>.
- ² European Committee for Standardization and European Electrotechnical Committee for Standardization (CEN/CENELEC) <https://www.cenelec.eu/>.
- ³ EUR-Lex (2012), *Regulation 1025/2012 on European standardisation*, <https://eur-lex.europa.eu/eli/reg/2012/1025/oj/eng>.
- ⁴ EUR-Lex, (2024), *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence*, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.
- ⁵ 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>.
- ⁶ European Commission, *The AI Continent Action Plan*, <https://digital-strategy.ec.europa.eu/en/library/ai-continent-action-plan>.
- ⁷ EUR-Lex (2012), *Regulation 1025/2012 on European standardisation*, <https://eur-lex.europa.eu/eli/reg/2012/1025/oj/eng>.
- ⁸ European Commission, *First Draft of the General-Purpose AI Code of Practice published, written by independent experts*, 14 November 2024, <https://digital-strategy.ec.europa.eu/en/library/first-draft-general-purpose-ai-code-practice-published-written-independent-experts>.
- ⁹ European Commission, *Second Draft of the General-Purpose AI Code of Practice published, written by independent experts*, 19 December 2024 <https://digital-strategy.ec.europa.eu/en/library/second-draft-general-purpose-ai-code-practice-published-written-independent-experts>.
- ¹⁰ EUR-Lex, (2024), *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence*, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.
- ¹¹ European Commission, *Third Draft of the General-Purpose AI Code of Practice published, written by independent experts*, 11 March 2025, <https://digital-strategy.ec.europa.eu/en/library/third-draft-general-purpose-ai-code-practice-published-written-independent-experts>.
- ¹² Cabrera et al. "Human Rights are Universal, Not Optional: Don't Undermine the EU AI Act with a Faulty Code of Practice", *Tech Policy Press*, 28 March 2025 <https://www.techpolicy.press/human-rights-are-universal-not-optional-dont-undermine-the-eu-ai-act-with-a-faulty-code-of-practice/>.
- ¹³ European Commission (n.d.), *The General-Purpose AI Code of Practice*, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>.
- ¹⁴ Proton, "Tech fines tracker", <https://proton.me/tech-fines-tracker>.
- ¹⁵ Winecoff, Amy, "Artificial Sweeteners: The Dangers of Sycophantic AI", *Tech Policy Press*, 14 May, 2025, <https://www.techpolicy.press/artificial-sweeteners-the-dangers-of-sycophantic-ai/>.
- ¹⁶ Booth, Robert, "Teen killed himself after months of encouragement from ChatGPT; lawsuit claims", *The Guardian*, 27 August 2025. <https://www.theguardian.com/technology/2025/aug/27/chatgpt-scrutiny-family-teen-killed-himself-sue-open-ai>.
- ¹⁷ Ebers, Martin, "When Guidance Becomes Overreach: How the forthcoming Code of Practice Threatens to Undermine the EU's AI Act", *Verfassungsblog*, 8 April 2025, <https://verfassungsblog.de/when-guidance-becomes-overreach-gpai-codeofpractice-aiact/>.
- ¹⁸ Kroet, Cynthia, "Generative AI code contradicts copyright law, industry says", *Euronews*, 28 January 2025, <https://www.euronews.com/next/2025/01/28/generative-ai-code-contradicts-copyright-law-industry-says>.
- ¹⁹ Marcus, Gaia, "Ada Lovelace Institute responds to final General-Purpose AI Code of Practice", *Ada Lovelace Institute*, 15 July 2025, <https://www.adalovelaceinstitute.org/news/ada-responds-to-gpai-code-of-practice/>.
- ²⁰ Kroet, Cynthia, "Meta top official warns penalising US companies will not make EU competitive", *Euronews*, 5 February 2025, <https://www.euronews.com/next/2025/02/05/meta-top-official-warns-penalising-us-companies-will-not-make-eu-competitive>.
- ²¹ Volpicelli, Gian et al., "Trump Administration Pressures Europe to Ditch AI Rulebook", *Bloomberg*, 25 April 2025, <https://www.bloomberg.com/news/articles/2025-04-25/trump-administration-p pressures-europe-to-reject-ai-rulebook>.
- ²² Corporate Europe Observatory, "Coded for privileged access: How Big Tech weakens rules on advanced AI", 30 April 2025, <https://corporateeurope.org/en/2025/04/coded-privileged-access>.
- ²³ Cerulus, Laurens et al., "Ranked: The 10 most intensely lobbied EU laws", *Politico Europe*, <https://www.politico.eu/article/european-parliament-law-legislation-lobbying-technology-sustainability-amendments/>.
- ²⁴ Bertuzzi, Luca, "EU Commission eyes pausing AI Act's entry into application", *Mlex*, 26 May 2025, <https://www.mlex.com/mlex/articles/2344845/eu-commission-eyes-pausing-ai-act-s-entry-into-application>.
- ²⁵ Council of the European Union, *Simplification: Council gives final green light on the 'Stop-the-clock' mechanism to boost EU competitiveness and provide legal certainty to businesses*, 14 April 2025, <https://www.consilium.europa.eu/en/press/press-releases/2025/04/14/simplification-council-gives-final-green-light-on-the-stop-the-clock-mechanism-to-boost-eu-competitiveness-and-provide-legal-certainty-to-businesses/>.
- ²⁶ Bertuzzi, Luca, "Poland puts pausing enforcement of the AI Act on EU ministers' table", *Mlex*, 3 June 2025, <https://www.mlex.com/mlex/articles/2348475/poland-puts-pausing-enforcement-of-the-ai-act-on-eu-ministers-table>.
- ²⁷ Moreau, Claudie, "Lobbyists make last push to derail AI Code of Practice", *Euractiv*, 30 June 2025, <https://www.euractiv.com/section/tech/news/lobbyists-make-last-push-to-derail-ai-code-of-practice/>.
- ²⁸ European Commission (n.d.), *The General-Purpose AI Code of Practice*, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>.
- ²⁹ Schaefer, Peder, "Hiring struggles are plaguing the EU AI Office", *Transformer*, 15 September 2025. <https://www.transformernews.ai/p/eu-is-struggling-to-hire-ai-act-office-safety-unit>.
- ³⁰ European Commission, (n.d.), *Signatories of the AI Pact*, <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai#ecl-inpage-Signatories-of-the-AI-Pact>.
- ³¹ Kroet, Cynthia, and Hess, Amandine, "Meta won't sign the EU's AI Code, but who will?", *Euronews*, 23 July 2025, <https://www.euronews.com/my-europe/2025/07/23/meta-wont-sign-eus-ai-code-but-who-will>.
- ³² European Commission, "Commission launches consultation to develop guidelines and Code of Practice on transparent AI systems", 4 September 2025. <https://digital-strategy.ec.europa.eu/en/news/commission-launches-consultation-develop-guidelines-and-code-practice-transparent-ai-systems>.

With the strategic support of



Co-funded by the European Union