

July 23, 2025

EU Commission Publishes Its Code of Practice for General Purpose AI: What You Need to Know

On July 10, 2025, the European Commission (the “**Commission**”) published its Code of Practice for General Purpose AI (the “**Code**”)¹ relating to the European Union’s Regulation 2024/1689, also known as the EU AI Act (the “**AI Act**”)². The Code follows three draft iterations and a delay from its initially anticipated May 2025 publishing date. The Code is a voluntary framework intended to help its signatories (“**Signatories**”) comply with their obligations under the AI Act when providing general-purpose AI (“**GPAI**”) models (as defined in the AI Act). The Code does not create legally binding obligations and is instead stated to serve “*as guidance to help providers meet their existing obligations under the AI Act without creating new ones, extending existing ones, or imposing additional burdens*” until relevant standards are published for GPAI providers. This alert summarises the key provisions of the Code and identifies practical takeaways for providers of GPAI models and other stakeholders in the AI value chain.

Key Takeaways

1. **More to come.** The Commission and Member States need to approve the Code before it comes into effect, although no changes to the text are expected. In addition, the Commission has subsequently complemented the Code by publishing its finalised guidelines on GPAI models³, ahead of AI Act GPAI obligations coming into force on August 2, 2025. In our next alert, we will summarise the key takeaways from these guidelines which aim to provide clarity in respect of a number of key concepts of the AI Act (including the definition of GPAI model), and how the European Commission’s AI Office (the “**AI Office**”) will support providers to comply with the AI Act. While the guidelines are likely to be useful, particularly for borderline cases, model developers may nevertheless wish to use the Code as a useful reference for compliance obligations (especially given lack of certainty as to whether future regulatory investigations will use the Code as a baseline).
2. **Potential safe harbour?** The AI Office stated in its Q&A (an updated version of which was published with the Code)⁴ that it will work closely, particularly with Signatories, to ensure that models can be brought to the EU market “without delay” in compliance with the Code. For Signatories who are not fully compliant with the Code upon signing it, the AI Office will not enforce the requirements of the AI Act in the short term. Indeed, the AI Office has stated that it will treat those Signatories that comply with the Code as complying with the relevant obligations under the AI Act. This will provide comfort to GPAI providers who would like the benefits of being a Signatory but who are also concerned about their ability to comply from day one.

¹ Available here: <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>

² Available here: <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

³ Available here: <https://digital-strategy.ec.europa.eu/en/library/guidelines-scope-obligations-providers-general-purpose-ai-models-under-ai-act>

⁴ Available here: <https://digital-strategy.ec.europa.eu/en/faqs/questions-and-answers-code-practice-general-purpose-ai>

3. **Lightening of the load for developers.** Although safeguarding copyright-holders' rights has been a key consideration in the evolution of the Code, the final text has generally shown a movement towards a principles based set of obligations on GPAI providers that choose to become Signatories, rather than prescriptive obligations regarding use of copyright works and/or obligations to inform rightsholders of use of their works. However, GPAI providers who are scraping data online should carefully consider the measures they take in relation to rights reservations to text and data mining ("**TDM**") and monitoring use of content from websites that the AI Office has published as being high risk to ensure they do not trip over the clearest Code obligations. Given the political and regulatory focus on these issues, they may well become an area of focus for EU regulatory enforcement actions.
4. **Guidance, but limited detail, on systemic risk framework.** Detail on what Signatories should include in their systemic risk assessment frameworks has been stripped back from the final Code, with increased weight placed on general principles. The Commission may be hoping the collaboration principle of the Code will lead to a market-led consensus on requirements for these frameworks, but for now there is a lack of clarity. Providers of GPAI with systemic risk should carefully consider the requirements of the Code to identify what should be implemented as a matter of best practice, especially given the standalone obligations under the AI Act for such AI models.

Summary of the Code

Copyright

Consistent with the previous iteration of the Code, the final copyright chapter establishes measures to maintain a copyright policy, restrict crawling of pirated or infringing content, respect TDM reservations under Article 4(3) of Directive 2019/790 (the "**Copyright Directive**"), restrict copyright-infringing outputs and allow copyright-related complaints.

- **Copyright policy.** Signatories must create an internal policy setting out how they comply with EU copyright law, and are also "encouraged" to make publicly available a summary of this policy. The Code does not prescribe any further contents, although the AI Office will be providing a template copyright policy in due course. As such, this currently appears to be a fairly loose requirement on GPAI providers.
- **TDM reservations.** Although the Code recognises the widely-used Robots Exclusion Protocol (robots.txt) as a means for rights-holders to express TDM reservations under the Copyright Directive, it requires that Signatories also identify and comply with other "appropriate" machine-readable protocols (in addition to the Robots Exclusion Protocol) adopted by international or European standardisation organisations and other state-of-the-art (but technically implementable) measures widely adopted by rights-holders. This has been limited to solely when data is scraped though, with the obligation to comply as part of "crawling" removed. In addition, Signatories must take "appropriate" (changed from "reasonable" in the last iteration of the Code) measures to provide information about their web crawlers and steps they take to comply with TDM reservations (also enabling rightsholders to be automatically notified when that information changes). Although this provides a clear principle, there is a remaining lack of clarity for the industry on how relevant protocols will be determined, identified or agreed upon, and whether it is technologically possible to comply with all such protocols.
- **Copyright infringement and complaints.** Signatories must take a number of steps to prevent infringing copyright work, including:
 - Implementing appropriate and proportionate technical safeguards to prevent reproduction of copyright materials (including prohibiting any such infringing uses in model terms and policies). This marks a change from the "reasonable efforts" requirement under the previous iteration of the Code;
 - Not circumventing technological measures designed to prevent or restrict *unauthorised acts* in respect of such works (which the Code notes may be broader than paywalls, e.g. IP-based geographic restrictions). The reflects a more significant obligation compared to the previous iteration of the Code, which only prohibited Signatories from circumventing technological measures (such as paywalls) designed to prevent or restrict *access to* copyright works. The wider reference to *unauthorised acts* provides clarity, for example, that Signatories should not circumvent measures, website terms or technological means preventing *copying* (even though *access* may be allowed);
 - Excluding from their crawling activities those websites that persistently and repeatedly make infringing copyright material available on a commercial scale as recognised by courts or public authorities in the EU and European Economic Area. This reflects an increase from the "reasonable efforts" obligation in the previous iteration of the Code, indicating a tougher stance from the AI Office on use of pirated content by GPAI providers (although with the AI Office making the

task easier by agreeing to publish a dynamic list of such websites). Interestingly, the final version of the Code removed the obligation to take steps to understand the origins of information mined other than via web-crawling, reducing the burden on GPAI model developers in relation to offline source of information; and

- Designating a point of contact for affected rights-holders to submit complaints regarding alleged copyright infringement and act on complaints in a diligent, non-arbitrary manner “within a reasonable time”.

While many GPAI providers will already have processes in place to implement these commitments, it is clear that the AI Office is reducing the ability for Signatories to state that they were not aware of such allegedly infringing sources or content.

Transparency

The transparency chapter sets out the documentation requirements for GPAI providers. This chapter has only undergone minor changes from the previous iteration of the Code, in all cases to soften obligations on Signatories.

- Signatories are required to document, and keep updated, a range of information in relation to each GPAI model. The Code includes a “friendly” form⁵ to guide Signatories, which may also be of interest to non-Signatory GPAI providers seeking to comply with Article 53(1) AI Act. Model documentation must be retained for each version of the GPAI model for 10 years from the date the relevant version has been placed on the market.
- The friendly form differentiates between information that must be made available, upon request, to the AI Office and that which must be provided to downstream providers. A more limited and high-level set of information is required to be provided to downstream providers, in particular not including distribution channels, design parameters for training, details on selection, volume and verification of data sources, computational resources and energy consumption. The Code does not mandate any disclosure of this information to the public generally, although encourages Signatories to consider whether any such information should be publicly disclosed (including to comply with obligations under Art.54(1)(d) AI Act to publish a summary of content used for GPAI training).
- In addition to the specific information set out in the model documentation, the Code imposes additional obligations to: (i) provide the AI Office and relevant national authorities *any additional information* necessary for them to fulfil their tasks under the AI Act (including to assess compliance of high-risk AI systems built on GPAI models, where the system and model providers are different); and (ii) to provide additional information, within 14 days of a request from a downstream provider, to enable them to have a good understanding of the capabilities of the GPAI model relevant to its integration into downstream providers’ AI systems. Unlike in previous versions of the Code, all disclosure requirements are subject to protection of the Signatory’s IP rights and trade secrets.

Safety and Security

This section applies to any GPAI that has been determined to have systemic risk (which under the AI Act are GPAI models with high impact capabilities (including computation floating point operations greater than 10^{25}), or which are otherwise designated as such by the AI Office, in each case, which create a specific risk that has a significant impact on the EU market due to reach or actual or reasonably foreseeable negative effects on public health, safety, security, fundamental rights or society as a whole, that can be propagated across the AI value chain). The Code further describes systemic risks as including “*lowering barriers for the development of chemical or biological weapons, or risks related to loss of control over the model*”.

This section has been significantly reduced compared to previous iterations. In particular, the commitments and obligations on Signatories have been reduced, and replaced with a more principles- and outcomes-based approach.

- **Framework for Continuous and Contextual Risk Assessment.** The Code requires Signatories to establish and provide to the AI Office a state-of-the-art ‘Safety and Security Framework’. This framework must include: (i) a description and justification of the trigger points along the model lifecycle at which the Signatories will conduct “lighter-touch model evaluations”, including during development; (ii) information related to the Signatories’ determination of whether systemic risk is considered acceptable (e.g., a description and justification of criteria for accepting that risk); (iii) a description of how responsibility for systemic risk assessment is allocated within the organisation and to third parties; and (iv) a description of the process by which Signatories will update the framework. In implementing frameworks, Signatories must undertake ongoing assessment of the risk posed by a GPAI model across the AI value chain, including through the mentioned lighter-touch evaluations and ongoing post-market monitoring of real-world

⁵ Available here: <https://ec.europa.eu/newsroom/dae/redirection/document/118118>

use by end users. However, before placing a model on the market Signatories must conduct a full systemic risk assessment by: (a) identifying systemic risks through analysis of potential risk sources (i.e., factors which may give rise to systemic risks) against the model itself (including post-market monitoring, serious incidents and near misses); (b) analysing identified systemic risks qualitatively and quantitatively using a wide range of sources, including internal analysis of the training data and external analysis of sources such as literature and expert or lay interviews or panels; (c) determining whether the systemic risks are acceptable, according to the Signatory's criteria and risk tiers; and (d) implementing safety and security mitigations, discussed below (and then conducting the assessment again).

- **Safety and Security Mitigations and Measures.** For any models with systemic risk that are put on the EU market, Signatories are required to deploy safety mitigations and measures to ensure those risks are and will remain acceptable. Those mitigations must be at least state-of-the-art and supplemented by external independent audits and undertake security reviews to verify the Signatory's risk controls are in line with industry best practices. Safety mitigations should be sufficiently robust for the risk identified (including any third party pressure such as jailbreaking or fine-tuning attacks) and take into account the release and distribution strategy of the model – examples include cleaning training data, staging access to the model and offering tools to users to mitigate risks. Signatories must also implement an adequate level of cybersecurity protection for their models and physical infrastructure to ensure that systemic risks that could arise from unauthorised releases, unauthorised access, and/or model theft are acceptable, with a defined goal specifying the threat actors its security mitigations are intended to protect against and accordingly implementing mitigations. There are exemptions to security mitigations if the relevant model's capabilities are inferior to the capabilities of at least one model for which the parameters are publicly available for download.
- **Model Reporting.** Model providers (other than small and medium enterprises) are required to provide the AI Office with a 'Safety and Security Model Report'. Among other items, this must include information on the model (including its architecture, capabilities, and propensities), systemic risk assessments, and mitigation processes and measures. That information is provided on a confidential basis and is required prior to placing a model on the market and when there is any material change based on continuous assessment (within a reasonable time of such changes, unless such change results from a deliberate update to the model in which case the updated report should be provided prior to the model being placed on the market). The Safety and Security Model Report should essentially summarise the risk assessment and mitigations mentioned above, and justify why the model and systemic risks are acceptable in the Signatory's view. There are also obligations to give the AI Office information to understand how the systemic risk landscape may change over time (including as the model develops post-market) and to reflect the external evaluations and reviews noted above. The reporting obligations are all subject to confidentiality requirements on the AI Office.
- **Incident Reporting.** All Signatories are obliged to report (without admitting fault) to the AI Office: (i) any serious and irreversible disruptions to the management or operation of critical infrastructure (or a reasonably likely causal relationship between the model and such disruption), no later than two days after becoming aware; (ii) serious cybersecurity breaches, including the (self)-exfiltration of model weights and cyberattacks, no later than five days after becoming aware; (iii) a death of a person (or a reasonably likely causal relationship between the model and such death), no later than 10 days after becoming aware; and (iv) serious harm to a person's mental and/or physical health, an infringement of any fundamental rights and/or serious harm to property or the environment (or a reasonably likely causal relationship between the model and such incidents), no later than 15 days after becoming aware. Reports should provide root cause analysis, recommendations to the AI Office and national competent authorities in response to the incident (and detail what the Signatory has done in response) and technical detail on the model and the chain of events leading to the incident (although there is an acknowledgement that this information may only be known after the incident has occurred and some information may be lost in the process). Final reports containing all available information are due for any incident 60 days after resolution of the incident and, to the extent incidents are not resolved, iterative reports are required to be made every four weeks after the initial report. These obligations will require Signatories to implement robust data and information records and logging processes.
- **Risk Allocation and Cooperation.** Signatories commit to cooperate with the AI Office, other Signatories, downstream users and modifiers of their systems, and civil society and academia in relation to their Code obligations, and allocate responsibilities internally (including between HR, finance, knowledge and IT teams) and along the AI value chain to best manage systemic risk. Although these sections are some of the most high-level, Signatories will need to consider multiple stakeholders to ensure safe and secure model development.

Conclusion

With this final version of the Code, the Commission will hope to have struck an appropriate balance between the interests of developers, deployers and copyright owners, to attract broad uptake of the Code whilst enhancing protections for rightsholders. Interestingly, in the same manner as the UK Government with its final version of the Data (Use and Access) Act 2025⁶, the Commission has currently not clearly responded to calls by rightsholders for GPAI providers to disclose the exact materials and information used to train the GPAI models, limiting the requirements to internal policies on compliance with EU copyright laws and public summaries of the same (including under Art. 53 AI Act). It will be interesting to see whether and how this develops in the public summary templates that the AI Office has committed to release to give more clarity on how GPAI models should comply with the AI Act in relation to information on copyright works (including publication of training data sources and TDM reservations under the Copyright Directive).

Additionally, although voluntary, the Code provides an insight into the Commission's views on the steps required for leading GPAI providers and, as such, may indicate what will be considered in any regulatory investigation in the EU once the GPAI obligations come into force in August 2025 (although see more on this in our upcoming alert on the new GPAI model guidelines). There will be much focus on the extent and tone of the AI Office's enforcement of the AI Act's GPAI obligations in the context of the Code (following its suggested grace period for good faith non-compliance by Signatories), which may play a key role in determining uptake across the market. As such, GPAI providers are likely to balance the scope of the Code, its required commitments against current practice and the potential benefits of streamlined compliance, against the risks of publicly committing to commitments that may go beyond the text of the AI Act itself, to assess whether to sign up to the Code, treat it as general guidance on good practice or take an alternative path altogether. The initial responses from potential Signatories have been mixed.

* * *

This memorandum is not intended to provide legal advice, and no legal or business decision should be based on its content. Questions concerning issues addressed in this memorandum should be directed to:

Jonathan H. Ashtor

+1-212-373-3823

jashtor@paulweiss.com

John P. Carlin

+1-202-223-7372

jcarlin@paulweiss.com

Katherine B. Forrest

+1-212-373-3195

kforrest@paulweiss.com

Anna R. Gressel

+1-212-373-3388

agressel@paulweiss.com

Henrik Morch

+32 2 884 0802

hmorch@paulweiss.com

John Patten

+44-20-7367-1684

jpatten@paulweiss.com

Georgina Hoy

+44-20-7601-8743

ghoy@paulweiss.com

Alex Zapalowski

+44-20-7367-1697

azapalowski@paulweiss.com

Associates Edmund Berney, Charlie Burrell, Scott Caravello and Ali Fazeli-Nia contributed to this Client Memorandum.

⁶ Available here: <https://www.legislation.gov.uk/ukpga/2025/18>