


EU AI Office Clarifies Key Obligations for AI Models Becoming Applicable in August



CONTRIBUTORS



Laura De Boel



Maneesha Mithal



Karol Piwonski



Roberto Yunquera
Sehwan

ALERTS

April 29, 2025

On April 22, 2025, the EU Commission's AI Office published draft [guidelines](#) to clarify the obligations in the EU AI Act for providers of general-purpose AI models (guidelines). These obligations will be applicable to AI models released in the EU market after August 2, 2025. The guidelines are currently open for public consultation, and the AI Office invited stakeholders to provide feedback by May 22, 2025, using this [form](#). The AI Office plans to adopt a final version of the guidelines before August 2025. Below, we present the main takeaways from the guidelines.

Background

The EU Artificial Intelligence Act (AI Act) was adopted last year, and it is becoming applicable in phases. A ban on certain AI practices already became applicable on February 2, 2025 (as we [explain here](#)). The next wave of requirements will be the obligations on general-purpose AI models (GPAI), which will become applicable in August 2025. In addition, the AI Office has begun issuing guidance to clarify certain provisions of the AI Act, including the [definition of AI systems](#) and the [prohibited AI practices](#) (which we cover [here](#)).

The AI Act imposes specific obligations on GPAI models, which are defined as models that i) display significant generality, ii) are capable of competently performing a wide range of distinct tasks, and iii) can be integrated into a variety of downstream AI systems or applications. Large generative AI models typically qualify as GPAI. Providers of GPAI models must draw up technical documentation, publish a summary of the data used for training, and implement a policy for compliance with EU copyright rules. The AI Act also imposes additional requirements for the most powerful GPAI models, considered to be GPAI with "systemic risk" (or GPAI-SR). To learn more about the AI Act's rules for GPAI models, please see our [AI Act FAQs](#).

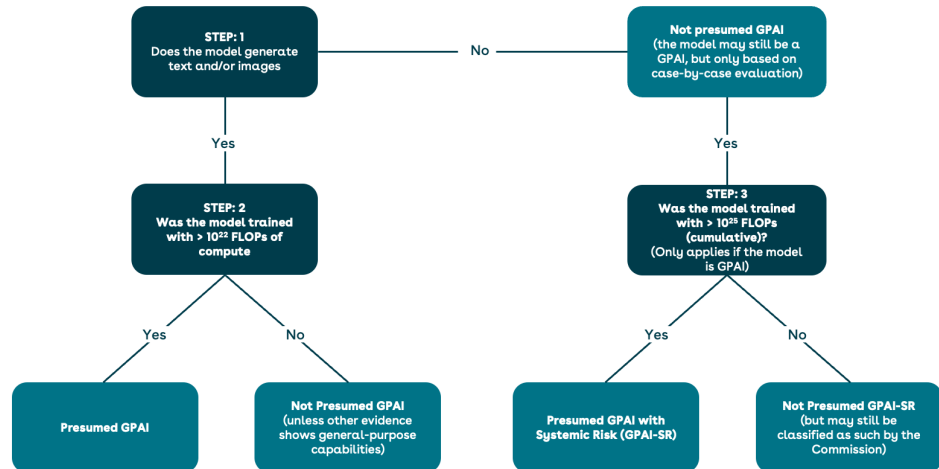
The guidelines clarify the AI Act's provisions on GPAI and GPAI-SR. They complement the voluntary Code of Practice on GPAI (CoP). Companies that develop GPAI models can voluntarily adhere to the CoP to demonstrate compliance with the AI Act. The AI Office has been coordinating the drafting of the CoP. The final version of the CoP is expected to be published in May.

Qualification of AI Model as GPAI

The guidelines provide a rebuttable presumption according to which generative AI models that were trained using compute greater than 1022 floating point operations (FLOPs) qualify as GPAI. This presumption is rebuttable, meaning that a model whose training compute exceeds 1022 FLOPs could still avoid the classification as GPAI if there is evidence that it does not display significant generality or that it cannot perform a wide range of distinctive tasks. For example, a model that can only be used for transcription of speech would not qualify as GPAI (narrow task), regardless of its training

compute. Conversely, a model may qualify as GPAI even where the training compute threshold is not met, if there is evidence that it displays sufficient generality and capabilities.

The presumption laid down in the guidelines for GPAI models is not mentioned in the AI Act itself. In contrast, the AI Act contains a different presumption for models exceeding the higher threshold of 1025 FLOPs in compute. Models that exceed this threshold laid down in the AI Act are presumed to qualify as GPAI-SR and are subject to additional obligations (e.g., risk assessment obligations). Please consult the decision tree below which summarizes these different presumptions:



The guidelines recognize two approaches to estimating compute: by tracking Graphics Processing Unit (GPU) usage (hardware-based approach), or by counting the expected number of FLOPs based on the model's architecture (architecture-based approach). The choice of the appropriate approach is left to the providers. The guidelines indicate that providers should estimate training compute before starting the pre-training run. They should notify the AI Office if the estimated compute exceeds the threshold for classifying the model as GPAI-SR (i.e., 1025 FLOPs).

Application of GPAI Rules When Fine-Tuning an AI Model

The guidelines clarify the application of the AI Act to the fine-tuning of a model (including a third-party model):

- *The guidelines lay down thresholds to identify when modifying a model can lead to a distinct new model.* If a GPAI model is modified, and the modification involves more than a third of the compute required for the model to be classified as GPAI (i.e., more than roughly 3 x 1021 FLOPs), then the modified model will be considered a new model, and compliance with the GPAI requirements will have to be considered separately. For instance, this could be relevant for cases where a GPAI provider develops a new GPAI model by fine-tuning an existing pre-trained model. If the fine-tuning uses significant compute exceeding the above thresholds, the modified GPAI model would need to separately comply with the GPAI requirements (e.g., the provider would need to draw up separate technical documentation for the new model).
- *Companies who conduct significant modifications on a third-party GPAI model would also be subject to GPAI requirements with respect to their modification.* The guidelines indicate that if a company licenses a GPAI model from a third-party and conducts significant modifications or fine-tuning (i.e., exceeding the compute threshold of roughly 3 x 1021 FLOPs described above), then the company responsible for the modification will need to separately assess compliance with the GPAI requirements. The guidelines clarify that this company would be subject to the GPAI requirements only with respect to the modification or fine-tuning it conducted (i.e., not with respect to the base model developed by a third-party, which is aligned with Recital 109 of the AI Act), and provided the resulting model qualifies as GPAI. This is relevant for cases where a company licenses an AI model from a third-party provider and significantly fine-tunes the model before integrating the resulting GPAI into its products and services. If the modifications meet the above thresholds, the company who conducted those modifications will need to comply with the AI Act with respect to those modifications (e.g., will need to draw up technical documentation explaining the modifications conducted, and will need to publish a summary of the training data used for the modifications).

- *If the fine-tuning results in a model that exceeds the thresholds for “systemic risk,” then the company responsible for the modification becomes subject to the GPAI-SR requirements.* If a company fine-tunes a GPAI model and the modified model ends up exceeding the thresholds for systemic risk (i.e., 1025 FLOPs, as shown in the decision tree above), then the company responsible for the fine-tuning would become subject to the GPAI-SR requirements. This means that the company would be required to conduct risk assessments and model evaluations, and notify the AI Office, amongst other GPAI-SR requirements. This could be relevant for companies that license third-party models that are close to meeting the GPAI-SR thresholds (e.g., close to 1025 FLOPs) and they conduct fine-tuning that brings the model above the threshold.

Next Steps

As the date for applicability of GPAI requirements approaches, we recommend that companies that develop or use AI take steps to map out their exposure to AI technologies and how the AI Act could apply to them. As first steps, companies could consider the following:

- *Establishing an AI governance structure:* Companies that develop or use AI should establish a dedicated AI governance team or function responsible for mapping out use of AI within the company and assessing the applicability of the EU AI Act.
- *Mapping AI use cases:* Companies should comprehensively map and document all AI technologies currently in use, being developed, or procured. This facilitates identifying any models that could qualify as GPAI or GPAI-SR under the AI Act, as well as any AI systems that could be prohibited or high-risk.
- *Developing and implementing internal AI policies:* Companies should develop internal AI policies to clearly define escalation and approval procedures for AI use cases (e.g., who should authorize the use of AI within a company, what other stakeholders within the company should be consulted before deploying AI tools, who should analyze the possible implications under the AI Act). Internal AI policies are the cornerstone of an AI governance program, allowing an organization to manage AI in a responsible and efficient manner.
- *Reviewing contract terms:* Companies should review their contract terms with their AI suppliers or with their customers (if their products include AI). Contracts should ensure sufficient rights for companies to train or use AI in the way that they intend and should ensure that the company receives all the information it needs from AI suppliers to comply with its own requirements under the AI Act, and make the assessments discussed in the guidelines.
- *Enforcement and cooperation with the AI Office:* The guidelines remind that obligations imposed on GPAI providers start to apply on August 2, 2025. However, providers of GPAI models placed on the market before August 2, 2025, have an additional two years to ensure compliance with the AI Act (i.e., until August 2, 2027). The guidelines explain what it means to “place a model on the market” and this includes making the model available via an API or software library, integrating it into a chatbot available via a web interface, into a mobile app, or into a product or service made available on the market. The AI Office recognizes that some providers may face difficulties in ensuring timely compliance and invites providers to proactively inform the AI Office in case they plan to launch a GPAI model after August 2, 2025, to discuss the company’s compliance strategy.

If a company develops or uses models that qualify as GPAI, we recommend reviewing the guidelines and considering participating in the public consultation to submit any suggestions before the final draft is published.

For more information on how to ensure your AI systems and models comply with the EU AI Act, please contact [Cédric Burton](#), [Laura De Boel](#), [Yann Padova](#), or [Nikolaos Theodorakis](#) from Wilson Sonsini’s Data, Privacy, and Cybersecurity practice.

Wilson Sonsini’s AI Working Group assists clients with AI-related matters. Please contact [Laura De Boel](#), [Maneesha Mithal](#), [Manja Sachet](#), or [Scott McKinney](#) for more information.

Roberto Yunquera Sehwaní and Karol Piwonski contributed to the preparation of this alert.