

General FAQ

Why do we need rules for general-purpose AI models?

AI promises huge [benefits](#) to our economy and society. General-purpose AI models play an important role in that regard, as they can be used for a variety of tasks and therefore form the basis for a range of downstream AI systems, used in Europe and worldwide.

The [AI Act](#) aims to ensure that general-purpose AI models are safe and trustworthy.

To achieve that aim, it is crucial that providers of general-purpose AI models possess a good understanding of their models along the entire AI value chain, both to enable the integration of such models into downstream AI systems and to fulfil their obligations under the AI Act. As explained in more detail below, providers of general-purpose AI models must draw up and provide technical documentation of their models to the AI Office and downstream providers, must put in place a copyright policy, and must publish a training content summary. In addition, providers of general-purpose AI models posing systemic risks, which may be the case either because they are very capable or because they have a significant impact on the internal market for other reasons, must notify the Commission, assess and mitigate systemic risks, perform model evaluations, report serious incidents, and ensure adequate cybersecurity of their models.

In this way, the AI Act contributes to safe and trustworthy innovation in Europe.

What are general-purpose AI models?

The AI Act defines a general-purpose AI model as “an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications” (Article 3(63)).

The Recitals to the AI Act further clarify which models should be deemed to display significant generality and to be capable of performing a wide range of distinct tasks.

According to Recital 98, “whereas the generality of a model could, inter alia, also be determined by a number of parameters, models with at least a billion of parameters and trained with a large amount of data using self-supervision at scale should be considered to display significant generality and to competently perform a wide range of distinctive tasks.”

Recital 99 adds that “large generative AI models are a typical example for a general-purpose AI model, given that they allow for flexible generation of content, such as in the form of text, audio, images or video, that can readily accommodate a wide range of distinctive tasks.”

Note that significant generality and ability to competently perform a wide range of distinctive tasks may be achieved by models within a single modality, such as text, audio, images, or video, if the modality is flexible enough. This may also be achieved by models that were developed, fine-tuned, or otherwise modified to be particularly good at a specific task.

The AI Office intends to provide further clarifications on what should be considered a general-purpose AI model, drawing on insights from the Commission’s Joint Research Centre, which is currently working on a scientific research project addressing this and other questions.

What are general-purpose AI models with systemic risk?

Systemic risks are risks of large-scale harm from the most advanced (i.e. state-of-the-art) models at any given point in time or from other models that have an equivalent impact (see Article 3(65)). Such risks can manifest themselves, for example, through the lowering of barriers for chemical or biological weapons development, unintended issues of control over autonomous general-purpose AI models, or harmful discrimination or disinformation at scale (Recital 110). The most advanced models at any given point in time may pose systemic risks, including novel risks, as they are pushing the state of the art. At the same time, some models below the threshold reflecting the state of the art may also pose systemic risks, for example, through reach, scalability, or scaffolding.

Accordingly, the AI Act classifies a general-purpose AI model as a general-purpose AI model with systemic risk if it is one of the most advanced models at that point in time or if it has an equivalent impact (Article 51(1)). Which models are considered general-purpose AI models with systemic risk may change over time, reflecting the evolving state of the art and potential societal adaptation to increasingly advanced models. Currently, general-purpose AI models with systemic risk are developed by a handful of companies, although this may also change over time.

To capture the most advanced models, the AI Act initially lays down a threshold of 10^{25} floating-point operations (FLOP) used for training the model (Article 51(1)(a) and (2)). Training a model that meets this threshold is currently estimated to cost tens of millions of Euros ([Epoch AI, 2024](#)). The AI Office will continuously monitor technological and industrial developments and the Commission may update the threshold to ensure that it continues to single out the most advanced models as the state of the art evolves by way of delegated act (Article 51(3)). For example, the value of the threshold itself could be adjusted, and/or additional thresholds introduced.

To capture models with an impact equivalent to the most advanced models, the AI Act empowers the Commission to designate additional models as posing systemic risk, based on criteria such as number of users, scalability, or access to tools (Article 51(1)(b), Annex XIII).

The AI Office intends to provide further clarifications on how general-purpose AI models will be classified as general-purpose AI models with systemic risk, drawing on insights from the Commission’s Joint Research

Centre which is currently working on a scientific research project addressing this and other questions.

What is a provider of a general-purpose AI model?

The AI Act rules on general-purpose AI models apply to providers placing such models on the market in the Union, irrespective of whether those providers are established or located within the Union or in a third country (Article 2(1)(a)).

A provider of a general-purpose AI model means a natural or legal person, public authority, agency or other body that develops a general-purpose AI model or that has such a model developed and places it on the market, whether for payment or free of charge (Article 3(3)).

To place a model on the market means to first make it available on the Union market (Article 3(9)), that is, to supply it for distribution or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge (Article 3(10)). Note that a general-purpose AI model is also considered to be placed on the market if that model's provider integrates the model into its own AI system which is made available on the market or put into service, unless the model is (a) used for purely internal processes that are not essential for providing a product or a service to third parties, (b) the rights of natural persons are not affected, and (c) the model is not a general-purpose AI model with systemic risk (Recital 97).

What are the obligations for providers of general-purpose AI models?

The obligations for providers of general-purpose AI models apply from 2 August 2025 (Article 113(b)), with special rules for general-purpose AI models placed on the market before that date (Article 111(3)).

Based on Article 53 of the AI Act, providers of general-purpose AI models must document technical information about the model for the purpose of providing that information upon request to the AI Office and national competent authorities (Article 53(1)(a)) and making it available to downstream providers (Article 53(1)(b)). They must also put in place a policy to comply with Union law on copyright and related rights (Article 53(1)(c)) and draw up and make publicly available a sufficiently detailed summary about the content used for training the model (Article 53(1)(d)).

The General-Purpose AI Code of Practice should provide further detail on these obligations in the sections dealing with transparency and copyright (led by Working Group 1).

Based on Article 55 of the AI Act, providers of general-purpose AI models with systemic risk have additional obligations. They must assess and mitigate systemic risks, in particular by performing model evaluations,

keeping track of, documenting, and reporting serious incidents, and ensuring adequate cybersecurity protection for the model and its physical infrastructure.

The General-Purpose AI Code of Practice should provide further detail on these obligations in the sections dealing with systemic risk assessment, technical risk mitigation, and governance risk mitigation (led by Working Groups 2, 3, and 4 respectively).

If someone open-sources a model, do they have to comply with the obligations for providers of general-purpose AI models?

The obligations to draw up and provide documentation to the AI Office, national competent authorities, and downstream providers (Article 53(1)(a) and (b)) do not apply if the model is released under a free and open-source license and its parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available. This exemption does not apply to general-purpose AI models with systemic risk (Article 53(2)). Recitals 102 and 103 further clarify what constitutes a free and open-source license and the AI Office intends to provide further clarifications on questions concerning open-sourcing general-purpose AI models.

By contrast, providers of general-purpose AI models with systemic risk must comply with their obligations under the AI Act regardless of whether their models are open-source. After the open-source model release, measures necessary to ensure compliance with the obligations of Articles 53 and 55 may be more difficult to implement (Recital 112). Therefore, providers of general-purpose AI models with systemic risk may need to assess and mitigate systemic risks before releasing their models as open-source.

The General-Purpose AI Code of Practice should provide further detail on what the obligations in Articles 53 and 55 imply for different ways of releasing general-purpose AI models, including open-sourcing.

An important but difficult question underpinning this process is that of finding a balance between pursuing the benefits and mitigating the risks from the open-sourcing of advanced general-purpose AI models: open-sourcing advanced general-purpose AI models may indeed yield significant societal benefits, including through fostering AI safety research; at the same time, when such models are open-sourced, risk mitigations are more easily circumvented or removed.

Do the obligations for providers of general-purpose AI models apply in the Research & Development phase?

Article 2(8) specifies that the AI Act “does not apply to any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market or put into service.”

At the same time, many of the obligations for providers of general-purpose AI models (with and without systemic risk) explicitly or implicitly pertain to the Research & Development phase of models intended for but prior to the placing on the market. For example, this is the case for the obligations for providers to notify the Commission that their general-purpose AI model meets or will meet the training compute threshold (Articles 51 and 52), to document information about training and testing (Article 53), and to assess and mitigate systemic risk (Article 55). In particular, Article 55(1)(b) explicitly specifies that “providers of general-purpose AI models with systemic risk shall assess and mitigate possible systemic risks at Union level, including their sources, that may stem from the development (...) of general-purpose AI models with systemic risk.”

In any case, the AI Office expects discussions with providers of general-purpose AI models with systemic risk to start early in the development phase. This is consistent with the obligation for providers of general-purpose AI models that meet the training compute threshold laid down in Article 51(2) to “notify the Commission without delay and in any event within two weeks after that requirement is met or it becomes known that it will be met” (Article 52(1)). Indeed, training of general-purpose AI models takes considerable planning, which includes the upfront allocation of compute resources, and providers of general-purpose AI models are therefore able to know if their model will meet the training compute threshold before the training is complete (Recital 112).

The AI Office intends to provide further clarifications on this question.

If someone fine-tunes or otherwise modifies a model, do they have to comply with the obligations for providers of general-purpose AI models?

General-purpose AI models may be further modified or fine-tuned into new models (Recital 97). Accordingly, downstream entities that fine-tune or otherwise modify an existing general-purpose AI model may become providers of new models. The specific circumstances in which a downstream entity becomes a provider of a new model is a difficult question with potentially large economic implications, as many organisations and individuals fine-tune or otherwise modify general-purpose AI models developed by another entity. The AI Office intends to provide further clarifications on this question.

In the case of a modification or fine-tuning of an existing general-purpose AI model, the obligations for providers of general-purpose AI models in Article 53 should be limited to the modification or fine-tuning, for example, by complementing the already existing technical documentation with information on the modifications (Recital 109). The obligations for providers of general-purpose AI models with systemic risk in Article 55 may be limited in similar ways. The General-Purpose AI Code of Practice could reflect differences between providers that initially develop general-purpose AI models and those that fine-tune or otherwise modify an existing model.

Note that regardless of whether a downstream entity that incorporates a general-purpose AI model into an AI system is deemed to be a provider of the general-purpose AI model, that entity must comply with the relevant AI Act requirements and obligations for AI systems.

What is the General-Purpose AI Code of Practice?

Based on Article 56 of the AI Act, the [General-Purpose AI Code of Practice](#) should detail the manner in which providers of general-purpose AI models and of general-purpose AI models with systemic risk may comply with their obligations under the AI Act. The AI Office is facilitating the drawing-up of this Code of Practice, with four working groups chaired by independent experts and involving nearly 1000 stakeholders, EU Member States representatives, as well as European and international observers.

More precisely, the Code of Practice should detail at least how providers of general-purpose AI models may comply with the obligations laid down in Articles 53 and 55. This means that the Code of Practice can be expected to have two parts: one that applies to providers of all general-purpose AI models (Article 53), and one that applies only to providers of general-purpose AI models with systemic risk (Article 55). Another obligation that may be covered by the Code of Practice is the obligation to notify the Commission for providers of general-purpose AI models that meet or are expected to meet the conditions listed in Article 51 for being classified as a general-purpose AI model with systemic risk (Article 52(1)).

What is not part of the Code of Practice?

The Code of Practice should not address inter alia the following issues: defining key concepts and definitions from the AI Act (such as “general-purpose AI model”), updating the criteria or thresholds for classifying a general-purpose AI model as a general-purpose AI model with systemic risk (Article 51), outlining how the AI Office will enforce the obligations for providers of general-purpose AI models (Chapter IX Section 5), and questions concerning fines, sanctions, and liability.

These issues may instead be addressed through other means (decisions, delegated acts, implementing acts, further communications from the AI Office, etc.).

Nevertheless, the Code of Practice may include commitments by providers of general-purpose AI models to document and report additional information, as well as to involve the AI Office and third parties throughout the entire model lifecycle, in so far as this is considered necessary for providers to effectively comply with their obligations under the AI Act.

Do AI systems play a role in the Code of Practice?

The AI Act distinguishes between AI systems and AI models, imposing requirements for certain AI systems (Chapters II-IV) and obligations for providers of general-purpose AI models (Chapter V). While the provisions of the AI Act concerning AI systems depend on the context of use of the system, the provisions of the AI Act concerning general-purpose AI models apply to the model itself, regardless of what is or will be its ultimate use. The Code of Practice should only pertain to the obligations in the AI Act for providers of general-purpose AI models.

Nevertheless, there are interactions between the two sets of rules, as general-purpose AI models are typically integrated into and form part of AI systems. If a provider of the general-purpose AI model integrates a general-purpose AI model into an AI system, that provider must comply with the obligations for providers of general-purpose AI models and, if the AI system falls within the scope of the AI Act, must comply with the requirements for AI systems. If a downstream provider integrates a general-purpose AI model into an AI system, the provider of the general-purpose AI model must cooperate with the downstream provider of the AI system to ensure that the latter can comply with its obligations under the AI Act if the AI system falls within the scope of the AI Act (for example by providing certain information to the downstream provider).

Given these interactions between models and systems, and between the obligations and requirements for each, an important question underlying the Code of Practice concerns which measures are appropriate at the model layer, and which need to be taken at the system layer instead.

How does the Code of Practice take into account the needs of start-ups?

The Code of Practice should set out its objectives, measures and, as appropriate, key performance indicators (KPIs) to measure the achievement of its objectives. Measures and KPIs related to the obligations applicable to providers of all general-purpose AI models should take due account of the size of the provider and allow simplified ways of compliance for SMEs, including start-ups, that should not represent an excessive cost and not discourage the use of such models (Recital 109). Moreover, the KPIs related to the obligations applicable to providers of general-purpose AI models with systemic risk should reflect differences in size and capacity between various providers (Article 56(5)), while ensuring that they are proportionate to the risks (Article 56(2)(d)).

When will the Code of Practice be finalised?

After the publication of the first draft of the Code of Practice, it is expected that there will be three more drafting rounds over the coming months. Thirteen Chairs and Vice-Chairs, drawn from diverse backgrounds in computer science, AI governance and law, are responsible for synthesizing submissions from a multi-stakeholder consultation and discussions with the Code of Practice [Plenary](#) consisting of around 1000 stakeholders. This iterative process will lead to a final Code of Practice which should reflect the various submissions whilst ensuring a convincing implementation of the legal framework.

What are the legal effects of the Code of Practice?

If approved via implementing act, the Code of Practice obtains general validity, meaning that adherence to the Code of Practice becomes a means to demonstrate compliance with the AI Act. Nevertheless, compliance with the AI Act can also be demonstrated in other ways.

Based on the AI Act, additional legal effects of the Code of Practice are that the AI Office can enforce adherence to the Code of Practice (Article 89(1)) and should take into account commitments made in the Code of Practice when fixing the amount of fines (Article 101(1)).

How will the Code of Practice be reviewed and updated?

While the first draft of the Code of Practice does not yet contain details on its review and updating, further iterations of the draft, and any implementing act adopted to approve the final Code of Practice, can be expected to include this information.

Which enforcement powers does the AI Office have?

The AI Office will enforce the obligations for providers of general-purpose AI models (Article 88), as well as support governance bodies within Member States in their enforcement of the requirements for AI systems (Article 75), among other tasks. Enforcement by the AI Office is underpinned by the powers given to it by the AI Act, namely the powers to request information (Article 91), conduct evaluations of general-purpose AI models (Article 92), request measures from providers, including implementing risk mitigations and recalling the model from the market (Article 93), and to impose fines of up to 3% of global annual turnover or 15 million Euros, whichever is higher (Article 101).