



EUROPEAN
COMMISSION

Brussels, 18.7.2025
C(2025) 5045 final

ANNEX

ANNEX

to the

Communication to the Commission

**Approval of the content of the draft Communication from the Commission – Guidelines
on the scope of the obligations for general-purpose AI models established by Regulation
(EU) 2024/1689 (AI Act)**

Contents

1. Background and objectives of the Commission guidelines	3
2. General-purpose AI models.....	5
2.1. When is a model a general-purpose AI model?	5
2.2. The lifecycle of a general-purpose AI model	8
2.3. When is a general-purpose AI model a general-purpose AI model with systemic risk?10	
2.3.1. Classification	10
2.3.2. Notification	11
2.3.3. Procedure for contesting classification	12
2.3.4. Designation	14
2.3.5. Procedure for contesting designation.....	15
3. Providers placing on the market general-purpose AI models	15
3.1. When is an actor a provider placing on the market a general-purpose AI model?.....	16
3.1.1. Examples of providers of general-purpose AI models.....	16
3.1.2. Examples of placing on the market of general-purpose AI models.....	17
3.1.3. Considerations for general-purpose AI models integrated into AI systems.....	17
3.2. Downstream modifiers as providers of general-purpose AI models	18
3.2.1. Downstream modifiers becoming providers of general-purpose AI models	20
3.2.2. Downstream modifiers becoming providers of general-purpose AI models with systemic risk	20
4. Exemptions from certain obligations for certain models released as open-source	21
4.1. Scope of the exemptions.....	21
4.2. Conditions for the exemptions to apply	22
4.2.1. Conditions on the licence.....	22
4.2.2. Lack of monetisation	24
4.2.3. Public availability of parameters, including the weights, the information on the model architecture, and the information on model usage	25
5. Enforcement of the obligations for providers of general-purpose AI models.....	25
5.1. Effects of adhering to a code of practice assessed as adequate	25
5.2. Supervision, investigation, enforcement, and monitoring of Chapter V of the AI Act .	27
5.3. Compliance after entry into application of Chapter V of the AI Act on 2 August 2025	28
6. Review and update of the Commission guidelines	30
Annex: Training compute of general-purpose AI models	31
A.1. What should be estimated?	31
A.2. How should training compute be estimated?.....	32

A.2.1. Hardware-based approach.....	33
A.2.2. Architecture-based approach.....	33
A.3. Training compute of models with approximately one billion parameters.....	34

1. Background and objectives of the Commission guidelines

- (1) Regulation (EU) 2024/1689 of the European Parliament and the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending certain regulations ('AI Act')¹ entered into force on 1 August 2024. Its aim is to promote innovation in and uptake of artificial intelligence ('AI') in the Union, while ensuring a high level of protection of health, safety, and fundamental rights, including democracy and the rule of law.
- (2) General-purpose AI models play a significant role in the innovation in and uptake of AI in the Union. This is because they can be used for a variety of tasks and be integrated into a wide array of downstream AI systems. Therefore, providers of general-purpose AI models have a particular role and responsibility along the AI value chain, including regarding downstream providers who need a good understanding of the models and their capabilities. This understanding is important for downstream providers to be able to integrate such models into their products, and to fulfil their obligations under the AI Act.
- (3) For this reason, the AI Act lays down proportionate transparency measures, including an obligation to draw up and keep up-to-date documentation to be made available to downstream providers and, upon request, to the AI Office and the national competent authorities (Article 53(1), points (a) and (b), AI Act). Providers of general-purpose AI models that are released under a free and open-source licence are exempted from these transparency-related requirements under certain conditions, unless they are general-purpose AI models with systemic risk.
- (4) Given that general-purpose AI models are often trained on vast amounts of text, images, videos and other data, and that this training data may include copyright protected content, the AI Act also requires providers of general-purpose AI models to put in place a policy to comply with Union copyright law (Article 53(1), point (c), AI Act), and to make publicly available a summary of the content used for the training (Article 53(1), point (d), AI Act).
- (5) General-purpose AI models may also present systemic risks. These are risks which, are specific to the most advanced general-purpose AI models and can have a significant impact on the Union market (Article 3(65) AI Act). Under the AI Act, general-purpose AI models trained with a cumulative amount of computational resources greater than 10^{25} floating-point operations ('FLOP', see Article 3(67) AI Act) are presumed to be general-purpose AI models with systemic risk (Article 51(2) AI Act). The Commission may also designate general-purpose AI models as general-purpose AI models with systemic risk based on the criteria in Annex XIII AI Act. Providers of general-purpose AI models with systemic risk are subject, in addition to the obligations for providers of general-purpose AI models, to obligations aimed at assessing and mitigating those risks. These include conducting model evaluations, reporting serious incidents and ensuring an adequate level of cybersecurity protection.
- (6) Under the AI Act, the Commission has the exclusive competence to enforce compliance of providers of general-purpose AI models, with and without systemic

risk, with their obligations laid down in Chapter V of the AI Act, the implementation of which is entrusted to the European Artificial Intelligence Office ('AI Office', see Article 3(47) AI Act).

- (7) Under Article 96(1) AI Act, the Commission must develop guidelines on the practical implementation of the AI Act. These guidelines focus on the scope of the obligations for providers of general-purpose AI models laid down in the AI Act, in light of their imminent entry into application on 2 August 2025. They solely concern the AI Act and are not applicable to other Union laws, including on liability.
- (8) More precisely, these guidelines address four key topics: (i) general-purpose AI models (Section 2); (ii) providers placing on the market general-purpose AI models (Section 3); (iii) exemptions from certain obligations for providers of general-purpose AI models released as open-source (Section 4); and (iv) compliance with the obligations for providers of general-purpose AI models (Section 5). In doing so, these guidelines should in particular help actors along the AI value chain to determine: (i) whether their model is a general-purpose AI model; (ii) whether they are the provider placing on the market that general-purpose AI model; (iii) whether they are exempt from the obligations for providers of general-purpose AI models; and (iv) what to expect regarding the Commission's enforcement of compliance with these obligations, in particular during the period immediately after their entry into application on 2 August 2025.
- (9) These guidelines are not binding for providers of general-purpose AI models; an authoritative interpretation of the AI Act may only be given by the Court of Justice of the European Union ('CJEU'). Nevertheless, these guidelines set out the Commission's interpretation and application of the AI Act, on which it will base its enforcement action. This facilitates providers' compliance with their obligations and contributes to the effective implementation of the AI Act. Still, a case-by-case assessment will always be necessary to account for the specifics of each individual case. These guidelines will also be reviewed and updated as necessary, in particular in light of technological, societal, and market developments (Section 6).
- (10) These guidelines are complementary to other initiatives of the Commission contributing to the proper implementation of the obligations for providers of general-purpose AI models under the AI Act, notably to the codes of practice drawn up in line with Article 56 AI Act. The guidelines concern the scope of the obligations for providers of general-purpose AI models laid down in the AI Act and their implementation (including the role of codes of practice in demonstrating compliance). By contrast, the codes of practice lay down specific measures which providers of general-purpose AI models may implement to comply with their obligations laid down in Articles 53(1) and 55(1) AI Act. These guidelines are also complementary to the template provided by the Commission for the public summary of the content used to train general-purpose AI models under Article 53(1), point (d), AI Act.
- (11) These guidelines take into account input gathered by the Commission through a public multi-stakeholder consultation, which was launched on 22 April 2025 and

closed on 22 May 2025. An earlier version of these guidelines was presented to the Member States via the European Artificial Intelligence Board ('Board') on 30 June 2025. These guidelines also take into account input from the pool of experts, put into place by the Commission's Joint Research Centre ('JRC') to provide the AI Office with advice on the categorisation of AI models as general-purpose AI models and general-purpose AI models with systemic risk, as well as other experts.

2. General-purpose AI models

- (12) This section of the guidelines focuses on the concept of a 'general-purpose AI model'. It clarifies when the Commission considers a model to be a general-purpose AI model by providing an indicative criterion (Section 2.1). It also clarifies the notion of the 'lifecycle' of a general-purpose AI model and its implications for the obligations for providers of general-purpose AI models (Section 2.2). Further, it clarifies when a general-purpose AI model is a general-purpose AI model with systemic risk under Articles 51 and 52 AI Act (Section 2.3).

2.1. When is a model a general-purpose AI model?

- (13) Article 3(63) 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, except AI models that are used for research, development or prototyping activities before they are placed on the market'*. This definition lists in a general manner factors that determine whether a model is a general-purpose AI model. Nevertheless, it does not set out specific criteria that potential providers can use to assess whether their model is a general-purpose AI model. Yet it is important to provide conditions that are simple to check to limit the burden on the many actors that must assess whether they are providers of general-purpose AI models.
- (14) Given the wide variety of capabilities and use cases for general-purpose AI models, it is not feasible to provide a precise list of capabilities that a model must display and tasks that it must be able to perform in order to determine whether it is a general-purpose AI model.
- (15) Instead, the Commission's approach is to use the amount of computational resources used to train the model ('training compute', see paragraph 115 for how this notion should be understood in these guidelines) measured in FLOP as well as the modalities of the model to set an indicative criterion for when it considers a model to be a general-purpose AI model. This approach builds on recital 98 AI Act which specifies that *'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’, as well as on recital 99 AI Act which clarifies that ‘[l]arge 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’.

- (16) Training compute has the advantage of combining number of parameters and number of training examples into a single number that is reasonably straightforward for providers to estimate. This number is typically proportional to the number obtained by multiplying these two numbers, allowing a single threshold to be set rather than separate thresholds for model size and training data size. While training compute is an imperfect proxy for generality and capabilities, the Commission considers setting an indicative criterion which includes a training compute threshold to be the most suitable approach at present. Nevertheless, the Commission’s approach may change in the future as technology and the market evolve. Furthermore, the Commission will continue to investigate the availability of other criteria that could be used to assess generality and capabilities with relative ease, especially for smaller actors.¹
- (17) Based on the considerations above, an indicative criterion for a model to be considered a general-purpose AI model is that its training compute is greater than 10^{23} FLOP and it can generate language (whether in the form of text² or audio³), text-to-image or text-to-video.
- (18) This threshold corresponds to the approximate amount of compute typically used to train a model with one billion parameters on a large amount of data. While such models may be trained with varying amounts of compute depending on how much data is used, 10^{23} FLOP is typical for models trained on large amounts of data as of the time of writing (see the examples in Annex A.3)⁴.
- (19) The modalities are chosen based on the fact that models trained to generate language – be it via text or speech (as a type of audio) – are able to use language to communicate, store knowledge, and reason. No other modality confers such a wide range of capabilities. Consequently, models that generate language are typically more capable of competently performing a wider range of tasks than other models. Although models that generate images or video typically exhibit a narrower range of capabilities and use cases compared to those that generate language, such models may nevertheless be considered to be general-purpose AI models. Text-to-image and text-to-video models are capable of generating a wide range of visual outputs,

¹ In the future, if deemed appropriate, the Commission may take into account benchmarks to determine whether a model is a general-purpose AI model.

² The Commission understands the modality of ‘text’ to include ‘code’.

³ The Commission understands the modality of ‘audio’ to include ‘speech’.

⁴ This value for the threshold was updated from the value proposed in the public consultation for these guidelines. Previously, the reference examples in Annex A.3 contained several examples of models trained with less than 10^{23} FLOP. However, these models were trained for research purposes only by non-commercial actors (e.g. academic groups or research non-profits), and so they were deemed unrepresentative of models that should be in scope of the AI Act.

which enables flexible content generation that can readily accommodate a wide range of distinct tasks.

- (20) If a general-purpose AI model meets the criterion from paragraph 17 but, exceptionally, does not display significant generality or is not capable of competently performing a wide range of distinct tasks, it is not a general-purpose AI model. Similarly, if a general-purpose AI model does not meet that criterion but, exceptionally, displays significant generality and is capable of competently performing a wide range of distinct tasks, it is a general-purpose AI model.

Example of a model in scope

- A model is trained on a broad range of natural language data (i.e. text) curated and scraped from the internet and other sources (as is currently typical for language models) using 10^{24} FLOP.
 - The criterion from paragraph 17 indicates that the model should be a general-purpose AI model because it can generate text and its training compute is greater than 10^{23} FLOP. Training on a broad range of natural language further indicates that the model should display significant generality and should be capable of competently performing a wide range of distinct tasks. Therefore, the model likely is a general-purpose AI model.

Examples of models out of scope

- A model is trained on natural language data, using 10^{22} FLOP, and is not able to competently perform a wide range of distinct tasks.
 - While the model can generate text, its training compute is not greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should not be a general-purpose AI model. The fact that the model is not capable of performing a wide range of distinct tasks confirms that it is not a general-purpose AI model.
- A model is trained specifically for the task of transcribing speech to text, using 10^{24} FLOP.
 - The model can generate text and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (transcribing speech), it is not actually a general-purpose AI model.
- A model is trained specifically for the task of generating speech from text, using 10^{24} FLOP.
 - The model can generate speech and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (generating speech from text), it is not actually a general-purpose AI model.
- A model is trained specifically for the task of increasing the resolution of images, using 10^{24} FLOP.

- The model can generate images and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (upsampling images), it is not actually a general-purpose AI model.
- A model is trained specifically for the task of filling in damaged or missing parts of images, using 10^{24} FLOP.
 - The model can generate images and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (inpainting images), it is not actually a general-purpose AI model.
- A model is trained specifically for playing chess or video games, using 10^{24} FLOP.
 - While the model's training compute is greater than 10^{23} FLOP, if it cannot generate text, speech, text-to-image, and text-to-video, the criterion from paragraph 17 indicates that the model should not be a general-purpose AI model. If the model can only competently perform a narrow set of tasks (playing chess or video games), this fact confirms that it is not a general-purpose AI model.
- A model is trained specifically for modelling weather patterns or physical systems, using 10^{24} FLOP.
 - While the model's training compute is greater than 10^{23} FLOP, if it cannot generate text, speech, text-to-image, and text-to-video, the criterion from paragraph 17 indicates that the model should not be a general-purpose AI model. If the model can only competently perform a narrow set of tasks (modelling weather patterns or physical systems), this fact confirms that it is not a general-purpose AI model.
- A model is trained specifically for generating music, based on text prompts including lyrics, using 10^{24} FLOP.
- (21) The model can generate speech (in the form of song lyrics) and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (generating music, based on text prompts including lyrics) it is not actually a general-purpose AI model.
- A model is trained specifically for generating sound effects, using 10^{24} FLOP. It can generate language (speech), but the quality of the content of the speech generated is low (e.g. the content is not consistent over the course of a brief conversation).
 - The model can generate language and its training compute is greater than 10^{23} FLOP. Therefore, the criterion from paragraph 17 indicates that the model should be a general-purpose AI model. However, if the model can only competently perform a narrow set of tasks (sound effects) it is not actually a general-purpose AI model.

2.2. The lifecycle of a general-purpose AI model

- (22) Providers of general-purpose AI models with systemic risk are expected to 'continuously assess and mitigate systemic risks', including by 'taking appropriate

measures along the entire model's lifecycle' (recital 114 AI Act). They must also *'ensure an adequate level of cybersecurity protection for the model and its physical infrastructure, if appropriate, along the entire model lifecycle'* (recital 115 AI Act). The notion of the 'lifecycle' of a model thus plays an important role for the obligations for providers of general-purpose AI models with systemic risk.

- (23) The iterative and interlinked process through which a provider may develop a 'model', for example through techniques such as distillation, quantisation, or merging of model weights, makes it difficult to clearly delineate a model and its lifecycle. In light of this challenge, the Commission understands the notion of 'model', and consequently its 'lifecycle', in a broad sense. In practice, the Commission considers the lifecycle of a general-purpose AI model to begin at the start of the large pre-training run.⁵ Any subsequent development of the model downstream of this large pre-training run performed by the provider or on behalf of the provider, whether before or after the model has been placed on the market, forms part of the same model's lifecycle rather than giving rise to new models. Different considerations apply if another actor modifies the model (see Section 3.2). A model is thus considered to be the same model along its entire lifecycle, i.e. throughout its development, market availability, and use. In particular, different stages of the development of a model are not considered to constitute different models.
- (24) This approach to the lifecycle of a general-purpose AI model has the implications described below for the obligations for providers of general-purpose AI models.
- The documentation required under Article 53(1), points (a) and (b), AI Act must be drawn up for each model placed on the market and kept up to date throughout its entire lifecycle.
 - The copyright policy required under Article 53(1), point (c), AI Act must be applied throughout the entire lifecycle of each of the provider's relevant models. Providers may choose to develop one policy and apply it to all their relevant models..
 - The summary of the content used for training required under Article 53(1), point (d), AI Act must be drawn up and made publicly available for all the provider's models placed on the market. The template that is currently being prepared by the AI Office will specify circumstances under which the summary must be updated.
 - The systemic risk assessment and mitigation required under Article 55(1) AI Act must be carried out continuously for each model throughout its entire lifecycle. This may involve a combination of ongoing measures, as well as more comprehensive and thorough measures at regular intervals and before key decisions in the lifecycle of the model. These measures should be effective, adequate, and proportionate to the systemic risks. The governance

⁵A large pre-training run is understood as the foundational training run conducted on a large amount of data to build the model's general capabilities, which may take place after smaller experimental training runs, and which may be followed by fine-tuning for specialisation or other post-training enhancements.

mitigations required under Article 55(1) AI Act must be applied throughout the entire lifecycle of each of the provider's relevant models. For some mitigations, providers may be able to develop them once and then apply them to all their relevant models.

2.3. When is a general-purpose AI model a general-purpose AI model with systemic risk?

- (25) 'General-purpose AI models with systemic risk' form a special class of general-purpose AI models. Providers of such models are subject to additional obligations concerning the assessment and mitigation of the 'systemic risks' presented by these models, under Articles 52 and 55 AI Act. The AI Act defines a 'systemic risk' as *'a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain'* (Article 3(65) AI Act).

2.3.1. Classification

- (26) Under Article 51(1) AI Act, a general-purpose AI model is classified as a general-purpose AI model with systemic risk if it meets either of the following two conditions:
- it has 'high-impact capabilities', namely *'capabilities that match or exceed those recorded in the most advanced models'* (Article 3(64) AI Act);
 - *'based on a decision of the Commission, ex officio or following a qualified alert from the scientific panel, it has capabilities or an impact equivalent to those set out in [the preceding point] having regard to the criteria set out in Annex XIII'.*
- (27) From the moment when a general-purpose AI model meets either of the two conditions above, the model is classified as a general-purpose AI model with systemic risk and its provider must comply with the relevant obligations.
- (28) Whether a given general-purpose AI model has high-impact capabilities should be *'evaluated on the basis of appropriate technical tools and methodologies, including indicators and benchmarks'*, under Article 51(1), point (a), AI Act. These tools and methodologies are to be further specified by the Commission through adoption of delegated acts (see paragraph 28). Notwithstanding future adoption of such delegated acts, Article 51(2) AI Act states that a general-purpose AI model is *'presumed to have high-impact capabilities pursuant to paragraph 1, point (a), when the cumulative amount of computation used for its training measured in floating point operations is greater than 10^{25} '.* This is because the cumulative amount of computation used for the training of a general-purpose AI model ('cumulative training compute', see paragraph 116 for how this notion should be

understood in these guidelines) measured in FLOP is considered to be a relevant metric for identifying high-impact capabilities (see recital 111 AI Act).

- (29) Under Article 51(3) AI Act, the Commission is empowered to adopt delegated acts, in line with Article 97 AI Act, to adjust the thresholds set out in Article 51(1) and (2) AI Act. The Commission is also empowered to introduce additional benchmarks and indicators when necessary, to account for technological advancements – such as improvements in algorithms or greater hardware efficiency – so that these thresholds continue to reflect the state of the art.

2.3.2. Notification

- (30) When a general-purpose AI model has met, or it becomes known that it will meet a requirement leading to the presumption that the model has high-impact capabilities,⁶ the provider must notify the Commission in line with Article 52(1) AI Act (see also recital 112 AI Act). The notification must happen *‘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) AI Act).
- (31) In particular, a notification may be required *before* training is complete, if the provider can reasonably foresee that the requirement that leads to the presumption of the model having high-impact capabilities is reasonably likely to be met. In this regard, recital 112 AI Act specifies that *‘training of general-purpose AI models takes considerable planning which includes the upfront allocation of compute resources and, therefore, providers of general-purpose AI models are able to know if their model would meet the threshold before the training is completed.’* Since the ‘planning’ and ‘upfront allocation of compute resources’ take place before the start of the large pre-training run, providers should estimate the cumulative amount of training compute that they will use before starting this run (see the Annexes A.1 and A.2 to these guidelines for how to estimate training compute). Following this estimation, providers should proceed in the following way.
- If the estimated value meets the threshold laid down in Article 51(2) AI Act, providers should notify the Commission *‘without delay and in any event within two weeks’*, under Article 52(1) AI Act.
 - If the estimated value does not meet the threshold laid down in Article 51(2) AI Act, providers should in any case closely monitor their actual and expected compute usage over the course of training and throughout the model’s entire lifecycle. This is to ensure that they are able to know if and when the cumulative amount of training compute has met or will meet the

⁶ In the absence of a delegated act amending the thresholds or supplementing the benchmarks and indicators listed in Article 51(1) and (2) AI Act, the only possible such requirements are the requirement that the model meets or will meet the threshold laid down in Article 52(1) AI Act, and the requirement that the model is the result of a modification of a general-purpose AI model with high-impact capabilities that meets or will meet the threshold laid down in paragraph 60.

threshold, and that they notify the Commission accordingly under Article 52(1) AI Act.

- (32) Under Article 52(1) AI Act, the notification should include *‘the information necessary to demonstrate that the relevant requirement has been met’*. In the case where the notification has been triggered because the model meets or will meet the threshold laid down in Article 51(2) AI Act, or because the model is the result of a modification of a general-purpose AI model with high-impact capabilities that meets or will meet the threshold laid down in paragraph 60,⁷ this information should cover in particular:
- the amount of compute estimated by the provider that has triggered the requirement to notify, reported in FLOP and with two significant figures (e.g. 2.3×10^{25} FLOP);
 - a description of the approach used to estimate this amount of compute, including approaches used to make approximations where precise information is not available.

2.3.3. Procedure for contesting classification

- (33) When a provider notifies the Commission under Article 52(1) AI Act, they *‘may present, with its notification, sufficiently substantiated arguments to demonstrate that, exceptionally, although it meets that requirement, the general-purpose AI model does not present, due to its specific characteristics, systemic risks and therefore should not be classified as a general-purpose AI model with systemic risk’* (Article 52(2) AI Act). This provision acknowledges that while general-purpose AI models that meet the threshold laid down in Article 51(2) AI Act or that have high-impact capabilities typically present systemic risks, there may be cases where this is not so.
- (34) Providers may present arguments aimed at demonstrating that their model does not or will not present systemic risks by demonstrating that their model does not or will not have high-impact capabilities, namely capabilities that *‘match or exceed those recorded in the most advanced models’* (Article 3(64) AI Act). This is because systemic risks are *‘specific to high-impact capabilities’* (Article 3(65) AI Act). Such an approach would be appropriate in cases where the notification has been triggered because of the *presumption* that the model has high-impact capabilities (for example, because the model has met or will meet the threshold laid down in Article 51(2) AI Act), but the provider deems that the presumption is not well-founded for its model. The Commission notes that the burden of adducing evidence that the presumption deriving from the fulfilment of the quantitative thresholds should not apply should be borne by that provider.

⁷ In the absence of a delegated act amending the thresholds or supplementing the benchmarks and indicators listed in Article 51(1) and (2) AI Act, these two requirements are the only requirements that trigger the obligation to notify the Commission under Article 52(1) AI Act.

- (35) As part of justifying why their model does not or will not have high-impact capabilities and therefore does not or will not present systemic risks, providers should include information available to them at the time of notification about the model's achieved or anticipated capabilities, including in the form of actual or forecasted benchmark results (for example based on scaling analyses). To complement this information and allow the Commission to properly assess whether the model does not or will not have high-impact capabilities and therefore does not or will not present systemic risks, providers are strongly advised to include any other information that may have a bearing on the model's capabilities, such as model architecture, number of parameters, number of training examples, data curation and processing techniques, training techniques, input and output modalities, expected tool use, and expected context length.
- (36) Upon receiving a notification with arguments for why the model does not present systemic risks, the Commission will assess the arguments and decide whether to accept or reject them, in line with the rules of procedure of the European Commission⁸ and established principles of EU law. In particular, in the case of rejection, in line with Article 41 of the Charter of Fundamental Rights of the European Union, providers will have the right to be heard before the decision is taken (through receiving a draft decision letter in advance), and the right to have access to their file, while respecting the legitimate interests of confidentiality and of professional and business secrecy. The Commission's decision will include reasons for why the arguments submitted by the provider under Article 52(2) AI Act have been accepted or rejected.
- (37) In the case where the provider's arguments are aimed at rebutting the presumption that the model has high-impact capabilities and therefore does not present systemic risks, the Commission will assess whether the provider has presented sufficiently substantiated arguments manifestly calling into question this presumption.
- (38) High-impact capabilities are, by definition, *'capabilities that match or exceed those recorded in the most advanced models'* (Article 3(64) AI Act). Moreover, they determine whether a general-purpose AI model is classified as a general-purpose AI model with systemic risk. For this reason, the level of capabilities beyond which a model is considered to have 'high-impact capabilities', or equivalently to be amongst the 'most advanced models', is expected to change over time, as capabilities increase and risks evolve. In particular, the Commission does not consider it to refer to a *fixed* level of capabilities.
- (39) In its assessment of whether the model is amongst the most advanced models at the time of notification, the Commission will take into account the extent to which the cumulative training compute of the model is indicative of the model being amongst these models. In this regard, the extent to which the cumulative training compute of the model exceeds the threshold laid down in Article 51(2) AI Act will also be taken into account. The Commission will also assess any other elements beyond

⁸ [Rules of Procedure of the European Commission | EUR-Lex.](#)

cumulative training compute which influence the achieved or expected capabilities of the model, including forecasted or achieved benchmark scores.

- (40) In the case where the provider's arguments are aimed at demonstrating that their model does not present systemic risks despite having high-impact capabilities, the Commission considers that arguments that a model does not present systemic risks because of mitigations already or planned to be implemented are not suitable grounds for a model being excluded from classification as a general-purpose AI model with systemic risk. In those cases, the model still poses systemic risks which must continuously be assessed and mitigated. Nevertheless, the provider could use these mitigations as part of the systemic risk mitigation that they must carry out under Article 55(1) AI Act.
- (41) The Commission's decision to accept or reject the provider's arguments has the legal effects set out below.
 - If the Commission decides to reject the arguments, the model is confirmed to be a general-purpose AI model with systemic risk. The provider is subject to the obligations for providers of general-purpose AI models with systemic risk from the moment when the model meets the condition laid down in Article 51(1), point (a), AI Act. In particular, this means that presenting arguments with the notification does not suspend providers' obligations as providers of general-purpose AI models with systemic risk.
- (42) If the Commission decides to accept the arguments, the general-purpose AI model is no longer classified as a general-purpose AI model with systemic risk. The provider is no longer subject to the obligations for providers of general-purpose AI models with systemic risk from the moment when it is informed of the acceptance decision. This acceptance decision will require the provider to renotify the Commission if there is or it becomes known that there will be a substantial change in any of the facts on which the acceptance decision was based or it becomes known that the acceptance decision was based on substantially incomplete, incorrect or misleading information. In either case, the provider must provide the Commission with all relevant information, and the Commission may turn its acceptance decision into a rejection decision, in which case the model becomes a general-purpose AI model with systemic risk. Moreover, an acceptance decision does not preclude the model later being classified as a general-purpose AI model with systemic risk via designation by the Commission based on the criteria set out in Annex XIII AI Act (see Section 2.3.4).

2.3.4. Designation

- (43) In addition to the above mechanism whereby a general-purpose AI model is classified as a general-purpose AI model with systemic risk through having high-impact capabilities, the AI Act also lays down a designation mechanism.
- (44) Specifically, under Article 52(1), Article 51(1), point (b), and Article 52(4) AI Act, the Commission may designate a general-purpose AI model as a general-purpose AI

model with systemic risk on its own initiative (*ex officio*) or following a qualified alert from the scientific panel, as provided for in Article 90(1), point (a), AI Act.

- (45) The designation can occur:
- under Article 52(4) AI Act, if the Commission concludes that a model has capabilities or an impact equivalent to high-impact capabilities based on the criteria set out in Annex XIII AI Act, or
 - under Article 52(1) AI Act, if the provider of a general-purpose AI model meeting the condition referred to in Article 51(1), point (a), AI Act failed to notify the Commission, in breach of its obligation under Article 52(1) AI Act. In this case, the provider may be fined under Article 101 AI Act.
- (46) In both cases, the provider must comply with the obligations for providers of general-purpose AI models with systemic risk from the moment the model is classified as a general-purpose AI model with systemic risk. In the first case, this is the moment when it is informed of the designation decision. In the second case, this is the moment when the model meets the condition laid down in Article 51(1), point (a), AI Act.

2.3.5. Procedure for contesting designation

- (47) Under Article 52(5) AI Act, providers of general-purpose AI models that have been designated by the Commission as general-purpose AI models with systemic risk under Article 52(4) AI Act may, at the earliest six months after designation, submit a reasoned request for the Commission to reassess the designation. In doing so they should provide *‘objective, detailed and new reasons that have arisen since the designation decision’* for why their model no longer presents systemic risks (Article 52(5) AI Act). Where the designation is maintained by the Commission following its reassessment, providers may request a further reassessment at the earliest six months after that decision.

3. Providers placing on the market general-purpose AI models

- (48) This section of the guidelines focuses on the concepts of the ‘provider’ and ‘placing on the market’ of a general-purpose AI model. It aims to offer clarity on when an actor along the AI value chain must comply with the obligations for providers of general-purpose AI models under the AI Act. First, it clarifies through examples which actor the Commission considers to be the ‘provider’ of a general-purpose AI model and when the Commission considers a general-purpose AI model to be ‘placed on the market’ (Section 3.1). It then addresses the case of a general-purpose AI model being integrated into an AI system (Section 3.2). Finally, it clarifies when the Commission considers an actor modifying a general-purpose AI model to be the provider of a general-purpose AI model, and thus to have to comply with the relevant obligations (Section 3.3).

3.1. When is an actor a provider placing on the market a general-purpose AI model?

- (49) Article 3(3) AI Act defines a ‘provider’ of a general-purpose AI model as *‘a natural or legal person, public authority, agency or other body that develops ... a general-purpose AI model or that has ... a general-purpose AI model developed and places it on the market ... under its own name or trademark, whether for payment or free of charge’*. In turn, Article 3(9) AI Act defines a ‘placing on the market’ of a general-purpose AI model as *‘the first making available of ... a general-purpose AI model on the Union market’*, while Article 3(10) AI Act defines the ‘making available on the market’ of a general-purpose AI model as *‘the supply of ... a general-purpose AI model for distribution or use on the Union market in the course of a commercial activity, whether in return for payment or free of charge’*. Recital 97 AI Act further clarifies that *‘[g]eneral-purpose AI models may be placed on the market in various ways, including through libraries, application programming interfaces (APIs), as direct download, or as physical copy.’*
- (50) In line with Article 2(1), point (a), AI Act, an actor that places a general-purpose AI model on the Union market may become the provider of that general-purpose AI model and be subject to the obligations for providers of general-purpose AI models under the AI Act *‘irrespective of whether they are established or located within the Union or in a third country’*. To facilitate compliance, providers established or located within a third country must appoint an authorised representative established in the Union before placing a model on the Union market (Article 54 AI Act).

3.1.1. Examples of providers of general-purpose AI models

- (51) The following list gives insights into which actor should be considered to be the provider of the general-purpose AI model in various situations:
- If actor A develops a general-purpose AI model and places it on the market, then actor A is the provider.
 - If actor A has a general-purpose AI model developed on its behalf by actor B and actor A places that model on the market, then actor A is the provider.
 - If actor A develops a general-purpose AI model and uploads it to an online repository hosted by actor C, then actor A is the provider.
 - If a collaborative or consortium has a general-purpose AI model developed for it by different individuals or organisations and places the model on the market, then usually the coordinator of the collaborative or the consortium is the provider. Alternatively, the collaborative or the consortium might be the provider. This must be assessed on a case-by-case basis.
- (52) The above examples should be interpreted in line with the relevant provisions of the AI Act (especially Article 3(3) AI Act) and require case-by-case assessment.

3.1.2. Examples of placing on the market of general-purpose AI models

- (53) The following list of examples gives insights into when a general-purpose AI model should be considered to be placed on the market, building on the examples given in recital 97 AI Act:
- a general-purpose AI model is made available for the first time on the Union market via a software library or package;
 - a general-purpose AI model is made available for the first time on the Union market via an application programming interface (API);
 - a general-purpose AI model is uploaded for the first time to a public catalogue, hub, or repository for direct download on the Union market;
 - a general-purpose AI model is made available for the first time on the Union market as a physical copy;
 - a general-purpose AI model is made available for the first time on the Union market via a cloud computing service;
 - a general-purpose AI model is made available for the first time on the Union market by being copied onto a customer's own infrastructure;
 - a general-purpose AI model is integrated into a chatbot made available for the first time on the Union market via a web interface;
 - a general-purpose AI model is integrated into a mobile application made available for the first time on the Union market via app stores;
 - a general-purpose AI model is used for internal processes that are essential for providing a product or service to third parties or that affect the rights of natural persons in the Union.
- (54) The above examples should be interpreted in line with the general guidance for the definition of 'placing on the market' provided in the Blue Guide⁹ and the relevant provisions in the AI Act (Articles 3(9) and (10), and recital 97 AI Act) and require case-by-case assessment.

3.1.3. Considerations for general-purpose AI models integrated into AI systems

- (55) According to recital 97 AI Act, *'AI models are typically integrated into and form part of AI systems'*. Furthermore, *'[the] rules for general-purpose AI models and for general-purpose AI models that pose systemic risks, ... should apply also when these models are integrated or form part of an AI system.'* The following paragraphs describe special cases in which a general-purpose AI model that is integrated into an AI system should be considered to have been placed on the market, in addition to the examples in Section 3.1.2.

⁹ Commission notice, The 'Blue Guide' on the implementation of EU product rules 2022 (Text with EEA relevance) 2022/C 247/01.

- (56) First, as specified in recital 97 AI Act, *‘[w]hen the provider of a general-purpose AI model integrates an own model into its own AI system that is made available on the market or put into service, that model should be considered to be placed on the market and, therefore, the obligations in this Regulation for models should continue to apply in addition to those for AI systems.’*
- (57) Second, when an upstream actor develops or has developed a general-purpose AI model and makes the model available for the first time to a downstream actor on the Union market, the model should be considered to have been placed on the market and the upstream actor to be the provider of the model. The upstream actor therefore has to comply with the obligations for providers of general-purpose AI models. The downstream actor integrating the model into an AI system and placing the system on the Union market or putting it into service in the Union may be the provider of the system. In this case the downstream actor has to comply with the applicable requirements and obligations for AI systems laid down in the AI Act.
- (58) Third, when an upstream actor develops or has developed a general-purpose AI model and makes the model available for the first time to a downstream actor outside the Union market, and the downstream actor integrates the model into an AI system which it places on the Union market or puts into service in the Union, the following considerations apply.
- To ensure that all general-purpose AI models used in the Union comply with the requirements of the AI Act, the model should be considered to be placed on the Union market once the system into which it is integrated is placed on the Union market or put into service in the Union.
- (59) The upstream actor should then be considered the provider of the model, unless the upstream actor has excluded, in a clear and unequivocal way, the distribution and use of the model on the Union market,¹⁰ including its integration into AI systems that are intended to be placed on the Union market or put into service in the Union. If the upstream actor has done this, the downstream actor that integrates the model into a system and places the system on the Union market or puts it into service in the Union should be considered the provider of the model.

3.2. Downstream modifiers as providers of general-purpose AI models

- (60) According to recital 97 AI Act, *‘[general-purpose AI] models may be modified or fine-tuned into new models’*.¹¹ In particular, downstream actors (distinct from the original provider and not acting on its behalf) may modify a general-purpose AI model (with or without integrating it into an AI system).¹² Nevertheless, the AI Act

¹⁰ See definition of ‘making available on the market’ in Article 3(10) AI Act.

¹¹ The Commission considers ‘fine-tuning’ to be one way of ‘modifying’ a general-purpose AI model.

¹² Whether it is the original provider or a downstream actor who modifies the general-purpose AI model must be assessed on a case-by-cases basis. An important factor in this assessment may be who has the control over the model’s weights, for example, in case of fine-tuning via API.

does not specify the conditions under which downstream modifiers should be considered the providers of the modified general-purpose AI models.

- (61) The Commission deems that it is not necessary for every modification of a general-purpose AI model to lead to the downstream modifier being considered the provider of the modified general-purpose AI model. This is in line with the Blue Guide which states that *‘a product which has been subject to important changes or overhauls aiming to modify its original performance, purpose or type may be considered as a new product’*.
- (62) Instead, the Commission considers a downstream modifier to become the provider of the modified general-purpose AI model only if the modification leads to a significant change in the model’s generality, capabilities, or systemic risk.
- (63) Based on the considerations above, an indicative criterion for when a downstream modifier is considered to be the provider of a general-purpose AI model is that the training compute used for the modification is greater than a third of the training compute of the original model (see paragraph 115 for how ‘training compute’ should be understood in these guidelines).
- (64) However, if the downstream modifier cannot be expected to know this value (for example because it has not been communicated by the provider of the original model) and cannot estimate it (see the Annexes A.1 and A.2 to these guidelines for how to estimate training compute) then the threshold should be replaced as follows. If the original model is a general-purpose AI model with systemic risk, the threshold should be replaced with a third of the threshold for a model being presumed to have high-impact capabilities (i.e. currently 10^{25} FLOP, see Article 51(2) AI Act). Otherwise, it should be replaced with a third of the threshold for a model being presumed to be a general-purpose AI model (i.e. currently 10^{23} FLOP, see Section 2.1).
- (65) This threshold is based on the expectation that a model that is modified with this amount of compute will display a significant change which warrants the downstream modifier being subject to the obligations for providers or general-purpose AI models and potentially of general-purpose AI models with systemic risk:
 - The change in model properties may be relevant for the Commission, national competent authorities, or downstream providers, and justifies the downstream modifier being subject to the transparency obligations in Article 53(1), points (a) and (b), AI Act.
 - A modification that meets this threshold can also be expected to have used a significant amount of data, which may be relevant for the obligations laid down in Article 53(1), points (c) and (d), AI Act, namely, the copyright policy and the public summary of the content used for training.
 - In cases where the original model is a general-purpose AI model with systemic risk, the modified model can be expected to present significantly different systemic risk compared to the original model. In particular, the Commission deems that the original provider cannot reasonably foresee the change in systemic risk posed by such modifications in its systemic risk

assessment and mitigation. This justifies the downstream modifier being subject to the obligations in Article 55 AI Act.

- (66) The threshold is relative to the amount of compute used to train the original model. This is because, as explained in paragraph 16, the amount of compute used to train a model is typically proportional to the number obtained by multiplying the number of its parameters with the number of its training examples. Modifications typically involve training the model on some amount of additional data. This means that more compute is required to perform a modification with a fixed amount of data on large models compared to small models. In other words, the amount of compute needed to perform a given modification on a model is typically proportional to the number of parameters of the model. This justifies setting a relative threshold for determining when a modification should lead to the downstream modifier becoming a provider. This approach aims to ensure equal treatment of models regardless of the number of their parameters. In particular, it should not disincentivise modifying models with fewer parameters, since although the threshold for such models is lower, modifying such models also uses less compute.
- (67) While currently few modifications may meet the criterion set out in paragraph 60, the number of downstream modifiers that become providers of general-purpose AI models may increase over time as the compute used to modify models increases. The criterion is thus primarily forward-looking, and in line with the risk-based approach of the AI Act. Therefore, the Commission's approach may change in the future as technology and the market evolve.

3.2.1. Downstream modifiers becoming providers of general-purpose AI models

- (68) According to recital 109 AI Act, *'in the case of a modification or fine-tuning of a model, the obligations for providers of general-purpose AI models should be limited to that modification or fine-tuning, for example by complementing the already existing technical documentation with information on the modifications, including new training data sources, as a means to comply with the value chain obligations provided in this Regulation.'* That is, the documentation required by Article 53(1), points (a) and (b), AI Act is limited to information on the modification, while the copyright policy required by Article 53(1), point (c), AI Act and the summary of the content used for training required by Article 53(1), point (d), AI Act are limited to the data used as part of the modification.
- (69) A downstream modifier who becomes the provider of a general-purpose AI model must also comply with Article 54 AI Act.

3.2.2. Downstream modifiers becoming providers of general-purpose AI models with systemic risk

- (70) If a downstream actor modifies a general-purpose AI model that has been classified as a general-purpose AI model with systemic risk, in such a way that they become

the provider of the modified general-purpose AI model, then the resulting model is presumed to have high-impact capabilities as well. It is therefore considered to be a general-purpose AI model with systemic risk based on Article 51(1), point (a), AI Act.

- (71) In this case, the downstream modifier must comply with the obligations for providers of general-purpose AI models with systemic risk. In particular, the provider must notify the Commission in line with Article 52(1) AI Act, providing the information specified in paragraph 31.

4. Exemptions from certain obligations for certain models released as open-source

- (72) In principle, all providers of general-purpose AI models must comply with the obligations set out in Articles 53 and 54 AI Act. Nevertheless, Articles 53(2) and 54(6) AI Act lay down exemptions from some of these obligations for *‘providers of AI models that are released under a free and open-source licence that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available’* as long as the model is not a general-purpose AI model with systemic risk. This section of the guidelines first describes the obligations to which these exemptions apply (Section 4.1), and then clarifies the conditions that must be fulfilled for these exemptions to apply (Section 4.2).

4.1. Scope of the exemptions

- (73) For providers meeting the conditions referred to in Section 4.2, the following obligations do not apply, unless their model is a general-purpose AI model with systemic risk:
- Article 53(1), point (a), AI Act: the obligation to *‘draw up and keep up-to-date the technical documentation of the model, including its training and testing process and the results of its evaluation, which shall contain, at a minimum, the information set out in Annex XI for the purpose of providing it, upon request, to the AI Office and the national competent authorities’*;
 - Article 53(1), point (b), AI Act: the obligation to *‘draw up, keep up-to-date and make available information and documentation to providers of AI systems who intend to integrate the general-purpose AI model into their AI systems’*;
 - Article 54 AI Act, which concerns the obligation to appoint an authorised representative for providers of general-purpose AI models established in third countries.
- (74) The rationale behind these exemptions is twofold. The first is that models released under a free and open-source licence *‘can contribute to research and innovation in the market and can provide significant growth opportunities for the Union economy’*

(recital 102 AI Act). The second is that *‘[g]eneral-purpose AI models released under a free and open-source licence should be considered to ensure high levels of transparency and openness if their parameters, including the weights, the information on the model architecture, and the information on model usage are made publicly available’* (recital 102 AI Act), which justifies exemption specifically from the transparency related obligations. Nevertheless, if the model is a general-purpose AI model with systemic risk, the fact that the model is released under a free and open-source licence is not sufficient to exempt its provider from complying with these obligations under the AI Act.

- (75) The release of general-purpose AI models under a free and open-source licence does not necessarily reveal substantial information on the data used for training or modifying a model, nor on how compliance with copyright law was ensured. For this reason, providers of general-purpose AI models that meet the conditions referred to in Section 4.2 are not exempted from the obligation laid down in Article 53(1), point (c), AI Act to put in place a policy to comply with Union copyright law. This obligation includes identifying and complying with a reservation of rights expressed under Article 4(3) of Directive (EU) 2019/790 of the European Parliament and of the Council. For the same reason, they are not exempted from the obligation laid down in Article 53(1), point (d), AI Act to produce a summary about the content used for training.

4.2. Conditions for the exemptions to apply

4.2.1. Conditions on the licence

- (76) To qualify for the exemptions referred to in Section 4.1, a general-purpose AI model must in particular be released under a *‘free and open-source licence that allows for the access, use, modification, and distribution of the model’* (Articles 53(2) and 54(6) AI Act).
- (77) The term ‘licence’ in this context should be understood as the granting of permissions related to the model under specified terms and conditions. ‘Free and open-source’ should be understood as a form of licensing that makes use of copyright to allow wide dissemination of the model and incentivise further developments.
- (78) Recital 102 AI Act clarifies that the licence must allow the model to be ‘openly shared’ and users to be able to ‘freely access, use, modify and redistribute’ the model ‘or modified versions thereof’. If one of these rights (i.e. rights to access, use, modify, and redistribute) is missing, the licence cannot be considered free and open-source under the AI Act. In this case the provider is not exempt from the obligations referred to in Section 4.1.
- (79) The Commission understands access as the right of anyone interested to freely obtain the model without any payment requirements or other restrictions. However, reasonable safety and security measures, such as user verification processes, may

be implemented provided that they do not unfairly discriminate against persons, for example on the basis of their country of origin.

- (80) The Commission understands usage, in this context, to mean that the licence guarantees that the original provider will not use their intellectual property rights to restrict the use of the model, or charge for its use. The licence should allow the model to be used without any restrictions besides restrictions ensuring attribution and distribution of the model or derivatives under the same or comparable terms. These terms could include an obligation triggered when the model is distributed, requiring it to be distributed under the same (or compatible) copyright terms.
- (81) The Commission understands modification as the right of anyone to freely make alterations to the model without any payment requirements or other restrictions.
- (82) The Commission understands distribution as the right of those who access, use, or modify the general-purpose AI model to freely distribute it onwards, subject to limited conditions. Such limited conditions often consist only in crediting the author(s) and retaining their copyright notice, i.e. attribution. Therefore, the recipient of the licenced model should be able to modify it and redistribute the resulting derivative work under different licence terms, including closed-source proprietary terms, without prejudice to any restrictions acceptable under paragraph 80.
- (83) Examples of restrictions that would disqualify a licence from meeting these criteria include:
 - limitations to non-commercial or research-only use (e.g. ‘you may use the model solely for non-commercial, non-revenue generating, research purposes’);
 - prohibitions on distributing the model or its components (e.g. ‘you may use the model, but you may not distribute the model’);
 - usage restrictions triggered by user scale thresholds (e.g. requiring additional licensing if monthly active users exceed a certain number);
 - requirements to obtain separate commercial licences for specific use cases.
- (84) Aside from the rights described above, open-source licences consist of various terms and conditions. Users can use, modify, and distribute general-purpose AI models under free and open-source licences, in compliance with the terms and conditions of the licence. These can include crediting the original creators, respecting the terms of distribution, and making any modifications or improvements available under the same or comparable licence terms. While the licence must generally permit use for any purpose, licensors may include specific, safety-oriented terms that reasonably restrict usage in applications or domains where such use would pose a significant risk to public safety, security, or fundamental rights. Such restrictions must be proportionate to the risks being mitigated, and based on objective, non-discriminatory criteria that do not unfairly target specific categories of users or other actors.

4.2.2. Lack of monetisation

- (85) The concept of monetisation also plays a crucial role in determining whether the exemptions apply to general-purpose AI models, as clarified by recital 103 AI Act. According to this recital, in order for the exemptions to apply, no monetary compensation should be required in exchange for access, use, modification, and distribution of the AI model. In this context, monetisation should be understood as encompassing not only the provision of the model against a price but also other types of monetisation strategies.
- (86) The following are scenarios that the Commission considers to be forms of monetisation:
- the model is provided under a dual licensing model or a comparable approach that allows free academic use but requires payment for commercial usage or use over a certain scale.
 - technical support and other services that are indistinguishably linked to the model itself or its security, and without which the model would not work or be accessible, provided against payment.
 - users must purchase support, training, and maintenance services to access the model.
 - the model is exclusively hosted by the provider on its individual hosting platform or website that requires users to pay for access, including when users can freely access the platform or website but are served with paid advertisements.
- (87) While fully recognising that the protection of personal data is a fundamental right and that therefore personal data cannot be considered a commodity or monetised, access, use, modification, or redistribution of the model requiring the collection or otherwise processing of personal data should be treated in same manner as monetisation strategies, unless that processing is exclusively and strictly limited to the security of the model and without any commercial or financial gain. In any event, the processing of personal data should comply with the rules under Union data protection law.
- (88) By contrast, the following are scenarios that the Commission would not consider, for the purposes of applying the exemptions, to be forms of monetisation.
- The model is provided together with paid services that do not affect the usability or free usage of the model and that are purely optional. This could cover, for example, a business model that offers commercial services unrelated to the open-source licence.
 - Paid services or support are made available alongside the model, without any purchase obligation, as long as the model's usage and free and open access are guaranteed. Such services or support could include premium versions of the model with advanced features or additional tools, update systems, extensions, or plug-ins that help users work with the open-source model or extend its functionality.

- (89) The monetisation aspect is considered absent for transactions between microenterprises.

4.2.3. Public availability of parameters, including the weights, the information on the model architecture, and the information on model usage

- (90) For the exemptions referred to in Section 4.1 to apply, the model's '*parameters, including the weights, the information on the model architecture, and the information on model usage*' must be '*made publicly available*' (Articles 53(2) and 54(6) AI Act).
- (91) This information should be made publicly available in a format, and with a degree of clarity and specificity, that enables access, usage, modification, and distribution of the model.
- (92) Information about the model's usage should as a minimum include information about the model's input and output modalities, capabilities, and limitations. Information about the model's usage should also include the technical means (e.g. instructions for use, infrastructure, tools) required for the model to be integrated into AI systems, which may include the appropriate configuration for the intended use cases, where applicable. This information ensures that downstream providers of AI systems, developers, and users can use the model for practical applications.

5. Enforcement of the obligations for providers of general-purpose AI models

- (93) The obligations for providers of general-purpose AI models under the AI Act enter into application on 2 August 2025 (Article 113, point (b), AI Act) and the Commission, acting through the AI Office, is entrusted with supervising and enforcing these obligations (Article 88(1) AI Act). This section of the guidelines first clarifies the effects of signing and implementing a code of practice assessed as adequate (Section 5.1), then provides some insights into the AI Office's mandate for supervising, investigating, enforcing, and monitoring Chapter V of the AI Act (Section 5.2), and finally clarifies the Commission's expectations for compliance after entry into application of Chapter V of the AI Act on 2 August 2025 (Section 5.3).

5.1. Effects of adhering to a code of practice assessed as adequate

- (94) Providers of general-purpose AI models can demonstrate compliance with the obligations in Articles 53(1) and 55(1) AI Act by adhering to a code of practice that is assessed as adequate by the AI Office and the Board (Article 56(6) AI Act). Any opt-out from chapters of the code of practice results in losing the benefits of facilitating the demonstration of compliance in that respect. While providers can also demonstrate compliance with these obligations through alternative adequate means, adherence to a code of practice that is assessed as adequate is a

straightforward way of demonstrating compliance (Articles 53(4) and 55(2) AI Act). For providers of general-purpose AI models that adhere to a code of practice that is assessed as adequate, the Commission will focus its enforcement activities on monitoring their adherence to the code of practice. Providers of general-purpose AI models that sign a code of practice that is assessed as adequate will be transparent about the measures they implement to comply with the AI Act. They will therefore benefit from increased trust from the Commission and other stakeholders.

- (95) Providers of general-purpose AI models that do not adhere to a code of practice that is assessed as adequate are expected to demonstrate how they comply with their obligations under Chapter V of the AI Act via other adequate means and will have to report the measures they have implemented to the AI Office. Furthermore, such providers are expected to explain how the measures they implement ensure compliance with their obligations under the AI Act, for instance by carrying out a gap analysis that compares the measures they have implemented with the measures set out by a code of practice that is assessed as adequate. They may also be subject to a larger number of requests for information and requests for access to conduct model evaluations throughout the entire model lifecycle because the AI Office will have less of an understanding of how they are ensuring compliance with their obligations under the AI Act and will typically need more detailed information, including about modifications made to general-purpose AI models throughout their entire lifecycle.
- (96) The Commission may take into account commitments implemented in line with a code of practice that is assessed as adequate as a mitigating factor when fixing the amount of fines, depending on the specific circumstances (Article 101(1) AI Act).
- (97) Providers of general-purpose AI models without systemic risk may sign and implement commitments relevant for providers of general-purpose AI models with systemic risk (Article 56(7) AI Act).
- (98) Commitments implemented in line with a code of practice that is assessed as adequate will only become relevant for assessing compliance with the AI Act when the obligations for providers of general-purpose models laid down in Chapter V of the AI Act enter into application on 2 August 2025 (Article 113, point (b), AI Act).
- (99) The Commission may also approve a code of practice via implementing act, thereby giving it general validity within the Union (Article 56(6) AI Act). If a code of practice cannot be finalised by 2 August 2025, the Commission may adopt common rules via implementing act, which would be applicable to all providers of general-purpose AI models and general-purpose AI models with systemic risk (Article 56(9) AI Act).
- (100) A code of practice is a temporary tool for demonstrating compliance with the AI Act (see Articles 53(4) and 55(2) AI Act) until harmonised standards are developed and approved (Article 40 AI Act). As opposed to adherence to a code of practice, compliance with harmonised standards grants a presumption of conformity with the corresponding obligations under the AI Act (see Articles 53(4) and 55(2) AI Act). If, after a standardisation request, harmonised standards are not developed in time

or in a satisfactory manner, the Commission may adopt common specifications via an implementing act (Article 41 AI Act).

5.2. Supervision, investigation, enforcement, and monitoring of Chapter V of the AI Act

- (101) The AI Office is tasked with supervising, investigating, enforcing, and monitoring the obligations for providers of general-purpose AI models laid down in Chapter V of the AI Act (Articles 88 and 89 AI Act) and monitoring the compliance of AI systems based on general-purpose AI models with the obligations under the AI Act when the providers of the model and of the system are the same (Article 75(1) AI Act). The following paragraphs clarify the supervision, investigation, enforcement, and monitoring activities of the AI Office regarding providers of general-purpose AI models.
- (102) The AI Office will take a collaborative, staged, and proportionate approach. The AI Office encourages close informal cooperation with providers during the training of their general-purpose AI models to facilitate compliance and ensure timely market placement, in particular for providers of general-purpose AI models with systemic risk. Furthermore, the AI Office expects proactive reporting by providers of general-purpose AI models with systemic risk, whether as part of commitments under a code of practice or as part of alternative means of demonstrating compliance. Finally, the AI Office expects providers to collaborate and actively engage with the AI Office at the later stages of potential formal proceedings conducted in line with the AI Office's powers under the AI Act (for example, through structured dialogues).
- (103) Such proactive reporting is without prejudice to the obligation to keep track of, document, and report relevant information about 'serious incidents' under Article 55(1), point (c), AI Act. The AI Office considers that this obligation covers serious cybersecurity breaches related to the model or its physical infrastructure, including the (self-)exfiltration of model parameters and cyberattacks, due to their possible implications for the obligations provided for in Article 55(1), points (b) and (d), AI Act. Apart from this, the AI Office considers a 'serious incident' in the context of Chapter V AI Act as any incident or malfunctioning of a general-purpose AI model that directly or indirectly leads to any of the events listed in the corresponding definition for AI systems in Article 3(49), points (a) to (d), AI Act.
- (104) In assessing compliance, including with respect to measures related to cybersecurity, accountability, and governance processes, the AI Office may take into account, as applicable, whether the provider has implemented relevant technical standards, in particular, if available, the harmonised standards published in the *Official Journal of the European Union*.
- (105) Under Article 89(1) AI Act, the AI Office is tasked with monitoring whether providers of general-purpose AI models effectively comply with their obligations under the AI Act. Such monitoring includes monitoring adherence to codes of practice. Under Article 89(2) AI Act, downstream providers have the right to lodge a complaint alleging an infringement of the AI Act by providers of general-purpose

AI models. Such a complaint must be duly reasoned and indicate at least the information listed in Article 89(2), points (a) to (c), AI Act.

- (106) The Commission's enforcement of the obligations for providers of general-purpose AI models is underpinned by the powers given to it under the AI Act. These are the powers to: (i) request information (Article 91 AI Act); (ii) conduct evaluations of general-purpose AI models (Article 92 AI Act); (iii) request measures from providers, including implementing mitigation measures and recalling the model from the market (Article 93 AI Act); and (iv) impose fines of up to 3% of global annual turnover or EUR 15 million, whichever is higher (Article 101 AI Act); starting on 2 August 2026. More detailed acts will follow to further specify how these powers will be implemented, in particular the implementing acts under Articles 92(6) and 101(6) AI Act.
- (107) The Commission can take decisions under Article 52 AI Act (such as the acceptance or rejection of arguments as well as the designation of general-purpose AI models as general-purpose AI models with systemic risk) from 2 August 2025 (Article 113, point (b), AI Act).
- (108) The Commission will ensure the confidentiality of the data obtained in the context of supervision, investigation, enforcement, and monitoring activities to protect, in particular, intellectual property rights, confidential business information or trade secrets of natural or legal persons, and public safety and security, in line with Article 78 AI Act.

5.3. Compliance after entry into application of Chapter V of the AI Act on 2 August 2025

- (109) The obligations for providers of general-purpose AI models laid down in Chapter V of the AI Act apply from 2 August 2025 (Article 113, point (b), AI Act). However, under Article 111(3) AI Act, *'[p]roviders of general-purpose AI models that have been placed on the market before 2 August 2025 shall take the necessary steps in order to comply with the obligations laid down in this Regulation by 2 August 2027.'* This covers such models throughout their entire lifecycle (see Section 2.2).
- (110) The Commission recognises that in the months following the entry into application of the obligations for providers of general-purpose AI models on 2 August 2025, providers of general-purpose AI models placed on the market *before* 2 August 2025 may face various challenges to comply with their obligations under the AI Act by 2 August 2027. The AI Office is therefore dedicated to supporting providers in taking the necessary steps to comply with their obligations by 2 August 2027.
- (111) In particular, providers of general-purpose AI models placed on the market before 2 August 2025 are not required to conduct retraining or unlearning of models, where it is not possible to do this for actions performed in the past, where some of the information about the training data is not available, or where its retrieval would cause the provider disproportionate burden. Such instances must be clearly disclosed and justified in the copyright policy and in the summary of the content used for training.

- (112) The Commission also recognises that in the months following the entry into application of the obligations for providers of general-purpose AI models on 2 August 2025, providers of general-purpose AI models placed on the market *after* 2 August 2025 may need time to adapt their policies and procedures to ensure they are compliant with their obligations under the AI Act. The AI Office is therefore dedicated to supporting providers in taking the necessary steps to comply with their obligations and invites providers to immediately and proactively contact the AI Office to ensure they are taking the right steps to ensure compliance.
- (113) In particular, providers who, on 2 August 2025, have trained, are in the process of training, or are planning to train a general-purpose AI model with a view to placing it on the market after 2 August 2025, and who anticipate difficulties in complying with the obligations for providers of general-purpose AI models, especially those with systemic risk, should proactively inform the AI Office regarding how and when they will take the necessary steps to comply with their obligations. In the specific case where a provider has not placed on the market a general-purpose AI model with systemic risk before 2 August 2025, the Commission will give particular consideration to their challenging situation, in particular to allow a timely placing on the market.
- (114) Furthermore, the Commission plans to take into account during enforcement that the ecosystem of external model evaluators is still developing, and that initially the methodologies used by such evaluators may be inconsistent. The AI Office will monitor the situation and, if necessary, work to set out a common approach for external model evaluators, for instance, by involving the network of evaluators and other subject-matter experts.
- (115) Finally, the AI Act acknowledges that immediately after 2 August 2025, providers may not be able to fully implement certain measures laid down in a code of practice that is assessed as adequate or alternative adequate measures. In the first year from 2 August 2025 onwards, the Commission cannot take any enforcement actions because its enforcement powers only enter into application on 2 August 2026. This lack of enforcement does not put into question the applicability of the obligations laid down in Chapter V of the AI Act from 2 August 2025 onwards (Article 113, point (b), AI Act). In particular, providers who, on 2 August 2025, have trained, are in the process of training, or are planning on training a general-purpose AI model with systemic risk with a view to placing the model on the market after 2 August 2025, are expected to notify the Commission without delay and in any event within two weeks after 2 August 2025 under Article 52(1) AI Act and to cooperate with the AI Office in view of ensuring full compliance. From 2 August 2026 onwards, the Commission will enforce with fines compliance with all obligations by providers of general-purpose AI models who are not fully compliant on that date, taking into account all relevant aspects mentioned in Article 101 AI Act.

6. Review and update of the Commission guidelines

- (116) The Commission will review these guidelines as soon as it becomes necessary in the light of practical experience gained by providers of general-purpose AI models in implementing their obligations, and technological, societal, and market developments in this area. This also includes any relevant experience from enforcement actions and interpretations given by the CJEU on the obligations for providers of general-purpose AI models and other provisions of the AI Act examined in these guidelines. During its review, the Commission may decide to withdraw or amend these guidelines. The Commission encourages providers of general-purpose AI models, national market surveillance authorities (through the Board), the advisory forum, the scientific panel, the research community, and civil society organisations to contribute to this process by responding to future calls for public consultation, workshops, or other opportunities for input.

Annex: Training compute of general-purpose AI models

- (117) Article 51(2) AI Act lays down a threshold for the compute used to train a general-purpose AI model that can lead to the model being classified as a general-purpose AI model with systemic risk. Sections 2.1 and 3.2 also set thresholds involving training compute. To know whether a model meets any of these thresholds, potential providers must estimate the relevant amount of compute used for training. This Annex to the guidelines provides guidance on how they may do so. Annex A.1 specifies what potential providers of general-purpose AI models and of general-purpose AI models with systemic risk should include when estimating training compute, while Annex A.2 describes two widely used approaches that potential providers may use. Annex A.3 estimates the training compute for a number of models trained with approximately one billion parameters to illustrate the formulae in Annex A.2 and to justify the threshold that forms part of the indicative criterion for when a model is a general-purpose AI model (Section 2.1).

A.1. What should be estimated?

- (118) In these guidelines, ‘training compute’ of a general-purpose AI model refers to either:
- the total amount of compute directly contributing to parameter updates in the model if the model is not a general-purpose AI model with systemic risk;
 - the cumulative amount of compute used to train the model if, and for the purpose of assessing whether, the model is a general-purpose AI model with systemic risk.
- (119) Under the AI Act, regarding how ‘cumulative training compute’ should be understood, recital 111 AI Act specifies that *‘the cumulative amount of computation used for training includes the computation used across the activities and methods that are intended to enhance the capabilities of the model prior to deployment, such as pre-training, synthetic data generation and fine-tuning’*. Given the wide variety of such ‘activities and methods’, and their rapid evolution, it is not feasible to provide a precise and exhaustive list of what should and should not be included. Instead, the Commission considers that providers should, as a general rule, account for all compute that contributed or will contribute to the model’s capabilities. In particular, any compute directly contributing to parameter updates in the model should be included.
- (120) If the model is trained on synthetic data that is not publicly accessible, the forward passes used to generate the data, including discarded data, should be included in the estimation of the cumulative training compute. For example, if 100 samples were generated and only the top 10 samples were selected for training, the compute used to generate all 100 samples should be counted since the compute used to generate all 100 samples was necessary to create the selected 10.
- (121) Where a model has been created by combining model weights, for example via model weight merging or model weight averaging, or through integrating pre-

existing model weights, for example through weight initialisation, the training compute used to train the combined model weights should be included in the estimation of the cumulative training compute of the model.

(122) By contrast, the following are examples of compute which need not be included in the estimation of the cumulative training compute. This list may change as technology evolves:

- compute used to generate synthetic data which is publicly accessible, since it may be indistinguishable from other publicly accessible data;
- compute spent on purely diagnostic processes that do not contribute to enhancing model capabilities, such as model evaluations or red-teaming;
- compute whose expenditure contributes to enhancing model capabilities only through lessons learnt by humans, such as compute spent on exploratory research projects that result in more efficient training techniques, or on experiments in synthetic data generation that fail and whose results are discarded;
- compute used to train parent model(s) used in distillation (whether or not the parent model coincides with the concerned model);
- compute used to train auxiliary model(s) such as value functions or reward models (whether or not the auxiliary model coincides with the concerned model);
- recomputation of activations to save memory.

A.2. How should training compute be estimated?

- (123) Providers may choose any method to estimate the relevant amount of training compute, so long as the estimated amount is, in the providers' best judgement, accurate within an overall error margin of 30% of the reported estimate. This is to account for the difficulties providers may face in obtaining precise estimates of their model's cumulative training compute, for example, when estimates are made before training of the model is complete. Providers are nevertheless expected to document the assumptions made in making their estimations, including the method of estimation, and the associated uncertainties.
- (124) Providers may choose to estimate the relevant amount of training compute by tracking graphics processing unit ('GPU') usage (hardware-based approach) or by estimating operations directly based on the relevant model's architecture (architecture-based approach), as appropriate to what is being estimated.
- (125) Regardless of the method of estimation, all operations should be counted equally, independently of floating-point precision.

A.2.1. Hardware-based approach

- (126) To estimate training compute using the hardware-based approach, the potential provider of a general-purpose AI model or general-purpose AI model with systemic risk should first estimate:
- the number N of GPUs or other hardware units used or expected to be used;
 - the total duration of use L , achieved or expected (measured in seconds);
 - the peak theoretical performance H of the GPUs or other hardware units used or expected to be used (measured in FLOP/second), calculated via weighted average if different types of GPUs or hardware units are used with different peak theoretical performances (where the weights are determined by the number of hardware units of each type);
 - the average percentage of GPU utilisation U achieved or expected to be achieved for the duration of use.
- (127) If all hardware units N are used for the entirety of the duration L , the corresponding amount of training compute C (measured in FLOP) can then be calculated according to the following formula:
- $$C = N \cdot L \cdot H \cdot U.$$
- (128) If different numbers of hardware units are used for different periods of time, the formula above can be applied to each period of time for which a constant number of hardware units is used (provided that the total time matches the total duration of training). The resulting estimates can then be summed up to yield the total amount of training compute.
- (129) Annex A.3 includes example calculations using this approach.

A.2.2. Architecture-based approach

- (130) This approach consists in directly estimating the number of operations performed based on the relevant model's architecture.
- (131) For example, essentially all general-purpose AI models or potentially general-purpose AI models today are based on neural networks that are trained through a succession of forward and backward passes. For such models, the corresponding amount of training compute C (measured in FLOP) is the product of the number of full passes (one full pass is the combination of a forward and backward pass) made during training with the total number of operations performed in a full pass, i.e.

$$C = \text{Number of forward passes made during training} \cdot \text{number of operations/full pass}.$$

- (132) For some dense large¹³ language models based on the transformer architecture, the amount of training compute C may be approximated by the following formula:

$$C \approx 6 \cdot P \cdot D,$$

where P refers to the total number of model parameters and D refers to the total number of training tokens used for training¹⁴.

- (133) Annex A.3 includes example calculations using this approach.

A.3. Training compute of models with approximately one billion parameters

- (134) The threshold that forms part of the indicative criterion for when a model is a general-purpose AI model (Section 2.1) has been informed by the estimated amount of compute used in practice to train widely used models with approximately one billion parameters. Although the models and their providers are not referred to by name, information relevant to this estimate is shared below.

- (135) The estimated compute used to train the investigated models is as follows:
- Model A (a language model with 3.8 billion parameters): $7.5 \cdot 10^{22}$ FLOP.
 - Model B (an image diffusion model with 1 billion parameters): 10^{23} FLOP.
 - Model C (a language model with 600 million parameters): $1.3 \cdot 10^{23}$ FLOP
 - Model D (a language model with 1.5 billion parameters): $1.6 \cdot 10^{23}$ FLOP
 - Model E (a language model with 1.7 billion parameters): $3.7 \cdot 10^{23}$ FLOP
 - Model F (a language model with 2.5 billion parameters): $1.8 \cdot 10^{23}$ FLOP
 - Model G (a language model with 2.6 billion parameters): $1.5 \cdot 10^{23}$ FLOP
 - Model H (a language model with 1 billion parameters): $6.5 \cdot 10^{23}$ FLOP

- (136) The values above have been obtained via the following calculations, which also serve to illustrate the two estimation methods described in Annex A.2.

- (137) **Model A. Architecture-based approach:** Model A is a language model with a transformer decoder architecture and 3.8 billion parameters, trained on 3.3 trillion tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 3.8 \cdot 10^9 \cdot 3.3 \cdot 10^{12} \approx 7.5 \cdot 10^{22} \text{ FLOP.}$$

- (138) **Model B. Hardware-based approach:** Model B is a diffusion model for image generation. According to its model card it was trained for a total of 200,000 GPU hours on 40GB A100 GPUs. This implies a product NL of number of GPUs N and total training duration of L equal to $NL = 200,000 \cdot 3,600$ seconds. Assuming a peak performance of 300 TFLOP/second and GPU utilisation of 50%, the approximation from Annex A.2.1 yields:

$$C = N \cdot L \cdot H \cdot U \approx 10^{23} \text{ FLOP.}$$

¹³ A model based on a dense transformer architecture should be considered large if its number of parameters is significantly larger than one twelfth of the number of tokens in the input context. This ensures that the contribution to the training compute coming from the context-dependent compute is negligible in comparison to the contribution coming from the non-embedding compute and which equals $6PD$ (see for example the reference in footnote 16).

¹⁴ See for example Kaplan, J., et al, *Scaling laws for neural language models*, arXiv preprint 2001.08361, 2020.

- (139) **Model C.** *Architecture-based approach:* Model C is a language model with a transformer decoder architecture and 600 million parameters, trained on 36 trillion tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 0.6 \cdot 10^9 \cdot 36 \cdot 10^{12} \approx 1.3 \cdot 10^{23} \text{ FLOP.}$$

- (140) **Model D.** *Architecture-based approach:* Model D is a language model with a transformer decoder architecture and 1.54 billion parameters, trained on 18 trillion tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 1.54 \cdot 10^9 \cdot 18 \cdot 10^{12} \approx 1.663 \cdot 10^{23} \text{ FLOP.}$$

- (141) **Model E.** *Architecture-based approach:* Model E is a language model with a transformer decoder architecture and 1.7 billion parameters, trained on 36 trillion tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 1.7 \cdot 10^9 \cdot 36 \cdot 10^{12} \approx 3.7 \cdot 10^{23} \text{ FLOP.}$$

- (142) **Model F.** *Architecture-based approach:* Model F is a language model with a transformer decoder architecture and 2.5 billion parameters, trained on 12 trillion tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 2.5 \cdot 10^9 \cdot 12 \cdot 10^{12} \approx 1.8 \cdot 10^{23} \text{ FLOP.}$$

- (143) **Model G.** *Architecture-based approach:* Model G is a language model with a transformer decoder architecture and 2.6 billion parameters, trained on 61 billion tokens for 164 epochs, i.e. a total of about 10 trillion training tokens. The approximation from Annex A.2.2 yields:

$$C \approx 6 \cdot P \cdot D = 6 \cdot 2.6 \cdot 10^9 \cdot 10 \cdot 10^{12} \approx 1.5 \cdot 10^{23} \text{ FLOP.}$$

- (144) **Model H.** *Hardware-based approach:* Model H is a language model. According to its model card, Model H was trained for a total of 370,000 GPU hours on 80GB H100 GPUs. Assuming a peak performance of 989 TFLOP/second and GPU utilisation as previously, the approximation from Annex A.2.1 yields:

$$C = N \cdot L \cdot H \cdot U = 6.5 \cdot 10^{23} \text{ FLOP.}$$