WE ARE SHAPING THE INTERNET.
YESTERDAY.TODAY.BEYOND TOMORROW.

eco

POSITIONPAPER

on the third draft of the General-Purpose AI Code of Practice

Berlin, 28.03.2025

The European economy can benefit significantly from the potential of artificial intelligence and the technology is widely considered to be one of the major technologies of the future. In order to fully exploit this potential, a practical and legally sound regulation is needed that excludes undesirable aspects while at the same time enabling innovation. A Code of Practice is currently being developed to support providers of general-purpose AI models (GPAI models) in fulfilling their obligations under the AI Act. The purpose of the Code is twofold: firstly, to specify the provisions for providers of these models, and secondly, to give specific recommendations on how they can be fulfilled. The Code of Practice has the potential to serve as a valuable instrument in ensuring compliance with the provisions of the AI Act in a practical and streamlined manner, thereby contributing to the effective governance of AI-related activities.

However, for this to be successful, it is essential to consider a number of aspects during the design process. The current third draft incorporates simplifications and enhancements over the second version, while retaining critical elements. From the perspective of the Internet industry, it is important that the Code of Practice does not deviate from its original purpose. Consequently, it is imperative that the Code of Practice does not encompass any requirements that exceed the provisions stipulated within the AI Act. This is problematic for two reasons. Firstly, the AI Act already has a high level of regulatory depth, which potentially imposes high bureaucratic hurdles on providers of AI models. Secondly, the AI Act already represents a democratically agreed balance of the various interests. Consequently, any proposed amendments or supplements to the Code of Practice must be initiated within the framework of a legislative procedure. Furthermore, it is imperative that the guidelines adhere to the principle of simplification, a central tenet declared by the newly constituted Commission as a pivotal objective for the present legislature. This principle should not be contravened.

In detail, the eco Association of the Internet Industry has the following comments on the third draft of the Code of Practice:

1. Scope of application

To be useful to providers of GPAI models, it is important to clearly define the scope of the code. The AI Act stipulates that deployers of models must adhere to the requirements for providers if they utilise a model themselves and implement adjustments through fine-tuning. However, the extent to which fine-tuning of models is considered a modification is not specified. This issue remains unaddressed in the Code of Practice. Due to the adaptation of an existing model a

significant proportion of companies were to come under the provisions of the AI Act, therefore it is possible that the regulations could potentially affect a considerably higher number of businesses than originally anticipated. Consequently, the Internet industry contends that fine-tuning should not generally be regarded as a substantial modification. At the very least, the Commission is required to provide further clarification so that companies can better assess which requirements apply to them.

In this context, the definition of open source models is also unclear. The term 'fully open source model' is included in the section on security without being clearly defined. It is essential to exercise caution and ensure that the terminology employed aligns with the provisions of the AI Act and the other sections of the code. Any potential distinctions should be elucidated with precision. For instance, it remains ambiguous whether a model excluded from certain use cases by contract is considered a 'fully open source model'.

2. Transparency requirements

A number of transparency requirements are applicable to providers of GPAI models in accordance with the AI Act. These requirements are partially specified in this third draft of the Code of Practice. A notable enhancement over previous iterations is the incorporation of a model card, a feature that was previously absent. This iteration of the Code of Practice is notable for its provision of a detailed overview of the information required and the extent to which it must be provided. It also outlines what information providers must make available to the national competent authorities, the AI Office and downstream providers. This template is generally a beneficial instrument, particularly for SMEs, in assessing the scope and type of data to be provided. The limitation of the required words is also beneficial in clarifying that these are generally rather short descriptions, which is to be welcomed in the context of practical implementation.

However, it is necessary to ensure that only data that is truly needed and does not contain sensitive business secrets is made available, especially to downstream providers. This is evidenced by the ongoing challenges associated with specifying the files utilised for training purposes. While a certain degree of transparency is generally acknowledged as necessary, the information required in the model card, particularly with regard to the sub-categories of data used and the websites utilised, is of a highly technical nature and is overly detailed. This creates a risk that sensitive business secrets could be part of the information required. Consequently, it is imperative that only the information deemed indispensable by downstream providers is conveyed to them. It is evident that the current version of the model card imposes obligations that are not comprehensible in this respect. This is particularly evident in the section pertaining to additional information. A thorough examination is therefore required to ascertain the necessity for the provision of information on the system architecture, and whether these obligations can at least be defined with greater precision.

It is also important to ensure that the information required can be clearly determined by the provider of GPAI models and does not exceed the provisions of

the AI Act. This is not always the case, for example, with regard to the required information on the energy consumption of the models. For instance, it proves challenging for providers to accurately measure the energy consumption of inference computing. It is crucial that transparency requirements do not compromise the security of AI models. However, a detailed breakdown could potentially assist in identifying vulnerabilities in the models that could be exploited by malicious actors.

The new draft code introduces further obligations for providers of GPAI models. For instance, providers are now obliged to retain the information provided via model cards for a further 10 years after the shutdown of an AI model. The Internet Industry has expressed concerns that this period is disproportionate to the benefits derived from this information. Consequently, the Internet Industry has advocated for the removal of this obligation, citing its excessive duration.

3. Risks

A fundamental element of the code is the management of potential risks associated with GPAI models. In the opinion of eco, this necessitates a close alignment with the provisions of the AI Act. It is acknowledged that a number of the risks to be addressed are difficult to operationalise or assess. Furthermore, it is imperative to ensure that providers of models are only required to exclude risks that are within their own comprehension. In principle, the third version of the Code of Practice represents a commendable advancement in this regard. From the Internet Industry's perspective, the draft's emphasis on 'high-impact capabilities' and its subsequent narrowing of risks without compromising the overarching objective of ensuring AI's security and reliability are commendable. Conversely, it is encouraging to note that Measure II.3.1 has addressed certain challenges associated with the definition and management of risks that are not easily delineated at the model level.

Nevertheless, the terms employed in this document are, in some cases, excessively broad, which has the potential to impede the capacity of companies to mitigate risk. An illustration of this is evident in the definition of harmful manipulation, which requires refinement to ensure a precise scope of application. Furthermore, the burden of risk minimisation remains disproportionately placed on the providers of AI models. The issue of manipulation is most often addressed at the level of an AI system. Given the inherent inability of providers to foresee every potential use case, it is impractical for them to fully mitigate the risk of manipulation due to the model. The Internet Industry has expressed the opinion that a shift to the level of the AI systems would be more effective in this regard, where this is applicable.

4. Safety requirements

The AI Act imposes a series of obligations on providers of GPAI models with systemic risks, with the aim of ensuring the security of these models. Specifically, the Code of Practice sets out requirements for the testing of models. However, eco believes that some of these requirements are problematic and should be reconsidered in order to enable a practicable design. For instance, the general

framework adequacy assessment, which is to be prepared in accordance with Article 55 of the AI Act, is required to be submitted to the AI Office within five working days. In eco's opinion, this deadline is impractical and should be extended. In order to prevent further bureaucratic burdens for companies, the deadline for the model-specific adequacy assessment should also be extended. It is difficult to understand why model providers should be required to submit such a report before the end of the development process.

The third draft also continues to provide for an obligation for model providers to have their models tested for security by third parties. To this end, external auditors are to be granted extensive insight into the model and its functionality. However, the Internet Industry has expressed concerns that this could potentially compromise trade secrets and the security of the models. Furthermore, the wording in the Code exceeds the provisions set out in the AI Act. Recital 114 merely stipulates the option of security testing by external experts, not an obligation to do so. This discrepancy must be addressed and reflected in the Code. It is important to note that the model providers themselves have the necessary expertise to carry out the evaluation of the models internally, especially as it is also in their own interest to minimise risks. However, it must be acknowledged that the prerequisite of employees in model evaluation teams, as outlined in Measure II.4.11, for instance, through the attainment of a PhD, could pose a significant challenge for SMEs in particular. Consequently, it is recommended that the qualification requirements for employees in model evaluation teams be rendered more practically oriented.

5. Copyright

The code will also address the provisions on copyright. From the perspective of the Internet Industry, it is essential to emphasise that the provisions on copyright should not go beyond what is laid down in the AI Act or the Copyright Directive. Furthermore, it is equally important that it remains feasible and attractive to train and develop AI models in Europe, so that Europe does not lose touch with this important emerging technology. The third draft of the code still requires a number of adjustments. For instance, the providers of AI models are still required to fulfil very granular transparency obligations with regard to the training data used, which in some cases is not technically feasible.

The Internet Industry has expressed the view that rights holders should retain the right to object to the use of their works. This option is also provided for in Article 5 of the Copyright Directive and the AI Act. The implementation of such protocols should be straightforward, with robots.txt protocols having demonstrated their efficacy in this regard and enjoying widespread acceptance among providers of models. It is further recommended that other protocols should only be declared binding in dialogue with the providers of models, with a view to ensuring that implementation is as error free and consistent as possible. The Code of Practice also stipulates that a rights holder's objection to a use must not have a negative impact on the findability of the works in search engines.

However, this would constitute an overreach of the scope of the AI Act, which does not mention this and should therefore be deleted. A similar argument can be made

for the provision that obliges model providers to make reasonable efforts to mitigate 'memorisation' of training content as well as to prohibit copyright infringing uses. The absence of such an obligation within the AI Act constitutes a fundamental contradiction with the overarching objective of establishing alignment with both the AI Act itself and the prevailing international approaches.

It is necessary to make clear in the Code of Practice that the EU legislation applies only within the EU and does not apply to the training of models outside the Union. Ensuring legal certainty for providers of models and their availability in Europe is imperative. The Internet Industry's perspective is that this principle should also apply to the information on protected content not web-crawled by the Signatory mentioned in Measure I.2.4. The extension of these provisions beyond the geographical boundaries of the EU would also exceed the scope of the AI Act.

It is also crucial to emphasise that the data utilised for model training is not retained. Furthermore, the obligations stipulated in Measure I.2.5 should be addressed at the level of the AI systems themselves, as it is generally not at the level of the models that protected works are created.

6. Summary

The third draft of the Code of Practice has been revised and contains certain enhancements in comparison to previous iterations. Nevertheless, in order to facilitate its utilisation as a pragmatic and innovation-promoting instrument in the implementation of the AI Act, further adjustments are imperative, particularly to avert the potential for the obligations stipulated herein to exceed the provisions of the AI Act, thereby engendering heightened uncertainty and compliance costs for companies. Specifically, eco considers the following amendments to be necessary:

- Clear definition of the scope of application
  In order to provide companies with clarity and legal certainty, it is essential to define the scope of application of the Code with greater precision. Of particular note is the ambiguity surrounding the categorisation of fine-tuning as a substantial modification to an AI model, a decision which would consequently impose regulatory obligations on users. Given that the development of new models may not have any impact on a significant number of companies, it is the adjustments to existing models that may have consequences for them. It is therefore possible that the provisions of the AI Act could have a greater impact on companies than was originally envisaged. Consequently, the Internet Industry has expressed its support for the exclusion of fine-tuning from the definition of significant changes, or at the very least, the establishment of explicit requirements. Furthermore, there is a necessity to harmonise the definition of open source models with the AI Act in order to avoid contradictions.

- Less bureaucratic approach to transparency obligations
  The third draft of the Code of Practice represents an enhancement in the form of the model card, which provides a clear articulation of the stipulated transparency requirements and offers a valuable orientation resource for

SMEs. The restriction of the model card to a limited number of words is a positive development, as it facilitates practical implementation. Nevertheless, there is a necessity for greater precision in the requirements for the disclosure of sensitive data. The disclosure of training data and system architecture is associated with potential risks, including the exposure of business secrets and compromised security. Furthermore, certain information obligations, such as the measurement of energy consumption, are not always feasible in practice. The stipulation that model card data be stored for a period of ten years can be regarded as excessive and should be reconsidered. It is therefore essential to establish clear and proportionate requirements in order to harmonise transparency and innovation.

- Defining practical safety standards
  Notwithstanding the presence of certain enhancements in the present version, significant terms such as 'harmful manipulation' are delineated in an overly ambiguous manner, thereby impeding the feasibility of practical implementation. Furthermore, the document inappropriately shifts the responsibility for risk mitigation to model providers, even though many risks only arise at the level of the AI system. The deadlines for model assessments are set to an unrealistically short timeframe, resulting in undue pressure being placed on companies. Of particular concern is the mandatory external assessment, which puts business secrets at risk and goes beyond the requirements of the AI Act. The evaluation team qualification requirements are too strict, creating unnecessary hurdles for SMEs. In the opinion of eco, there is still room for improvement in this area.

- Copyright in accordance with the CRD and the AI Act
  In the context of copyright provisions, the Code in question has been found to exceed the standards set out in both the AI Act and the Copyright Directive. This discrepancy necessitates immediate rectification. The detailed transparency obligations for training data that are still required are not technically feasible in many cases and represent an unnecessary hurdle. It is acknowledged that rights holders should retain the capacity to object to the utilisation of their works; however, this process should be facilitated by established, readily implementable protocols such as robots.txt. Furthermore, additional obligations, such as the 'memorisation' of training data or the prevention of the influence on the findability of copyrighted works, go beyond the scope of the AI Act and must be removed. Furthermore, it is imperative that EU legislation does not extend to model training outside the Union, in order to circumvent legal uncertainty. Finally, mandatory measures on the use of protected works should be placed at the level of AI systems rather than models. This necessitates a substantial degree of adaptation.

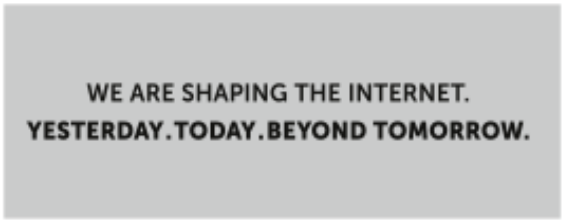

About eco: With approximately 1,000 member companies, eco (international.eco.de) is the leading Association of the Internet Industry in Europe. Since 1995, eco has been highly instrumental in shaping the Internet, fostering new technologies, forming framework conditions, and representing the interests of its members in politics and international forums. eco has offices based in Cologne, Berlin and Brussels. In its work, eco primarily advocates for a high-performance, reliable and trustworthy ecosystem of digital infrastructures and services.