22 May 2025

European Union AI Office
Rue de la Loi / Wetstraat 200
1049 Brussels
Belgium

RE: Comments of ACT | The App Association to the European Commission AI Office Regarding Guidelines for Providers of General-Purpose AI Models Under the AI Act

ACT | The App Association appreciates the opportunity to submit views to the European Commission AI Office in response to its request for comment on guidelines for providers of general-purpose artificial intelligence (GPAI) models under the AI Act.[1]

The App Association is a not-for-profit policy trade association for the small business technology developer community. Our members are small and medium-sized enterprises (SMEs) within the app ecosystem that engage with verticals in industries across the European Union (EU) and around the globe. We work with and for our members to promote a policy environment that rewards and inspires innovation while providing resources that help them raise capital, create jobs, and continue to build incredible technology. Today, the ecosystem the App Association represents – which we call the app economy – is valued at approximately 830 billion euros globally and is responsible for over 1.3 million jobs in the EU. App Association members create innovative software and hardware technology solutions and are at the forefront of incorporating artificial intelligence (AI) into their products and services.

Even more than other areas of technology, the AI sector is extraordinarily fast-moving, both in terms of the technology itself and the business models that companies create to leverage that technology. Much about the AI industry is unsettled, meaning that regulations crafted with one set of industry players in mind may have serious implications for others. With the release of the DeepSeek model earlier this year, we have already seen one high profile instance of a smaller, lesser-known entity rising quickly to compete with the extant dominant players in the AI marketplace. Therefore, SMEs cannot assume that they will not be providers of GPAIs in the near future, and so must consider the implications of AI Act requirements applying to them today.

We generally support and share the goal of supporting responsible and pro-innovation governance of advanced AI systems across their lifecycle. To support this effort, we urge for alignment with the following resources:

- The App Association's comprehensive AI Policy Principles: https://actonline.org/wpcontent/uploads/2023-11-16-ACT-AI-Policy-Principles-FINAL.pdf; and
- The App Association's 'AI Roles and Interdependencies Framework', which proposes clear definitions of stakeholders across the AI value chain, from development to distribution, deployment, and end use; and discusses roles for supporting safety, ethical use, and fairness for each of these important stakeholder groups that are intended to illuminate the

[1] https://digital-strategy.ec.europa.eu/en/news/commission-seeks-input-clarify-rules-general-purpose-ai-models.

interdependencies between these actors, thus advancing the shared responsibility concept: https://actonline.org/wp-content/uploads/ACT-AI-Roles-Interdependencies-Framework-final-text-May-2024-UK-English.pdf.

We appreciate the steps taken in the third draft to tailor expectations to the size and capacity of the GPAI provider, especially acknowledging that SMEs and startups often lack the same financial or technical resources as large corporations. However, this sentiment is not consistently reflected in the operational sections. We believe the third draft could be further improved in the following areas to better support the ability of SMEs to innovate in the AI marketplace:

- We are concerned about the proposed approach to compliance flexibility embedded in Measure II.4.5 of the Safety and Security section, which addresses rigorous model evaluations. It requires evaluations to meet standards comparable to those used in leading scientific journals or major machine learning conferences, a formidable bar, even for large tech players. SMEs, theoretically, are allowed to deviate from these standards if they lack the expertise or financial resources, but only on the condition that they formally notify the AI Office and negotiate an alternative approach. This presents several challenges:
    - Administrative complexity: Many SMEs may not have the legal or compliance resources to navigate these formal processes effectively.
    - Lack of clarity: The Code provides no specific benchmarks or examples for what might constitute an "acceptable alternative."
    - Uncertainty and inconsistency: With decisions left to case-by-case negotiations, the process could become unpredictable or opaque, especially as the AI Office is not yet fully operational, and its future capacity remains uncertain.

  This mechanism lacks a predictable structure for a fruitful implementation, resulting in an inconsistent application, legal uncertainty, and could even discourage SMEs from engaging with the Code, contradicting the overarching goal of promoting inclusive and widespread adherence to responsible AI practices.

- Several provisions of the guidelines, including those relating to incident reporting and monitoring, are unclear as to how a developer is meant to comply from a technical, real-world perspective. Without clear guidelines, SMEs in particular will find it difficult to innovate in the AI space.
- For SMEs, ease of understanding is essential for reducing compliance burdens. Many SMEs do not have dedicated compliance departments and have difficulty affording legal teams to monitor the latest complex rules. Therefore, to the extent the guidelines can be simplified by combining overlapping Measures and reducing reliance on prescriptive details, SMEs will be better able to comply.
- Some elements of the AI Act that the guidelines should explain further are not addressed, such as the "adequate level of detail for the summary of content used for training." Better explanation of such terms and any obligations created by them would SMEs ensure that they are in compliance with relevant requirements.
- SMEs are quite likely to be modifiers of GPAIs, so obligations created for modifiers and how they fit into the regulatory structure is of particular importance to our members. It is unclear, for example, how a modifier's obligations can only apply to "the extent of

their…modifications" in a practical sense, as there is not yet a standardized way to fully separate modifications from the underlying behavior of a model. It may also be unfeasible for SMEs to have access to the underlying training data or code history of a GPAI that they are modifying, depending on whether the model is open source and of what type, making it difficult to know for certain whether a model's behavior is the result of an SMEs modifications.

- The draft guidelines should clarify how compliance with their terms relates to compliance with EU copyright law. Though Objective III states that the Code is intended to assist covered model providers in "effectively comply[ing] with Union law on copyright and related rights," it is unclear whether adherence to the guidelines' suggested copyright compliance policy would actually insulate an SME with an unsophisticated compliance apparatus from claims of non-compliance with EU copyright law. Because a court would likely not find compliance with the Code to be a copyright "safe harbor," the language of the guidelines should not imply to SMEs that it could be.
- A number of aspects of the guidelines create heavier obligations on large developers. While it may seem that such size discrimination would benefit SMEs by giving them a comparative advantage, in reality many SMEs build on top of the products and services initially developed by larger entities. Unduly burdening large developers when entity size is not necessarily a risk factor for potentially harmful GPAI behavior could have downstream effects that harm SMEs as well.

Further specific feedback from the App Association on various sections of the Proposed Code include:

The Draft Code's Principles and Assumptions:

With respect to the draft's plan, principles, and assumptions, we recommend the following :

- **Principles 3a and 3c:** It's worth reconsidering whether "risk tiers" are the optimal framework for risk management in foundational principles. Risks often exist on a spectrum or across multiple dimensions, which may not fit neatly into discrete, sequential categories. Alternative models—such as varying levels or degrees of risk—could be explored to avoid the presumption that risks are always discretized. This perspective aligns with other parts of the document, such as the commitment to conducting systemic risk analysis with differing levels of depth and focus.
- **Principles 6 & 7:** From a technical standpoint, the terms "safe" and "safety" (used in phrases like "AI safety infrastructure," "ensuring the safety of…," "AI safety governance," and "safety outcomes") require clarification, especially since "safety" is treated as distinct from "human centric" and "trustworthy." The broad interpretation of "safety" could result in inconsistent application of the Code.
- **Final Paragraph:** As with earlier drafts, there remains ambiguity around what it means for those modifying GPAI models to have obligations limited to "the extent of their…modifications." Significant alterations to a model can fundamentally change its behavior, making it unclear which responsibilities—beyond documenting the original

model's training data and processes—would no longer apply to those making substantial changes.

- **Training Data Transparency:** The latest draft still does not address the need for an "adequate level of detail for the summary about the content used for training," even though this is required under Article 56(2)(b) of the EU AI Act. We strongly recommend including this information to ensure compliance with regulatory expectations.

The Draft Code's Preamble:

The App Association requests that the preamble of the Code explicitly state that following its guidelines does not demonstrate compliance with intellectual property or trade secret laws. The EU AI Act stands apart from intellectual property regulations. While the preamble currently notes that adherence to the Code is not proof of meeting AI Act requirements, it should also clarify that this does not equate to compliance with intellectual property, copyright, or trade secret protections.

We are also concerned that Objective III (Copyright) oversimplifies the issue by equating the text and data mining (TDM) exception with the entirety of copyright law, overlooking the courts' essential role in interpreting these laws. The AI Act requires GPAI providers to establish policies for compliance with EU copyright and related rights, especially concerning rights reservations under Article 4(3) of Directive (EU) 2019/790 (Art. 53(1)(c)). However, the Code only suggests possible policy measures, and legal compliance is ultimately determined by the courts. EU copyright law is complex and varies across Member States. The AI Act itself confirms that it does not affect the enforcement of copyright rules (Recital 108) and that GPAI compliance is without prejudice to Union copyright law (Recital 109).

We recommend revising Objective III to clarify that the Code's aim is to help GPAI model providers meet the AI Act's requirements for copyright and related rights, while also protecting intellectual property and confidential business information. We also suggest adding language to ensure the Code is not interpreted in a way that conflicts with EU intellectual property laws, including copyright and trade secrets.

Finally, we note that the Safety and Security section (Recital f) already highlights the need to interpret commitments in line with the Code's objectives, especially Objective IV.

The Draft Code's Transparency Commitments for General-Purpose AI Models:

We offer the following reactions and suggestions to the Draft Code's transparency commitments for general-purpose AI models:

- **Intellectual Property and Confidentiality:** As emerging AI solution developers and deployers, we believe that the current model documentation requirements do not sufficiently address the protection of proprietary information, trade secrets, or confidential business data. The present draft lacks robust mechanisms for safeguarding sensitive material, especially regarding information disclosure obligations and guidelines for redacting confidential content. While the new references to Article 78 represent progress compared to earlier versions, and Measure 1.1.2 now acknowledges confidentiality protections, more explicit safeguards are needed.

- **Disclosure of Training Data Details:** The stipulation to reveal comprehensive information about datasets used for model training raises practical concerns. If a developer deems most or all of this information to be proprietary or a trade secret, there is ambiguity about how such claims will be evaluated and who will make these determinations. Clear, fair procedures are essential to resolve these issues without compromising competitive advantage.
- **Documentation and Energy Reporting:** Standardized methods for assessing AI energy consumption are still developing, making it difficult to ensure that energy usage reports are genuinely comparable or independently verifiable. Methodological differences will likely highlight inconsistencies rather than enable meaningful comparison. Achieving true comparability would require access to data—such as cooling energy use—that is often unavailable.

  Additionally, the term "withdrawn from the market" is vague, as older model versions may remain accessible or supported long after newer iterations are released. We propose that documentation be retained either until a model is replaced or for a fixed period (e.g., 10 years after being superseded).
- **Information Templates and Stakeholder Input:** The documentation refers to a forthcoming template from the AI Office for public disclosure of training data, as required by Article 53(1)(d). We seek clarification from the European Commission on whether this template is distinct from the Model Documentation Form included in the current materials. If the template is still under development, we request an opportunity for small business stakeholders to provide feedback before finalization.
- **Scope of Information Sharing:** It remains unclear what constitutes "additional information" that providers are expected to share with downstream users. We recommend limiting disclosures strictly to what is necessary, to avoid unnecessary exposure of sensitive business information.
- **Alignment with AI Office Templates and Copyright Policy:** The model documentation form should be harmonized with the AI Office's forthcoming summary template, ensuring that requirements for transparency about training data are consistent and do not inadvertently expose intellectual property or trade secrets. The January guidance from the AI Office indicated more detailed expectations—including copyright and trade secret protections—that are not yet reflected in the current draft, raising concerns about possible misalignment and lack of opportunity for meaningful input.
- **Incorporating IP and Trade Secret Protections:** We urge that the documentation process explicitly recognize and protect intellectual property, trade secrets, and confidential business information. For example, the form should provide clear instructions on how to redact or withhold sensitive details, and clarify whether information shared with the AI Office will also be disclosed to other market participants, including potential competitors.
    - **Data Collection and Curation**: Where methodologies for data gathering and curation are proprietary, guidance should be provided on how to safeguard this information, and under what circumstances it may be withheld from downstream users.
    - **Licensing Information**: For providers not operating under open-source licenses, sharing license terms may reveal confidential business arrangements. Guidance is

needed on how to redact sensitive terms, especially when different clients receive customized agreements.

- **Design Specifications**: The EU AI Act requires disclosure of model and training process design, but this should not override protections for trade secrets. The documentation should make clear how to identify and protect such information in compliance with Article 78.

- **Practicality and Relevance of Requested Information:** Some of the information currently requested—such as the exact number of data points, training duration, system architecture, or parameter relevance—may be difficult to provide, especially for small businesses using varied data sources or continuous learning approaches. We recommend that every information request in the documentation form be clearly justified by a specific need and that the form allow flexibility where providing certain details is impractical or of limited value.

The Draft Code's Copyright Provisions for General-Purpose AI Models:

As emerging AI startups and small businesses, we recognize the complexity of ongoing policy debates in the U.S. and the technical nuances of EU copyright law. While we respect that some matters remain unresolved, we urge policymakers to reference Recital 109 of the AI Act within Recital (b). This would clarify that requirements for proportionality—specifically, those that consider the scale of providers and offer streamlined compliance paths for SMEs and startups—do not override EU copyright regulations. Alternatively, we propose evaluating whether Recital (b) is necessary at all, as it's unclear why obligations should differ based on provider size.

Commitment I.2: Copyright Policy Guidance: We see value in clear instructions regarding how copyright-related risks might factor into "systemic risk" assessments. Guidance would help us understand which policy measures may be relevant and how to address them appropriately.

Measure I.2.2: Responsible Web Crawling

- We suggest renaming this measure to: "Minimize use of unlawfully obtained copyright-protected content during web crawling."
- The opening sentence could be revised to: "To help limit the use of works and protected content acquired without proper authorization by web crawlers…" The current language overstates the certainty of compliance, given ongoing debates about what constitutes "lawful access." The Code should acknowledge these ambiguities, while still encouraging avoidance of clearly illicit practices (such as bypassing technical protections or using content from piracy sources).

Measure I.2.3: Observing Rights Reservations: Some elements, such as encouraging participation in standardization, are more suitable for introductory or explanatory sections rather than as binding commitments.

Measure I.2.5: Reducing Copyright-Infringing Outputs:

- References to "free and open-source licenses" do not reflect the language of the EU AI Act. We suggest citing Article 53(2) and specifying "and not GPAISR" at the sentence's end, particularly if copyright infringement risks could lead to a model being classified under GPAISR.

- We also recommend considering a parallel commitment to address copyright risks in training datasets, acknowledging the complexity and ongoing nature of these discussions.

The Draft Code's Safety and Security Provisions for Providers of General-Purpose AI Models with Systemic Risk:

As small business AI innovators, we appreciate the intent behind the proposed safety and security provisions but believe that the current approach requires significant refinement to be practical, effective, and inclusive of diverse stakeholders like ourselves.

First, it is essential to broaden the understanding of risk beyond a narrow focus on "capabilities." Risk management must consider evolving use patterns, newly identified limitations, and emerging evidence about how AI systems interact with people and institutions. Systemic risks can arise not only from what AI can do but also from what it cannot do—misunderstandings about its abilities can be just as dangerous. Therefore, shifting the focus from "capabilities" to "behaviors" offers a more comprehensive and realistic foundation for managing risks.

We also urge clearer definitions and scope for the frameworks being proposed. Calling the framework a "Safety and Security Framework" is misleading since systemic risks encompass more than just safety or security concerns. Renaming it to "Systemic Risk Management Framework" would better reflect its purpose and help organizations, especially small businesses, align their policies effectively. Moreover, providers should be allowed to build on and adapt existing frameworks they already use, rather than being forced into rigid new structures.

The current emphasis on rigid risk tiers and acceptance criteria is problematic. There is no broad expert consensus that static tiers are the best way to categorize or manage risk. Instead, acceptance criteria should be flexible and evidence-based, accommodating different approaches depending on the context and nature of the AI system. Additionally, risk mitigation should not be limited to technical fixes; organizational and procedural measures are equally vital and should be recognized.

Forecasting future AI capabilities or risk timelines is inherently speculative and unreliable. Mandating such forecasts risks fostering false confidence or unnecessary work that distracts from meaningful risk management. We recommend removing these forecasting requirements entirely. Similarly, requirements for transparency into external inputs in decision-making add little value and extend beyond the legal scope.

Updating risk management frameworks should be driven by actual need and evidence, not by a vague obligation to "improve." Small businesses, in particular, benefit from clear, achievable expectations that allow them to focus resources where they matter most.

Regarding risk assessment throughout the AI model lifecycle, it is important to recognize that many model providers do not control how their models are deployed or used after release. Expectations for post-deployment assessments should be realistic and tailored accordingly. Distinguishing between standalone model developers and those who deploy models for specific applications would help clarify responsibilities and reduce undue burdens.

Many of the proposed measures are overly prescriptive and detailed, limiting the flexibility that innovation requires. Ambiguous terms like "effective compute" and "model elicitation" need clearer definitions or should be removed to avoid confusion. Risk identification and analysis should be based on demonstrable evidence that a model increases systemic risk beyond existing alternatives, rather than speculative or broad lists of potential risks.

Evaluations should consider the broader context in which models operate, not just isolated technical tests. Quantitative risk estimates are valuable only when grounded in rigorous methodology. External evaluations should be optional, especially for open-source models already accessible to the community.

Security provisions must be comprehensive, covering all relevant assets and aligned with established cybersecurity standards like ISO 27001 or the NIST Cybersecurity Framework. Overly detailed or duplicative security requirements should be streamlined to focus on effective protection rather than exhaustive checklists.

Documentation and reporting obligations should be reasonable and proportional. Excessive demands for continuous updates, detailed incident reporting, or algorithmic improvement disclosures place undue strain on small businesses without clear benefits. Instead, a straightforward feedback mechanism for addressing risks is more practical.
Incident response should be a shared responsibility across the AI value chain, with clear roles and alignment to international standards such as those from the OECD. Notification timelines should be flexible (e.g., phrases like "without undue delay" are preferable to rigid deadlines).

External assessments should be required only when internal expertise is insufficient, and post-market external reviews should remain voluntary to avoid unnecessary burdens. Some proposed commitments duplicate existing legal protections and can be removed to reduce complexity.

The Draft Code's Safety and Security Section Glossary and Systemic Risk Taxonomy (Appendix 1):

While we recognize the importance of identifying and managing risks associated with AI, the existing framework feels insufficiently grounded in evidence and often lacks clear justification for the risks it highlights. Many of the risks prioritized appear speculative, making it difficult for resource-constrained small businesses to understand which threats are genuinely pressing and which are more hypothetical.

- **Appendix 1.1:** One of our key observations is that some of the most critical risks—such as privacy violations, the creation and distribution of nonconsensual intimate imagery, and child sexual abuse material—are not given the prominence they deserve. Instead, they are relegated to secondary categories without clear rationale, creating confusion about how these risks should be addressed. The division between primary and secondary risk types is ambiguous, which complicates compliance efforts for small teams.

  We also find that certain risks, especially those involving cyber, chemical, biological, radiological, and nuclear threats, are framed in ways that don't fully acknowledge the human role involved. These risks often arise from human-machine interactions rather than AI capabilities alone, raising questions about whether they truly fit the definition of

systemic risk under the AI Act. This distinction matters because it influences how responsibility and oversight should be allocated.

The way manipulation is defined in the taxonomy is another area of concern. Labeling manipulation as any instance where an AI causes someone to make a decision they wouldn't have otherwise made is overly broad. After all, influencing decisions is a fundamental part of effective communication. If this definition is applied strictly, it risks categorizing nearly all general-purpose AI as harmful, even when many users benefit from the same interactions.

The category of "loss of control" is particularly vague and seems to borrow from speculative narratives about existential risks. Without clear boundaries, it's unclear whether everyday technologies—like self-driving cars or language models inadvertently revealing private data—fall under this umbrella. This ambiguity makes it challenging for small innovators to know how to design and deploy AI responsibly without fear of unintentionally triggering regulatory requirements.

- **Appendix 1.2:** We also see a lack of clarity around what constitutes systemic risk in autonomous AI research and development. Since most AI projects incorporate some degree of autonomy in their processes, the absence of clear criteria leaves innovators uncertain about when their work might be considered risky.

  Moreover, some sections of the taxonomy, such as "risks to society as a whole," are too broad and speculative. Including vague threats like risks to non-human welfare or financial system stability stretches the definition of systemic risk beyond what is reasonably foreseeable. We believe systemic risk should focus on tangible, evidence-backed harms to ensure regulations are practical and enforceable.

  It's important to note that many risks depend heavily on the context of AI deployment rather than the model's inherent design. For instance, risks arising from the way a model is integrated into applications or systems should be addressed differently, especially when the model provider also acts as the deployer.

- **Appendix 1.4:** Finally, several capabilities listed as sources of systemic risk, such as adaptive learning or forecasting, are described in overly broad terms. These capabilities cover a wide range of non-harmful uses, and without considering the degree of autonomy or context, even minimally autonomous models could be unfairly labeled as systemic risks.

  In light of these issues, we urge regulators to refine the taxonomy by removing ambiguous terms, clarifying definitions, and grounding risk categories in solid evidence and practical relevance. Doing so will help small AI innovators navigate compliance more effectively while continuing to drive innovation that benefits society.

The App Association appreciates the EC's consideration of the above views. We urge the Commission to continue to streamline the regulatory environment to bolster innovation in the EU technology sector.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

Chapin Gregor
Policy Counsel

ACT | The App Association
1401 K St NW (Ste 501)
Washington, DC 20005
202-331-2130