



September 18, 2024

Ms. Lucilla Sioli
Director
Artificial Intelligence Office
Directorate-General for Communications Networks, Content and Technology
European Commission
B-1049 Brussels
Belgium

Re: Response to the Consultation on the AI Act: Trustworthy General-Purpose AI

Dear Director Sioli:

The U.S. Chamber of Commerce (“Chamber”) welcomes the opportunity to express our views on the European Commission’s consultation on the Code of Practice (“CoP”) for General-Purpose Artificial Intelligence (“GPAI”) models under the Artificial Intelligence (“AI”) Act.

The Chamber is the world’s largest business advocacy organization, promoting free enterprise and representing businesses of all sizes and sectors of the economy. In Europe, we collaborate closely with our partner organizations at AmCham EU and other American Chambers of Commerce across the 27 member states, and with our counterparts at BusinessEurope and other member-state business organizations.

We welcome the opportunity to comment on the proposed Code of Practice for GPAI models, focusing on key areas such as transparency, copyright-related rules, and risk management. Our feedback aims to support the Commission in creating a balanced and flexible framework that mitigates risks while maximizing the innovative potential of AI technologies.

General Principles

To support the objectives of the AI Act, the CoP for GPAI models must uphold four key principles. First, the CoP should be clear and well-focused within the parameters established by the AI Act to ensure legal certainty and encourage widespread adoption. A well-defined scope will help stakeholders understand their obligations and foster compliance.

Second, the CoP should align with existing standards and best practices, such as those from the G7 Hiroshima AI Code of Conduct, while remaining within the

framework of the EU AI Act. This alignment can help make swift progress in drawing up the CoP while also limiting the risk of fragmentation or access barriers to the Single Market and supporting European innovation.

Third, the CoP should be anchored in the technical realities of rapidly developing technologies. It should avoid favoring specific solutions and remain adaptable to developments, allowing flexible ways to demonstrate compliance. This approach will ensure that the CoP remains relevant and effective as AI technologies evolve.

Finally, the CoP will only achieve its objectives if it is created with a deep understanding of AI technologies' workings, possibilities, limitations, and risks. The drawing-up process should center around the contributions of AI model providers, as envisaged by Article 56(3) AI Act. Engaging AI experts will ensure that the CoP is grounded in practical knowledge and experience.

The Process for the Drawing-Up of the CoP

Given the technical nature of the obligations under the Code and the short timeframe for its development, the EU Artificial Intelligence Office (AIO) should focus on the contributions of GPAI model providers and deployers to develop the CoP.

Drafting the Code must be guided by the principle of proportionality and should focus on what is appropriate and necessary to achieve the purpose of the AI Act. Overly burdensome requirements could stifle innovation and hinder the adoption of AI technologies. As noted in Mario Draghi's recent report, "The Future of European Competitiveness," the complexity and potential conflicts between the AI Act and the General Data Protection Regulation (GDPR) can "undermine developments in the field of AI by EU industry actors."¹

To ensure that the first draft of the Code is fit for purpose, the AIO should plan sufficient time for workshops with GPAI model providers before the kick-off plenary. These workshops will allow for in-depth discussions and feedback, ensuring that the CoP addresses practical concerns and challenges and presents technically-sound regulatory responses. These workshops should focus on areas where participation of GPAI model providers is crucial for the CoP's success, such as technical documentation, risk assessment and mitigation, transparency and explainability, copyright and data usage, and evaluation and testing.

¹ European Commission. (2024). *The future of European competitiveness: In-depth analysis and recommendations* (p. 79). Retrieved from https://commission.europa.eu/document/download/ec1409c1-d4b4-4882-8bdd-3519f86bbb92_en?filename=The%20future%20of%20European%20competitiveness_%20In-depth%20analysis%20and%20recommendations_0.pdf.

Safeguards must be in place for these workshops to allow GPAI model providers and other stakeholders to share confidential information with the AIO and the working group chairs. Balancing transparency with confidentiality will foster trust and encourage open communication.

While independent experts can play a helpful coordinating role, the AIO must be closely involved in the drawing-up process to ensure the Code does not become detached from the practicalities of oversight. The AIO's involvement will ensure that the CoP is both practical and enforceable and that it does not go beyond the letter of the AI Act or diverge due to specific interests.

Technical Documentation for GPAI Models

We appreciate the efforts within the CoP to simplify technical documentation requirements for GPAI model providers under Article 53 of the AI Act. Clear documentation standards will help ensure that AI models are developed and used responsibly. Finding the right level of detail in this documentation is crucial to balance the need for transparency with the need to protect sensitive information and avoid security vulnerabilities.

The CoP should consider adopting model cards to establish a shared understanding of the AI models released by GPAI providers. This now widely adopted concept can provide essential information about the model's capabilities, limitations, and intended uses, promoting transparency and accountability, without undermining AI providers' legitimate interest.

Disclosures should be oriented around the intended purpose/audiences and be adaptable to the evolving AI landscape. This approach will ensure that the information provided is relevant and useful to stakeholders.

The CoP should avoid documentation that imposes undue compliance burdens on GPAI model providers and rightsholders. Streamlined and practical documentation requirements will encourage compliance without stifling innovation. In particular, the Code should provide clear guidance on what constitutes "keeping technical documentation up to date." While initial documentation at launch is crucial, continuous model improvements that do not fundamentally change the model's nature should not necessitate extensive updates. Requiring exhaustive documentation for every minor change would overwhelm stakeholders and increase undue compliance burdens for GPAI providers.

The CoP should focus on GPAI model providers' responsibilities and acknowledge the limitations they face in gathering information about downstream uses. Recognizing these limitations will help create realistic and achievable standards.

The CoP also should support fair processes for information requests from the AI Office. Clear and balanced processes will ensure that information requests are handled efficiently and transparently.

Transparency and Copyright-Related Rules

The EU maintains a long history of robust intellectual property (IP) protection that furthers both innovation and creativity, as evidenced by existing copyright laws. The Chamber's International IP Index illustrates the strength of the EU's current copyright framework, with the U.S., Singapore, UK, and EU leading the rankings in the copyright-related indicators.

In implementing the AI Act's goal to foster AI adoption and development in Europe, the CoP must not create new or contradict existing laws. Rather, it should encourage ongoing discussions and innovation, allowing flexibility in an ever-changing technological and creative environment. The Chamber looks forward to working with the European Commission to achieve this goal and ensure the CoP enables the EU to retain its global leadership on IP policy.

Systemic Risk: Risk Taxonomy, Assessment, and Mitigation

The CoP should focus evaluations on systemic risks and safety testing in internationally agreed-upon risk domains, such as those outlined in the G7 Principles and the June 2023 Voluntary White House Commitments on Artificial Intelligence. These domains include critical chemical, bio, radio, nuclear (CBRN) capabilities, critical cybersecurity capabilities, critical self-replication/self-proliferation (including tool-use) capabilities, and bias and discrimination. A clear and comprehensive risk taxonomy will help ensure that evaluations are focused on the most significant risks.

High-quality data sets are a critical precursor for accurate and robust outputs. In that vein, we believe it is crucial that providers undertake activities that can help to ensure that training data is secure, accurate, relevant, complete, and consistent, and that bias is mitigated to the extent possible. This may include, for example, monitoring and curating empirical datasets ingested for deviations, drifts, anomalies, and bias. Deployers should also undertake activities to ensure that other input data is secure, accurate, relevant, complete, and consistent, with bias mitigated to the extent possible.

The CoP should promote emerging procedural best practices but avoid being over-specific on evaluation methodologies where there are open scientific questions. Guidance should be outcome-focused, flexible, and promote procedural best practices. This will allow for continuous improvement and adaptation as new evaluation methods are developed.

The CoP should emphasize the importance of AI model providers' responsibility and governance structures without mandating prescriptive rules. Governance structures should include internal channels for thorough review and assessment of evaluation results to ensure comprehensive oversight and promote informed decision-making. Flexible governance structures will allow organizations to tailor their risk mitigation strategies to their specific needs and contexts. For example, the Code must take into account unique roles and responsibilities across the AI value chain, as well as challenges in accessing information about downstream applications.

Reviewing and Monitoring the CoP

The multi-stakeholder consultation seeks input on the timing of reviews and adaptations of the CoP for GPAI models to ensure it reflects the state of the art. The CoP should avoid favoring specific solutions and remain adaptable to developments, allowing flexible ways to demonstrate compliance. The AIO should provide as much implementation lead time as possible ahead of the initial August 2025 deadline, and for any new compliance timelines thereafter. Regular reviews and updates will ensure that the CoP remains relevant and effective as AI technologies and practices evolve.

Key Performance Indicators (KPIs)

The AI Office should refrain from setting prescriptive KPIs, especially in areas such as risk mitigation measures and capability evaluations. The lack of scientific research and standards on AI measurement creates challenges for providers in assessing the effectiveness of their risk mitigations based on a globally accepted objective benchmark.

GPAI models are designed to be adapted to a wide range of uses, making it infeasible to anticipate all possible risks. The rapidly evolving nature of technology also poses challenges in identifying new mitigation strategies as new risks emerge. A flexible approach to KPIs will allow for continuous improvement and adaptation.

Technical Documentation

The Code will set out details on how to provide information about the data used for model training, testing, and validation, including the type and provenance of data

and curation methodologies. To safeguard trade secrets and foster innovation, providers should be required to give only high-level descriptions of data types, without needing a time period or disclosing privacy-related information (which would be covered under other legislation, notably GDPR). For example, providers should be able to reference data categories (e.g., audio, video, images, text), notable characteristics (e.g. language), and whether data was publicly available online, acquired from third parties, and/or synthetically generated. Providers should also be encouraged to explain their rationale in the treatment of data, particularly regarding inclusion/exclusion choices. Another solution would be to lean into current industry work into data provenance and cross-industry metadata standards which brings transparency to the origin and use of empirical curated datasets. Visibility on provenance not only increases efficiency, accuracy, and quality assurance but also leads to more informed decision making, speeding up innovation.

Developing standardized principles for measuring compute will be key to ensure that safety measures are applied proportionately. Verifying the exact final amount of compute is challenging due to the iterative nature of model development. The approach to measuring compute will evolve over time as technology changes.

The Frontier Model Forum – an industry non-profit dedicated to advancing the safe development and deployment of frontier AI system – recommends principles for measuring and reporting training compute, including treating all operations equally, context-dependent approaches, excluding recomputations, and focusing on systems trained end-to-end. Standardized principles will help ensure consistency and comparability in reporting computational resources.

Currently, there are no standardized practices to document a model's known or estimated energy consumption. More research and standardization are needed. Instead of prematurely mandating potentially inaccurate energy consumption measurements, the Code should promote development of robust, standardized methodologies for assessing the energy efficiency of AI models and refrain from setting prescriptive requirements. This will ensure that energy-related information is reliable and useful, ultimately contributing to more sustainable AI practices. Addressing energy consumption is important for sustainability, and it requires careful consideration and development of appropriate standards.

Conclusion

The U.S. Chamber of Commerce appreciates the opportunity to contribute to the development of the Code of Practice for GPAI models. We believe that a well-crafted CoP, grounded in the principles of clarity, international alignment, technical realities, and expertise, will support the responsible development and use of AI

technologies. By addressing key areas including transparency and risk assessment, the CoP can foster innovation while ensuring safety and trustworthiness.

Engaging industry stakeholders and policymakers is crucial for the CoP's success. Continuous dialogue will keep the CoP relevant and effective, helping to identify emerging challenges and opportunities for timely updates and improvements. We appreciate the Commission's consideration of our recommendations and welcome the opportunity to elaborate on them in the weeks ahead.

Sincerely,



Tom Quaadman
Senior Vice President
Economic Policy
U.S. Chamber of Commerce



Marjorie Chorlins
Senior Vice President
European Affairs
U.S. Chamber of Commerce