Multi-stakeholder Consultation on Trustworthy General Purpose AI
Submission of the Software & Information Industry Association to the European AI Office

18 September 2024

On behalf of the Software & Information Industry Association (SIIA), we appreciate the opportunity to contribute to the European AI Office's consultation on trustworthy general purpose AI (GPAI). SIIA is the principal trade association for companies in the information industry. Our membership of nearly 400 companies, many of whom have operations in the EU, reflects the broad and diverse landscape of digital content providers and users in academic publishing, education technology, financial information, software, platforms, data analytics, and information services. Our members include upstream and downstream AI designers, developers, and deployers of AI systems across various environments. As the leading voice in the software and digital content industries, we recognize the significant role that this consultation will play in shaping the first General-Purpose AI Code of Practice (CoP).

We appreciate the AI Office's commitment to a multi-stakeholder process that includes perspectives across industry, civil society, academia, and public authorities. SIIA is eager to contribute our perspective with the aim of helping ensure that the CoP promotes the ethical use of AI while also supporting continued technological innovation.

The CoP Should Promote Transparency Measures Aligned with Existing Best Practices (Comments on Section 1)

We provide the following comments on Section 1, Parts A, B, and D. We are not at this time addressing issues relating to copyright.

Transparency is essential for responsible AI because it helps to engender trust among those using AI tools and provides a means to assess whether AI developers and deployers have implemented appropriate policies, procedures, and governance structures to advance AI that mitigates the risk of harm. The CoP must balance an interest in transparency against important countervailing policy objectives such as protecting trade secrets and exposing models to increased security risk.

The CoP should endorse transparency practices that will have a meaningful impact on downstream providers, downstream users, and oversight authorities. We believe this should be grounded in a robust understanding of what model providers are doing to advance transparency; any requirements beyond current state-of-the-art practices should be viewed

with caution. By doing so, the CoP can help to advance best practices across industry and align practices to improve interoperability.

We recommend focusing on model cards as the bedrock of transparency under the CoP. We further recommend that the CoP adopt an approach to model cards that is sufficiently flexible to address the needs of different audiences and the intended purposes of the models. To this end, it is important to recognize the challenges inherent in assessing all downstream uses (and misuses) of GPAI models. Any requirements should reflect the information that model providers reasonably have access to.

Lastly, we recommend that the CoP avoid requiring measures that remain subject to the development of formal technical standards. Providing a degree of flexibility to incorporate guidance that comes out of formal standards development processes will help to advance alignment across the AI industry and across borders.

The CoP Should Advance International Alignment in Risk Taxonomy, Assessment, and Mitigation (Comments on Section 2)

We recommend that the CoP align as much as possible with emerging international best practices for categorizing, assessing, and mitigating risk associated with GPAI models. The CoP should provide guidance that aligns with the work of the National Institute of Standards and Technology in the United States, as well as he G7 Hiroshima AI Code of Conduct, the Voluntary Commitments that the White House negotiated with leading AI model providers, and the emerging AI Safety Institute Network, including guidance from the UK AI Safety Institute (AISI) and AISIs in other nations. Focusing the CoP evaluations on systematic risk and safety testing in these internationally agreed-upon risk domains can also improve the accessibility of the CoP for a diverse set of model providers.

This approach will provide a solid foundation for the CoP and also further promote international alignment. That alignment is necessary for fostering innovation, interoperability, and promoting responsible AI development on a global scale. By aligning on best practices, countries can collectively shape the future of AI in a way that balances innovation while also collaborating to meet the shared goal of mitigating risks. This also reduces the compliance burden of companies operating in multiple jurisdictions that are already relying on existing benchmarks.

In addition to aligning with the risk taxonomies in the instruments and from the organizations noted above, we recommend that the CoP build in a concept of marginal risk, which compares the risk associated with a GPAI model against the risk associated with the next best alternative(s).

We also recommend that the CoP provide sufficient flexibility to accommodate developments in measurement science (which remains in its infancy for advanced AI models) and technical standards. The CoP should provide a framework for assessing and mitigating risk that can adapt as the state of the art around risk assessment and mitigation develops.

Lastly, while we concur with the need for model providers to implement best practices in responsible AI, including around governance, we recommend that the CoP take care to avoid creating compliance burdens - especially those that are not realizable or those that would impose weighty compliance burdens that outweigh the potential benefits from making GPAI models available to users.

The CoP Should Provide a Flexible Framework Grounded in Providers' Expertise (Comments on Section 3)

We believe the success of the CoP will be reflected in the ability of providers to comply with requirements while continuing to advance innovation and make GPAI models available to business, government, and consumer users. To that end, as noted above, we recommend that the CoP be grounded in the best practices currently in use by model providers. While practices differ, many leading model providers have a significant interest in advancing best practices because doing so facilitates model adoption, downstream development, and overall trust in the models' capabilities. These best practices are also reflective of technical expertise and a robust understanding of each model's capabilities and risk profile.

We recommend that the first CoP reflect existing best practices and standards (where available) that rely on the state of the art. Further iterations of the CoP can draw on experience under the first CoP and consider augmentation if performance indicators suggest a need to do so.