Recommendations from BSA | The Software Alliance on the Draft Basic Law for the Promotion of Responsible AI

April 11, 2024

General Comments

BSA | The Software Alliance (**BSA**)[1] appreciates the leadership of the Project team on the Evolution and Implementation of AI (**AI Project Team**), launched under the Digital Society Promotion Headquarters of the Liberal Democratic Party, and the AI Project Team's ongoing efforts to promote responsible AI. We support the objectives of the draft "Basic Law for the Promotion of Responsible AI" (**Draft Proposal**), presented by the Working Group members of the AI Project Team. The Draft Proposal aims to maximize the benefits of the sound development of AI while minimizing the risks to fundamental human rights and the interests of the public. BSA and its members are eager to support the AI Project Team to achieve these goals. As private sector involvement is explicitly stated in the proposed "co-regulation" model, we look forward to having constructive discussions on the specifics of the Draft Proposal.

BSA is the leading advocate for the global software industry. BSA members are at the forefront of developing cutting-edge services — including AI — and their products are used by businesses across every sector of the economy.[2] For example, BSA members provide tools including cloud storage and data processing services, customer relationship management software, human resource management programs, identity management services, cybersecurity services, and collaboration software and systems. BSA members are on the leading edge of providing AI-enabled products and services. As a result, they have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible AI.

In this submission, we urge the AI Project Team to engage with BSA, our members, and other interested stakeholders in further developing any legislative framework regarding AI. Our recommendations, described in more detail below, include:

- Adopting ***consistent and internationally recognized approaches to AI governance***.
- ***Avoiding legal inconsistencies*** within the legislative framework.
- Adopting a ***risk-based approach*** that focuses attention on high-risk uses of AI systems.

[1] BSA's members include: Adobe, Alteryx, Altium, Amazon Web Services, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, Cloudflare, CNC/Mastercam, Dassault, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, Hubspot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Nikon, Okta, Oracle, PagerDuty, Palo Alto Networks, Prokon, Rockwell, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Workday, Zendesk, and Zoom Video Communications, Inc.

[2] See BSA | The Software Alliance, Artificial Intelligence in Every Sector, available at https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf

- Ensuring the ***application of responsibilities to appropriate entities*** based on their role in developing or deploying AI.
- Ensuring obligations are ***reasonable and proportionate*** and avoiding imposition of overly prescriptive and burdensome requirements.
- Facilitating ***multi-stakeholder engagement*** during legislative and regulatory processes.
- ***Avoiding implying a preference or requirement for third-party verification***.

As policymakers around the world are developing regulatory approaches to AI, the global nature of today's technology ecosystem demands coordinated policy responses to foster innovation. We encourage countries to pursue interoperability through multistakeholder dialogue, developing a shared vision for a risk-based policy approach for addressing common AI challenges and advancing norms around responsible AI governance (e.g., risk-based approaches and proportionate and role-based responsibilities along the AI value chain). Global partners should also agree on common AI terminology and taxonomy to enable innovators the flexibility to adopt the technology for beneficial applications with confidence. As such, we recommend the Draft Proposal reflect such a harmonized approach. AI is a global technology, developed and used across borders and the product of many international collaborations. Setting out a globally coherent governance framework for AI is important to address risks that transcend borders and allow for international collaboration on the development and use of the technology. Governance frameworks should align with international best practices and prioritize interoperability with other frameworks globally.

Focus on High-Risk Uses and Ensure Global Interoperability

As a threshold matter, we recommend more clearly aligning the definition of "AI System" with internationally recognized definitions. Specifically, we suggest adopting the OECD's definition of an AI system:

> "An AI system is a machine-based system that, for explicit or implicit objectives infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."[3]

Given that AI systems are developed and deployed in an international context, definitions that apply to AI should operate across different jurisdictions to facilitate and promote further widescale adoption and use of AI technologies. Using an accepted and internationally recognized definition of AI, such as the OECD's, will facilitate the international alignment of Japan's policies, promoting dialogue, adoption, and compliance with the proposed AI legislation.

BSA supports regulatory frameworks that apply guardrails around high-risk uses of AI. AI can be used in a wide array of contexts, and policymakers should focus on those uses that pose the greatest risks to consumers. AI systems are used in a wide range of scenarios that do not present such risks, from detecting and lowering background noise on a video call to optimizing manufacturing production. For low-risk systems — like an AI system used to predict the types of fonts used in a document — additional obligations are not necessary. But

[3] Updates to the OECD's definition of an AI system explained, November 29, 2023, at https://oecd.ai/en/wonk/ai-system-definition-update

for high-risk systems, developers and deployers should implement measures to assess and mitigate risks.

The Draft Proposal seeks to focus on foundation models or frontier models, but we urge caution in defining these concepts and imposing new obligations, given that the capabilities and underlying technology continue to evolve, and a risk-based approach is currently more suitable for addressing uses that cause the most significant harm. Further, as you continue to consider these issues, it is important for the Government to align, to the extent possible, definitions and regulatory approaches with those of other countries to minimize unnecessary policy fragmentation and to promote international interoperability. This will assist AI developers and deployers, regulators, and consumers to maximize the benefits that will come from this technology, which will frequently be offered internationally, while identifying and minimizing risks. It also allows Japan's regime to benefit from and build on the extensive effort and thought leadership offered by AI legislative developments in other parts of the world. Further, it helps to ensure that any legislative framework is flexible and adaptable so that concepts and definitions can be updated as the technology evolves.

As described in further detail below, BSA supports AI developers and deployers implementing risk-management programs, impact assessments, and internal testing protocols for high-risk uses of AI. AI developers, including those creating foundation models, should provide information about model capabilities, limitations, testing, and security along the AI value chain based on the level of risk involved, to assist deployers and other entities in the AI ecosystem to better understand, identify, and address issues that may emerge in particular high-risk uses of the AI solution. And the Government of Japan should remain active in international discussions on AI governance to ensure that Japan's approach is well represented as further consensus develops in this fast- moving and consequential technological and regulatory landscape.

Avoid Legal Inconsistencies

We recommend avoiding legislation that would disregard the principle of technological neutrality and could lead to legal inconsistencies with existing statutes. Instead, any proposed legislation should focus on filling regulatory gaps. Specifically, many AI systems are already regulated by existing laws. Any AI legislation should keep that in mind and avoid imposing duplicative, conflicting, or unnecessary new requirements on operators of AI systems that are already subject to existing legal requirements.

Adopt a Risk-Based Approach

We strongly support a risk-based approach to AI policies that focuses on use cases that create high risks to individuals. For example, high-risk applications of AI include those that determines an individual's eligibility and results in the provision or denial of housing, employment, credit, education, access to physical places of public accommodation, healthcare, or insurance. The Draft Proposal describes safety verifications for an AI system used in particularly high-risk areas. However, we recommend requirements in the Draft Proposal be limited to high-risk use cases rather than "areas". The benefits, harms, and policy considerations around different applications of AI vary greatly.

For example, an AI system might be used by a healthcare provider to facilitate scheduling, address billing issues, or otherwise assist in routine administrative tasks. These generally would be "low-risk" uses of AI in the healthcare sector, while AI systems involved in determining eligibility for health insurance reimbursement or for a particular treatment might be considered "higher-risk" uses. As such, we recommend against creating compliance obligations in low risk use cases, which also include background blurring on video calls, autocorrect, email spam filters, web search engines, and TV show recommendations. Requiring compliance obligations for such low-risk technology could result in substantially slowing down business activities and not providing a meaningful benefit to consumers that are expecting business activities to be performed using well-accepted, widely used technology.

Ensure Balanced and Proportionate Allocation of Responsibilities Among AI Actors in the AI Ecosystem

AI systems can be used in an extraordinarily broad range of scenarios, and the risks arising from an AI system can vary greatly across specific use cases. The Draft Proposal focuses on compliance requirements for developers of foundation models, but this may not be sufficient to achieve the legislative goals. As the AI value chain is diverse and complex, we recommend that the Draft Proposal allocates responsibilities to the entities best placed to comply with them.

Because a developer is the entity that designs, codes, or produces an AI system, and a deployer is the entity that uses an AI system, these two organizations will have different roles in identifying and mitigating potential risks. Moreover, the two types of organizations will have access to different types of information — and will be positioned to take different steps to mitigate potential risks. For example, developers that design an AI system are well-positioned to have access to information about the type of data used to train the AI system, the system's known limitations, and its intended use cases. In contrast, a deployer using an AI system is well-positioned to have access to information regarding the specific ways in which it uses that system that impacts consumers. Any policies focused on supporting AI accountability should reflect these different roles and assign obligations accordingly.

Organizations may also take on other roles, such as integrating an existing AI model into the organization's products and services. Any obligations placed on these organizations should similarly reflect their role in integrating the AI system into the organization's products and services.

Avoid Prescriptive Transparency/Reporting Requirements

The Draft Proposal describes that designated advanced AI foundation model developers will be required to establish a system to implement seven obligations, including third-party vulnerability verification and public disclosure of AI capabilities and limitations. Under the requirement, the designated developers will also need to periodically report to the Government or a third-party institution (i.e., the AI Safety Institute) on the status of compliance with the obligations.

While we support the goal of the Draft Proposal to provide guardrails against associated risks through these obligations, it is important that regulations of foundation models are commensurate with the models' risks and capabilities. As such, foundation model developers should provide information about model capabilities, limitations, testing, and security along the AI value chain based on the level of risk involved. For high-risk uses of AI systems, we encourage robust testing and evaluation for safety, security, accuracy, and harmful bias. However, it is important to understand that existing technical standards for AI testing are nascent and should be developed consistent with longstanding voluntary, market-driven, and consensus-based approaches to standards development.

We encourage the use of watermarks or other disclosure methods for AI-generated content. However, we are concerned that broad reporting requirements could result in an inundation of documentation review for regulators and may require companies to disclose proprietary or confidential information related to development or use of AI systems.

Promote Effective Co-Regulation

The Draft Proposal describes the so-called "monitoring and review" framework as the core of co-regulation. Compared to mandatory intervening measures, co-regulation is considered a regulatory approach suitable for rapidly changing business environments to facilitate quick and flexible responses among stakeholders. On the other hand, if the co-regulation is applied arbitrarily, there is a risk that the dialogue may function as an excessive administrative burden and inadvertently stifle routine activities of business operators. As such, we recommend considering the following for the implementation of co-regulation:

- Verify carefully whether the involvement of the government is necessary to provide explanations to the public and respect the autonomy of each company. Do not require companies to implement specific initiatives that may not effectively control risks development or deployment for particular AI systems.
- Upon checking with or addressing inquiries to a Specified Advanced AI Foundation Model Developer, the supervising authority must provide sufficient explanation for why the requested information is relevant to the purpose of, and is in line with, the basic principles of the legislation.
- Specified Advanced AI Foundation Model Developers should be able to withhold information that may contain trade secrets. The supervising authority requesting information should handle any provided information as confidential.

Avoid Imposing External Safety Verification and Detection and Reporting of Vulnerabilities by Third Parties

The seven obligations in the Draft Proposal include conducting internal and external safety verification and using third parties for vulnerability detection and reporting. We agree that safety verification and detection are keys to identifying risks. However, we advise against suggesting that organizations always should conduct external testing. There are circumstances where an organization may elect to perform external testing. However, internal testing — which can be performed by a team of employees that is independent from the team tasked with developing an AI system — can identify and mitigate risks without creating concerns about sharing trade secrets, information that could jeopardize information or network security, and other proprietary information that will arise in external testing. As a result, we recommend focusing on internal testing and removing independent external testing in the obligation.

The Draft Proposal also includes detection and reporting of vulnerabilities by third parties, which may imply external audits. We caution against external audits, as current auditable standards for AI are not yet mature. There are few existing procedures or best practices for companies to either: (1) choose a reputable company capable of auditing an AI system, or (2) determine what standards any such auditing company should apply. Although organizations such as the International Organization for Standardization (ISO) have issued several AI-related standards, many remain under development by different bodies. Consequently, currently there is a lack of sufficient voluntary consensus-based standards addressing AI systems. Without common standards, the quality of audits will vary significantly because different audits may measure against different benchmarks, undermining the goal of obtaining an evaluation based on an objective benchmark. Furthermore, while BSA understands the need to promote transparency, we recommend not requiring business operators to publish audited results, which include confidential or proprietary information. This could disincentivize companies from voluntarily undertaking a rigorous review of their AI systems. For these reasons, external audits are not an appropriate solution to achieving transparency over AI governance, and we recommend removing this from the obligations.

Conclusion

BSA and our members look forward to working with the AI Project Team to support its goal of developing effective safeguards for AI. In addition to sharing this recommendation, we would appreciate the opportunity for a dialogue to better understand the intention of the Draft Proposal and discuss how we can further assist in the effort.