April 9, 2024

The Honorable Rebecca Bauer-Kahan
Assembly Privacy and Consumer Protection Committee
1020 N Street
Room 162
Sacramento CA 95814

Dear Chair Bauer-Kahan:

BSA | The Software Alliance appreciates the opportunity to share insights from the enterprise software sector on artificial intelligence (AI) generally and AB 3211 and AB 3050. BSA is the leading advocate for the global software industry.[1] BSA members are at the forefront of developing cutting edge services and their products are used by businesses of all sizes across every sector of the economy. AI is much more than robots, self-driving vehicles, or social media; it is used by companies large and small to create and improve the products and services they provide to consumers, to streamline their internal operations, and to enhance their capacity to make data-informed decisions. BSA members are on the leading edge of providing businesses-to-business tools that help companies leverage the remarkable benefits of AI.[2]

As leaders in the development of enterprise AI, BSA members have unique insights into the technology's tremendous potential to further spur digital transformation in the private and public sectors and the policies that can best support the responsible use of AI, especially high-risk uses of AI. BSA's views are informed by our recent experience with members developing BSA Framework to Build Trust in AI,[3] a risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle. Built on a vast body of research and informed by the experience of leading AI developers, the BSA Framework outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and highlights corresponding risk mitigation best practices. BSA's extensive experience has helped us identify effective policy solutions for addressing AI risks.

[1] BSA's members include: Adobe, Alteryx, Asana, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, Hubspot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Okta, Oracle, PagerDuty, Palo Alto Networks, Prokon, 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.
[3] *See* BSA | The Software Alliance, Confronting Bias: BSA's Framework to Build Trust in AI, *available at* https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

There has been tremendous innovation in AI, but it can also exacerbate risks of misinformation. Transparency about AI-generated content is key to ensuring responsible AI. When examining AI, we believe policymakers should focus on high-risk uses of AI, meaning AI that is used to make consequential decisions about an individual's eligibility for important services and benefits. In crafting legislation, policymakers should (1) clearly define the types of companies and types of decisions subject to the legislation, (2) require both developers and deployers of high-risk AI systems to conduct impact assessments, recognizing that the content of those assessments should be different based on whether a company is a developer or a deployer, and (3) require both AI developers and AI deployers to adopt risk management programs, to ensure companies have policies and personnel in place to identify and mitigate risks across the lifecycle of an AI system.

While we support the intention of promoting transparency in AI, AB 3211 and AB 3050 present several concerns, as described in more detail below.

I. AB 3211

Among other provisions, AB 3211 places several obligations on "generative AI system providers," meaning organizations or individuals that develop AI systems or substantial components thereof. The bill requires generative AI system providers to:

(1) place watermarks on content created by an AI system;

(2) provide downloadable software tools or online services available to all "large online platforms" and the public to determine whether content was created on the provider's system;

(3) conduct third-party testing of watermarks;

(4) only distribute generative AI systems that do not allow for removal of the system's watermarking functionality; and

(5) report AI system vulnerabilities and failures to the Department of Technology and certain other generative AI system providers within 24 hours of discovery and notify affected parties (reports to the Department of Technology would be made publicly available, unless a limited exception applies).

The bill also broadly prohibits companies from distributing products and services that have the capacity to remove watermarks from "synthetic content," meaning content significantly modified or generated by algorithms. The legislation's obligations on generative AI system providers present at least three distinct concerns:

First, the bill does not define "content," and "synthetic content" is defined broadly and includes text, meaning any content created by a generative AI system provider will require a watermark. Requirements to include watermarks or other disclosure methods for AI-generated content that consists of text are overly broad, since such a requirement would encompass all content generated by AI systems.

Second, the requirements for generative AI system providers to provide downloadable software tools or online services to determine whether content was created on the provider's system and to test watermarks through third-party experts are impractical and overly burdensome, particularly considering the bill's broad treatment of content.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Third, the notification system envisioned in the bill is unworkable. The bill does not define "vulnerability" or "failure" of an AI system, creating an extremely low threshold for triggering the legislation's reporting and notification requirements. Furthermore, the bill assumes generative AI system providers know all the parties affected by an AI system, which is often not the case, since a different company likely deploys the AI system. Also, reporting vulnerabilities and failures of AI systems, even if clearly defined, within 24 hours of discovery is unreasonable. Such a requirement would divert critical time and efforts from mitigating vulnerabilities or failures to fulfilling reporting and notification responsibilities. These concerns are compounded by the bill's requirement that reports to the Department of Technology be made public, unless a limited exception applies. Not only could these reports include confidential information (including trade secrets and IP-protected information), but they may also create roadmaps for bad actors to exploit known vulnerabilities.

II. AB 3050

AB 3050 prohibits "AI-generating entit[ies]" from creating "covered AI-generated materials" unless the material includes a watermark that meets the Department of Technology's standards, among other requirements. This obligation presents at least two distinct concerns:

First, "AI-generated material" is defined to include text, meaning virtually any AI-generated content that meets the definition of "covered AI-generated material" will require a watermark.

Second, granting the Department of Technology the authority to issue watermarking standards creates uncertainty for businesses (many of which are already working toward consensus-based standards) and consumers (who are likely to be confused if a range of overlapping and potentially conflicting standards emerge to address the same issues). We recommend the bill be revised to encourage consensus-based, industry-led standards, such as those developed by the Content Authenticity Initiative (CAI).

III. BSA's Principles for Promoting Transparency in AI

Our recommendations focus on the aspects of these bills addressing transparency and consumer-facing disclosures. BSA supports the following principles for promoting transparency in AI technologies.

a. Encouraging the Use of Watermarks or Other Disclosure Methods for AI-Generated Content

Disclosures, including watermarks, can help consumers tell whether content is human- or AI-generated. This can be helpful in preventing misinformation. Encouraging the use of watermarks or other disclosure methods for AI-generated content can help address this concern. Governments can play an important role in promoting transparency in AI technologies, including through watermarks or other disclosures.

b. Promoting the Coalition for Content Provenance and Authenticity Standard

BSA supports the CAI's efforts to promote the open Coalition for Content Provenance and Authenticity standard for content authenticity and provenance. This standard will help consumers decide what content is trustworthy and promote transparency around the use of AI. In conjunction with watermarking, the CAI approach provides secure, indelible provenance. BSA recommends California promote the CAI's efforts, which are consensus-based and built on industry best

practices.

c. Disclosing When Consumers Are Interacting with AI

Consumers should know when they are interacting with AI depending on the circumstances and context of use. For example, chatbots should disclose that consumers are interacting with AI instead of a human. This type of transparency is important in building trust in AI systems, and in educating consumers about a company's use of AI.

* * *

Thank you for allowing us to provide the enterprise software sector's perspective. We welcome the opportunity to serve as a resource and further engage with you or a member of your staff on these important issues.

Sincerely,

Meghan Pensyl
Director, Policy

cc: Assembly Privacy and Consumer Protection Committee

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org