



BSA | The Software Alliance Comments on the New York City Department of Consumer and Worker Protection's Proposed Rules Implementing Local Law 144 of 2021 Regarding Automated Employment Decision Tools

October 24, 2022

BSA | The Software Alliance appreciates the opportunity to submit comments on the proposed regulations implementing Local Law 144 of 2021—New York City's ordinance on automated employment decision tools ("Ordinance"). BSA is the leading advocate for the global software industry before governments and in the international marketplace. Our members are enterprise software companies that create the technology products and services that power other businesses.¹ They provide trusted tools that help other businesses innovate and grow, including cloud storage services, customer relationship management software, human resources management programs, identity management services, and collaboration software.

BSA members are on the leading edge of providing businesses — in every sector of the economy — with innovative software services, including Artificial Intelligence (AI). As leaders in the development of enterprise AI systems, BSA members have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

While the adoption of AI can unquestionably be a force for good, we also recognize the significant risks to society if this technology is not developed and deployed responsibly. BSA agrees that when AI is used in ways that could unlawfully discriminate, the public should be assured that such systems have been thoroughly vetted to identify and mitigate risks associated with unintended bias.

A. BSA's Framework to Build Trust in AI

This has been an area of particular focus for BSA and our member companies over the last several years. In June 2021, BSA released *Confronting Bias: BSA's Framework to Build Trust in AI*, which sets forth a risk management approach for confronting concerns about bias.² The Framework was informed by leading AI developers and outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and corresponding best practices for mitigating those risks. AI is used in so many different contexts that only a flexible, risk management approach will be successful. There are a variety of AI

¹ BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, CrowdStrike, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Intel, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Prokon, PTC, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

² *Confronting Bias: BSA's Framework to Build Trust in AI*, available at <https://ai.bsa.org/confronting-bias-bsas-framework-to-build-trust-in-ai>.

development and deployment models, and the Framework recognizes that the appropriate allocation of risk management responsibilities will vary depending on the type of system, including who develops the algorithm, trains the model, and ultimately deploys the system.

BSA supports the overarching goal of the proposed regulations, which is to ensure high-risk uses of AI are subject to safeguards. One crucial safeguard that promotes responsible uses of AI systems is ensuring that companies that develop or use high-risk AI systems establish a comprehensive approach for performing impact assessments. Impact assessments are widely used in a range of other fields—from environmental protection to data protection—as an accountability mechanism that promotes trust by demonstrating that a system has been designed in a manner that accounts for the potential risks it may pose to the public. The purpose of an impact assessment is to establish organizational processes to guide the development and use of high-risk systems by requiring internal stakeholders to identify the risks that a system may pose, quantify the degree of harm the system could generate, and document any steps that have been taken to mitigate those risks to an acceptable level. By establishing a process for personnel to document key design choices and their underlying rationale, impact assessments are an important transparency and accountability mechanism.

B. The Department of Consumer and Worker Protection’s Proposed Regulations

The proposed regulations address the need for safeguards that promote responsible uses of AI by focusing on how companies are to implement bias audits under the Ordinance. As the proposed regulations recognize, one critical aspect of operationalizing such audits is identifying a set of individuals who can conduct them. Importantly, the proposed regulations appropriately recognize that internal stakeholders who are not involved in the development or use of an automated employment decision tool are competent to conduct a bias audit. Acknowledging the independence of internal stakeholders will incentivize companies to implement multiple layers of independent review, which will enhance trust in the use of these systems and create safeguards that function in practice.

We have recommendations, however, about three aspects of the proposed regulations:

First, we suggest deleting “or modify” in the definition of “automated employment decision tool” in Section 5-300. The definition refers to relying “solely” on simplified outputs “with no other factors considered” and use of simplified outputs to “override” conclusions derived from other factors including human decision-making. As illustrated by these descriptions, the intent of this definition appears to be identifying circumstances that could “substantially assist or replace discretionary decision making.” This is a relatively high threshold. Outputs that do not rise to this level, such as merely contributing to or modifying, should be excluded from the scope of this defined term.

Second, we suggest adding a notation explaining that the examples in Section 5-301 for the bias audit are not exhaustive of all scenarios. Along the same lines, we recommend revising the two examples in this section to recognize a broader range of possible scenarios regarding the appropriate entity to conduct a bias audit. There are myriad development and deployment models, and multiple stakeholders may play varying roles depending on the circumstances. Specifically, we recommend offering an illustrative example without inadvertently suggesting that the types of entities described in the example will always have the same roles.

Further, in the examples for both subsections (a) and (b) of this section, we recommend replacing “historical data” with “test data” because there may not always be historical data if

the organization has not previously used the technology. In addition, historical data may not be available in many scenarios—such as when that data is subject to privacy or security safeguards that prevent its dissemination. In addition, we recommend deleting “planned” from the example in subsection (b), as it inserts ambiguity regarding the timing of the use of the automated employment decision tool.

Third, we recommend omitting the requirement to publish the selection rates and impact ratios for all categories and instead require a summary statement on adverse impact. Although Section 20-871(2) of the Ordinance requires a “summary of the results” of the bias audit to be published, it does not call for the level of specificity contemplated by the proposed regulations. Publishing the specific information contemplated by the proposed regulations could inadvertently undermine the goals of the Ordinance. For example, it may discourage applicants from groups that are selected less frequently from applying to an organization at all, hampering efforts to attract a diverse workforce. Moreover, requiring the public disclosure of such specific information could disincentivize companies from conducting thorough audits to avoid possible results that may not be optimal. Accordingly, we recommend striking “the selection rates and impact ratios for all categories” in Section 5-302 and replacing it with “a statement on adverse impact.” We further recommend aligning the categories for the selection rates and impact ratios with the EEOC’s approach for disparate impact testing.

We thank you for the opportunity to comment on the proposed regulations and look forward to serving as a resource as you finalize the proposed regulations.