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

January 23, 2023

BSA | The Software Alliance appreciates the opportunity to submit comments on the revised proposed regulations implementing Local Law 144 of 2021 — New York City's ordinance on automated employment decision tools ("Ordinance").[1] BSA is the leading advocate for the global software industry. Our members are enterprise software companies that are on the leading edge of providing businesses — in every sector of the economy — with innovative technology products and services, including those powered by artificial intelligence (AI).[2] Our members 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. 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.

The success of AI products and services will be based on public trust and confidence in these technologies. To earn that trust, organizations that develop AI, and those that use AI, must do so responsibly and in a manner that accounts for the unique opportunities and risks the technology poses. Policymakers can enhance public confidence and trust by establishing a legal and regulatory environment that supports responsible innovation, including by creating safeguards that help prevent unlawful discrimination. BSA supports the overarching goal of the proposed regulations to ensure that AI systems have been thoroughly vetted to identify and mitigate risks associated with unintended bias. However, we have recommendations on four aspects of the proposed regulations — the independent auditor requirement, publication of bias audit results, the definition of automated employment decision tool, and the use of historical data — to better achieve this objective.

[1] BSA's comments on the initial set of proposed rules is available at https://www.bsa.org/files/policy-filings/10242022nyctools.pdf.

[2] BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, CrowdStrike, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Juniper Networks, 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.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

A. Independent Auditor

As the proposed regulations recognize, one critical aspect of operationalizing bias audits is identifying a set of individuals who can conduct them. Importantly, the earlier version of the proposed regulations appeared to recognize that internal personnel who are not involved in the development or use of an automated employment decision tool are competent to responsibly conduct a bias audit. Acknowledging the independence of internal stakeholders would incentivize companies to implement multiple layers of independent review as they develop and test their products, which would enhance trust in the use of these systems and create safeguards that function in practice.

The revisions to the proposed rules eliminate this flexibility, clarifying that no employees are considered independent and, therefore, may not serve as auditors. We support the aim of the bias audit but respectfully recommend an alternative approach for two reasons.

First, there is currently no consensus on AI auditing standards. Unlike other areas, such as privacy and cybersecurity, where standards underpin notable certifications, AI lacks this essential element to create trust and accountability in the marketplace. Standards bodies, such as the International Organization for Standardization, are only in the early phases of their AI projects. In addition, recognizing the lack of auditing tools to detect bias and discrimination, Stanford University recently launched an AI audit challenge to help with this effort. A Stanford professor, who is also the President and CEO of Sagewood Global Strategies, a technology policy and risk advisory firm, and the former Senior Director for Cyber Policy on the White House National Security Council, recently examined this issue and acknowledged the lack of auditable criteria, concluding that "[t]he AI audit ecosystem is immature, at best."[3]

Although the proposed rules specify how selection rates and impact ratios should be calculated, this guidance does not replace the need for broader universal, consensus-based standards developed in a multistakeholder process. Nor does it address the lack of available bias detection tools. Without common standards, companies can shop around for auditors based on their preferences for particular methods, criteria, and scope. In addition, the quality of audits will vary significantly and are likely to correlate with price, undermining efforts to establish common objective benchmarks.

Second, there are no professional organizations to govern or train third-party auditors for AI systems. Auditors typically have professional bodies that create baseline criteria and maintain ethical guidelines. SOC audits, for example, are conducted by CPAs and governed by the American Institute of Certified Public Accountants. In addition, educational bodies are in place in other fields — such as privacy — to train professionals. No such body exists for AI auditors.

In short, the AI auditing landscape is nascent, lacking common developed standards and professional oversight bodies. For these reasons, requiring third-party audits is not a feasible or optimal approach. Instead, we recommend that the Department of Consumer and Worker Protection (DCWP) reinstate the earlier definition of "independent auditor," which would have

[3] Andrew Grotto, Sagewood Global Strategies LLC, *Audit of AI Systems: Overview, Current Status, and Future Prospects*, 5 (Nov. 2022), *available at* https://www.regulations.gov/comment/FTC-2022-0053-0906.

allowed internal personnel to conduct the audit so long as they were not involved in the development of the AI system.

B. Bias Audit

As we highlighted in our initial comments, we also recommend omitting the requirement to publish the selection rates and impact ratios for all categories and instead require a summary statement on adverse impact. As an initial matter, the definition of impact ratio includes either selection rates or scoring rates, so there is no need to publish both the impact ratio and the selection rate. In addition, although Section 20-871(a)(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 required 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 a more flexible approach to strike "the selection rates and impact ratios for all categories" in Section 5-303 and replace it with "a statement on adverse impact."

We also note that it is unclear how employers should account for applicants who do not disclose their race, ethnicity, or gender. For example, if an employer has 1,000 applicants, but only 750 of them disclose those fields, the result of the bias audit will be statistically meaningless. Because the selection rate calculation requires an evaluation of the total number of applicants, the disparate impact result from the 750 people who report those fields will be inaccurate and incomplete. A selection rate calculation that excludes the additional 250 people who did not report those demographic fields would also render inaccurate results, as it doesn't reflect the racial and gender diversity that could exist in 25% of the applicant pool. As a result, a requirement that employers publish the results of these calculations on its website could lead to misleading or inaccurate information.

We further recommend aligning the categories for the selection rates and impact ratios with the EEOC's approach for disparate impact testing. The regulations contemplate that race, ethnicity, and sex will be evaluated both separately and intersectionally. However, the US Equal Employment Opportunity Commission (EEOC) does not require intersectional testing. Accordingly, we recommend that you revise the proposed regulations to require that these demographic categories only be tested separately, consistent with the well-established EEOC approach.

C. Definition of Automated Employment Decision Tool

We welcome the proposed regulations' revision to the definition of automated employment decision tool, which strikes "or modify." This change appropriately focuses the definition on circumstances that overrule human decision making.

The proposed regulations would benefit from clarification on one additional point relating to this definition. If an automated employment decision tool produces a simplified output, but there is functionality in the tool to override or disregard that score (e.g., the user has the ability to change the score or select candidates independent of the score), it is unclear whether the company needs to produce artifacts to show that its automated employment decision tool is not in scope because users are not solely relying on the simplified output. We

interpret the rules' silence on this issue as not requiring the production of such artifacts, but we would appreciate further clarifications in the proposed regulations.

D. Use of Historical Data to Conduct Bias Audits

The requirement in Section 5-302 to use historical data to conduct bias audits poses many practical challenges and raises additional concerns regarding the data privacy of employees and job candidates. First, it is unclear what circumstances render historical data unavailable, triggering the option to use test data. Second, the historical data may reside with multiple employers, and vendors may be legally and contractually prevented from collecting or accessing that data. Third, as written, employers would be required to provide sensitive employee and candidate information to third-party auditors. Yet DCWP's rules do not impose limitations on how that data is handled by these third-party firms. Indeed, third parties may conceivably reuse this sensitive personal information for commercial uses or even sell it to other third parties. For these reasons, we recommend that: (1) employers should not be required to use historical data to conduct bias audits; (2) the use of test data for vendor-initiated audits should be the default rule; and (3) employers should not be required to provide sensitive personal information of their employees and job candidates to third parties.

Finally, we note that the current enforcement date is April 15, 2023. However, without final regulations, organizations cannot undertake all the efforts necessary to comply. We urge you to allow sufficient time for compliance once the regulations have been finalized and, if necessary, postpone the enforcement date.

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

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org