Hire★Vue

October 24, 2022

New York City Department of
Consumer and Worker Protection
42 Broadway New York, NY 10004
Submitted online to https://rules.cityofnewyork.us/

Re: HireVue Comments on New York City Local Law 144 of 2021

To Whom It May Concern:

HireVue is pleased to submit comments to New York City Local Law 144 of 2021 and the proposed Rules issued on September 23, 2022 ("the Law") regarding the use of automated employment decisions tools, including the use of Artificial Intelligence ("AI") in hiring.

HireVue is a global software technology provider that has hosted more than 30 million interviews and 200 million chat-based candidate engagements for over 800 customers around the globe. At HireVue our mission is to change lives by connecting talent and opportunity. Our approach is to use science, technology and best practice to enable a hiring process that is more fair, inclusive, and equitable, and that allows candidates to showcase their potential.

HireVue technologies support both the candidate and employer interviewing experience in and around the New York City area, in a broad spectrum of industries. HireVue's comments on the Law are based on our extensive experience with the use of AI in the context of hiring.

HireVue knows that human decision making in hiring is not immune to bias - ranging from the existence of overt prejudice to subconscious bias. HireVue's commitment to mitigating human bias in hiring has been rooted from the beginning in its use of scientifically-validated methods implemented by our team of advanced degree industrial and organizational (IO) psychologists and data scientists. For example, HireVue's interviewing technology (with or without AI) uses standardized and structured interviews, i.e., the same set of questions in the same order for every candidate applying for a given role. In addition to being a best practice approach in candidate selection, this greatly enhances the candidate experience and ensures each candidate has the same opportunity to demonstrate their competencies for the job.

When HireVue first began developing, designing, and testing our algorithmic assessment solutions, we referenced over a century of study of scientifically-validated methodologies, decades of experience, and well-established Federal guidelines, best practices, and employment laws. HireVue's algorithmic technology evaluates candidates' job-related skills, without considering criteria often used by humans that are unrelated to job success, such as how well dressed one is, which university one attended, which keywords were packed into their application, demographic information, or educational background. Due to our focus on only knowledge and skills relevant to success in a specific job role, all candidates are evaluated in a consistent and objective manner

using HireVue's technology. Furthermore, we develop our technology to minimize bias at every design step - from sampling representative training data, including diverse subject matter experts in the development of our interview questions, to using state-of-the-art bias mitigation algorithms to mitigate bias from the technology's predictions during the design and development stage, before any applicant is assessed.

HireVue's ongoing commitments to transparency, accountability, continual improvement and engaging with stakeholders drives more equitable hiring to offer a faster, fairer and more flexible experience for candidates and employers alike. HireVue has accomplished this through its attention to the following key principles, underpinned by best practices:

1. **Open and transparent approach:** HireVue believes the only thing that should be evaluated during the recruitment process are a candidates job-related competencies in relation to knowledge, skills, and abilities associated with the particular job they are being considered for. Our tools and services use deep industrial and organizational (I/O) psychology expertise and science-backed methods to increase the speed at which employers can screen and hire candidates and which candidates can display their skills. At its core, HireVue's standardized and structured interview (with or without AI) reflects scientifically validated best practices. When structured interviews are used in conjunction with our AI-driven assessments, they are designed to evaluate candidates' job-related skills, without taking into account criteria proven *not* to determine job success, such as demographic information or educational background. Our AI-driven assessments do not replace humans - HireVue simply helps recruiters and talent acquisition teams do a better job, faster by focusing their efforts on candidates exhibiting job-related competencies. Most recently, rather than ignore what has been described as the "black box" nature of AI, HireVue published our industry's first AI Explainability Statement to provide transparency into the methods of designing, developing and training our algorithms.[1]

2. **Continuous product improvement**: HireVue tracks innovative research, in addition to conducting its own academic research and applies advancements to ensure our technology continually meets or exceeds industry best practices. As evidence of our constant drive for improvement, in 2020 we made the decision to remove all video inputs from our algorithms in favor of solely using Natural Language Processing (NLP) after research and stakeholder feedback indicated that improvements in the NLP technology provided more predictive power and that the non-verbal data features did not add incremental accuracy to our assessments. More recently, a member of HireVue's team published ground-breaking research[2] that found no differences in HireVue game-based assessment results of candidates on the autism spectrum vs. the general population, whereby traditional human hiring approaches might overlook or, worse, reject these candidates. In addition, HireVue's core development process includes bias mitigation, testing and feedback at the design stage

---

[1] https://www.hirevue.com/why-hirevue/ai-explainability-statement

[2] https://www.hirevue.com/blog/hiring/hiring-in-a-neurodiverse-world-autism-and-the-science-of-game-based-assessments

and ongoing review at the deployment stage. This is in addition to the fact that our AI is locked when deployed and not dynamically learning with un-curated data. Lastly, HireVue has also voluntarily submitted its design methods for review through the first-of-their-kind independent 3rd party audits and took the additional step to make the results of those audits publicly available. Our work with O'Neil Risk Consulting and Algorithmic Auditing (ORCAA) on an algorithmic audit is available on our website[3] and Landers Science Workforce LLC's evaluation of our application of IO Psychology to candidate selection can also be reviewed online.[4]

3. **Engaging with our stakeholders:** Beyond leveraging the diverse expertise of a multidisciplinary workforce to design and build our products, we regularly engage with various stakeholders to solicit input and feedback on our research, product development, and bias-mitigation practices. HireVue actively engages with its candidates, its customers, and with third-party organizations like Integrate Autism Employment Advisors, representing neurodiverse candidates to learn about their experience in using our technology. In addition, HireVue's expert advisory board consists of outside experts in relevant fields of IO psychology, legal, and AI ethics. Lastly, in 2019, HireVue established ethical AI principles and we continue to engage with stakeholders through our advisory boards to understand the challenges of both public and private employers.

These guiding principles have allowed HireVue to establish its industry-leading best practices around AI and uphold principles of minimizing discrimination in hiring. In alignment with the foregoing principles, HireVue supports well-designed legislation to address the risks of AI if it is not designed and deployed responsibly. HireVue gleaned significant learnings after subjecting our technology to third-party independent audits and is eager to share its insights on how to craft these in a meaningful way to accomplish the goals the Law is seeking to meet.

Specifically with respect to the Law, HireVue questions severe restrictions on good practices and beneficial innovation that pragmatically broaden access to employment in the name of attempting to mitigate algorithmic bias. The Law's "one size fits all" approach, i.e., failing to differentiate use cases, industries, type of AI (learning vs static), or *what* can reasonably or feasibly be audited in each development stage of an algorithm (design vs. deployment stage) will not necessarily effectuate what appears to be the intent of the Law.

We also call to the attention of the agency that some requirements of the Law may not be reasonably satisfied in the employment context, such as seemingly requiring an audit of the AEDT at the deployment stage when no live-candidate data would be available to audit prior to use, running analysis on data (race, gender, marital status, etc.) that an employer cannot require due to pre-existing employment laws and services providers may not naturally have access to due to data minimization concepts in various privacy laws.

---

[3] https://www.hirevue.com/blog/hiring/industry-leadership-new-audit-results-and-decision-on-visual-analysis

[4] https://www.hirevue.com/blog/hiring/independent-audit-affirms-the-scientific-foundation-of-hirevue-assessments

Lastly, with respect to any algorithm audit requirement, a thoughtful audit and any legislation which requires it should leave room for continuous development and improvement to build on the "good" and to identify *and* address concerning results. Based on its experience HireVue offers the following points for the council's consideration with respect to algorithmic audits:

1. **Audit criteria must be clearly defined.** Much like audit standards in industries like privacy, finance, etc., an audit of an AI tool should include reference to the relevant industry and legal standards against which the tool is tested. It should explain how a model works, its purposes and its limitations, and the data it relies upon to make decisions. HireVue follows, and in many ways exceeds, previously established employment frameworks such as the EEOC's Uniform Guidelines on Employee Selection Procedures (1978)[5] to statistically test and mitigate for adverse impact in our lab models during the design and development phase. We also perform adverse impact checks (at a minimum annually) for deployed models and adjust to mitigating bias if it is observed in use with a particular population. The most recent guidance for the Law seems to be at odds with well-established best practices regarding the process and methods of how to perform a disparate impact analysis and fails to require a specific sample size threshold to ensure statistical validity. In addition, the requirement to publish results does not come with any consideration as to the intended audience; most laymen do not have sufficient knowledge to interpret statistical tables and do not know what meets federal guidelines or is legally considered "fair" in the context of employment law.

2. **Audit at the design and development phase.** The focus of the audit should be at the outset of product development to ensure algorithmic tools are designed, developed, trained and tested – including steps to identify and mitigate bias – *before* deployment. The algorithm should also be periodically monitored after deployment to identify any unexpected results. While the Law expressly applies to employers, the rulemaking creates confusion on what obligations a service providers has and fails to distinguish between audits performed at the development vs deployment stage, particularly the AEDT has not been previously been used and there are no selection rates for employers to audit.

3. **HireVue, like other vendors, should be responsible for delivering an independent audit on the AI-based products we provide to our customers.** Vendors will differ in how their tools are developed and what sort of data they use, thus the way audits are conducted will not be universal. Audits must always consider the industry and context of where and how the AI is being used. In HireVue's case, it is against the EEOC Guidelines to deploy algorithms that treat job applicants differently from "day-to-day", thus we have chosen to only deploy static algorithmic models (after auditing them and testing against established frameworks in hiring). This means our algorithms are "locked" and do *not* learn or continually change from real-time uncurated and unfiltered customer data - as this would be unfair to the job applicants. This approach prevents the risk of bias unexpectedly creeping into our pre-tested models. We also suggest that Customers using HireVue's audited algorithmic tool should be able to rely on those audit results without needing to

---

[5] https://www.eeoc.gov/laws/guidance/employment-tests-and-selection-procedures

conduct a third-party audit of the deployed model–though the model should continue to be periodically monitored to validate the deployment and use of an AI tool in their particular setting, e.g. with live candidate data that the employer would usually (but not always) collect. This distinction and concept fail to be captured in the current Law.

Ongoing dialogue between appropriate stakeholders is the key to creating sensible legislation that protects candidates, companies, and innovation. HireVue welcomes legislation related to the potential impact of all hiring technologies as it encourages transparency and improves the process and fairness for all candidates. Legislation like this demands that all vendors meet the high standards that we have supported since the beginning.

Should the Council have questions about this letter or our business, HireVue is absolutely ready and willing to engage in a dialogue with the Council and happy to share our perspective as the leader in responsible hiring solutions.

Sincerely,

Anthony Reynolds
Anthony Reynolds
Chief Executive Officer

Lindsey Zuloaga
Lindsey Zuloaga
Chief Data Scientist

Nathan Mondragon
Nathan Mondragon
Chief I/O Psychologist