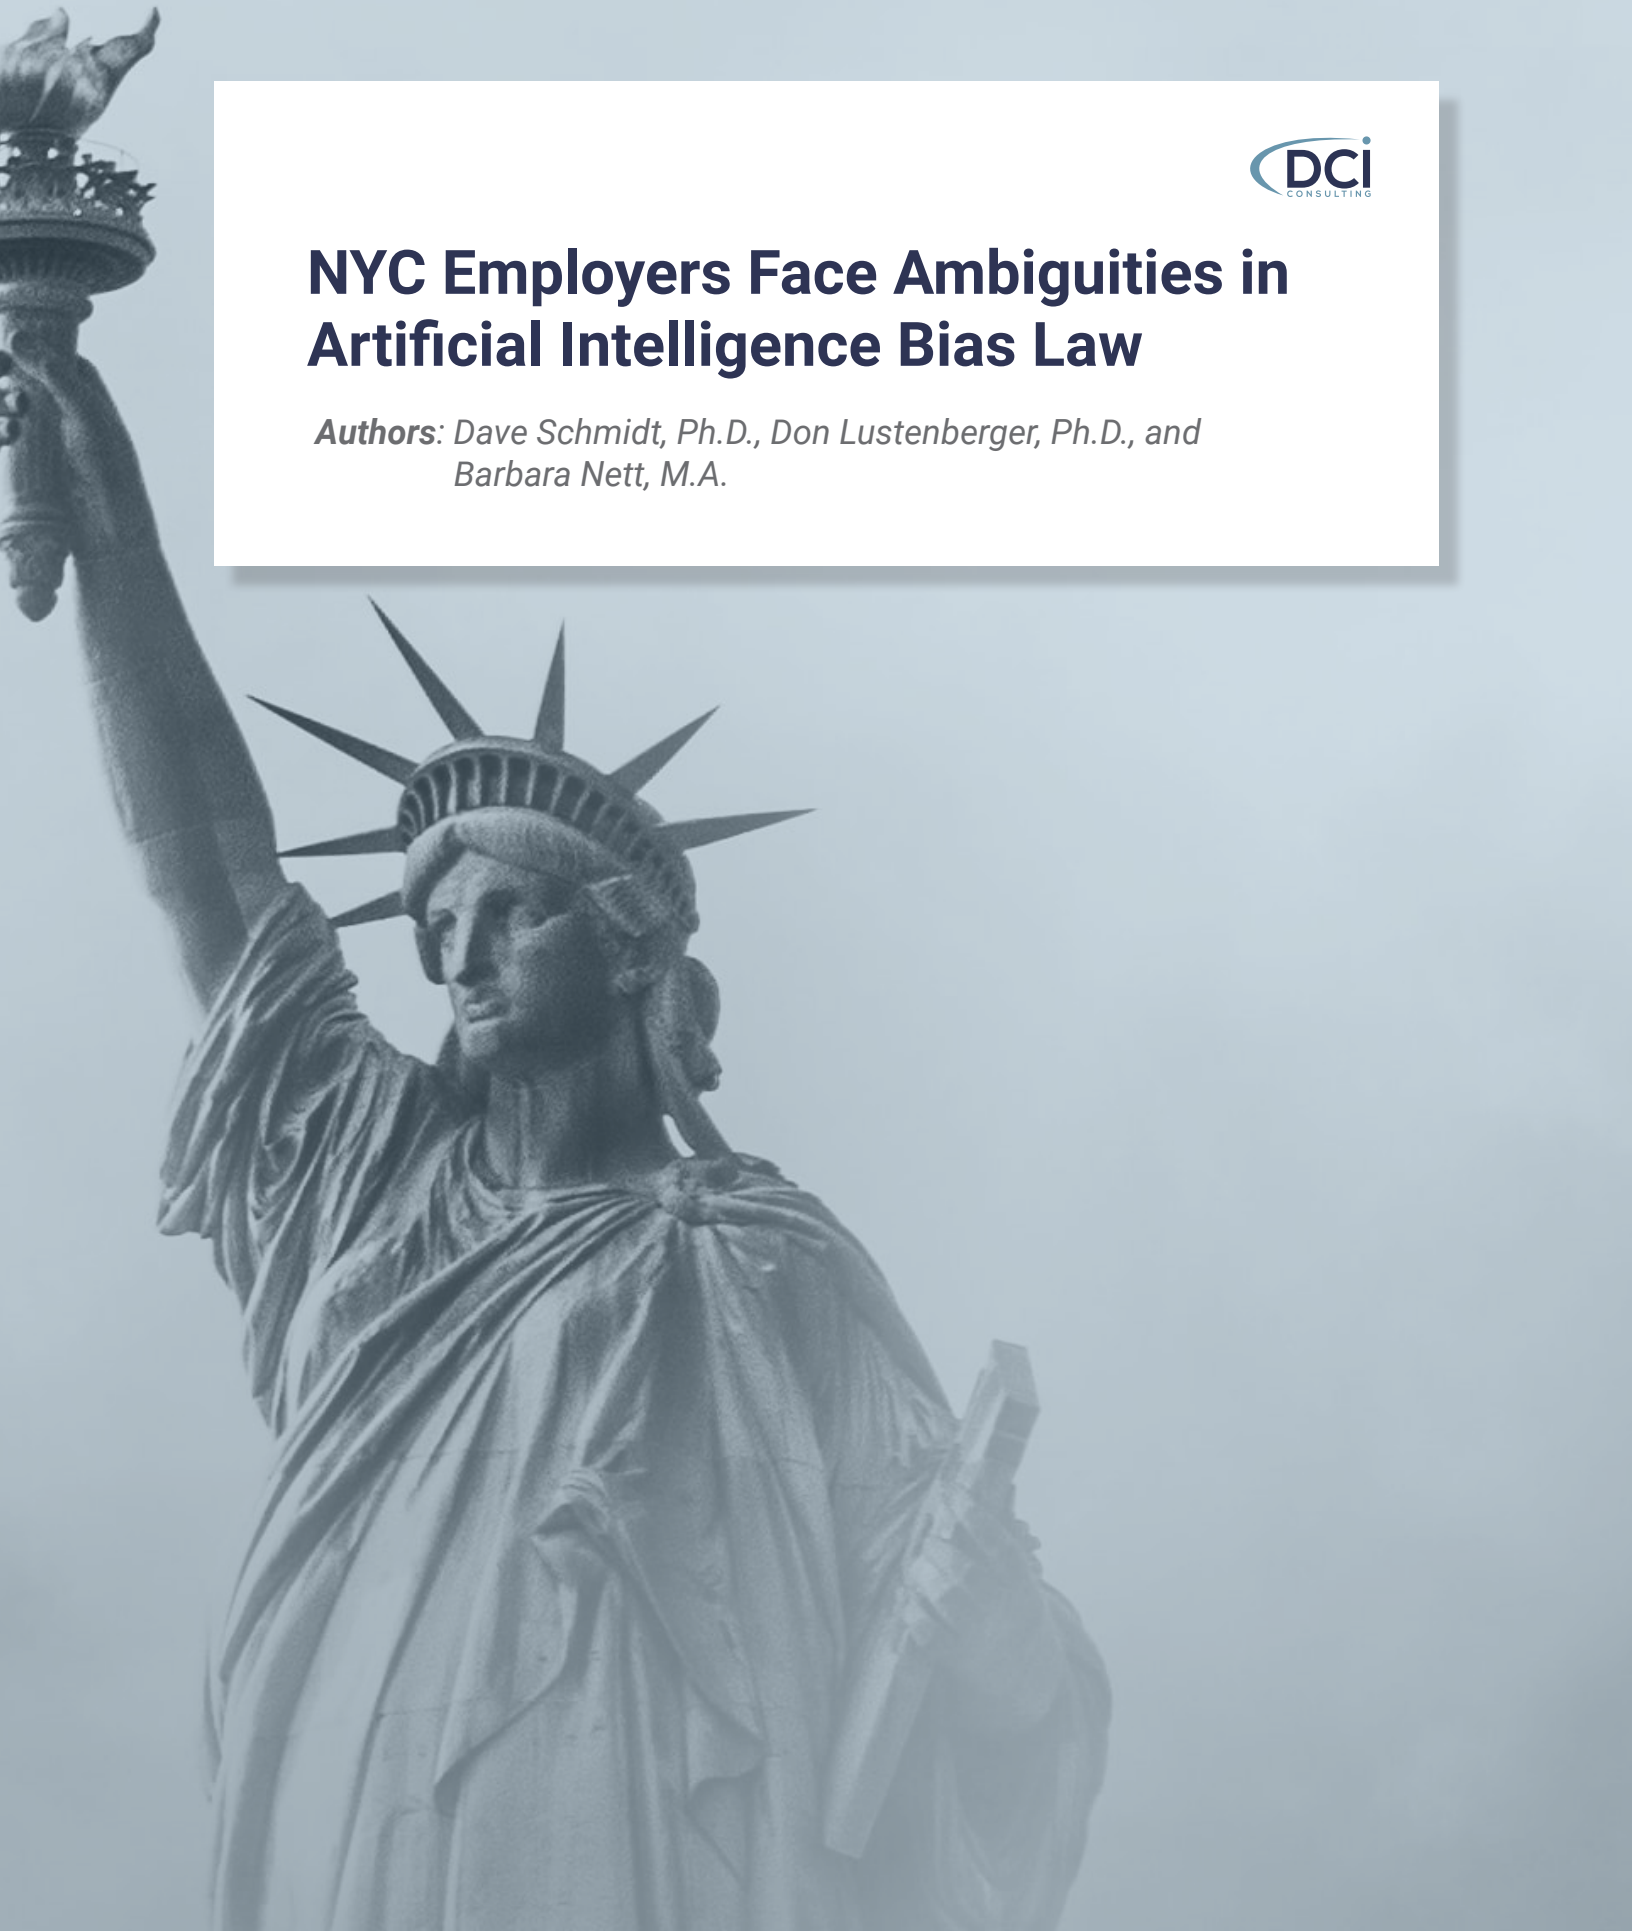

NYC Employers Face Ambiguities in Artificial Intelligence Bias Law

Authors: *Dave Schmidt, Ph.D., Don Lustenberger, Ph.D., and Barbara Nett, M.A.*

NYC Employers Face Ambiguities in Artificial Intelligence Bias Law

Employers have less than five months to prepare for the Jan. 1, 2023, enactment of New York City's Local Law 144 of 2021,[1] which will regulate the use of automated employment decision tools, or AEDTs, in employee selection.

Public commentary solicited by the New York City Department of Consumer and Worker Protection in June of this year has underscored numerous ambiguities in the law. In what follows, we will share several outstanding questions about the law and highlight critical areas of ambiguity that we hope will be addressed by the city.

The New York City AEDT law, one of the broadest in this area passed to date, will require employers using selection tools that leverage artificial intelligence, machine learning, natural language processing or other automated algorithms to:

- Conduct an independent bias audit of each AEDT;
- Make the audit results publicly available on its website; and
- Notify candidates that the selection process includes an AEDT, what job qualifications and characteristics are being evaluated by the AEDT and that they are allowed to request an alternative selection process or accommodation.

Employers failing to meet audit or notice requirements are subject to fines — $500 for the first violation and $500 to $1,500 for each subsequent violation.

Who is and is not covered by the law?

Public commentary has highlighted confusion about one of the more glaring issues with the law. Which employers are subject to its regulations? This confusion is in part sown by references to both city residents and employers.

For instance, the law states that, "the term 'employment decision' means to screen candidates for employment or employees for promotion within the city,"[2] suggesting that it is applicable to employers with operations in the city. In the notices section of the law however, it states that it applies to a job candidate who "resides in the city."[3]

For many, there are multiple ways to interpret the law, as depicted by Figure 1. The authors interpret the wording to refer to the intersection of the two circles.

However, given the impact of this interpretation on the scope of the law, clarification from the city would be beneficial, and should be relatively easy to provide.

In absence of clarification, any employer with both operations and candidates in the city will definitely be required to comply and should prepare accordingly. Other employers with applicants who reside in the city may want to prepare just in case.

Figure 1. *Is the law applicable to the left circle, the right circle, or the intersection of the two circles?*

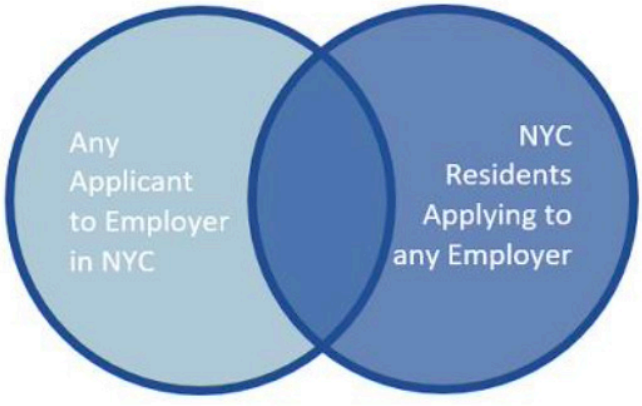

What selection practices are, or might be, considered automated employment decision tools?

Further related to scope, the definition provided by the law is quite broad and may encompass a much wider range of tools than intended:

The term "automated employment decision tool" means any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including a score, classification, or recommendation, that is used to substantially assist or replace discretionary decision making for making employment decisions that impact natural persons.[4]

This clearly covers tools using artificial intelligence methods; however, it may also extend to traditional selection tools — e.g., online multiple-choice tests, automatically scored application blanks — where statistical modeling or data analytics are applied to configure use for a particular employer or job.

Further, many selection tools are formulaically scored and are used to assist in making decisions about which candidates progress in the selection process. These tools have operated this way for decades, and it is unclear whether the intent of the law was to regulate these tools in addition to those leveraging more recent, sophisticated data-science approaches.

We hope that the city better defines these boundaries. Notwithstanding, it would be prudent for employers to ensure that any selection tools using an algorithm to generate scores, or to produce information considered in decision making, has been properly validated and is routinely monitored, in accordance with the Uniform Guidelines on Employee Selection Procedures.[5]

What qualifies as a bias audit?

Under the new law, employers must conduct a bias audit that minimally "assess[es] the tool's disparate impact" by race, ethnicity and sex.[6] While this sounds simple, there are complicating nuances that make this challenging.

First, disparate impact is typically examined within the context of job relevance — i.e., validation evidence — and the availability of suitable alternative selection procedures, as outlined by the Uniform Guidelines. Curiously, the law mentions only half of the equation by evaluating disparate impact absent job relatedness considerations.

Second, there are many methods and statistics available to evaluate disparate impact — results from different statistics may suggest similar or different conclusions.

Best practices are to consider multiple statistical indicators and examine the pattern of results in formulating conclusions.[7] These conclusions are rarely as simple as yes, there is disparate impact, or no, there is not.

Third, there is no indication in the law of whether these disparate impact analyses need to be conducted separately by requisition, job, or job family or separately by location.

Nor does the law indicate the time period upon which to base the disparate impact analyses — e.g., one month, one year. We simply know that an audit must be conducted for each AEDT that includes an analysis of disparate impact.

Fourth, it is not uncommon for employers to implement a new assessment using data gathered from nonlocal sources prior to having sufficient local applicant data for disparate impact analyses.

The law is not at all clear on whether existing evidence of subgroup differences for an AEDT collected from other organizations could stand in for audit purposes until local, applicant data are amassed. A tool cannot be evaluated for local disparate impact unless it is implemented in some way in that organization, even if as a pilot.

With such little guidance provided in this area, it is difficult for employers to know what the city will find acceptable. Perhaps they are leaving what qualifies as a sufficient audit up to the independent auditors themselves.

Who qualifies as an independent auditor?

There is also very little clarity in the law on who qualifies as an independent auditor to conduct the bias audits. The law states that the bias audit should be "an impartial evaluation by an independent auditor."[8]

It is unlikely that internal audits will be sufficient to meet this requirement. It is also unlikely that a vendor can be considered an independent auditor of an AEDT it develops, markets and sells. Beyond that, it is unclear what constitutes a qualified independent auditor.

Auditors might range from those who are relatively inexperienced and inexpensive to those with extensive knowledge and expertise in AEDTs, third-party reviews, validation research, disparate impact analyses and compliance with the uniform guidelines.

Will the city certify individuals or firms as auditors? That seems unlikely. Thus, employers will want to carefully consider the qualifications of any potential auditors.

What audit results must employers make available?

In addition to having a bias audit conducted by an independent auditor, the law states that a summary of the audit results must be "made publicly available on the website of the employer." [9]

However, there is no guidance on the form of this audit summary, nor the level of detail required. Summaries could range from a simple sentence concluding disparate impact does or does not exist to a multipage summary describing contextual factors, methods and statistics used, disparate impact results on a continuum, evidence of job-relevance and so on.

Guidance or examples from the city would be beneficial to help employers understand what will satisfy this requirement.

Who is entitled, under the notices section, to request an alternative process or accommodation?

In stipulating that employers must provide information about the AEDT to candidates, the law casually states that candidates may "request an alternative selection process or accommodation." [10]

Consideration of accommodations for individuals with disabilities is a standard step in the assessment design phase.

Federal law, of course, entitles individuals with disabilities to request accommodation to a selection procedure that allows for a more accurate measure of their capabilities — as opposed to a less accurate measure that includes noise from the assessment mode.[11]

This may be challenging for tools using artificial intelligence methods; however, it is important to ask the AEDT developers how some very common accommodation requests can be handled.

From the text of the law, however, it is not entirely clear that an individual must have a disability to request an accommodation or an alternative selection process; it reads as if anyone can request these.

Moreover, providing an alternative, non-AEDT selection process for those simply not wanting to be assessed by an AEDT is not at all a common request and raises numerous questions related to operational logistics and measurement integrity, among others.

Further guidance from the city around the intent of this language is needed and should be relatively easy to provide.

How much detail about an AEDT must employers provide candidates?

The law's notices section requires an AEDT have algorithmic transparency and explainability — how data are used, what is measured, how this relates to job requirements — but it is not clear how much information employers will need to share with candidates, except that:

- Candidates must be informed of the "job qualifications and characteristics that such automated employment decision tool will use in the assessment of such candidate or employee," [12] and
- Employers must disclose "information about the type of data collected for the automated employment decision tool, the source of such data." [13]

In some instances, these requirements could pose challenges for employers using AEDTs with so-called black-box algorithms, the inner workings of which can be particularly complex and difficult to explain.

In other instances, employers and vendors may be hesitant to divulge detailed information about precisely what an AEDT assesses and how. These dynamics may lead some employers to provide overly vague or uninformative information to candidates about an AEDT.

Additional guidance is needed to help employers understand what information they need to provide candidates about an AEDT, and whether they should invest resources in either simplifying the AEDT algorithms or leveraging additional advanced data-science approaches to sufficiently explain them.

What should employers be doing to prepare for the law's enactment?

We hope the city provides additional guidance on many of these issues. However, assuming the city provides no additional guidance we suggest that employers prepare on multiple fronts for when this law takes effect in January 2023.

First, employers should evaluate their employment decision tools to determine what may be in scope. Clearly those tools that use artificial intelligence, machine learning or natural language processing are subject to the law. But employers will probably also want to evaluate other tools that use data analytics, regression or other algorithms for scoring.

Second, employers or their vendors will need to be able to understand and communicate how the tool works — what data these tools are using, what the tools are measuring and how those data and measures relate to the jobs in question.

Third, for each AEDT, employers should be prepared to furnish their independent auditor with a dataset of candidate scores and demographic subgroup information for analysis purposes. These scores may need to be organized in some meaningful way by requisition, job family, location and so on.

Employers may also want to have on hand information about how each AEDT was designed and evidence of the job-relatedness validity or predictive accuracy.

Fourth, in advance of the law's enactment, employers should be:

- Engaging their AEDT vendors or developers to discuss transparency, explainability and accommodations or alternatives;
- Involving internal partners, including business, compliance and legal, to ensure they are ready to help where needed; and
- Identifying a short list of potential independent auditors.

Conclusion

The landscape for selection tools using artificial intelligence is evolving — and this will likely continue. Illinois[14] and Maryland[15] have laws on the use of artificial intelligence in job interviews. California[16] and Washington, D.C.,[17] have proposed legislation on automated tools used in employment decisions. Federal agencies have committed to strengthening guidance and regulations on the use of artificial intelligence in selection practices.[18]

While the New York City AEDT law is currently unique in the breadth of its requirements, this will not be the case for long. Employers would be wise to consider a broader lens when thinking about audits.

The New York City AEDT law is a useful step forward in regulating the use of artificial intelligence methods in recruiting and selection processes. However, its effectiveness will be limited if some of these outstanding questions are not addressed prior to the law being enacted.

---

[1] https://legistar.council.nyc.gov/View.ashx?M=F&ID=10399761&GUID=F99584B7-57C8- 469E-9637-46A0E780690E.
[2] See Section § 20-870 Definitions, Employment decision.
[3] See Section § 20-871, paragraph b and paragraph b1.
[4] Definition provided in Section § 20-870 Definitions, Automated employment decision tool.
[5] Equal Employment Opportunity Commission, Civil Service Commission, Department of Labor, and Department of Justice. (1978). Uniform guidelines on employee selection procedures. Federal Register, 43(166), 38295-38309. Technical requirements are contained in Section 14; documentation requirements are contained in Section 15.
[6] Definition provided in Section § 20-870 Definitions, Bias audit.
[7] Morris, S. B., & In Dunleavy, E. M. (2017). Adverse impact analysis: Understanding data, statistics, and risk. New York, NY: Routledge.
[8] See Section § 20-870, Definitions in the paragraph on bias audits
[9] See Section § 20-871, paragraph a2.
[10] See Section § 20-871, paragraph b1.
[11] Americans With Disabilities Act of 1990. Public Law 101-336. 108th Congress, 2nd session (July 26, 1990).
[12] See Section § 20-871, paragraph b2.
[13] See Section § 20-871, paragraph b3.
[14] https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=4015&ChapterID=68.
[15] https://legiscan.com/MD/text/HB1202/id/2169556.
[16] https://www.dfeh.ca.gov/wp-content/uploads/sites/32/2022/03/AttachB- ModtoEmployRegAutomated-DecisionSystems.pdf.
[17] https://oag.dc.gov/sites/default/files/2021-12/DC-Bill-SDAA-FINAL-to-file-.pdf.
[18] For example, In October 2021 the EEOC announced an initiative aimed at ensuring that artificial intelligence and other emerging tools used in employment decisions comply with federal civil right laws; Algorithmic Accountability Act of 2022; NIST in March 2022 released special publication 1270 on managing bias in artificial intelligence; guidance from the EEOC in May 2022 on artificial intelligence as it relates to the Americans with Disabilities Act; guidance from the DOJ in May 2022 on algorithms, artificial intelligence, and disability discrimination in hiring.