

New York City Adopts Final Regulations on Use of AI in Hiring and Promotion, Extends Enforcement Date to July 5, 2023

By Jim Paretti, Niloy Ray, Eli Freedberg, and Ellie McPike on April 13, 2023

- NYC has issued final rules to implement an ordinance regulating employer use of automated employment decision tools (AEDTs) in hiring and promotions.
- Employers have until July 5, 2023 to determine whether they use an AEDT to make employment decisions, and if so, commission an independent bias audit; publish a summary of the results; provide notice to applicants and employees of the tool's use and functioning; and provide notice that affected individuals may request an accommodation or alternative selection process.

After several rounds of public comment and revision, on April 5, 2023 New York City published final regulations implementing its first-in-the-nation ordinance that regulates the use of AI-driven hiring tools (Local Law 144 of 2021, or "NYC 144"). Initial regulations relating solely to the penalty schedule were issued a year ago, on April 27, 2022. New York City previously deferred enforcement of the law pending publication of final regulations. Now that the final rules have been published, enforcement of the law will begin on July 5, 2023. This gives employers less than three months to determine whether they are utilizing any AI-driven hiring or screening tools and, if so, take prescribed steps to ensure compliance.

NYC 144 prohibits employers or employment agencies from using an automated employment decision tool (AEDT) to make an employment decision unless the tool is audited for bias annually; the employer publishes a public summary of the audit; and the employer provides certain notices to applicants and employees who are subject to screening by the tool. NYC 144 *likely* applies only

to New York City residents who are either candidates for hire in, or employees up for promotion to positions in, New York City, although neither the statute nor any regulation states this limitation clearly.

Does the Employer Use an AEDT to Make an Employment Decision?

The threshold question is whether an employer is using an AEDT to make a covered employment decision and thus subject to the audit and notice requirements of the law. To answer this question, the employer must look to the technology that comprises the tool, and how and for what purposes the tool is used. These terms sometimes overlap and are redundant; for ease, we have laid out a simplified 3-step framework for analysis below.

The statute defines *AEDT* to be:

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

To determine whether a tool fits within the definition of AEDT, employers should ask themselves three questions: (1) does the technology used by the tool fall within the definition of AEDT? (2) is the tool being used in a way that brings it within the definition of AEDT? and (3) is the tool being used to make a covered employment decision within the definition of AEDT and NYC 144? The most efficient approach for employers will often be to consider these questions in reverse order, since this will spotlight situations to which the law clearly does not apply.

Does the tool make a covered employment decision?

The fairly narrow definition of covered “employment decisions” is the simplest starting point for analysis.¹ The statute provides that the only “employment decisions” it covers involve hiring and promotion. Decisions regarding compensation, termination, workforce planning, labor deployment, benefits, workforce monitoring, and likely even performance evaluations, are beyond the reach of the law.

The regulations also limit the scope of the law to candidates who have actually applied for a specific job. This means that the law does not apply to tools used to identify potential candidates who have not yet applied for a position.

It is also important to note that the regulations restrict application of the law to hiring or promotion decisions that “screen” candidates (for hiring) or employees (for promotion) by determining whether they “should be selected or advanced” in the process. Thus, for example, where an employer uses a tool to sort applicants into three tiers (“highly qualified,” “qualified” and “less qualified”), and as a practical matter does not advance any candidate who ranks “less qualified,” that use almost surely falls within the scope of the law. Conversely, if an employer uses a tool to assess a candidate’s attention to fine detail, but that assessment is not the most important or determinative factor, but rather just one of many considered in deciding (in practice, not just theoretically) whether the candidate advances, that use may fall outside the definition of AEDT and thus the law.

Is the employer using the tool in a way that brings it within the definition of AEDT?

If the tool is in fact going to be part of a hiring or promotion screening decision, the next consideration is how output from the tool will be used in making that decision. An AI tool falls within scope of NYC 144 only when it issues a “simplified output,”² and where that output is used to “substantially assist or replace discretionary decision making.” The final regulations clarify that a covered tool’s output – which may be a score, classification or recommendation – must be used in one of three ways:

- as the sole criterion in making the employment decision, with no other factors considered;
- as a criterion that is given more determinative weight than any other criterion; or
- to overrule conclusions derived from other factors including human decision-making.

Any other use of an AI tool’s output – for instance, where that ranking, score or classification is merely an equal part of the decision calculus for an individual – would appear to exempt that tool from the definition of an AEDT under the statute. So, for example, where an employer uses a tool to provide a ranking of candidates but does not rely solely on that ranking to advance the candidate, afford the ranking the most weight in its decision, or allow the ranking to overrule a human decisionmaker, that use probably falls outside the definition of AEDT and is not subject to the law’s requirements. Conversely, if the employer uses a tool to rank 100 resumes, and advances only the first 20 (rejecting the other 80) solely on the basis of that ranking, such use does likely fall within the definition of “substantially assist or replace” human decision-making.

Determining the weight afforded an AI tool’s output will ultimately become very fact-specific, and employers that determine their AI tool is exempt from coverage under NYC 144 because it is not used to “substantially” assist or replace human decision-making, may have those determinations

challenged. An employer using scores or ranking in its selection process will want to ensure and be able to demonstrate through adequate documentation that, in practice as well as on paper, its model falls outside any of the covered uses.

Does the tool's underlying technology fall within the definition of AEDT?

Per NYC Admin. Code 20-870, a tool is covered by NYC 144 only if it is derived from "machine learning, statistical modeling, data analytics, or artificial intelligence," which term is defined, in the final regulations, to mean a group of mathematical, computer-based techniques:

- that generate a prediction, meaning an expected outcome for an observation, such as an assessment of a candidate's fit or likelihood of success, or that generate a classification, meaning an assignment of an observation to a group, such as categorizations based on skill sets or aptitude; **and**
- for which a computer at least in part identifies the inputs, the relative importance placed on those inputs, and, if applicable, other parameters for the models in order to improve the accuracy of the prediction or classification.

Together, these two criteria describe machine learning-based predictive AI tools. In other words, despite the purportedly wider technological reach indicated by the term "machine learning, statistical modeling ...," this definition restricts the law's reach to those predictive tools that employ the specific algorithmic techniques that comprise machine learning. Chief among these are the staple machine-learning characteristics of feature extraction and training: the algorithm, to greater or lesser extent, identifies the inputs (aka "features") to be used in the model, and then, by "training" on datasets, the algorithm determines the optimal weight to place on each feature. Contrast this to a human-designed statistical model, where the specifications are all set by a human, and the distinction indicates that the human product falls outside the scope of NYC 144.

As a practical matter, this means, for example, that where a test is designed by a human who identifies the necessary inputs and their relative importance, the tool is likely outside the definition of AEDT and thus not within the scope of the law. For example, if an employer has a human programmer design a test to determine whether a candidate is likely to perform well in a position, or has the necessary skills for a promotion, and the test measures criteria X, Y, and Z (determined by the employer), and weights X as the most important criterion, the test is likely *not* an AEDT under the definition in the New York City law. Conversely, if the tool is programmed to autonomously adjust the weighting of criteria based on prior data sets and outputs, such that the basis for its output may "evolve" to give more weight to criterion Y than X, it may well be a covered AEDT.

NYC 144 Applies – Now What?

If the tool meets all three criteria – a machine-learning tool whose use has a specific determinative impact on the screening out of a job applicant or employee up for promotion – and thus falls within the scope of NYC 144's requirements, employers must take four steps: commission an independent bias audit; publish a summary of that audit's results; provide notice to applicants and employees of the tool's use and functioning; and provide notice that individuals subject to the tool may request an accommodation or alternative selection process.

Bias Audit Per 6 RCNY 5-301. NYC 144 requires employers to conduct a bias audit and publish a summary of the audit (and its distribution date) clearly and conspicuously *before* its use. As a practical matter, if a covered AEDT is already in use, a bias audit should be completed as quickly as possible. After the initial bias audit the law requires such an audit to be performed annually. Bias audits must be completed by "Independent Auditors," defined as objective individuals or groups who are not and have not been involved in the use, development, or distribution of the AEDT; have not at any point during the audit been employed by the employer, vendor, or AEDT developer; and who have no direct financial interest or material indirect financial interest in the use of the AEDT or the vendor that developed or distributed it.

The content of bias audits includes calculations of the selection rate for each category (or scoring rate, where the tool issues scores instead of classifications/groupings), including sex categories, race/ethnicity categories, and intersectional categories, as well as the impact ratio of each category. The final regulations have added an additional requirement for bias audits: publication of the number of individuals whose race and gender are "unknown" and thus excluded from the calculations. Additionally, an independent auditor may exclude a category that represents less than 2% of the data being used in the bias audit if the independent auditor provides a justification for the exclusion; in those instances, the number of applicants and "scoring rate" or "selection rate" for the excluded category must be included in the public summary.³

Note that the final regulations identify two sorts of data that may be used in a bias audit: "Historical Data" and "Test Data." "Historical Data" are, essentially, the results of an employer's prior use of the AEDT being audited. The regulations note that Historical Data "must" be used for the bias audit unless the employer can explain why there are insufficient Historical Data for the independent auditor's use. Only given such an explanation may an employer rely on a bias audit using "Test Data," which is defined as "not Historical Data" and thus represents other, alternative data sets – such as independent data sets that can be run through the AEDT to determine results. The goal of this distinction appears to be to encourage organizations to share the results of their use of an AEDT with the tool vendor or independent auditor, so that a single bias audit, conducted using data from multiple employers, can be used by each of those contributing employers.

Prior to the use of an AEDT, the employer must post in a clear and conspicuous manner on the employment section of its website the date of the most recent bias audit; a summary of the results, including the source and explanation of the data used; and the selection (or scoring) rates and impact ratios for all categories; and the distribution date of the AEDT. This requirement may be met by way of an active hyperlink to a website containing the required information, so long as such link is clearly identified as linking to the results of the bias audit. The employer must also list the type of data collected for the AEDT, the source of such data, and the organization's data retention policy. This requirement, too, may be met by clear and conspicuous posting on the employer's careers website.

Notice to Employees Per 6 RCNY 5-304. NYC 144 further requires employers to provide candidates for employment or promotion with notice 10 business days before use of the tool: (1) that an AEDT is being used in assessing and evaluating the candidate; (2) the job qualifications and characteristics the AEDT will use in its analysis; (3) if not disclosed elsewhere on its website, the AEDT's data source, type, and the employer's data retention policy; and (4) that a candidate may request an alternative selection process or accommodation. This notice may be accomplished by a single clear and conspicuous notice on the employment section of its website, in a job posting, or by way of U.S. mail or e-mail to the candidate. In practice, we anticipate that a website link or policy posting will be the preferred way to meet this requirement, given the relative ease of adding such linked content to existing career websites.

Alternative Selection Process or Accommodation Per NYC Admin. Code 20-871(b)(1); 6 RCNY 5-304(a). NYC 144 requires that employers offer either an alternative selection process (ASP) or reasonable accommodation to those employees or applicants who request it. The latter appears to be an ADA access requirement, while the ASP (if not mere surplusage) would be a non-ADA opt-out request mechanism. However, and critically, the regulations make clear that "[n]othing [in these regulations] requires an employer or employment agency to provide" an ASP. In sum, it would appear that employers are not required to do more than provide a mechanism for this non-ADA opt-out request, and there is apparently no requirement that these non-ADA opt-outs be granted. Notably, the law does not impose any standard or decisional method for decisions about whether to grant an ASP in either circumstance.

Penalties Per NYC Admin. Code 20-872; 6 RCNY 6-81. Per the law and initial regulations, any person who violates the law is liable for a civil penalty of \$375 for the first violation (and each additional violation occurring on the first day of the first violation). Each subsequent violation (or default, if the violation is left uncorrected) occurring after the first day will result in a civil penalty of

at least \$500 and not more than \$1,500. Each failure to meet applicant notice requirements constitutes a separate violation. Failure to meet the “bias audit” requirements results in separate, daily violations.

The final regulations go a long way toward clarifying the requirements of a short and broadly written law. That said, the law’s definitional criteria (both legal and technological) are complex, and employers using automated screening tools that may fall under the law would be well advised to consult counsel to discuss whether they are covered and if so what they may be required to do to ensure compliance.

¹ This third prong is where the law and regulations become a bit redundant and a lot confusing (another reason to address it first). The law prohibits use of AEDTs to screen candidates absent bias audit and notice steps—but, by definition, a tool that does not screen candidates is not an AEDT and thus already beyond scope. Deciding, first, whether a tool in fact “screens” individuals, cuts through this confusion.

² The final regulations clarify that a “simplified output” means a “prediction or classification [which] may take the form of a score (e.g., rating a candidate’s estimated technical skills), tag or categorization (e.g., categorizing a candidate’s resume based on key words, assigning a skill or trait to a candidate), recommendation (e.g., whether a candidate should be given an interview), or ranking (e.g., arranging a list of candidates based on how well their cover letters match the job description).” In practice, we expect most AI-driven tools used by employers will meet this criterion.

³ The 2% is to ensure the data is statistically sound and is a true representation of the whole.

Information contained in this publication is intended for informational purposes only and does not constitute legal advice or opinion, nor is it a substitute for the professional judgment of an attorney.

© 2023 Littler Mendelson P.C.

Littler Mendelson is part of the international legal practice Littler Global which operates worldwide through a number of separate legal entities. Attorney Advertising.