



How To Prepare Your Business For New Regulations Covering AI In HR

beamery.com

Private & Confidential - © Beamery Inc. All rights reserved

Introduction

With more and more businesses, across industries, adopting artificial intelligence (“AI”) and other complex algorithms, lots of decisions that used to be made by humans are now being made (either completely or partly) by algorithms or models. Examples of automated decision-making (“ADM”) include determining whether an email is spam or not, which adverts get displayed for someone on the internet, or whether a payment is potentially fraudulent. Organizations are seeing huge efficiency gains from automating previously manual tasks, and for many automations, the quality of the outcome is significantly improved.

12% of participants in a pilot by [World Economic Forum \(WEF\)](#) reported that their organization was already using AI-based HR tools, with another **9%** making concrete plans to do so. The largest group, **43%**, were in initial exploration stages, while 21% had no plans and **15%** were not sure.

For HR practitioners, the most common use case for ADM is creating shortlists of jobs or candidates: helping to determine who gets an interview, a job, or a promotion. AI-powered solutions are growing in popularity: as ways to reduce the time and cost associated with the hiring process (e.g. sorting through large amounts of resumes and applications); to improve hiring outcomes, and recruit the best talent in an increasingly competitive landscape. Beyond hiring, [AI can also bring benefits](#) when it comes to personalizing employee career paths, increasing retention, and improving the employee experience.

66 HR is challenging work. It is often executed with limited resources, and current practices are not always effective or fair. HR is also important work because people’s jobs are critical for their lives and livelihoods, and quality, reliable talent fuels the success of employers. The promises offered by AI-based HR tools to do this work faster, better and possibly even fairer, therefore, has caught the attention of many organizations, governments and individuals.



GROWING COMPLIANCE CONCERNS

Lawmakers and regulators are starting to pay closer attention to the use of AI across sectors, especially where its application may impact individuals' financial or other life opportunities. Clearly, employment is a key area, targeted for increased regulation. The question is: are individuals going to lose out on work opportunities due to the use of AI to screen and select candidates? Can AI and ADM be applied fairly, with appropriate concern for the security and privacy of applicants?

51% of respondents in the [Beamery Talent Index Fifth Edition](#) thought all AI algorithms were adhering to ethical standards, either some or all of the time. **25%** said either rarely or not at all, while **24%** were unsure.

But there were concerns about how bias can creep into algorithms: **38%** said AI can unintentionally discriminate against certain candidates.

The growing popularity of AI-powered hiring tools has led to concerns from regulators that they may be inadvertently discriminating, for example, by:



Suggesting candidates who are similar to those making up the company's current leadership team, which is not itself diverse;



Penalizing job candidates with "work gaps" in their resumes, which could lead to a bias against those with childcare or other caring responsibilities; or



Supporting automated "games" that are unfairly difficult for people with disabilities, as part of an evaluation process, even though they could do the job with some reasonable accommodations

66 the high-stakes role that HR decisions play in people's lives increases the importance of both the opportunities and the perils of using AI. The optimists and pessimists of AI in HR, therefore, often have a common interest in ensuring a better and fairer workplace.



The NYC AEDT Law

There is a broad trend, at a state, national, and supranational level, towards regulating how AI is used in employment processes. At the municipal level, last November, the [New York City Council passed a bill](#) that regulates employers' and employment agencies' use of "automated employment decision tools" in making employment decisions. The new law takes effect on 1st January 2023, with final guidance being drafted at the time of publication.

WHAT IS AN AUTOMATED EMPLOYMENT DECISION TOOL?

The broad definition of an AEDT covers "any computational process . . . that issues simplified output . . . that is used to substantially assist or replace discretionary decision making for making employment decisions."

The proposed regulations clarify that to "substantially assist or replace discretionary decision making" means to:



Rely only on a simplified output (e.g., score, tag, classification, ranking, etc.) with no other factors considered.



Use a simplified output as one of a set of criteria, where the output is weighted more than any other criteria.



Use a "simplified output to overrule or modify conclusions derived from other factors including human decision-making."


(A "simplified output" is defined as a prediction or classification that can take the form of a score, tag or categorization, recommendation, or ranking.)

So this isn't just about 'AI' or indeed full automation. A tool that, for example, produces a score from a collection of candidate-related data (ranking people 'automatically') will probably qualify, even if a human is looking at the scoring before making any decisions. It will likely cover many game-based interview tests, as well as some resume review tools and automated personality assessments.

It doesn't have to be a machine entirely making decisions – it just has to be influencing them.

THE REQUIREMENTS

Employers have to comply with some new requirements, including: a yearly bias audit; and a summary of the results of such an audit made publicly available; providing candidates with notice of automated processing; and allowing candidates or employees to request an alternative process or accommodation.



This law addresses the perceived bias risk of AI or ADM: is the ‘machine’ selecting candidates for interviewing, or filtering resumes, in a biased way? Is it inadvertently discriminating by gender, race or ethnicity, or some other protected characteristic? And do people know they are being reviewed by a machine, at least in part? Can they request an accommodation?

Let’s look at the requirements of the NYC AEDT law more closely.

i. Conduct an Independent, Annual Bias Audit

The AEDT says that employers must conduct annual, independent audits to assess “disparate impact”. An independent auditor is defined as “a person or group that is not involved in using or developing an AEDT that is responsible for conducting a bias audit of such AEDT.” It is not totally clear what level of independence is required of the auditor, including whether the auditor may be otherwise employed by the AEDT’s vendor.

Disparate impact relates to whether results from any automated employment decision tools are skewed based on protected characteristics such as gender, race or ethnicity, although the term is not defined within the proposed new law.

What is required in a bias audit depends on how the AEDT is used in the employment process. When an AEDT selects individuals to move forward in the recruitment process, or classifies them into groups, the bias audit must calculate (for each category i.e. race, ethnicity, sex):



the selection rate (the number of individuals in the category moving forward or assigned a classification, divided by the total number of individuals in the category who applied for a position or were considered for promotion),



the impact ratio (defined as either (1) the selection rate for a category divided by the selection rate of the most selected category, or (2) the average score of all individuals in a category, divided by the average score of individuals in the highest scoring category.)

Where an AEDT classifies individuals into groups, the calculations for the selection rate and impact ratio must be performed for each such classification. However, when an AEDT scores applicants or candidates, the audit must calculate the average score for individuals in each category, and calculate the impact ratio for each category.

An independent auditor conducting a bias audit must compare selection rates for each EEO-1 category and compare those rates to the most selected or favored category.

The “selection rate” and “impact ratio” for each category must be reported to the Equal Employment Opportunity Commission (EEOC) pursuant to the EEOC Employer Information Report EEO-1 (race, ethnicity and sex).

ii. Provide Audit Disclosures

Before using an AEDT, employers must publish in a “clear and conspicuous manner” on their careers site or jobs section:



the date of the most recent bias audit of the AEDT



a summary of the results of the audit, including selection rates and impact ratios for all categories



the distribution date of the AEDT, or the date the employer began using a specific AEDT.

It seems employers can meet this requirement by including a hyperlink to an external website that contains the required information. The required information must remain posted for at least six months after last using the AEDT for an employment decision.

iii. Notify Candidates or Employees

Employers must provide candidates (who are New York City residents) notice – of at least 10 days – that an automated employment decision tool will be used to assess their employment or candidacy, and the qualities or characteristics assessed. It must also notify them that they can request an “alternative selection procedure or accommodation.” This would suggest employers should get into compliance by mid-December to avoid any interruptions in their NYC hiring processes. The currently proposed guidance on the regulations also clarifies the law’s requirement that employers provide information regarding the type of data collected, and their data retention policy. Employers may satisfy this requirement by posting the information to their website’s careers or jobs section, or by providing it in writing via mail or email. If the notice is not included on the website, an employer must post instructions on how to make a written request for such information on its website.

iv. Provide an Accommodation or Alternative Selection Process

The NYC law requires that companies provide candidates or employees residing in New York City a way to request an alternative selection process or accommodation; it does not, however, say what these alternatives or accommodations should be, or whether anyone who requests this will be entitled to go through an “alternative selection process” – meaning at this time accommodation is not mandated.

MORE INFORMATION ON THE AEDT LAW

What we have is simply step one regarding compliance with the NYC AEDT law: regulators are now undertaking the process to provide further guidance on compliance.



Candidates: To be considered a “candidate”, the individual must have made “affirmative steps” – that is, by submitting information or other items to be covered under the law.



Penalties: On June 6th, the NYC Department of Consumer and Worker Protection held a hearing on the AEDT Law and shared a penalty schedule. Companies that are found to not be in compliance will face penalties of \$375 for a first violation, \$1,350 for a second violation, and \$1,500 for a third violation and any subsequent violations.



Violations: What constitutes a ‘violation’? For each day that a company uses a covered tool that doesn’t comply with the law, and for each day they fail to provide the required notice, there is a separate violation.



Enforcement: In terms of enforcement, New York City’s Corporation Counsel (who brings suits on behalf of the government) is the body most likely to bring proceedings to enforce the AEDT. The New York City Commission on Human Rights may also enforce this new law. These are two groups we will watch closely as the regulations roll out. The AEDT does not include a private right of action, but it does not stop private plaintiffs from bringing civil actions related to a company’s practices or automated tools (including, for example, discrimination claims).



Scope: When it comes to jurisdiction, it is as yet unclear whether the AEDT applies to all employers and all jobs based in New York City. Companies located in New York City, which are hiring or promoting City residents, for jobs that are located in the City (using “automated employment decision tools” to “replace” or “substantially assist” decision-making in hiring or promotions), are certainly covered.

The full scope of the AEDT is likely to be clarified in the forthcoming rulemaking process and we hope to have more information to share in due course.

Other Existing & Emerging Regulations

NYC is not the only place releasing guidance/regulations around AI (or automation) that is being used in a hiring context. In the US, two states have AI laws in effect: Illinois and Maryland. Both focus on the use of AI in a video (interview) context, and are concerned with avoiding bias in the employment process. Other states have proposals under consideration, while larger jurisdictions are assessing the right approach to this topic.

Illinois: Requirements around disclosure and data destruction, on employers using video-based AI in interviews (in effect).

Originally the regulation was about making sure interviewees had consented to the use of AI. Following amendments in January 2022, employers who rely solely on AI in video interviews now have to provide annual reports on demographic data, outlining who is and isn't afforded an interview, and who is hired, with reference to race/ethnicity.

Washington DC: Requirements around algorithmic audits (proposed).

A lot of people are watching to see what happens in DC. If passed, the D.C. bill would prohibit use of algorithmic tech to determine eligibility for employment based on someone's race, color, religion, familial status, source of income or other information in a way that discriminates against them "or otherwise makes important life opportunities unavailable to an individual or class of individuals."

They would also have to produce annual audits around disparate impact, and share reporting to the office of the Attorney General (with an audit trail covering 5 years). This law would enforce lots of obligations on employers; there is a public hearing in September.

Maryland: Prohibiting facial recognition in pre-employment job interviews without the applicant's written consent (in effect).

California: Prohibiting employers from screening out applicants on the basis of a protected characteristic through use of AI (proposed).

EEOC: Warning about use of AI leading to unlawful employment discrimination, and recommending that employers provide candidates with notice about a tool's use and that employers offer accommodations or exemptions to candidates (non-binding guidance).

EU: Placing regulatory and disclosure requirements on "high-risk" AI systems, likely including employment tools (proposed). This would impose onerous obligations on employers in the European Union.

The [UK privacy watchdog](#) is also set to probe whether employers using artificial intelligence in their recruitment systems could be discriminating against ethnic minorities and people with disabilities.

All of the legislation means that, for large enterprises in particular, choosing technology partners will mean looking for a flexible solution that meets compliance requirements in different geographies.

LOOKING TO THE FUTURE

Longstanding principles around avoiding discrimination will likely be applied to employers' use of AI at the local, state, and national level.

The AEDT law approach, requiring audits and the provision of notice and disclosures, may well gain favor in the United States, at least among states and large municipalities.

Globally, a risk-based framework like that in the EU may become a preferred approach among national governments: categorizing AI systems in hiring processes as "high risk" and subjecting the use of these to strict rules.




Tips for Staying Compliant

Broadly speaking, AI application providers – the vendors providing the solution, such as candidate matching, resume screening or automated interviewing tools – should focus on taking steps to ease customers' compliance burdens. Employers also should train hiring managers and HR personnel regarding the law's lengthy requirements, including its publication and notice provisions.

LEAN ON YOUR VENDORS

As an employer, you will want to create a cross-functional taskforce for ensuring compliance with AI regulations: including HR as subject matter experts, legal and compliance, and probably the technology function, and we recommend you include your vendors/service providers.

Providers (responsible vendors) should be able to tell you which of their tools are covered by the AEDT law and other similar laws and regulations, and how they are complying and can help you comply. The AEDT law (and others like it) will present similar responsibilities as GDPR (or CCPA): the data processor/data controller relationship is mirrored here. While the legal burden is on the employer, they will expect their suppliers to help them with compliance.



i. Testing & Bias Audit

In order to comply with the NYC law and others, it will be necessary to undertake an independent bias audit. Unless you have built and are running your own AI-powered system, vendors will need to support their customers here. This is similar to how things work in information security: if a supplier was handling data on your behalf, you would have questions; now it will be similar for AI tools. Vendors will have to provide bias audits in addition to ISO certifications or SOC audits.

Your vendor should be able to share the results of an independent bias audit (from within the last 12 months) along with results of 'disparate impact' testing and other related testing that looks at potential bias within their data, models and AI interfaces.

ii. Transparency & Notice: Updated Career Sites & Explainability Statement

The proposed guidance allows the required notices to be given to candidates who reside in New York City by including the notice on the employer's careers or jobs section online, in a job posting, or via email or mail. An employer can satisfy the notice requirements for current employees by including it in a written policy, procedure or job posting, or providing it in person, or via email or mail. Many employers are used to posting a candidate privacy statement: these disclosures can be incorporated therein.

Your vendor should have an Explainability Statement, and build in functionality to help you easily fulfill disclosure and notice obligations.

We assume that Explainability Statements (comparable with Privacy Statements) will be published online in relation to any given AEDT, and will be part of the process when someone applies for a role, is added to a talent pool, or is using internal mobility tools as an employee. This would be shared with the candidate at the same time as a privacy statement – to explain what the tools do. It's not clear yet where such a statement would have to live, what exact form it should take, and when in the process it should be shared with candidates. It must of course be written in plain language, so it is understandable to the general public.

At Beamery, we have a publicly published [Explainability Statement](#), which customers can link to in their application process.

'Explainable AI' ensures that users always understand how candidate recommendations have been made: they can see the components that were weighed up in the decision making process, and what elements impacted the recommendations the most. AI is most effective with [human-in-the-loop guidance](#) – in short, it's the human, not the AI, that makes the final decision. The AI element makes the recruitment faster and more efficient, but does not replace humans.

iii. Product: In-App Explainability

Your vendor should make it easy for users (i.e. HR professionals using the AEDT) to understand how AI is working within the product.

As someone using these tools (on behalf of an employer), you might not even know that they use AI; why you've been sent a set of resumes; or why someone has been recommended for an open vacancy. In-app explainability (the background to recommendations) is not a requirement under the law but we think it is very important.

At Beamery, we want to make it clearer in the product why particular candidates were surfaced. Our technology does not make decisions but it does help steer them, so it's important to understand the rationale behind recommendations (e.g. not only relevant skills but also seniority level, industry context, or expected proficiency).

iv. Consent, Preferences & Accessibility

There are already a lot of rules in place that help protect and accommodate individuals with disabilities as consumers and on the job market. For example, the Americans with Disabilities Act ([ADA](#)) and the [European Accessibility Act](#) have all sorts of important regulations for companies to follow regarding hiring those with disabilities, and how businesses can best serve their customers who have disabilities (i.e. making sure their website graphics have alt text for those who are visually impaired).

Vendors should also be considering how they handle 'accommodation' as it pertains to the AEDT law. A preference center will enable candidates (those being assessed by the employer, or individuals interested in the employer) to manage their preferences in regards to AI use, and request an alternative selection process or accommodation directly. This actually goes beyond the AEDT law: GDPR compliance can be handled through this mechanism as well.



Questions To Ask AEDT Vendors

i. What are you doing to mitigate bias in your AI?

AI is trained using data sets, and if there is bias in the data, the outcomes from the AI will be biased too. Your vendor should be able to point to actions they have taken to reduce the chance of recommendations being biased.

ii. Do you have an AI Explainability Statement?

One of the fears people have around AI is that it is hard to understand: that decisions or recommendations could be made without anyone understanding the rationale. Related to the reduction of bias, if an AI-powered solution is not explainable, there is a real chance of outcomes being unfair. A document explaining how the AI works is a great start.

iii. What third parties have you been working with to audit your models and your training data?

The new rules mean that companies that use “automated employment decision tools” have to conduct a yearly bias audit; and a summary of the results of such an audit must be made publicly available. Does your supplier offer this? Who is auditing their AI and what are the results?

iv. How do you always ensure a human understands when they are interacting with AI?

Employers affected by the NYC law also have to provide candidates with notice of automated processing; and allow candidates or employees to request an alternative process or accommodation. Does your vendor provide a way for candidates/employees to be informed as to AI being used, how it is being used, and can they manage their preferences? Further, do the employers’ users of the tool know they are using an AI-powered tool that may steer their decision making?

v. What are your AI ethical principles?

AI works best with human involvement: it is not there to replace TA/HR people, but to save time on manual processes, and enhance the work they do. Your vendor should have principles in place aligned to your own hiring principles.



Beamery uses AI to power software that helps reduce bias in the hiring process by focusing on skills, and to enable us to deliver on Beamery's mission: to create equal access to work. Our AI helps recruiters identify which candidates to contact based on their skills and titles, and helps candidates see what career paths might be a good match for them based on the skills and interests. The use of AI can help enable the reduction of bias across Talent Lifecycle Management and enable compliance through explainability: we can provide transparent assistance to users in helping them make effective and informed decisions.

This document and all its contents is the property of Beamery Inc. THIS DOCUMENT IS FOR INFORMATION PURPOSES ONLY. DIAGRAMS, PERCENTAGES AND/OR OUTCOMES USED IN THIS DOCUMENT ARE FOR ILLUSTRATION PURPOSES ONLY. INDIVIDUAL OUTCOMES VARY BY CUSTOMER. THE BEAMERY CUSTOMER IS SOLELY RESPONSIBLE FOR ITS USE OF THE BEAMERY TECHNOLOGY. BEAMERY WILL NOT BE RESPONSIBLE FOR ANY LIABILITY, LOSS OR DAMAGE OF ANY KIND RESULTING FROM OR CONNECTED WITH THE USE OF THIS DOCUMENT.

© Beamery Ltd 2022. All rights reserved

