

Artificial Intelligence Discrimination Prevention & Compliance Game Plan



Technology moves faster than law. This forces businesses to fit old regulations to novel circumstances. Artificial intelligence (AI) is the perfect example. AI is digital but the laws that govern its use are analog. One of those analog laws is the one that bans employers from discriminating on the basis of race, religion, sex, etc. Potential bias within AI algorithms brings discrimination laws into play.

The liability risks of [AI bias](#) are more than speculative. In a novel case, iTutorGroup Inc. shelled out \$365,000 to settle [age discrimination](#) claims for using software programmed to reject female job applicants over age 55 and male applicants over age 60. Before the ink on the settlement had even dried, a group of 100 rejected job applicants filed a class action lawsuit accusing software vendor Workday of using AI screening tools with algorithms wired to exclude older African Americans and individuals with disabilities. While both of these cases involved U.S. companies, it's only a matter of time before Canadian employers become the target for AI discrimination lawsuits.

Bottom Line: If you use AI for hiring or other HR functions, you need to guard against [algorithmic discrimination](#) risks. Here's a 7-step Game Plan.

1. Use Due Diligence in Selecting AI Vendors & Products

The name of the game is to ensure that whatever AI technology you use for hiring isn't programmed to limit or exclude groups of job applicants on the basis of protected characteristics. And if the AI technology you use comes from third-party vendors, you need to select vendors and products that you can trust.

Compliance Strategy: Perform due diligence before signing a contract with a vendor who designs or implements an AI-based hiring tool. Specifically, HR should work together with IT and the legal department to vet the product and its design and features for potentially discriminatory outcomes. Key things to confirm about the product:

- The vendor performs regular end-to-end testing to ensure that algorithms are similarly predictive across and don't disadvantage any protected groups.
- It allows for detection of discriminatory outcomes.

- The vendor takes steps to correct disparities it identifies.

2. Include Liability Protections in Vendor Contract

Having completed your due diligence, negotiate a contract that makes the AI vendor you select responsible for any discrimination or other liability issues arising out of your use of the product for hiring purposes that are the vendor's fault.

Compliance Strategy: Such provisions may include:

- An allocation of risk clause.
- A vendor warranty that the product is and will remain compliant with human rights and other laws.
- An indemnification clause requiring the vendor to pay all fines, damages, losses, and other costs you incur for discrimination or other legal actions arising from use of the product to perform hiring functions.

3. Guard Against Disability Discrimination Caused by AI Screen Out

Human rights laws require employers to make reasonable accommodations, such as specialized equipment, alternative testing, and even changes to the job itself, to ensure individuals with disabilities equal opportunity to apply and receive fair consideration for a job. The danger is that algorithmic decision-making hiring tools may screen out disabled applicants without taking reasonable accommodations into account. The result is to exclude disabled applicants who are qualified for the job when provided with reasonable accommodations.

Example: Chatbot software that screens out applicants for cashier jobs requiring standing for long periods may reject an applicant who uses a wheelchair who'd be entitled to the reasonable accommodation of lowering the cash register so that it can be operated from a sitting position.

Compliance Strategy: Ensure that your AI tools measure abilities and qualifications truly necessary for the job—even for those who are entitled to on-the-job reasonable accommodation. Steer clear of algorithmic decision-making tools that don't directly measure but instead make inferences about necessary abilities and qualifications based on characteristics that are correlated with them. For example, a tool for a job requiring the ability to analyze data may rate that ability by measuring the similarity between an applicant's personality and the typical personality of successful data analysts. Result: The tool would reject an applicant who's great at analyzing data but who, due to a disability, has a personality that's far from the norm for successful data analysts.

4. Beware of Hiring Processes that Aren't Accessible to Disabled Applicants

AI hiring technology may not be [accessible](#) to applicants with particular kinds of disabilities. For example, "gamified" tests that use video games to measure applicants' abilities, personality traits, and other qualities may exclude people with visual impairments compromising their ability to play the game. So, a policy of requiring a particular score, such as 90% on a gamified assessment of memory, would exclude a blind applicant with a good memory and is in all other ways perfectly capable of performing the job.

Compliance Strategy: Ask AI vendors if the tool was developed with disabled people in mind. Specific questions:

- Does the tool ask job applicants illegal questions that are likely to elicit information about a disability?
- Is the tool's user interface, if any, accessible to as many individuals with disabilities as possible?
- Does the tool present materials to applicants in alternative formats and, if so, which ones?
- Are there any disabilities for which the tool can't provide accessible formats, (in which case you might have to provide reasonable accommodations)?
- Did the vendor take steps to determine whether use of the algorithm disadvantages the disabled, such as assessing whether any traits or characteristics the tool measures are correlated with certain disabilities?

You also need some kind of alternative method for rating job applicants in case the current AI evaluation process you're using is inaccessible or otherwise unfairly disadvantages someone with a disability.

5. Be Transparent About How You Use AI

Hiring discrimination lawsuits are often the product of miscommunication and misunderstanding about how the process works. So, being transparent about your AI use may reduce risk of being sued. It might also save you time and money by discouraging unqualified people from applying.

Compliance Strategy: Provide job applicants information about any AI algorithmic decision-making or screening tools you use, including:

- The traits or characteristics the tool measures.
- The methods it uses to measure those tools and traits.
- The disabilities or other protected characteristics, if any, that might potentially lower assessment results or cause applicants to be screened out.

6. Let Applicants Know that Reasonable Accommodations Are Available

Defuse reasonable accommodations issues by proactively addressing them early in the hiring process.

Compliance Strategy: Notify applicants that reasonable accommodations, including alternative formats and tests, extended deadlines, etc., are available to individuals with disabilities. Provide clear instructions on how to request accommodations and establish a process for responding quickly so that requesters have ample time to be considered for the job before it's filled and you won't have to choose between meeting your accommodation responsibilities and unduly delaying the hiring process.

7. Monitor Your AI Hiring Selection Rates for Potentially Discriminatory Outcomes

Monitor AI hiring selection rates by race, sex, disability, age, and other protected characteristics and analyze the data to identify and mitigate discriminatory outcomes resulting from your use of the tool. **Explanation:** Selection rate refers to the proportion of applicants that got a yes from AI tool to move forward in the hiring process, such as applicants the AI selected to interview. To get the selection rate of a group, you divide the number of people selected by the total number of applicants in that group. So, if 100 women apply for a position and 40 are selected for an interview, the selection rate for women is 40%.

Compliance Strategy: Monitor AI hiring decisions for potential bias by comparing selection rates of different groups. There are two basic approaches:

Option 1. The Four-Fifths Rule: The unofficial rule is that discrimination may be present when the selection rate for a protected group is less than 80% as compared to non-protected groups. **Example:** An algorithm used for a personality test selects Black applicants at a rate of 30% and White applicants at a rate of 60%. The resulting 50% selection rate for Blacks versus Whites ($30/60 = 50\%$) would raise a [racial discrimination](#) red flag because it's lower than 4/5 (80%) of the rate at which White applicants were selected.

Option 2. AI Bias Audits: A more precise method is to periodically perform what's called an [AI bias audit](#) that uses specific metrics measuring historical data that employers keep track of relating to their real-life use of the AI tool to determine whether such use results in disproportionately negative outcomes against certain protected groups. The metrics vary depending on whether the AI tool is a regression system that generates a continuous score or a classification system that provides a "Yes/No" or other binary output. Employer use of AI bias audits, which are often performed by an independent third party, will probably become mandatory in Canada in the not-too-distant future the way they have in some parts of the US, including New York City.

- **Discrimination Risks of Using AI for Hiring**

- **Advantages of Using AI for Hiring**

- **Discrimination Risks of Using AI for Hiring**

Human rights laws ban employers from making hiring decisions based on an applicant's race, sex, age, religion, national origin, disability and other protected characteristics. Discrimination doesn't have to be deliberate or direct. Employers are liable for employment practices that appear neutral on their face but have the effect of discriminating against protected groups. This is true even if "adverse effect" discrimination is the result of using an AI product created by a third-party vendor.

The problem with AI hiring platforms is that they may be wired to consider job applicants' protected characteristics in carrying out their functions. While this can be done deliberately, it may also happen as a result of automated systems operations that are designed not to comply with discrimination laws but to make the particular function, e.g., advertising, screening or evaluation, more efficient in meeting the user's objectives. There are 2 major technical factors that contribute to AI bias:

Training Data: The decision-making processes that AI systems develop are based on training data. If that data either overrepresents or underrepresents certain groups, the system is likely to make biased decisions. **Example:** Facial recognition software that overrepresents white people may result in less accurate and racially biased facial recognition of people of colour. Training data may also be mislabeled in a way that disadvantages protected groups.

Programming Errors: Data and algorithms built into AI may incorporate the subtle prejudices of the humans who create them, e.g., assumptions about race or gender based on indicators like income or vocabulary. So, relying on AI as a decision-making tool, particularly for hiring, can expose your company to discrimination liability

risks.

Example: Amazon got rid of an AI-based recruitment program after discovering that its algorithm skewed against women. The model was programmed to vet candidates by observing patterns in resumes submitted to the company over a 10-year period. Most of the candidates in the training set were men. As a result, the AI taught itself that male candidates were preferred over female candidates.

• Advantages of Using AI for Hiring

AI technologies can help improve recruiting, [pre-employment screening](#) and retention in multiple ways, including for:

Creating More Targeted Job Descriptions: HR can use augmented writing platforms that analyze databanks of job postings and internal demographics to write focused job descriptions that are more likely to appeal to the kinds of applicants the employer is seeking to hire. target audience of the hiring employer.

Creating Job More Targeted Job Ads: AI advertising platforms typically enable advertisers to target the audience for their ads by gender, age, income, location, interests, activities, connections and other categories based on online usage and other personal data. Categorization tools come in different forms, such as drop-down menus, toggle buttons, search boxes or maps. For example, the ad placement interface may display a toggle button prompting the advertiser to select “men” or “women” as the potential audience.

Job Applicant [Screening](#): Keyword-matching software enable employers to sort through and screen out resumes and applications quickly and easily; more sophisticated products make it possible to compare applicants and their likelihood of succeeding in the job based on specific factors, which may include factors that the program itself has determined are indicative of likely success.

Job Applicant Communications: You can use interactive Chatbot products for basic tasks such as asking and answering applicants’ simple questions, scheduling interviews and sending feedback and reminders, which frees up HR to spend more quality time connecting with qualified applicants that the firm wants to recruit.

Evaluating Job Applicants: AI tools can be used at the interview stage to assess a job candidate’s competency, cognitive skills, personality traits and/or “fit” with the organization and its culture.