



The Future of Privacy Forum (FPF) is a non-profit organization that serves as a catalyst for privacy leadership and scholarship, advancing principled data practices in support of emerging technologies. Learn more about FPF by visiting fpf.org.

HIGHLIGHTS

1. Clearly distinguishes between developer and deployer responsibilities regarding AI and hiring
2. Prohibits secret use of AI tools to hire, terminate, and take other actions that have consequential impacts
3. Provides for testing to ensure AI hiring tools are fit for their intended purpose and assessed for bias
4. Highlights the importance of inclusive data sets in testing and bias assessment
5. Bars the use of AI hiring tools in a manner that harmfully discriminates, and urges the implementation of anti-discrimination protections that go beyond current law as needed
6. Warns against using facial characterization and emotion inference technologies in the hiring process absent public disclosures supporting the tools' efficacy, fairness, and fitness for purpose
7. Recognizes the sensitivity of hiring and employment data, and includes heightened privacy and security protections
8. Uses an AI governance framework informed by the NIST AI Risk Management Framework
9. Advises organizations against claiming AI hiring tools are "bias-free"
10. Requires that AI hiring tools be designed and operated with informed human oversight and engagement

TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
INTRODUCTION	2
BEST PRACTICES	3
1. Non-Discrimination	3
2. Responsible AI Governance	5
3. Transparency	7
4. Data Security and Privacy	9
5. Human Oversight	10
6. Alternative Review Procedures	10
APPENDIX A: DEFINITIONS	11
APPENDIX B: RELEVANT U.S. LAWS AND RULES	12
ENDNOTES	13

EXECUTIVE SUMMARY

Organizations are increasingly using AI tools as a part of their hiring and employment decisions (“AI tools”). These tools can help match candidates with relevant opportunities and inform organizations’ decisions about who to recruit, hire, and promote. More broadly, they can also help candidates discover and describe their skills, find new opportunities that match their experience, and suggest steps to position them for career growth. However, AI tools present risks that, if not addressed, can impact job candidates and hiring organizations, and pose challenges for regulators and other stakeholders. Like all hiring tools, AI tools are imperfect, and the stakes are high for decisions impacting employment.

The best practices set forth below provide guidance regarding AI tools that are used in ways that have consequential impacts on employment relationships, including consequential impacts in recruiting, hiring, promotion, or termination of an employment relationship. Specifically, the best practices provide guidance for organizations that develop and/or deploy AI tools in the employment context, addressing issues related to non-discrimination, responsible AI governance, transparency, data security and privacy, human oversight, and alternative review procedures.

Our key recommendations include:

- » Developers and deployers should have clearly defined responsibilities regarding AI hiring tools’ operation and oversight;
- » Organizations should not secretly use AI tools to hire, terminate, and take other actions that have consequential impacts;
- » AI hiring tools should be tested to ensure they are fit for their intended purpose and assessed for bias;
- » AI tools should not be used in a manner that harmfully discriminates, and organizations should implement anti-discrimination protections that go beyond laws and regulations as needed;
- » Organizations should not use facial characterization and emotion inference technologies in the hiring process absent public disclosures supporting the tools’ efficacy, fairness, and fitness for purpose;
- » Organizations should implement AI governance frameworks informed by the NIST AI Risk Management Framework;
- » Organizations should not claim that AI hiring tools are “bias-free;” and
- » AI hiring tools should be designed and operated with informed human oversight and engagement.

INTRODUCTION

In the last few years, organizations in the U.S. have incorporated artificial intelligence (AI) tools into their hiring and employment practices at an unprecedented pace. The use of automated technology in the workplace can result in faster hiring for employers, increased access to diverse candidates and a broader pool of applicants, and greater access to hiring tools for small to mid-sized businesses. For candidates, automated technology can help match their skills to a broader variety of roles and identify new potential career paths.

Nevertheless, the use of AI in workplace decision-making comes with the potential for serious negative impacts absent mitigation of risks. Organizations need to implement meaningful guardrails to ensure the responsible and ethical use of these systems. Vendors and employers must deal with risks of bias, which can permeate the entirety of employment processes. When properly designed and utilized, AI can and must process vast amounts of personal data fairly and ethically, and can create more equitable access for people with disabilities, and people from underrepresented, marginalized and multi-marginalized communities.¹ Because of the complexity of both issues and equities within the hiring and employment space, there is a clear need for best practices to guide the use of AI tools and employment for U.S. organizations.

Organizations using AI to make employment decisions need to acknowledge the risks and opportunities these tools can pose, especially to marginalized or underrepresented communities. One's employment directly impacts one's socioeconomic status; thus, landing a good job can change outcomes for individuals and their loved ones for generations.

Developers and deployers of AI in the employment context have an ethical responsibility to center humans as they develop AI tools, and use them in a manner compliant with civil rights, employment, and privacy laws, as well as the highest ethical standards.

The Future of Privacy Forum's *Best Practices for AI and Workplace Assessment Technologies* ("Best Practices") provides a policy framework for (1) non-discrimination, (2) responsible AI governance, (3) transparency, (4) privacy and data security, (5) human oversight, and (6) alternative review procedures. While existing law applies to the use of AI tools, and best practices are not legal advice, best practices are needed because the field of AI governance is still maturing. These best practices are meant to inform the broader AI governance field and advance the conversation about the specific use case of AI in employment, which currently lacks comprehensive benchmarks for best practices. As AI regulatory requirements, frameworks, and technical standards continue to mature, these best practices may be updated.

The Best Practices are informed by leading frameworks, including the EEOC's *Select Issues: Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964*,² NIST's *AI Risk Management Framework*, the *Civil Rights Principles for Hiring Assessment Technologies*,³ and the Data and Trust Alliance's initiative *Algorithmic Safety: Mitigating Bias in Workforce Decisions*.⁴

BEST PRACTICES



1 Non-Discrimination: AI tools should not be used in a manner that harmfully discriminates. Organizations should comply with anti-discrimination laws and regulations, and implement additional protections as needed.

- a) **Compliance with Civil Rights Laws:** The use of an AI tool with Consequential Impacts is governed by current anti-discrimination laws.⁵ Developers should assess an AI tool intended for use in contexts with Consequential Impacts in light of current laws and regulations. Deployers should ensure that their use of AI tools is compliant with non-discrimination laws.
- b) **Internal Testing:** Developers and Deployers should engage in internal testing that is aligned with current law.⁶
 - i) Developers that provide AI tools that have Consequential Impacts should, to the best of their ability, assess whether their tools are fit for the intended purpose and aligned with current legal standards. This includes, but is not limited to, testing for unintentional bias with respect to race, gender, sexual orientation, gender identity, disability, age, religion, socioeconomic status, and national origin. Testing should take into consideration the nuances of the established industry and legal standards as well as evolving testing practices.⁷
 - ii) Deployers of AI tools should:
 - (1) Ensure they understand whether the AI tool has Consequential Impacts, how the AI tool has been tested for discriminatory bias, and how the tool may be tested once deployed, consistent with existing legal obligations;
 - (2) Determine whether their implementation of the AI tools provided by Developers is likely to raise additional risks of discriminatory bias;
 - (3) If such risks are present, assess — to the best of their ability — whether their implementation of the tools is fit for the intended purpose and aligned with current legal standards; and
 - (4) Clearly articulate their processes to ensure that their use of AI tools that have Consequential Impacts is consistent with their obligations under anti-discrimination laws and is informed by EEOC guidance.
 - iii) To the extent possible in light of data protection, confidentiality, and other obligations, Developers should ensure that Deployers can access the data necessary to test for unintentional discrimination as it relates to the Deployer's implementation of AI tools that have Consequential Impacts, or alternatively, conduct such tests themselves for Deployers to use.⁸

- iv) Developers and Deployers should test their tools in a manner that is consistent with existing law and technical standards, which in most cases leverage disparate impact testing as the mechanism for testing for unlawful unintentional discrimination. Where there is insufficient data to test for unlawful unintentional discrimination, Developers and Deployers should be able to take the reasonable steps necessary to mitigate potential unlawful disparate impacts and to articulate steps as to how they did so. Similarly, for situations where standards and practices do not utilize disparate impact testing, Developers and Deployers should leverage appropriate alternative methods to the specific circumstances.⁹

c) Responsible Use Standards: Implement and develop responsible use standards for AI tools.

- i) Developers and Deployers should support the responsible use of AI tools that have Consequential Impacts and support the development of new standards as AI governance matures.
- ii) Developers and Deployers should exercise extreme caution before providing or using AI tools that have Consequential Impacts based on an analysis of an employee's or candidate's:
 - (1) emotional state, as inferred from their biometric data or other means;
 - (2) facial features or movements, body language, gait, tone of voice, vocal pitch, or pace of speech; or
 - (3) heart rate, respiration, or other bodily functions regulated by the autonomic nervous system.
- iii) If Developers and Deployers provide or use AI tools described in subsection (ii), they should:
 - (1) Publicly state the intended purpose(s) of the tools;
 - (2) Demonstrate or verify the factual analysis that supports claims about the tools' efficacy, fairness, and fitness for the intended purpose(s); and
 - (3) Publicly disclose a summary of this factual analysis in a manner that is easily available and understandable to Individuals who interact with the tools and to regulators.

d) Future Laws and Regulations: Compliance with legal requirements is a baseline expectation for organizations. Organizations should implement compliance frameworks to identify and promptly comply with relevant laws and regulations promulgated in the future.

2

Responsible AI Governance: Institute responsible AI governance practices.¹⁰

- a) **Responsible AI Practices:** Developers and Deployers should establish internal Responsible AI governance practices for AI tools that have Consequential Impacts. Governance frameworks should:
- i) Establish and document processes that map, measure, and manage reasonably foreseeable risks and harms to individuals that AI tools that have Consequential Impacts can pose;
 - ii) Address the full product life cycle of AI tools, recognizing the different roles and responsibilities of Developers and Deployers within the life cycle of a tool and relevant informational and technical limitations;
 - iii) Include an accountability and oversight structure that highlights potential risks and routes risks for analysis and potential mitigation;
 - iv) Include processes for ongoing feedback throughout the tool's life cycle that includes relevant internal and, as appropriate, external stakeholders. Such ongoing feedback should include, as appropriate: members of corporate leadership; engineers, developers, and other members of product teams; AI ethics practitioners; technical and legal experts; HR professionals with day-to-day responsibility for hiring, promotion, and termination processes; and experts in civil rights and employment law.
- b) **Organizational Roles and Responsibilities:** Developers and Deployers should establish clear roles and responsibilities for the teams who develop, manage, or implement AI tools that have Consequential Impacts, understanding that these may vary depending on whether the organization is a Developer, Deployer, or both.¹¹ Teams should build and implement responsible AI practices in concert. These roles typically include but are not limited to:
- i) **Role: Development**
 - (1) **Responsibility:** Develop products that seek to identify, quantify, and mitigate potential risks, including, but not limited to, risks to individuals, risks to communities and society, risks to the organization, and ethical risks. Those in development roles should be made aware of the potential for historical, sampling, and labeling bias in data sets generally, as well as any bias in the particular data sets they use, especially where the data sets use methods that attribute diverse characteristics (as compared to self-identification of individuals). Those in development roles should, where possible, take appropriate measures to detect and mitigate potential bias.
 - ii) **Role: Compliance**
 - (2) **Responsibility:** Mitigate risks arising from the use of AI tools. Individuals or teams tasked with compliance should clearly articulate the potential risks of an AI tool, identify the equity and legal considerations of the AI tool with respect to possible Consequential Impacts, and establish that the AI tool is fit for its intended purpose and consistent with legal and ethical standards. Compliance responsibilities should be operationally separate from the team developing or selling, implementing, or using the AI tool and those responsible for compliance activities should have a clear mechanism for identifying risks and communicating potential mitigation strategies to senior leadership.

iii) Role: Privacy

(3) **Responsibility:** Establish clear standards and processes that protect the data of Deployers and Individuals who use the AI tool. The privacy team should ensure that Personal Data is collected, used, and shared solely within the scope of the organization's privacy policy, and in a manner consistent with privacy laws. Organizations should pay special attention to ensuring that the use of Personal Data to train or improve the AI tool is documented, consistent with law and ethical standards, and does not create unreasonable risks of data disclosures to unauthorized individuals or entities.

c) Accountability: Ensure that there are internal mechanisms for risk management, monitoring, escalation, testing, and reporting and those systems are managed by a team or individual who verifies internal processes are followed.

i) Risk management: Developers and Deployers should have processes to manage both internal and external risks. Developers and Deployers should be able to consult, where appropriate, with outside experts to help mitigate risk, and flag practices that are deemed high risk.

ii) Internal escalation: Developers and Deployers should have internal structures that ensure that the AI tools are fit for purpose and allow for escalation if systems are deemed too high risk for use. Developers and Deployers should have the ability to, if need be, escalate new or emerging risks. Two ways to do that include both impact assessments and internal evaluations.

(1) Impact Assessments

(a) Developers and Deployers should conduct impact assessments¹² for their AI tools that have Consequential Impacts. Impact Assessments may include, but are not limited to:

- (i) Identification of the purpose and intended uses of the AI tool;
- (ii) The types of data used to train the AI tool;
- (iii) Mitigation measures for potential bias or other high-risk behaviors;
- (iv) Identified risks of the AI tool to individuals, communities, or society;
- (v) Testing structures that are in place for the AI tool post-deployment; and
- (vi) The mitigation structures that are in place if the AI tool engages in or is used for high-risk behavior.

(2) Monitoring Feedback

(a) Where possible, Developers should make it easy to provide feedback from Deployers and, as appropriate, Individuals, regarding issues and concerns related to the AI tools. Deployers should share feedback regarding issues where applicable.

(3) Internal Evaluations

- (a) Internal evaluations or other internal compliance structures that allow for independent feedback mechanisms should include:
 - (i) Places to test (or sandbox), validate, and recommend corrections to AI tools;
 - (ii) Structures that allow team members to escalate the AI tool's potential risks or harms; and
 - (iii) Structures that allow for evaluation criteria with respect to bias, risk, and accuracy of AI tools.
- iii) **Internal Controls on Data Sets:** Developers of AI tools should establish internal practices to ensure AI tools use the most diverse and representative data sets available in consideration of the tool's intended use. Internal controls should be in place to clarify what data should or should not be included in the use of the AI tool. Developers and Deployers should train product, privacy, and compliance staff to understand the bias risks in large-scale data sets and actively strive to include marginalized and underrepresented communities in their data sets.
- iv) **Testing:** Where possible, and without requiring access to data sets not native to Deployers, Developers should enable Deployers to test for bias when using AI tools that have Consequential Impacts. Discrimination is not the only form of bias that organizations need to address since there are different sources of bias. Other forms of bias such as formulation bias, historical bias, sampling bias, labeling bias, proxy bias, aggregation bias, deployment bias, and misuse bias can have a significant impact on both data sets and AI tools.¹³ Deployers should be able to test the AI tool using their own data sets or independent data sets.

3

Transparency: Provide disclosures to those who interact with and are impacted by the use of an AI tool with Consequential Impacts.

- a) **Baseline Principles:** Developers and Deployers each have important roles in ensuring that Individuals understand when — and to what extent — AI tools have Consequential Impacts, what alternative options are available to all Individuals, and what accommodations are available to Individuals with disabilities. Particular disclosures should be provided by the entity that is best positioned to develop the content of the disclosure and communicate it to Individuals.
- b) **Developer Transparency:** Developers are often best positioned to explain how an AI tool works, how it was trained, the limitations of its effectiveness (e.g., the tool is not a substitute for decision making by a human), and how a tool may be used or configured.
 - i) Developers providing AI tools that have Consequential Impacts should be clear about the choices available to Deployers and Individuals. Developers should provide Deployers with information about the ways in which the AI tool is fit for purpose, addresses bias, calculates risk, and attempts to limit harm. When feasible, Developers should provide impacted Individuals with this information as well.

- ii) When feasible, Developers should provide links to the Developer’s responsible AI practices and the Developer’s AI governance framework.
 - iii) When communicating about the capabilities of their AI tools, Developers should take care to substantiate their claims and avoid statements that their tools are “bias free.”¹⁴
- c) Deployer Transparency:** Deployers are often best positioned to explain to Individuals how they implement an AI tool and how an AI tool fits into the Deployer’s overall decision-making processes regarding use of AI tools that have Consequential Impacts. When applicable, Deployers should ensure that transparency flows downward by prominently displaying Developer disclosures to Individuals, including Developer disclosures that enable Individuals to understand that they are interacting with an AI tool, how it works (including, to the extent possible, the reasoning behind how the tool could have a Consequential Impact on the Individual), how it was trained, and what choices Individuals have when they interact with the tool. Deployers should be particularly mindful of disclosing to Individuals any Deployer practices that differ from Developer practices in relation to the deployed tool and its use.
- d) Specific Disclosures:** Developers and Deployers should, when applicable, make the following specific disclosures.
- i) Developers providing AI tools with Consequential Impacts should disclose to Deployers when applicable:
 - (1) the intended purposes of the AI tool;
 - (2) purposes for which the AI tool is not intended;
 - (3) known efficacy limits of the AI tool;
 - (4) how the AI tool was trained;
 - (5) whether the AI tool was assessed for potential discriminatory bias;
 - (6) whether the AI tool uses information from Deployers or Individuals to further train or otherwise improve the tool;
 - (7) how the AI tool is intended to be deployed;
 - (8) uses of the AI tool that are not intended;
 - (9) what choices the AI tool provides to Deployers regarding anti-discrimination, governance, transparency to Individuals, privacy, security, and human oversight; and
 - (10) what choices the AI tool provides to Deployers to communicate to Individuals about how they implement the tool, and how the tool fits into the Deployer’s overall decision-making processes regarding Consequential Impacts.
 - ii) Deployers using AI tools with Consequential Impacts should disclose to Individuals when applicable:
 - (1) the fact that Individuals are interacting with an AI tool;
 - (2) the intended use of the AI tool (e.g., to evaluate job candidates, make compensation decisions, or consider employees for promotion);
 - (3) how the AI tool was trained;
 - (4) how an AI tool may have a Consequential Impact and how the tool fits into the Deployer’s overall decision-making processes;
 - (5) the extent to which Individuals’ Personal Data is shared with third parties or used to train or improve the AI Tool; and
 - (6) what alternative options are available to all Individuals, and how individuals with disabilities may seek accommodations.
 - iii) Disclosures should be understandable by non-technical audiences, accessible, and accurate.¹⁵ Disclosures should not reveal private information about Individuals or confidential or trade secret information regarding Developers, Deployers, or other organizations.

4

Data Security and Privacy: Protect Personal Data while at rest and in transit and ensure the confidentiality and integrity of Personal Data when used in ways that have Consequential Impacts. Use Personal Data only consistent with individuals’ choices, legal obligations, and privacy policies.

- a) **Comprehensive Programs:** Developers and Deployers of AI tools with Consequential Impacts should maintain comprehensive privacy and security programs designed to protect Personal Data.
 - i) The programs should be reasonably designed to protect the security, privacy, confidentiality, and integrity of Personal Data against risks — such as unauthorized access or use, or unintended or inappropriate disclosure or breach — through the use of administrative, technical, and physical safeguards appropriate to the sensitivity of the information.
 - ii) The programs should address risks arising from the use of AI tools, including risks arising from the use of large language models and generative AI when applicable.

- b) **Data Security:** Developers and Deployers should implement data security practices that include meaningful safeguards against malicious or unauthorized access to Personal Data and confidential information or trade secrets exchanged between Developers and Deployers. Security tools, when possible, should account for both known threat vectors as well as knowable future threat vectors for which mitigation measures may be available. Developers should provide reasonable guidance on what should and, when appropriate, what should not be done with Developer tools.
 - i) Data security practices should include but not be limited to:
 - (1) secure storage of Personal Data;
 - (2) encryption of confidential digital records in transit;
 - (3) guarding against injection attacks into the model;
 - (4) data-use agreements;
 - (5) contractual obligations; and
 - (6) accountability measures (e.g. training, access controls, and logs).

- c) **Privacy:** The sensitive nature of employment decisions obligates Developers and Deployers to institute systems and structures around AI tools that ensure Personal Data is not used in a way inconsistent with legal obligations, the Fair Information Practice Principles (including the principle of data minimization), and with other protections that are appropriate to the data’s sensitivity.
 - i) Developers and Deployers should protect Personal Data when using AI tools and ensure that AI tools do not inadvertently expose Personal Data.
 - ii) Developers and Deployers should ensure that Personal Data is handled consistently with applicable laws, data agreements, and privacy policies. Developers and Deployers should ensure that Personal Data is not used or disclosed in a manner inconsistent with Individuals’ choices.
 - iii) Developers should give Deployers the means to understand how Developers use Personal Data obtained from Deployers. Developers should ensure that their use of Personal Data about Individuals obtained from Deployers complies with privacy and data security laws.

5

Human Oversight: AI tools with Consequential Impacts should be designed and operated with informed human oversight and engagement.

- a) **Human in the Loop:** AI tools should be designed by Developers and implemented by Deployers with an informed “human in the loop” to enhance the explainability, transparency, and accountability of processes that have Consequential Impacts. AI tools are best used to augment or enhance human decision-making for processes that have Consequential Impacts, not replace human decision-making.
- b) **Human Involvement:** AI tools should be operated by Deployers with human oversight to guard against unfair and illegal employment practices. Humans should be accountable for using AI tools in connection with employment and assess the impact of AI tools on individuals.
- c) **Human Oversight:** Effective human oversight of AI tools will be implemented differently for:¹⁶
 - i) Developers and Deployers;
 - ii) different Consequential Impacts; and
 - iii) various purposes and use cases for AI tools.
- d) **Clarifying Roles and Responsibilities for Processes Involving Consequential Impacts:** Developers have different responsibilities in the hiring process than Deployers. Developers should be responsible for creating AI tools that allow Deployers to have human oversight when Deployers are the ultimate decision-makers for employment outcomes.

6

Alternative Review Procedures

Developers should design AI tools that have Consequential Impacts with alternative review procedures in mind. Deployers may be legally required to provide Individuals with alternative review procedures, for example, by configuring the AI tool such that it reasonably accommodates Individuals with disabilities or allows for alternative tools and procedures altogether.¹⁷

- a) **Practicability:** Alternatives may not be practicable in certain AI use cases such as recommendations or search results.

Appendix A

DEFINITIONS

Artificial Intelligence¹⁸ (AI) — An AI system is a machine-based system that is capable of influencing the environment by producing an output (predictions, recommendations, or decisions) for a given set of objectives. It uses machine and/or human-based data and inputs to (i) perceive real and/or virtual environments; (ii) abstract these perceptions into models through analysis in an automated manner (e.g., with machine learning), or manually; and (iii) use model inference to formulate options for outcomes. AI systems are designed to operate with varying levels of autonomy.

Consequential Impacts — An activity that has a legal, dispositive, or other similarly significant impact on an individual’s employment status, including in recruiting,¹⁹ hiring, promotion, or termination of an employment relationship.

Deployer — An entity that uses an AI tool that has Consequential Impacts.

Developer — An entity that designs, codes, creates, or modifies a tool that produces or is intended to produce a Consequential Impact, whether for internal use or for use by third parties. A Deployer’s use of Developer-provided software options does not constitute “modification” of an AI tool.

Individual — An applicant, candidate, or employee.

Human in the Loop — Humans as a part of the design and operation process of AI systems or tools that are accountable to people.

Personal Data — Information that is related to an identified or identifiable person.

Appendix B

RELEVANT U.S. LAWS AND RULES

Anti-Discrimination Laws

Title VII

Title VII of the Civil Rights Act of 1964 (Title VII) prohibits employers from discriminating against someone based on race, color, religion, national origin, or sex.²⁰ Forms of discrimination may include “failure or refusal to hire or to discharge any individual,” “discrimination with respect to his compensation, terms, conditions, or privileges,” or “to limit, segregate, or classify employees.”²¹ While this law largely applies to intentional discrimination, employers may be liable for violating Title VII on a disparate impact theory of liability.²² Courts use a burden shifting framework to adjudicate disparate impact cases. A plaintiff must show that an employer uses a “particular employment practice that causes a disparate impact on the basis of race, color, religion, sex, or national origin. The respondent must then demonstrate that the challenged practice is job related for the position in question and consistent with business necessity.”²³ A plaintiff may also bring a case if they present an “alternative employment practice” and the employer refuses to adopt such alternative employment practice.²⁴ An alternative employment practice is another method available to an employer that is equally effective but less discriminatory.²⁵ Nearly 15 years after the Act was passed, the Pregnancy Discrimination Act of 1978 amended Title VII to prohibit sex discrimination on the basis of pregnancy, childbirth, or a medical condition related to pregnancy or childbirth.²⁶

Equal Pay Act of 1963

The Equal Pay Act of 1963 prohibits discrimination between employees on the basis of sex “by paying wages to employees in such establishment at a rate less than the rate at which he pays wages to employees of the opposite sex in such establishment for equal work on jobs the performance of which requires equal skill, effort,

and responsibility, and which are performed under similar working conditions.”²⁷ This does not include payments made pursuant to a seniority system, a merit system, a sales system, or a differential based on any other factor other than sex.²⁸

Americans with Disabilities Act

The American with Disabilities Act (ADA) prohibits discrimination against a qualified individual with a disability, by reason of such disability, regarding job application procedures, the hiring, advancement, or discharge of employees, employee compensation, job training, and other terms, conditions, and privileges of employment.²⁹ The Act further prevents public entities from discriminating against disabled individuals.³⁰

Age Discrimination in Employment Act of 1967

The Age Discrimination in Employment Act of 1967 (ADEA) makes it unlawful for an employer to “fail or refuse to hire or to discharge any individual or otherwise discriminate against any individual with respect to his compensation, terms, conditions, or privileges of employment, because of such individual’s age.”³¹ Furthermore, an employer may not “limit, segregate, or classify his employees in any way which would deprive or tend to deprive any individual of employment opportunities or otherwise adversely affect his status as an employee, because of such individual’s age.”³²

Title II of the Genetic Information Nondiscrimination Act of 2008

Title II of the Genetic Information Nondiscrimination Act of 2008 makes it illegal for employers to discriminate based on genetic information.³³ This adds genetic information to the list of characteristics, outlined in Title VII, that employers may not use to inform hiring decisions. Employers may not request, require, or purchase genetic information of an employee or a family member of an employee, except under certain circumstances.³⁴

ENDNOTES

- 1 Certain underrepresented, marginalized, and multi-marginalized communities have a long and demonstrable history of bias and discrimination in the workplace at large. See, e.g., *Race in the Workplace*, McKinsey & Company (February 2021), <https://www.mckinsey.com/featured-insights/diversity-and-inclusion/race-in-the-workplace-the-black-experience-in-the-us-private-sector>. The use of those terms in this document is both to include all of those recognized as part of a protected community under the law (for instance, those of different races (White, Black, Latino, Asian American and Pacific Islander, Indigenous, etc.), people with disabilities, members of the LGBTQIA+ community, women, or those who are pregnant), those people who intersect between those communities, and other communities who may receive legal protection in the future.
- 2 *Select Issues: Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964*, U.S. Equal Opportunity Employment Commission (July 2023), <https://www.eeoc.gov/select-issues-assessing-adverse-impact-software-algorithms-and-artificial-intelligence-used>.
- 3 *AI Risk Management Framework*, National Institute of Standards and Technology (January 2023), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
- 4 *Algorithmic Safety: Mitigating Bias in Workforce Decisions, Data and Trust Alliance* (December 2021), <https://dataandtrustalliance.org/our-initiatives/algorithmic-safety-mitigating-bias-in-workforce-decisions>.
- 5 *Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems*, U.S. Consumer Financial Protection Bureau, U.S. Department of Justice, U.S. Equal Employment Opportunity Commission, and U.S. Federal Trade Commission (May 2023), https://www.ftc.gov/system/files/ftc_gov/pdf/EEOC-CRT-FTC-CFPB-AI-Joint-Statement%28final%29.pdf.
- 6 There is an inherent tension between privacy and the need for more data to better train AI tools. Certain protected categories are not well represented in broader data sets, which can lead to inequitable outcomes for some workers. At the same time, protected category data is typically treated as highly sensitive and subject to heightened protections under privacy laws. Organizations welcome the input of policymakers and civil society organizations for ways to procure more representative and non-biased data and training sets.
- 7 For example, while disparate impact testing is a standard approach for assessing unintentional race and gender discrimination in hiring, it is not a standard approach for assessing bias based on certain protected characteristics, including disability and sexual orientation discrimination in hiring. Developers should engage experts so that their practices reflect these nuances in existing legal standards and so they modify AI tools or provide appropriate alternative assessment tools to mitigate discrimination risks. At the same time, developers may not have access to the data necessary to do this testing and may be reliant on Deployers who are in a better position to provide relevant data. In such cases, Developers and Deployers may need to cooperate to do such testing and effective mitigation. We note that disparate impact testing does not assess whether a particular tool is job-related and consistent with business necessity, even if it does have adverse impact.
- 8 For example, where the nature of a tool obligates Deployers to perform testing for unintentional discrimination, subject to privacy, confidentiality, and intellectual property considerations, Developers should ensure Deployers have the means to export all data fields necessary to perform such assessment.
- 9 The 4/5ths rule is defined as “[the] selection rate for any race, sex, or ethnic group which is less than four-fifths (4/5) (or eighty percent) of the rate for the group with the highest rate will generally be regarded by the Federal enforcement agencies as evidence of adverse impact, while a greater than four-fifths rate will generally not be regarded by Federal enforcement agencies as evidence of adverse impact.” (29 CFR § 1607.4) There is a circuit split and a variety of opinions with respect to the 4/5ths rule, with some jurisdictions considering the 4/5ths rule to be the bare minimum and some considering it to be a good standard. Having the 4/5ths rule as the baseline for disparate impact testing, and not the high water mark, is both a best practice and legally compliant in all jurisdictions.
- 10 Audits can be a valuable way to assess a governance program or digital tool to determine whether it meets a defined standard or criteria. Both Developers and Deployers, for example, engage in audits of their privacy and cybersecurity controls and often use third-party auditors to assess and communicate the robustness of these controls to stakeholders. The AI auditing field is currently maturing, as there are neither consensus technical standards nor a common set of criteria to audit against. Nor are there widely accepted professional standards that are binding on third-party auditors, which are necessary for auditing integrity and to address potential malfeasance. As noted in the U.S.’s recent AI National Research & Development Plan, credible consensus standards can and must be established, as they are necessary to address the “significant practical challenges” for AI audits to be viable at scale. Developers and Deployers should contribute to the development of such standards. Future iterations of the Best Practices may incorporate third-party auditing, if these essential preconditions are in place.
- 11 See BSA | The Software Alliance, “AI Developers and Deployers: An Important Distinction” (March 2023) <https://www.bsa.org/files/policy-filings/03162023aidevdep.pdf>.
- 12 See Meghan Chilappa and Dileep Srihari, *Impact Assessments: Supporting AI Accountability & Trust*, Access Partnership (March 20, 2023) <https://www.workday.com/content/dam/web/en-us/documents/legal/access-partnership-workday-impact-assessment-paper.pdf>
- 13 See *Confronting Bias: BSA’s Framework to Build Trust in AI*, BSA: The Software Alliance (June 8, 2021) <https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai>.
- 14 See Elisa Jillson, *Aiming for truth, fairness, and equity in your company’s use of AI*, Federal Trade Commission (April 19, 2021) <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>; Michael Atleson, *Keep your AI claims in check*, Federal Trade Commission (February 27, 2023) <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check>.

-
- 15 See Phillips et al, “NISTIR 8312: Four Principle of Explainable Artificial Intelligence” (September 2021) <https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8312.pdf>.
- 16 See *OECD Framework for the Classification of AI systems*, Organisation for Economic Co-operation and Development (February 22, 2022) https://www.oecd-ilibrary.org/science-and-technology/oecd-framework-for-the-classification-of-ai-systems_cb6d9eca-en;jsessionid=qffqX8g_5UI4LiDU6YKpTNd8lqeXZ4Wcgvop58we.ip-10-240-5-69 (AI in the lab” versus “AI in the field” distinction).
- 17 See EEOC, “The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees.” <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>.
- 18 See William M. Mac Thornberry National Defense Authorization Act for Fiscal Year 2021 Conference Report, p. 1164 (December 3, 2020) <https://www.congress.gov/116/crpt/hrpt617/CRPT-116hrpt617.pdf#page=1210>, Recommendation of the Council on Artificial Intelligence, Organisation for Economic Co-operation and Development (May 21, 2019) <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.
- 19 Recruiting may include a variety of activities. To the extent recruiting activities prevent a job seeker from applying to employment opportunities, those activities could be considered Consequential Impacts.
- 20 42 U.S.C. § 2000e.
- 21 42 U.S.C. § 2000e-2.
- 22 42 U.S.C. § 2000e-2(k)(1)(A).
- 23 42 U.S.C. § 2000e-2(k)(1)(A)(i).
- 24 42 U.S.C. § 2000e-2(k)(1)(A)(ii)(C).
- 25 *Alternative Employment Practice Legal Meaning*, Quimbee (last visited Oct. 18, 2022), <https://www.quimbee.com/keyterms/alternative-employment-practice>.
- 26 See Pub. L. No. 95-555, 92 Stat. 2076.
- 27 29 U.S.C. § 206(d).
- 28 *Id.*
- 29 See 42 U.S.C. § 12112.
- 30 See 42 U.S.C. § 12132.
- 31 29 USC § 623(a).
- 32 *Id.*
- 33 See 42 U.S.C. § 2000ff(a).
- 34 See 42 U.S.C. § 2000ff(b).

