STATE OF NEW YORK

1169

2025-2026 Regular Sessions

IN SENATE

January 8, 2025

Introduced by Sen. GONZALEZ -- read twice and ordered printed, and when printed to be committed to the Committee on Internet and Technology

AN ACT to amend the civil rights law and the executive law, in relation to the use of artificial intelligence systems

The People of the State of New York, represented in Senate and Assembly, do enact as follows:

Section 1. This act shall be known and may be cited as the "New York artificial intelligence act (New York AI act)".

§ 2. Legislative findings and intent. The legislature finds and declares the following:

(a) A revolution in artificial intelligence (AI) has advanced to the point that comprehensive regulations must be enacted to protect New Yorkers.

(b) Artificial intelligence is already an integral part of New Yorkers' daily lives. In the private sector, AI is currently in use in areas such as education, health care, employment, insurance, credit scoring, public safety, retail, banking and financial services, media, and more with little transparency or oversight. A growing body of research shows that AI systems that are deployed without adequate testing, sufficient oversight and robust guardrails can harm consumers and deny historically disadvantaged groups the full measure of their civil rights and liberties, thereby further entrenching inequalities. The legislature must act to ensure that all uses of AI, especially those that affect important life chances, are free from harmful biases, protect our privacy, and work for the public good.

(c) Safe innovation must remain a priority for the state. New York state is home to thousands of technology start-ups, many of which experiment with new applications of AI and which have the potential to find new ways to employ technology at the service of New Yorkers. The goal of the legislature is to encourage safe innovation in the AI sector by providing clear guidance for AI development, testing, and validation

EXPLANATION--Matter in *italics* (underscored) is new; matter in brackets [-] is old law to be omitted.

LBD04409-01-5

both before a product is launched and throughout the product's life cycle.

(d) New York must establish that the burden of responsibility of proving that AI products do not cause harm to New Yorkers will be shouldered by the developers and deployers of AI. While government and civil society must act to audit and enforce human rights laws around the use of AI, the companies employing and profiting from the use of AI must lead in ensuring that their products are free from algorithmic discrimination.

(e) Close collaboration and communication between New York state and industry partners is key to ensuring that innovation can occur with safeguards to protect all New Yorkers. This legislation will ensure that lines of communication exist and that there is clear statutory authority to investigate and prosecute entities that break the law.

(f) As new forms of AI are developed beyond what is currently technologically feasible, the goal of the legislature is to use this section as a guiding light for future regulations.

(g) Lastly, it is in the interest of all New Yorkers that certain uses of AI that infringe on fundamental rights, deepen structural inequality, or that result in unequal access to services shall be banned.

§ 3. The civil rights law is amended by adding a new article 8-A to read as follows:

ARTICLE 8-A

PROTECTIONS REGARDING USE OF ARTIFICIAL INTELLIGENCE

Section 85. Definitions.
86. Unlawful discriminatory practices.
86-a. Deployer and developer obligations.
86-b. Whistleblower protections.
87. Audits.
88. High-risk AI system reporting requirements.
89. Risk management policy and program.
89-a. Social scoring AI systems prohibited.
89-b. Enforcement.

§ 85. Definitions. The following terms shall have the following meanings:

1. "Algorithmic discrimination" means any condition in which the use of an AI system contributes to unjustified differential treatment or impacts, disfavoring people based on their actual or perceived age, race, ethnicity, creed, religion, color, national origin, citizenship or immigration status, sexual orientation, gender identity, gender expression, military status, sex, disability, predisposing genetic characteristics, familial status, marital status, pregnancy, pregnancy outcomes, disability, height, weight, reproductive health care or autonomy, status as a victim of domestic violence or other classification protected under state or federal laws. Algorithmic discrimination shall not include:

(a) a developer's or deployer's testing of their own AI system to identify, mitigate, and prevent discriminatory bias;

(b) expanding an applicant, customer, or participant pool to increase diversity or redress historical discrimination; or

(c) an act or omission by or on behalf of a private club or other establishment that is not in fact open to the public, as set forth in Title II of the federal Civil Rights Act of 1964, 42 U.S.C. section 2000a(e), as amended.

2. "Artificial intelligence system" or "AI system" means a machine-based system or combination of systems, that for explicit and implicit objectives, infers, from the input it receives, how to generate outputs

such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Artificial intelligence shall not include any software used primarily for basic computerized processes, such as anti-malware, anti-virus, auto-correct functions, calculators, databases, data storage, electronic communications, fire-wall, internet domain registration, internet website loading, network-ing, spam and robocall-filtering, spellcheck tools, spreadsheets, web caching, web hosting, or any tool that relates only to internal manage-ment affairs such as ordering office supplies or processing payments, and that do not materially affect the rights, liberties, benefits, safe-ty or welfare of any individual within the state.

3. "Auditor" shall refer to an independent entity including but not limited to an individual, non-profit, firm, corporation, partnership, cooperative, or association commissioned to perform an audit.

4. "Consequential decision" means a decision or judgment that has a material, legal or similarly significant effect on an individual's life relating to the impact of, access to, or the cost, terms, or availabili-ty of, any of the following:

(a) Employment, workers' management, or self-employment, including, but not limited to, all of the following:

(i) Pay or promotion;

(ii) Hiring or termination; and

(iii) Automated task allocation.

(b) Education and vocational training, including, but not limited to, all of the following:

(i) Assessment or grading, including, but not limited to, detecting student cheating or plagiarism;

(ii) Accreditation;

(iii) Certification;

(iv) Admissions; and

(v) Financial aid or scholarships.

(c) Housing or lodging, including rental or short-term housing or lodging.

(d) Essential utilities, including electricity, heat, water, internet or telecommunications access, or transportation.

(e) Family planning, including adoption services or reproductive services, as well as assessments related to child protective services.

(f) Health care or health insurance, including mental health care, dental, or vision.

(g) Financial services, including a financial service provided by a mortgage company, mortgage broker, or creditor.

(h) Law enforcement activities, including the allocation of law enforcement personnel or assets, the enforcement of laws, maintaining public order, or managing public safety.

(i) Government services.

(j) Legal services.

5. "Deployer" means a person, partnership, association or corporation that uses an AI system or commerce in the state of New York or provides an AI system for use by the general public in the state of New York. A developer may also be considered a deployer if its actions satisfy this definition.

6. "Deployer-employer" means a deployer that is an employer.

7. "Developer" means a person, partnership, or corporation that designs, codes, or produces an AI system, or creates a substantial change with respect to an AI system, whether for its own use in the state of New York or for use by a third party in the state of New York.

8. "Developer-employer" means a developer that is an employer.

9. "Employee" means an individual who performs services for and under the control and direction of an employer for wages or other remuneration, including former employees, or natural persons employed as independent contractors to carry out work in furtherance of an employer's business enterprise who are not themselves employers.

10. "Employer" means any person, firm, partnership, institution, corporation, or association that employs one or more employees.

11. "End user" means any individual or group of individuals that:

(a) is the subject of a consequential decision made entirely by or with the assistance of an AI system; or

(b) interacts, directly or indirectly, with the relevant AI system on behalf of an individual or group that is the subject of a consequential decision made entirely by or with the assistance of an AI system.

12. "High-risk AI system" means any AI system that, when deployed: (a) is a substantial factor in making a consequential decision; or (b) will have a material impact on the statutory or constitutional rights, civil liberties, safety, or welfare of an individual in the state.

13. "Software stack" means the group of individual software components that work together to support the execution of an AI system.

14. "Substantial change" means any (a) deliberate modification to an AI system that would result in material inaccuracies in the reports created under section eighty-eight of this article; or (b) unintentional and substantial change in the data that the AI system uses as input data.

15. "Substantial factor" means a factor that assists in making a consequential decision or is capable of altering the outcome of a consequential decision. "Substantial factor" includes, but is not limited to, any use of an AI system to generate any content, decision, prediction, or recommendation that is used as a basis, in whole or in part, to make a consequential decision regarding an end user.

§ 86. Unlawful discriminatory practices. It shall be an unlawful discriminatory practice:

1. for a developer or deployer to use, sell, or share a high-risk AI system or a product featuring a high-risk AI system that produces algorithmic discrimination; or

2. for a developer to use, sell, or share a high-risk AI system or a product featuring a high-risk AI system that has not passed an independent audit, in accordance with section eighty-seven of this article, that has found that the product does not in fact produce algorithmic discrimination.

§ 86-a. Deployer and developer obligations. 1. (a) Any deployer that employs a high-risk AI system for a consequential decision must inform the end user at least five business days prior to the use of such system for the making of a consequential decision in clear, conspicuous, and consumer-friendly terms, made available in each of the languages in which the company offers its end services, that AI systems will be used to make a decision or to assist in making a decision. The deployer must allow sufficient time and opportunity in a clear, conspicuous, and consumer-friendly manner for the consumer to opt-out of the automated process and for the decision to be made by a human representative. A consumer may not be punished or face any other adverse action for opting out of a decision by an AI system and the deployer must render a decision to the consumer within forty-five days.

(b) If a deployer employs a high-risk AI system for a consequential decision to determine whether to or on what terms to confer a benefit on

an end user, the deployer shall offer the end user the option to waive their right to advance notice of five business days under this subdivision.

(c) If the end user clearly and affirmatively waives their right to five business days' notice, the deployer shall then inform the end user at least one business day before the making of the consequential decision in clear, conspicuous, and consumer-friendly terms, made available in each of the languages in which the company offers its end services, that AI systems will be used to make a decision or to assist in making a decision. The deployer must allow sufficient time and opportunity in a clear, conspicuous, and consumer-friendly manner for the consumer to opt-out of the automated process and for the decision to be made by a human representative. A consumer may not be punished or face any other adverse action for opting out of a decision by an AI system and the deployer must render a decision to the consumer within forty-five days.

2. Any deployer that employs a high-risk AI system for a consequential decision must inform the end user within five days in a clear, conspicuous, and consumer-friendly manner if a consequential decision has been made entirely by or with assistance of an automated system. The deployer must then provide and explain a process for the end user to appeal the decision, which must at minimum allow the end user to (a) formally contest the decision, (b) provide information to support their position, and (c) obtain meaningful human review of the decision. A deployer must respond to an end user's appeal within forty-five days of receipt of the appeal. That period may be extended once by forty-five additional days where reasonably necessary, taking into account the complexity and number of appeals. The deployer must inform the end user of any such extension within forty-five days of receipt of the appeal, together with the reasons for the delay.

3. The deployer or developer of a high-risk AI system is legally responsible for quality and accuracy of all consequential decisions made, including any bias, algorithmic discrimination, and/or misinformation resulting from the operation of the AI system.

4. The rights and obligations under this section may not be waived by any person, partnership, association or corporation.

§ 86-b. Whistleblower protections. 1. Developer-employers and/or deployer-employers of high-risk AI systems shall not:

(a) prevent an employee from disclosing information to the attorney general, including through terms and conditions of employment or seeking to enforce terms and conditions of employment, if the employee has reasonable cause to believe the information indicates a violation of this article; or

(b) retaliate against an employee for disclosing information to the attorney general pursuant to this section.

2. An employee harmed by a violation of this article may petition a court for appropriate relief as provided in subdivision five of section seven hundred forty of the labor law.

3. Developer-employers and deployer-employers of high-risk AI systems shall provide a clear notice to all employees working on such AI systems of their rights and responsibilities under this article, including the right of employees of contractors and subcontractors to use the developer's internal process for making protected disclosures pursuant to subdivision four of this section. A developer-employer or deployer-employer is presumed to be in compliance with the requirements of this

subdivision if the developer-employer or deployer-employer does either of the following:

(a) at all times post and display within all workplaces maintained by the developer-employer or deployer-employer a notice to all employees of their rights and responsibilities under this article, ensure that all new employees receive equivalent notice, and ensure that employees who work remotely periodically receive an equivalent notice; or

(b) no less frequently than once every year, provides written notice to all employees of their rights and responsibilities under this article and ensures that the notice is received and acknowledged by all of those employees.

4. Each developer-employer and deployer-employer shall provide a reasonable internal process through which an employee may anonymously disclose information to the developer if the employee believes in good faith that the information indicates that the developer has violated any provision of this article or any other law, or has made false or materially misleading statements related to its safety and security protocol, or failed to disclose known risks to employees, including, at a minimum, a monthly update to the person who made the disclosure regarding the status of the developer's investigation of the disclosure and the actions taken by the developer in response to the disclosure.

5. This section does not limit protections provided to employees under section seven hundred forty of the labor law.

§ 87. Audits. 1. Prior to deployment of a high-risk AI system, six months after deployment, and at least every eighteen months thereafter for each calendar year a high-risk AI system is in use after the first post-deployment audit, every developer or deployer of a high-risk AI system shall cause to be conducted at least one third-party audit in compliance with the provisions of this section to ensure that the product does not produce algorithmic discrimination and complies with the provisions of this article. Regardless of final findings, the deployer or developer shall deliver all audits conducted under this section to the attorney general.

2. A deployer or developer may hire more than one auditor to fulfill the requirements of this section.

3. The audit shall include the following:

(a) an analysis of data management policies including whether personal or sensitive data relating to a consumer is subject to data security protection standards that comply with the requirements of section eight hundred ninety-nine-bb of the general business law;

(b) an analysis of the system accuracy and reliability according to each specified use case listed in the entity's reporting document filed by the developer or deployer under section eighty-eight of this article;

(c) disparate impacts and a determination of whether the product produces algorithmic discrimination in violation of this article by each intended and foreseeable identified use as identified by the deployer and developer;

(d) analysis of how the technology complies with existing relevant federal, state, and local privacy and data privacy laws; and

(e) an evaluation of the developer's or deployer's documented risk management policy and program required under section eighty-nine of this article for conformity with subdivision one of such section eighty-nine.

4. The attorney general may promulgate further rules as necessary to ensure that audits under this section assess whether or not AI systems produce algorithmic discrimination and otherwise comply with the provisions of this article.

5. The independent auditor shall have complete and unredacted copies of all reports previously filed by the deployer or developer under section eighty-eight of this article.

6. An audit conducted under this section shall be completed in its entirety without the assistance of an AI system.

7. (a) An auditor shall be an independent entity including but not limited to an individual, non-profit, firm, corporation, partnership, cooperative, or association.

(b) For the purposes of this article, no auditor may be commissioned by a developer or deployer of an AI system if such entity has already been commissioned to provide any auditing or non-auditing service, including but not limited to financial auditing, cybersecurity auditing, or consulting services of any type, to the commissioning company in the past twelve months.

(c) Fees paid to auditors may not be contingent on the result of the audit and the commissioning company shall not provide any incentives or bonuses for a positive audit result.

8. The attorney general may promulgate further rules to ensure (a) the independence of auditors under this section, and (b) that teams conducting audits incorporate feedback from communities that may foreseeably be the subject of algorithmic discrimination with respect to the AI system being audited.

§ 88. High-risk AI system reporting requirements. 1. Every developer and deployer of a high-risk AI system shall comply with the reporting requirements of this section. Regardless of final findings, reports shall be filed with the attorney general prior to deployment of a high-risk AI system and then annually, or after each substantial change to the system, whichever comes first.

2. Together with each report required to be filed under this section, developers and deployers shall file with the attorney general a copy of the last completed independent audit required by this article and a legal attestation that the high-risk AI system: (a) does not violate any provision of this article; or (b) may violate or does violate one or more provisions of this article, that there is a plan of remediation to bring the high-risk AI system into compliance with this article, and a summary of such plan of remediation.

3. Developers of high-risk AI systems shall file with the attorney general a report containing the following:

(a) a description of the system including:

(i) a description of the system's software stack;

(ii) the purpose of the system;

(iii) the system's use to end users; and

(iv) reasonably foreseeable uses outside of the current or intended uses;

(v) how the system should be used or not used;

(b) the intended outputs of the system and whether the outputs can be or are otherwise appropriate to be used for any purpose not previously articulated;

(c) the methods for training of their models including:

(i) any pre-processing steps taken to prepare datasets for the training of a model underlying a high-risk AI system;

(ii) datasheets comprehensively describing the datasets upon which models were trained and evaluated, how and why datasets were collected, how that training data will be used and maintained going forward through the development cycle; and

(iii) steps taken to ensure compliance with privacy, data privacy, data security, and copyright laws;

(d) detailed use and data management policies;

(e) any other information necessary to allow the deployer to understand the outputs and monitor the system for compliance with this article;

(f) any other information necessary to allow the deployer to comply with the requirements of subdivision four of this section; and

(g) for any high-risk AI system that is a substantial factor in making a consequential decision:

(i) a detailed description of the proposed uses of the system, including what consequential decisions the system will support;

(ii) a detailed description of the system's capabilities and any developer-imposed limitations, including capabilities outside of its intended use, when the system should not be used, any safeguards or guardrails in place to protect against unintended, inappropriate, or disallowed uses, and testing of any such safeguards or guardrails;

(iii) an internal risk assessment including documentation and results of testing conducted to identify all reasonably foreseeable risks related to algorithmic discrimination, accuracy and reliability, privacy and autonomy, and safety and security, as well as actions taken to address those risks, and subsequent testing to assess the efficacy of actions taken to address risks; and

(iv) whether the system should be monitored, and if so, how such system should be monitored.

4. Deployers of high-risk AI systems shall file with the attorney general a report containing the following:

(a) a description of the system including:

(i) a description of the system's software stack;

(ii) the purpose of the system;

(iii) the system's use to end users; and

(iv) reasonably foreseeable uses outside of the current or intended uses;

(b) the intended outputs of the system and whether the outputs can be or are otherwise appropriate to be used for any purpose not previously articulated;

(c) assessment of the relative benefits and costs to the consumer given the system's purpose, capabilities, and probable use cases;

(d) whether the deployer collects revenue or plans to collect revenue from use of the high-risk AI system, and if so, how it monetizes or plans to monetize use of the system; and

(e) for any high-risk AI system that is a substantial factor in making a consequential decision:

(i) a detailed description of the proposed uses of the system, including what consequential decisions the system will support;

(ii) whether the system is designed to make consequential decisions itself or whether and how it supports consequential decisions;

(iii) a detailed description of the system's capabilities and any deployer-imposed limitations, including capabilities outside of its intended use, when the system should not be used, any safeguards or guardrails in place to protect against unintended, inappropriate, or disallowed uses, and testing of any such safeguards or guardrails;

(iv) an assessment of the relative benefits and costs to the consumer given the system's purpose, capabilities, and probable use cases;

(v) an internal risk assessment including documentation and results of testing conducted to identify all reasonably foreseeable risks related

to algorithmic discrimination, accuracy and reliability, privacy and autonomy, and safety and security, as well as actions taken to address those risks, and subsequent testing to assess the efficacy of actions taken to address risks; and

(vi) whether the system should be monitored, and if so, how such system should be monitored.

5. The attorney general shall:

(a) promulgate rules for a process whereby developers and deployers may request redaction of portions of reports required under this section to ensure that they are not required to disclose sensitive and protected information; and

(b) maintain an online database that is accessible to the general public with reports, redacted in accordance with this subdivision, and audits required by this article which shall be updated biannually.

6. For high-risk AI systems which are already in deployment at the time of the effective date of this article, developers and deployers shall have eighteen months from such effective date to complete and file the reports and independent audit required by this article.

§ 89. Risk management policy and program. 1. Each developer or deployer of high-risk AI systems shall plan, document, and implement a risk management policy and program to govern development or deployment, as applicable, of such high-risk AI system. The risk management policy and program shall specify and incorporate the principles, processes, and personnel that the deployer uses to identify, document, and mitigate known or reasonably foreseeable risks of algorithmic discrimination covered under subdivision one of section eighty-six of this article. The risk management policy and program shall be an iterative process planned, implemented, and regularly and systematically reviewed and updated over the life cycle of a high-risk AI system, requiring regular, systematic review and updates, including updates to documentation. A risk management policy and program implemented and maintained pursuant to this section shall be reasonable considering:

(a) The guidance and standards set forth in version 1.0 of the "Artificial Intelligence Risk Management Framework" published by the National Institute of Standards and Technology in the United States department of commerce, or the latest version of the "Artificial Intelligence Risk Management Framework" published by the National Institute of Standards and Technology if, in the attorney general's discretion, the latest version of the "Artificial Intelligence Risk Management Framework" published by the National Institute of Standards and Technology in the United States department of commerce is at least as stringent as version 1.0;

(b) The size and complexity of the developer or deployer;

(c) The nature, scope, and intended uses of the high-risk AI system developed or deployed; and

(d) The sensitivity and volume of data processed in connection with the high-risk AI system.

2. A risk management policy and program implemented pursuant to subdivision one of this section may cover multiple high-risk AI systems developed by the same developer or deployed by the same deployer if sufficient.

3. The attorney general may require a developer or a deployer to disclose the risk management policy and program implemented pursuant to subdivision one of this section in a form and manner prescribed by the attorney general. The attorney general may evaluate the risk management policy and program to ensure compliance with this section.

§ 89-a. Social scoring AI systems prohibited. No person, partnership, association or corporation shall develop, deploy, use, or sell an AI system which evaluates or classifies the trustworthiness of natural persons over a certain period of time based on their social behavior or known or predicted personal or personality characteristics, with the social score leading to either or both of the following:

1. differential treatment of certain natural persons or whole groups thereof in social contexts which are unrelated to the contexts in which the data was originally generated or collected; or

2. differential treatment of certain natural persons or whole groups thereof that is unjustified or disproportionate to their social behavior or its gravity.

§ 89-b. Enforcement. 1. Whenever there shall be a violation of any provision of this article, an application may be made by the attorney general in the name of the people of the state of New York, to the supreme court having jurisdiction by a special proceeding to issue an injunction, and upon notice to the respondent of not less than ten days, to enjoin and restrain the continuance of such violation; and if it shall appear to the satisfaction of the court that the respondent has, in fact, violated this article, an injunction may be issued by the court, enjoining and restraining any further violations, without requiring proof that any person has, in fact, been injured or damaged thereby. In any such proceeding, the court may make allowances to the attorney general as provided in paragraph six of subdivision (a) of section eighty-three hundred three of the civil practice law and rules, and direct restitution. Whenever the court shall determine that a violation of this article has occurred, the court may impose a civil penalty of not more than twenty thousand dollars for each violation.

2. There shall be a private right of action by plenary proceeding for any person harmed by any violation of this article by any natural person or entity. The court shall award compensatory damages and legal fees to the prevailing party.

3. In evaluating any motion to dismiss a plenary proceeding commenced pursuant to subdivision two of this section, the court shall presume the specified AI system was created and/or operated in violation of a specified law or laws and that such violation caused the harm or harms alleged.

(a) A defendant can rebut presumptions made pursuant to this subdivision through clear and convincing evidence that the specified AI system did not cause the harm or harms alleged and/or did not violate the alleged law or laws. An algorithmic audit can be considered as evidence in rebutting such presumptions, but the mere existence of such an audit, without additional evidence, shall not be considered clear and convincing evidence.

(b) Where such presumptions are not rebutted pursuant to this subdivision, the action shall not be dismissed.

(c) Where such presumptions are rebutted pursuant to this subdivision, a motion to dismiss an action shall be adjudicated without any consideration of this section.

4. The supreme court in the state shall have jurisdiction over any action, claim, or lawsuit to enforce the provisions of this article.

§ 4. Section 296 of the executive law is amended by adding a new subdivision 23 to read as follows:

23. It shall be an unlawful discriminatory practice under this section for a deployer or a developer, as such terms are defined in section

eighty-five of the civil rights law, to engage in an unlawful discriminatory practice under section eighty-six of the civil rights law.

§ 5. This act shall take effect immediately.