STATE OF NEW YORK

1962

2025-2026 Regular Sessions

IN SENATE

January 14, 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 general business law, in relation to preventing the use of artificial intelligence algorithms to discriminate against protected classes

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

Section 1. Short title. This act shall be known and may be cited as the "New York artificial intelligence consumer protection act".

§ 2. The general business law is amended by adding a new article 45-A to read as follows:

ARTICLE 45-A

NEW YORK ARTIFICIAL INTELLIGENCE CONSUMER PROTECTION ACT

Section 1550. Definitions.
1551. Required documentation.
1552. Risk management.
1553. Technical documentation.
1554. Required disclosure.
1555. Preemption.
1556. Enforcement.

§ 1550. Definitions. For the purposes of this article, the following terms shall have the following meanings:

1. "Algorithmic discrimination":

(a) shall mean any condition in which the use of an artificial intelligence decision system results in any unlawful differential treatment or impact that disfavors any individual or group of individuals on the basis of their actual or perceived age, color, disability, ethnicity, genetic information, English language proficiency, national origin, race, religion, reproductive health, sex, veteran status, or other classification protected pursuant to state or federal law; and

(b) shall not include:

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

LBD01361-01-5

(i) the offer, license, or use of a high-risk artificial intelligence decision system by a developer or deployer for the sole purpose of:

(A) such developer's or deployer's self-testing to identify, mitigate, or prevent discrimination or otherwise ensure compliance with state and federal law; or

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

(ii) an act or omission by or on behalf of a private club or other establishment not open to the general public, as set forth in title II of the Civil Rights Act of 1964, 42 U.S.C. § 2000a(e), as amended.

2. "Artificial intelligence decision system" shall mean any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including any content, decision, prediction, or recommendation, that is used to substantially assist or replace discretionary decision making for making consequential decisions that impact consumers.

3. "Bias and governance audit" means an impartial evaluation by an independent auditor, which shall include, at a minimum, the testing of an artificial intelligence decision system to assess such system's disparate impact on employees because of such employee's age, race, creed, color, ethnicity, national origin, disability, citizenship or immigration status, marital or familial status, military status, religion, or sex, including sexual orientation, gender identity, gender expression, pregnancy, pregnancy outcomes, and reproductive healthcare choices.

4. "Consequential decision" shall mean any decision that has a material legal or similarly significant effect on the provision or denial to any consumer of, or the cost or terms of, any:

(a) education enrollment or education opportunity;

(b) employment or employment opportunity;

(c) financial or lending service;

(d) essential government service;

(e) health care service, as defined in section 42 U.S.C. § 324(d)(2), as amended;

(f) housing or housing opportunity;

(g) insurance; or

(h) legal service.

5. "Consumer" shall mean any New York state resident.

6. "Deploy" shall mean to use a high-risk artificial intelligence decision system.

7. "Deployer" shall mean any person doing business in this state that deploys a high-risk artificial intelligence decision system.

8. "Developer" shall mean any person doing business in this state that develops, or intentionally and substantially modifies, an artificial intelligence decision system.

9. "General-purpose artificial intelligence model":

(a) shall mean any form of artificial intelligence decision system that:

(i) displays significant generality;

(ii) is capable of competently performing a wide range of distinct tasks; and

(iii) can be integrated into a variety of downstream applications or systems; and

(b) shall not include any artificial intelligence model that is used for development, prototyping, and research activities before such artificial intelligence model is released on the market.

10. "High-risk artificial intelligence decision system":

(a) shall mean any artificial intelligence decision system that, when deployed, makes, or is a substantial factor in making, a consequential decision; and

(b) shall not include:

(i) any artificial intelligence decision system that is intended to:

(A) perform any narrow procedural task; or

(B) detect decision-making patterns, or deviations from decision-making patterns, unless such artificial intelligence decision system is intended to replace or influence any assessment previously completed by an individual without sufficient human review; or

(ii) unless the technology, when deployed, makes, or is a substantial factor in making, a consequential decision:

(A) any anti-fraud technology that does not make use of facial recognition technology;

(B) any artificial intelligence-enabled video game technology;

(C) any anti-malware, anti-virus, calculator, cybersecurity, database, data storage, firewall, Internet domain registration, Internet-web-site loading, networking, robocall-filtering, spam-filtering, spellchecking, spreadsheet, web-caching, web-hosting, or similar technology;

(D) any technology that performs tasks exclusively related to an entity's internal management affairs, including, but not limited to, ordering office supplies or processing payments; or

(E) any technology that communicates with consumers in natural language for the purpose of providing consumers with information, making referrals or recommendations, and answering questions, and is subject to an accepted use policy that prohibits generating content that is discriminatory or harmful.

11. "Intentional and substantial modification":

(a) shall mean any deliberate change made to:

(i) an artificial intelligence decision system that results in any new reasonably foreseeable risk of algorithmic discrimination; or

(ii) a general-purpose artificial intelligence model that:

(A) affects compliance of the general-purpose artificial intelligence model;

(B) materially changes the purpose of the general-purpose artificial intelligence model; or

(C) results in any new reasonably foreseeable risk of algorithmic discrimination; and

(b) shall not include any change made to a high-risk artificial intelligence decision system, or the performance of a high-risk artificial intelligence decision system, if:

(i) the high-risk artificial intelligence decision system continues to learn after such high-risk artificial intelligence decision system is:

(A) offered, sold, leased, licensed, given or otherwise made available to a deployer; or

(B) deployed; and

(ii) such change:

(A) is made to such high-risk artificial intelligence decision system as a result of any learning described in subparagraph (i) of this paragraph;

(B) was predetermined by the deployer, or the third party contracted by the deployer, when such deployer or third party completed the initial impact assessment of such high-risk artificial intelligence decision system pursuant to subdivision three of section one thousand five hundred fifty-two of this article; and

(C) is included in the technical documentation for such high-risk artificial intelligence decision system.

12. "Person" shall mean any individual, association, corporation, limited liability company, partnership, trust or other legal entity authorized to do business in this state.

13. "Red-teaming" shall mean an exercise that is conducted to identify the potential adverse behaviors or outcomes of an artificial intelligence decision system and how such behaviors or outcomes occur, and stress test the safeguards against such adverse behaviors or outcomes.

14. "Substantial factor":

(a) shall mean a factor that:

(i) assists in making a consequential decision;

(ii) is capable of altering the outcome of a consequential decision; and

(iii) is generated by an artificial intelligence decision system; and

(b) includes, but is not limited to, any use of an artificial intelligence decision system to generate any content, decision, prediction, or recommendation concerning a consumer that is used as a basis to make a consequential decision concerning such consumer.

15. "Synthetic digital content" shall mean any digital content, including, but not limited to, any audio, image, text, or video, that is produced or manipulated by an artificial intelligence decision system, including, but not limited to, a general-purpose artificial intelligence model.

16. "Trade secret" shall mean any form and type of financial, business, scientific, technical, economic, or engineering information, including, but not limited to, a pattern, plan, compilation, program device, formula, design, prototype, method, technique, process, procedure, program, or code, whether tangible or intangible, and whether stored, compiled, or memorialized physically, electronically, graphically, photographically, or in writing, that:

(a) derives independent economic value, whether actual or potential, from not being generally known to, or readily ascertainable by proper means by, other persons who can obtain economic value from its disclosure or use; and

(b) is the subject of efforts that are reasonable under the circumstances to maintain its secrecy.

§ 1551. Required documentation. 1. (a) Beginning on January first, two thousand twenty-seven, each developer of a high-risk artificial intelligence decision system shall use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination arising from the intended and contracted uses of a high-risk artificial intelligence decision system. In any enforcement action brought on or after such date by the attorney general pursuant to this article, there shall be a rebuttable presumption that a developer used reasonable care as required pursuant to this subdivision if:

(i) the developer complied with the provisions of this section; and

(ii) an independent third party identified by the attorney general pursuant to paragraph (b) of this subdivision and retained by the developer completed bias and governance audits for the high-risk artificial intelligence decision system.

(b) No later than January first, two thousand twenty-six, and at least annually thereafter, the attorney general shall:

(i) identify independent third parties who, in the attorney general's opinion, are qualified to complete bias and governance audits for the purposes of subparagraph (ii) of paragraph (a) of this subdivision; and

(ii) publish a list of such independent third parties available on the attorney general's website.

2. Beginning on January first, two thousand twenty-seven, and except as provided in subdivision five of this section, a developer of a high-risk artificial intelligence decision system shall make available to each deployer or other developer the following information:

(a) A general statement describing the reasonably foreseeable uses, and the known harmful or inappropriate uses, of such high-risk artificial intelligence decision system;

(b) Documentation disclosing:

(i) high-level summaries of the type of data used to train such high-risk artificial intelligence decision system;

(ii) the known or reasonably foreseeable limitations of such high-risk artificial intelligence decision system, including, but not limited to, the known or reasonably foreseeable risks of algorithmic discrimination arising from the intended uses of such high-risk artificial intelligence decision system;

(iii) the purpose of such high-risk artificial intelligence decision system;

(iv) the intended benefits and uses of such high-risk artificial intelligence decision system; and

(v) any other information necessary to enable such deployer or other developer to comply with the provisions of this article;

(c) Documentation describing:

(i) how such high-risk artificial intelligence decision system was evaluated for performance, and mitigation of algorithmic discrimination, before such high-risk artificial intelligence decision system was offered, sold, leased, licensed, given, or otherwise made available to such deployer or other developer;

(ii) the data governance measures used to cover the training datasets and examine the suitability of data sources, possible biases, and appropriate mitigation;

(iii) the intended outputs of such high-risk artificial intelligence decision system;

(iv) the measures such deployer or other developer has taken to mitigate any known or reasonably foreseeable risks of algorithmic discrimination that may arise from deployment of such high-risk artificial intelligence decision system; and

(v) how such high-risk artificial intelligence decision system should be used, not be used, and be monitored by an individual when such high-risk artificial intelligence decision system is used to make, or as a substantial factor in making, a consequential decision; and

(d) Any additional documentation that is reasonably necessary to assist a deployer or other developer to:

(i) understand the outputs of such high-risk artificial intelligence decision system; and

(ii) monitor the performance of such high-risk artificial intelligence decision system for risks of algorithmic discrimination.

3. (a) Except as provided in subdivision five of this section, any developer that, on or after January first, two thousand twenty-seven, offers, sells, leases, licenses, gives, or otherwise makes available to a deployer or other developer a high-risk artificial intelligence decision system shall, to the extent feasible, make available to such deployers and other developers the documentation and information relating to such high-risk artificial intelligence decision system necessary for a deployer, or the third party contracted by a deployer, to complete

an impact assessment pursuant to this article. The developer shall make such documentation and information available through artifacts such as model cards, dataset cards, or other impact assessments.

(b) A developer that also serves as a deployer for any high-risk artificial intelligence decision system shall not be required to generate the documentation and information required pursuant to this section unless such high-risk artificial intelligence decision system is provided to an unaffiliated entity acting as a deployer.

4. (a) Beginning on January first, two thousand twenty-seven, each developer shall publish, in a manner that is clear and readily available, on such developer's website, or a public use case inventory, a statement summarizing:

(i) the types of high-risk artificial intelligence decision systems that such developer:

(A) has developed or intentionally and substantially modified; and

(B) currently makes available to a deployer or other developer; and

(ii) how such developer manages any known or reasonably foreseeable risks of algorithmic discrimination that may arise from the development or intentional and substantial modification of the types of high-risk artificial intelligence decision systems described in subparagraph (i) of this subdivision.

(b) Each developer shall update the statement described in paragraph (a) of this subdivision:

(i) as necessary to ensure that such statement remains accurate; and

(ii) no later than ninety days after the developer intentionally and substantially modifies any high-risk artificial intelligence decision system described in subparagraph (i) of paragraph (a) of this subdivision.

5. Nothing in subdivisions two or four of this section shall be construed to require a developer to disclose any information:

(a) that is a trade secret or otherwise protected from disclosure pursuant to state or federal law; or

(b) the disclosure of which would present a security risk to such developer.

6. Beginning on January first, two thousand twenty-seven, the attorney general may require that a developer disclose to the attorney general, as part of an investigation conducted by the attorney general and in a form and manner prescribed by the attorney general, the general statement or documentation described in subdivision two of this section. The attorney general may evaluate such general statement or documentation to ensure compliance with the provisions of this section. In disclosing such general statement or documentation to the attorney general pursuant to this subdivision, the developer may designate such general statement or documentation as including any information that is exempt from disclosure pursuant to subdivision five of this section or article six of the public officers law. To the extent such general statement or documentation includes such information, such general statement or documentation shall be exempt from disclosure. To the extent any information contained in such general statement or documentation is subject to the attorney-client privilege or work product protection, such disclosure shall not constitute a waiver of such privilege or protection.

§ 1552. Risk management. 1. (a) Beginning on January first, two thousand twenty-seven, each deployer of a high-risk artificial intelligence decision system shall use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination. In

any enforcement action brought on or after said date by the attorney general pursuant to this article, there shall be a rebuttable presumption that a deployer of a high-risk artificial intelligence decision system used reasonable care as required pursuant to this subdivision if:

(i) the deployer complied with the provisions of this section; and

(ii) an independent third party identified by the attorney general pursuant to paragraph (b) of this subdivision and retained by the deployer completed bias and governance audits for the high-risk artificial intelligence decision system.

(b) No later than January first, two thousand twenty-seven, and at least annually thereafter, the attorney general shall:

(i) identify the independent third parties who, in the attorney general's opinion, are qualified to complete bias and governance audits for the purposes of subparagraph (ii) of paragraph (a) of this subdivision; and

(ii) make a list of such independent third parties available on the attorney general's web site.

2. (a) Beginning on January first, two thousand twenty-seven, and except as provided in subdivision seven of this section, each deployer of a high-risk artificial intelligence decision system shall implement and maintain a risk management policy and program to govern such deployer's deployment of the high-risk artificial intelligence decision system. The risk management policy and program shall specify and incorporate the principles, processes, and personnel that the deployer shall use to identify, document, and mitigate any known or reasonably foreseeable risks of algorithmic discrimination. The risk management policy shall be the product of an iterative process, the risk management program shall be an iterative process and both the risk management policy and program shall be planned, implemented, and regularly and systematically reviewed and updated over the lifecycle of the high-risk artificial intelligence decision system. Each risk management policy and program implemented and maintained pursuant to this subdivision shall be reasonable, considering:

(i) the guidance and standards set forth in the latest version of:

(A) the "Artificial Intelligence Risk Management Framework" published by the national institute of standards and technology;

(B) ISO or IEC 42001 of the international organization for standardization; or

(C) a nationally or internationally recognized risk management framework for artificial intelligence decision systems, other than the guidance and standards specified in clauses (A) and (B) of this subparagraph, that imposes requirements that are substantially equivalent to, and at least as stringent as, the requirements established pursuant to this section for risk management policies and programs;

(ii) the size and complexity of the deployer;

(iii) the nature and scope of the high-risk artificial intelligence decision systems deployed by the deployer, including, but not limited to, the intended uses of such high-risk artificial intelligence decision systems; and

(iv) the sensitivity and volume of data processed in connection with the high-risk artificial intelligence decision systems deployed by the deployer.

(b) A risk management policy and program implemented and maintained pursuant to paragraph (a) of this subdivision may cover multiple high-risk artificial intelligence decision systems deployed by the deployer.

3. (a) Except as provided in paragraphs (c) and (d) of this subdivision and subdivision seven of this section:

(i) a deployer that deploys a high-risk artificial intelligence decision system on or after January first, two thousand twenty-seven, or a third party contracted by the deployer, shall complete an impact assessment of the high-risk artificial intelligence decision system; and

(ii) beginning on January first, two thousand twenty-seven, a deployer, or a third party contracted by the deployer, shall complete an impact assessment of a deployed high-risk artificial intelligence decision system:

(A) at least annually; and

(B) no later than ninety days after an intentional and substantial modification to such high-risk artificial intelligence decision system is made available.

(b) (i) Each impact assessment completed pursuant to this subdivision shall include, at a minimum and to the extent reasonably known by, or available to, the deployer:

(A) a statement by the deployer disclosing the purpose, intended use cases and deployment context of, and benefits afforded by, the high-risk artificial intelligence decision system;

(B) an analysis of whether the deployment of the high-risk artificial intelligence decision system poses any known or reasonably foreseeable risks of algorithmic discrimination and, if so, the nature of such algorithmic discrimination and the steps that have been taken to mitigate such risks;

(C) A description of:

(I) the categories of data the high-risk artificial intelligence decision system processes as inputs; and

(II) the outputs such high-risk artificial intelligence decision system produces;

(D) if the deployer used data to customize the high-risk artificial intelligence decision system, an overview of the categories of data the deployer used to customize such high-risk artificial intelligence decision system;

(E) any metrics used to evaluate the performance and known limitations of the high-risk artificial intelligence decision system;

(F) a description of any transparency measures taken concerning the high-risk artificial intelligence decision system, including, but not limited to, any measures taken to disclose to a consumer that such high-risk artificial intelligence decision system is in use when such high-risk artificial intelligence decision system is in use; and

(G) a description of the post-deployment monitoring and user safeguards provided concerning such high-risk artificial intelligence decision system, including, but not limited to, the oversight, use, and learning process established by the deployer to address issues arising from deployment of such high-risk artificial intelligence decision system.

(ii) In addition to the statement, analysis, descriptions, overview, and metrics required pursuant to subparagraph (i) of this paragraph, an impact assessment completed pursuant to this subdivision following an intentional and substantial modification made to a high-risk artificial intelligence decision system on or after January first, two thousand twenty-seven, shall include a statement disclosing the extent to which the high-risk artificial intelligence decision system was used in a manner that was consistent with, or varied from, the developer's intended uses of such high-risk artificial intelligence decision system.

(c) A single impact assessment may address a comparable set of high-risk artificial intelligence decision systems deployed by a deployer.

(d) If a deployer, or a third party contracted by the deployer, completes an impact assessment for the purpose of complying with another applicable law or regulation, such impact assessment shall be deemed to satisfy the requirements established in this subdivision if such impact assessment is reasonably similar in scope and effect to the impact assessment that would otherwise be completed pursuant to this subdivision.

(e) A deployer shall maintain the most recently completed impact assessment of a high-risk artificial intelligence decision system as required pursuant to this subdivision, all records concerning each such impact assessment and all prior impact assessments, if any, for a period of at least three years following the final deployment of the high-risk artificial intelligence decision system.

4. Except as provided in subdivision seven of this section, a deployer, or a third party contracted by the deployer, shall review, no later than January first, two thousand twenty-seven, and at least annually thereafter, the deployment of each high-risk artificial intelligence decision system deployed by the deployer to ensure that such high-risk artificial intelligence decision system is not causing algorithmic discrimination.

5. (a) Beginning on January first, two thousand twenty-seven, and before a deployer deploys a high-risk artificial intelligence decision system to make, or be a substantial factor in making, a consequential decision concerning a consumer, the deployer shall:

(i) notify the consumer that the deployer has deployed a high-risk artificial intelligence decision system to make, or be a substantial factor in making, such consequential decision; and

(ii) provide to the consumer:

(A) a statement disclosing:

(I) the purpose of such high-risk artificial intelligence decision system; and

(II) the nature of such consequential decision;

(B) contact information for such deployer;

(C) a description, in plain language, of such high-risk artificial intelligence decision system; and

(D) instructions on how to access the statement made available pursuant to paragraph (a) of subdivision six of this section.

(b) Beginning on January first, two thousand twenty-seven, a deployer that has deployed a high-risk artificial intelligence decision system to make, or as a substantial factor in making, a consequential decision concerning a consumer shall, if such consequential decision is adverse to the consumer, provide to such consumer:

(i) a statement disclosing the principal reason or reasons for such adverse consequential decision, including, but not limited to:

(A) the degree to which, and manner in which, the high-risk artificial intelligence decision system contributed to such adverse consequential decision;

(B) the type of data that was processed by such high-risk artificial intelligence decision system in making such adverse consequential decision; and

(C) the source of such data; and

(ii) an opportunity to:

(A) correct any incorrect personal data that the high-risk artificial intelligence decision system processed in making, or as a substantial factor in making, such adverse consequential decision; and

(B) appeal such adverse consequential decision, which shall, if technically feasible, allow for human review unless providing such opportunity is not in the best interest of such consumer, including, but not limited to, in instances in which any delay might pose a risk to the life or safety of such consumer.

(c) The deployer shall provide the notice, statements, information, description, and instructions required pursuant to paragraphs (a) and (b) of this subdivision:

(i) directly to the consumer;

(ii) in plain language;

(iii) in all languages in which such deployer, in the ordinary course of such deployer's business, provides contracts, disclaimers, sale announcements, and other information to consumers; and

(iv) in a format that is accessible to consumers with disabilities.

6. (a) Beginning on January first, two thousand twenty-seven, and except as provided in subdivision seven of this section, each deployer shall make available, in a manner that is clear and readily available on such deployer's website, a statement summarizing:

(i) the types of high-risk artificial intelligence decision systems that are currently deployed by such deployer;

(ii) how such deployer manages any known or reasonably foreseeable risks of algorithmic discrimination that may arise from deployment of each high-risk artificial intelligence decision system described in subparagraph (i) of this paragraph; and

(iii) in detail, the nature, source and extent of the information collected and used by such deployer.

(b) Each deployer shall periodically update the statement required pursuant to paragraph (a) of this subdivision.

7. The provisions of subdivisions two, three, four, and six of this section shall not apply to a deployer if, at the time the deployer deploys a high-risk artificial intelligence decision system, and at all times while the high-risk artificial intelligence decision system is deployed:

(a) the deployer:

(i) has entered into a contract with the developer in which the developer has agreed to assume the deployer's duties pursuant to subdivisions two, three, four, or six of this section; and

(ii) does not exclusively use such deployer's own data to train such high-risk artificial intelligence decision system;

(b) such high-risk artificial intelligence decision system:

(i) is used for the intended uses that are disclosed to such deployer pursuant to subparagraph (iv) of paragraph (b) of subdivision two of section one thousand five hundred fifty-one of this article; and

(ii) continues learning based on a broad range of data sources and not solely based on the deployer's own data; and

(c) such deployer makes available to consumers any impact assessment that:

(i) the developer of such high-risk artificial intelligence decision system has completed and provided to such deployer; and

(ii) includes information that is substantially similar to the information included in the statement, analysis, descriptions, overview, and metrics required pursuant to subparagraph (i) of paragraph (b) of subdivision three of this section.

8. Nothing in this subdivision or subdivisions two, three, four, five, or six of this section shall be construed to require a deployer to disclose any information that is a trade secret or otherwise protected from disclosure pursuant to state or federal law. If a deployer withholds any information from a consumer pursuant this subdivision, the deployer shall send notice to such consumer disclosing:

(a) that the deployer is withholding such information from such consumer; and

(b) the basis for the deployer's decision to withhold such information from such consumer.

9. Beginning on January first, two thousand twenty-seven, the attorney general may require that a deployer, or a third party contracted by the deployer pursuant to subdivision three of this section, as applicable, disclose to the attorney general, as part of an investigation conducted by the attorney general, no later than ninety days after a request by the attorney general, and in a form and manner prescribed by the attorney general, the risk management policy implemented pursuant to subdivision two of this section, the impact assessment completed pursuant to subdivision three of this section; or records maintained pursuant to paragraph (e) of subdivision three of this section. The attorney general may evaluate such risk management policy, impact assessment or records to ensure compliance with the provisions of this section. In disclosing such risk management policy, impact assessment or records to the attorney general pursuant to this subdivision, the deployer or third-party contractor, as applicable, may designate such risk management policy, impact assessment or records as including any information that is exempt from disclosure pursuant to subdivision eight of this section or article six of the public officers law. To the extent such risk management policy, impact assessment, or records include such information, such risk management policy, impact assessment, or records shall be exempt from disclosure. To the extent any information contained in such risk management policy, impact assessment, or record is subject to the attorney-client privilege or work product protection, such disclosure shall not constitute a waiver of such privilege or protection.

§ 1553. Technical documentation. 1. Beginning on January first, two thousand twenty-seven, each developer of a general-purpose artificial intelligence model shall, except as provided in subdivision two of this section:

(a) create and maintain technical documentation for the general-purpose artificial intelligence model, which shall:

(i) include:

(A) the training and testing processes for such general-purpose artificial intelligence model; and

(B) the results of an evaluation of such general-purpose artificial intelligence model performed to determine whether such general-purpose artificial intelligence model is in compliance with the provisions of this article;

(ii) include, as appropriate, considering the size and risk profile of such general-purpose artificial intelligence model, at least:

(A) the tasks such general-purpose artificial intelligence model is intended to perform;

(B) the type and nature of artificial intelligence decision systems in which such general-purpose artificial intelligence model is intended to be integrated;

(C) acceptable use policies for such general-purpose artificial intelligence model;

(D) the date such general-purpose artificial intelligence model is released;

(E) the methods by which such general-purpose artificial intelligence model is distributed; and

(F) the modality and format of inputs and outputs for such general-purpose artificial intelligence model; and

(iii) be reviewed and revised at least annually, or more frequently, as necessary to maintain the accuracy of such technical documentation; and

(b) create, implement, maintain and make available to persons that intend to integrate such general-purpose artificial intelligence model into such persons' artificial intelligence decision systems documentation and information that:

(i) enables such persons to:

(A) understand the capabilities and limitations of such general-purpose artificial intelligence model; and

(B) comply with such persons' obligations pursuant to this article;

(ii) discloses, at a minimum:

(A) the technical means required for such general-purpose artificial intelligence model to be integrated into such persons' artificial intelligence decision systems;

(B) the information listed in subparagraph (ii) of paragraph (a) of this subdivision; and

(iii) except as provided in subdivision two of this section, is reviewed and revised at least annually, or more frequently, as necessary to maintain the accuracy of such documentation and information.

2. (a) The provisions of paragraph (a) and subparagraph (iii) of paragraph (b) of subdivision one of this section shall not apply to a developer that develops, or intentionally and substantially modifies, a general-purpose artificial intelligence model on or after January first, two thousand twenty-seven, if:

(i) (A) the developer releases such general-purpose artificial intelligence model under a free and open-source license that allows for:

(I) access to, and modification, distribution, and usage of, such general-purpose artificial intelligence model; and

(II) the parameters of such general-purpose artificial intelligence model to be made publicly available pursuant to clause (B) of this subparagraph; and

(B) unless such general-purpose artificial intelligence model is deployed as a high-risk artificial intelligence decision system, the parameters of such general-purpose artificial intelligence model, including, but not limited to, the weights and information concerning the model architecture and model usage for such general-purpose artificial intelligence model, are made publicly available; or

(ii) the general-purpose artificial intelligence model is:

(A) not offered for sale in the market;

(B) not intended to interact with consumers; and

(C) solely utilized:

(I) for an entity's internal purposes; or

(II) pursuant to an agreement between multiple entities for such entities' internal purposes.

(b) The provisions of this section shall not apply to a developer that develops, or intentionally and substantially modifies, a general-purpose artificial intelligence model on or after January first, two thousand twenty-seven, if such general purpose artificial intelligence model performs tasks exclusively related to an entity's internal management

affairs, including, but not limited to, ordering office supplies or processing payments.

(c) A developer that takes any action under an exemption pursuant to paragraph (a) or (b) of this subdivision shall bear the burden of demonstrating that such action qualifies for such exemption.

(d) A developer that is exempt pursuant to subparagraph (ii) of paragraph (a) of this subdivision shall establish and maintain an artificial intelligence risk management framework, which shall:

(i) be the product of an iterative process and ongoing efforts; and

(ii) include, at a minimum:

(A) an internal governance function;

(B) a map function that shall establish the context to frame risks;

(C) a risk management function; and

(D) a function to measure identified risks by assessing, analyzing and tracking such risks.

3. Nothing in subdivision one of this section shall be construed to require a developer to disclose any information that is a trade secret or otherwise protected from disclosure pursuant to state or federal law.

4. Beginning on January first, two thousand twenty-seven, the attorney general may require that a developer disclose to the attorney general, as part of an investigation conducted by the attorney general, no later than ninety days after a request by the attorney general and in a form and manner prescribed by the attorney general, any documentation maintained pursuant to this section. The attorney general may evaluate such documentation to ensure compliance with the provisions of this section. In disclosing any documentation to the attorney general pursuant to this subdivision, the developer may designate such documentation as including any information that is exempt from disclosure pursuant to subdivision three of this section or article six of the public officers law. To the extent such documentation includes such information, such documentation shall be exempt from disclosure. To the extent any information contained in such documentation is subject to the attorney-client privilege or work product protection, such disclosure shall not constitute a waiver of such privilege or protection.

§ 1554. Required disclosure. 1. Beginning on January first, two thousand twenty-seven, and except as provided in subdivision two of this section, each person doing business in this state, including, but not limited to, each deployer that deploys, offers, sells, leases, licenses, gives, or otherwise makes available, as applicable, any artificial intelligence decision system that is intended to interact with consumers shall ensure that it is disclosed to each consumer who interacts with such artificial intelligence decision system that such consumer is interacting with an artificial intelligence decision system.

2. No disclosure shall be required pursuant to subdivision one of this section under circumstances in which a reasonable person would deem it obvious that such person is interacting with an artificial intelligence decision system.

§ 1555. Preemption. 1. Nothing in this article shall be construed to restrict a developer's, deployer's, or other person's ability to:

(a) comply with federal, state or municipal law;

(b) comply with a civil, criminal or regulatory inquiry, investigation, subpoena, or summons by a federal, state, municipal, or other governmental authority;

(c) cooperate with a law enforcement agency concerning conduct or activity that the developer, deployer, or other person reasonably and in good faith believes may violate federal, state, or municipal law;

(d) investigate, establish, exercise, prepare for, or defend a legal claim;

(e) take immediate steps to protect an interest that is essential for the life or physical safety of a consumer or another individual;

(f) (i) by any means other than facial recognition technology, prevent, detect, protect against, or respond to:

(A) a security incident;

(B) a malicious or deceptive activity; or

(C) identity theft, fraud, harassment or any other illegal activity;

(ii) investigate, report, or prosecute the persons responsible for any action described in subparagraph (i) of this paragraph; or

(iii) preserve the integrity or security of systems;

(g) engage in public or peer-reviewed scientific or statistical research in the public interest that:

(i) adheres to all other applicable ethics and privacy laws; and

(ii) is conducted in accordance with:

(A) part forty-six of title forty-five of the code of federal regulations, as amended; or

(B) relevant requirements established by the federal food and drug administration;

(h) conduct research, testing, and development activities regarding an artificial intelligence decision system or model, other than testing conducted pursuant to real world conditions, before such artificial intelligence decision system or model is placed on the market, deployed, or put into service, as applicable;

(i) effectuate a product recall;

(j) identify and repair technical errors that impair existing or intended functionality; or

(k) assist another developer, deployer, or person with any of the obligations imposed pursuant to this article.

2. The obligations imposed on developers, deployers, or other persons pursuant to this article shall not apply where compliance by the developer, deployer, or other person with the provisions of this article would violate an evidentiary privilege pursuant to state law.

3. Nothing in this article shall be construed to impose any obligation on a developer, deployer, or other person that adversely affects the rights or freedoms of any person, including, but not limited to, the rights of any person:

(a) to freedom of speech or freedom of the press guaranteed in:

(i) the first amendment to the United States constitution; and

(ii) section eight of the New York state constitution; or

(b) pursuant to section seventy-nine-h of the civil rights law.

4. Nothing in this article shall be construed to apply to any developer, deployer, or other person:

(a) insofar as such developer, deployer or other person develops, deploys, puts into service, or intentionally and substantially modifies, as applicable, a high-risk artificial intelligence decision system:

(i) that has been approved, authorized, certified, cleared, developed, or granted by:

(A) a federal agency, including, but not limited to, the federal food and drug administration or the federal aviation administration, acting within the scope of such federal agency's authority; or

(B) a regulated entity subject to supervision and regulation by the federal housing finance agency; or

(ii) in compliance with standards that are:

(A) established by:

(I) any federal agency, including, but not limited to, the federal office of the national coordinator for health information technology; or

(II) a regulated entity subject to supervision and regulation by the federal housing finance agency; and

(B) substantially equivalent to, and at least as stringent as, the standards established pursuant to this article;

(b) conducting research to support an application:

(i) for approval or certification from any federal agency, including, but not limited to, the federal food and drug administration, the federal aviation administration, or the federal communications commission; or

(ii) that is otherwise subject to review by any federal agency;

(c) performing work pursuant to, or in connection with, a contract with the federal department of commerce, the federal department of defense, or the national aeronautics and space administration, unless such developer, deployer, or other person is performing such work on a high-risk artificial intelligence decision system that is used to make, or as a substantial factor in making, a decision concerning employment or housing; or

(d) that is a covered entity, as defined by the health insurance portability and accountability act of 1996 and the regulations promulgated thereunder, as amended, and providing health care recommendations that:

(i) are generated by an artificial intelligence decision system;

(ii) require a health care provider to take action to implement such recommendations; and

(iii) are not considered to be high risk.

5. Nothing in this article shall be construed to apply to any artificial intelligence decision system that is acquired by or for the federal government or any federal agency or department, including, but not limited to, the federal department of commerce, the federal department of defense, or the national aeronautics and space administration, unless such artificial intelligence decision system is a high-risk artificial intelligence decision system that is used to make, or as a substantial factor in making, a decision concerning employment or housing.

6. Any insurer, as defined by section five hundred one of the insurance law, or fraternal benefit society, as defined by section four thousand five hundred one of the insurance law, shall be deemed to be in full compliance with the provisions of this article if such insurer or fraternal benefit society has implemented and maintains a written artificial intelligence decision systems program in accordance with all requirements established by the superintendent of financial services.

7. (a) Any bank, out-of-state bank, New York credit union, federal credit union, or out-of-state credit union, or any affiliate or subsidiary thereof, shall be deemed to be in full compliance with the provisions of this article if such bank, out-of-state bank, New York credit union, federal credit union, out-of-state credit union, affiliate, or subsidiary is subject to examination by any state or federal prudential regulator pursuant to any published guidance or regulations that apply to the use of high-risk artificial intelligence decision systems, and such guidance or regulations:

(i) impose requirements that are substantially equivalent to, and at least as stringent as, the requirements of this article; and

(ii) at a minimum, require such bank, out-of-state bank, New York credit union, federal credit union, out-of-state credit union, affiliate, or subsidiary to:

(A) regularly audit such bank's, out-of-state bank's, New York credit union's, federal credit union's, out-of-state credit union's, affiliate's, or subsidiary's use of high-risk artificial intelligence decision systems for compliance with state and federal anti-discrimination laws and regulations applicable to such bank, out-of-state bank, New York credit union, federal credit union, out-of-state credit union, affiliate, or subsidiary; and

(B) mitigate any algorithmic discrimination caused by the use of a high-risk artificial intelligence decision system, or any risk of algorithmic discrimination that is reasonably foreseeable as a result of the use of a high-risk artificial intelligence decision system.

(b) For the purposes of this subdivision, the following terms shall have the following meanings:

(i) "Affiliate" shall have the same meaning as set forth in section nine hundred twelve of the business corporation law.

(ii) "Bank" shall have the same meaning as set forth in section two of the banking law.

(iii) "Credit union" shall have the same meaning as set forth in section two of the banking law.

(iv) "Out-of-state bank" shall have the same meaning as set forth in section two hundred twenty-two of the banking law.

(v) "Subsidiary" shall have the same meaning as set forth in section one hundred forty-one of the banking law.

8. If a developer, deployer, or other person engages in any action under an exemption pursuant to subdivisions one, two, three, four, five, six, or seven of this section, the developer, deployer, or other person bears the burden of demonstrating that such action qualifies for such exemption.

§ 1556. Enforcement. 1. The attorney general shall have exclusive authority to enforce the provisions of this article.

2. Except as provided in subdivision six of this section, during the period beginning on January first, two thousand twenty-seven, and ending on January first, two thousand twenty-eight, the attorney general shall, prior to initiating any action for a violation of this section, issue a notice of violation to the developer, deployer, or other person if the attorney general determines that it is possible to cure such violation. If the developer, deployer, or other person fails to cure such violation within sixty days after receipt of such notice of violation, the attorney general may bring an action pursuant to this section.

3. Except as provided in subdivision six of this section, beginning on January first, two thousand twenty-eight, the attorney general may, in determining whether to grant a developer, deployer, or other person the opportunity to cure a violation described in subdivision two of this section, consider:

(a) the number of violations;

(b) the size and complexity of the developer, deployer, or other person;

(c) the nature and extent of the developer's, deployer's, or other person's business;

(d) the substantial likelihood of injury to the public;

(e) the safety of persons or property; and

(f) whether such violation was likely caused by human or technical error.

4. Nothing in this article shall be construed as providing the basis for a private right of action for violations of the provisions of this article.

5. Except as provided in subdivisions one, two, three, four, and six of this section, a violation of the requirements established in this article shall constitute an unfair trade practice for purposes of section three hundred forty-nine of this chapter and shall be enforced solely by the attorney general; provided, however, that subdivision (h) of section three hundred forty-nine of this chapter shall not apply to any such violation.

6. (a) In any action commenced by the attorney general for any violation of this article, it shall be an affirmative defense that the developer, deployer, or other person:

(i) discovers a violation of any provision of this article through red-teaming;

(ii) no later than sixty days after discovering such violation through red-teaming:

(A) cures such violation; and

(B) provides to the attorney general, in a form and manner prescribed by the attorney general, notice that such violation has been cured and evidence that any harm caused by such violation has been mitigated; and

(iii) is otherwise in compliance with the latest version of:

(A) the Artificial Intelligence Risk Management Framework published by the national institute of standards and technology;

(B) ISO/IEC 42001 of the international organization for standardization and the international electrotechnical commission;

(C) a nationally or internationally recognized risk management framework for artificial intelligence decision systems, other than the risk management frameworks described in clauses (A) and (B) of this subparagraph, that imposes requirements that are substantially equivalent to, and at least as stringent as, the requirements established pursuant to this article; or

(D) any risk management framework for artificial intelligence decision systems that is substantially equivalent to, and at least as stringent as, the risk management frameworks described in clauses (A), (B), and (C) of this subparagraph.

(b) The developer, deployer, or other person bears the burden of demonstrating to the attorney general that the requirements established pursuant to paragraph (a) of this subdivision have been satisfied.

(c) Nothing in this article, including, but not limited to, the enforcement authority granted to the attorney general pursuant to this section, shall be construed to preempt or otherwise affect any right, claim, remedy, presumption, or defense available at law or in equity. Any rebuttable presumption or affirmative defense established pursuant to this article shall apply only to an enforcement action brought by the attorney general pursuant to this section and shall not apply to any right, claim, remedy, presumption, or defense available at law or in equity.

§ 3. This act shall take effect on the two hundred seventieth day after it shall have become a law.