Katten

New Colorado AI Act Targeting 'Algorithmic Discrimination' Provides AI Compliance Lessons

Michael Justus[1]

Starting February 1, 2026, businesses must comply with requirements of the Colorado AI Act (the Act) (SB 205) if they use artificial intelligence (AI) tools to make "consequential" decisions about Colorado consumers' education, employment, financial or lending services, essential government services, health care, housing, insurance or legal services.

The new law focuses on addressing "algorithmic discrimination" by high-risk AI systems. But it also requires that any AI system that interacts with consumers (even if not high-risk) must disclose to consumers that they are interacting with an AI system, unless that would be obvious to a reasonable person.

The Colorado Attorney General has exclusive authority to enforce the Act. There is no private right of action.

Colorado Gov. Jared Polis suggested in his signing statement that he has "reservations" about the new law, and that the Act should be amended between now and its 2026 entry into force. The governor also called upon the federal government to enact legislation for "a cohesive federal approach."

Under the Act, algorithmic discrimination is any condition in which the use of an Artificial Intelligence System (AI System) results in a differential treatment to an individual or group of individuals based on their actual or perceived age, color, disability, ethnicity, genetic information, limited proficiency in the English language, national origin, race, religion, reproductive health, sex, veteran status or any other classification protected under Colorado law or federal law.

The Act applies when an AI System makes or is a "substantial factor" in making "consequential decisions." A consequential decision "has material, legal or similarly significant effect on the provision, denial or the cost of terms to any Colorado resident on (a) education enrollment/opportunity, (b) employment/employment opportunity, (c) financial/lending service, (d) essential government service, (e) health care services, (f) housing, (g) insurance or (h) a legal service." A substantial factor means that the system (i) assists in making a consequential decision, (ii) can alter the outcome of a consequential decision and (iii) is generated by an AI System.

Duties imposed by the new law differ for "developers" and "deployers" of AI Systems. A developer means an entity that either develops or intentionally and substantially modifies[2] an AI System. A deployer is an entity that uses an AI System in commerce but does not develop or substantially modify the system. Both developers and deployers have a duty to exercise care to protect consumers from algorithmic discrimination by AI Systems. If developers and deployers comply with their obligations under the Act, there is a rebuttable presumption that they used reasonable care. The detailed responsibilities for developers and deployers are summarized in **Table 1** below.

The Act also provides for an affirmative defense where a business discovers a violation, cures it and is otherwise in compliance with certain recognized AI compliance frameworks. The discovery, cure and compliance requirements to establish an affirmative defense under the Act are summarized in **Table 2** below.

The Act contains a variety of exemptions.[3] For example, it exempts small deployers[4] from many deployer obligations.

Businesses that interact with Colorado consumers should evaluate whether they are using or plan to use AI in the impacted service categories. If so, they should begin compliance planning sooner rather than later. Although the Act may change between now and 2026, as suggested by Gov. Polis — or even be preempted by intervening federal legislation — it provides useful guidance in the meantime. The risk-based framework and general testing, monitoring and disclosure requirements of the Act are similar to the requirements of the EU AI Act and other relevant legal and regulatory frameworks. Further, the Act's incorporation of the National Institute of Standards and Technology's (NIST) framework and other voluntary standards for purposes of establishing an affirmative defense underscores the value of existing voluntary guidance in compliance planning. Accordingly, the Colorado Act may be viewed as a helpful guidepost in building an AI compliance program, along with other legal, regulatory and voluntary frameworks.

More information about **Katten's Artificial Intelligence practice** is available here.

CONTACT

Michael R. Justus
+1.202.625.3575
michael.justus@katten.com

** Summer associate Geomy George contributed to this article.*

TABLE 1

	Documentation & Disclosure Requirements	Impact Assessment	Risk Management Policy and Program	Duty to Consumers
Developer	Provide certain documentation to downstream Developer or Deployer.[5] Provide certain statements to the public on the Developer's website.[6] Inform the Colorado Attorney General and downstream Developers and Deployers of any discovered risks of algorithmic discrimination[7]. The Colorado Attorney General may request information provided to the Deployer or made available on the Developer's website.	The Developer must provide the downstream Developer or Deployer adequate information about the AI model, the training data, and other information to facilitate their Impact Assessment.	No express statutory requirement.	Use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination arising from the intended and contracted use of the AI system.
Deployer	Provide notice to customers that they are interacting with an AI System.[8] Provide notice on website of the type of high-risk AI Systems and the foreseeable risks.[9] Inform the Colorado Attorney General of any discovered risks of algorithmic discrimination.	Annual impact assessment of the AI System. The impact assessment must follow various parameters defined by the Act[10]. The Impact Assessment may be conducted by the Deployer or a contracted third party. Records should be kept for 3 years.	Implement a risk management policy and program with certain required features[11].	Use reasonable care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination arising from the intended and contracted use of the AI system. Provide for an appeal of adverse decisions, including human review where technically feasible.

TABLE 2

To establish an affirmative defense under the Act, the Developer or Deployer must discover and cure a violation (Column 1) and otherwise be in compliance with an approved framework (Column 2). In other words, the developer or deployer must meet at least one requirement from each of Column 1 and Column 2 below.

1. Discovery and Cure	2. Compliance Framework
Discovery and cure based on feedback that the Developer, Deployer or other person encourages Deployers or users to provide to the Developer, Deployer, or another person;	In compliance with the latest version of the "Artificial Intelligence Risk Management Framework" published by NIST, and Standard ISO/IEC 42001 of the International Organization for Standardization;
Discovery and cure based on adversarial testing or red teaming, as defined by NIST; or	In compliance with another nationally or internationally recognized risk management framework for AI Systems, if the standards are substantially equivalent to or more stringent than the requirements of the Act; or
Discovery and cure based on an internal review process.	In compliance with any risk management framework that the Attorney General designates.

[1] Michael Justus is a Partner and head of Katten's AI Working Group.

[2] The Act defines intentional and substantial modification as a deliberate change made to the AI System that results in any new reasonably foreseeable risk of algorithmic discrimination.

[3] The Act also exempts AI Systems that perform a narrow procedural task, detect decision-making patterns, or find deviations from previous decision-making patterns and are not intended to replace or influence a previously completed human assessment without sufficient human review. Additionally, there is a list of exceptions for technologies that are not making a consequential decision or are not a substantial factor in making a consequential decision. The Act also provides certain exemptions permitting nondisclosure of trade secrets, information protected from disclosure by state or federal law, or information that would pose a security risk to the developer. Deployers invoking these nondisclosure exemptions should notify the consumer and provide a basis for withholding the information. The Act contains exemptions for AI systems within the purview of certain federal regulatory agencies.

[4] Exempted Deployers employ less than 50 full time employees and do not use their own data to train the system, provided certain other use and disclosure requirements are satisfied.

[5] Documentation to the downstream user must include: general statement describing the reasonably foreseeable uses and known harmful or inappropriate uses of the AI System; disclosures on summaries of the type of data used to train the AI System; limitations of the AI System including the risks of discrimination arising from the intended use; the purpose of the AI system; the intended benefits and uses; how the AI System was evaluated for the mitigation of the algorithmic discrimination before offered in commerce; the data governance measures used; the intended outputs of the AI system; the measures taken to mitigate foreseeable risks; guidance on usage; and any additional documentation which is reasonably necessary for the downstream user to understand the outputs and monitor the performance of the system for discrimination.

[6] The statement must contain the types of high-risk AI systems the Developer has made available to a Deployer or other Developer and the Developer's plans to manage the known or reasonably foreseeable risks of discrimination which could arise from any development or substantial modification of such systems. The statement must be updated and remain accurate.

[7] The newly discovered risks could be based on the Developer's own testing, or a credible report from a Deployer. The report should be made within 90 days of the risk being discovered in a format specified by the Attorney General.

[8] The Act requires any AI system which interacts with a customer to disclose that the customer is interacting with an AI System. In addition, the notice should inform the customer that the Deployer has used an AI System in process of a consequential decision regarding the customer before making the decision. The notice should also include a statement disclosing the purpose of the AI system, the nature of the decision, contact information, and the information to opt out of the processing of personal data. This notice should be directly to the consumer, in plain language, in all languages the Deployer ordinarily conducts business, and accessible to consumers with disabilities.

[9] The notice must include the types of high-risk systems which are currently deployed, the plan for the management of known and reasonably foreseeable risks of discrimination, the nature, source, and extent of information collected. This notice must be periodically updated.

[10] The impact assessment must include, at a minimum:

(1) a statement from the Deployer about the purpose, use cases, use context and benefits of the AI System,

(2) an analysis whether the deployment of the AI System poses any known or reasonably foreseeable risks of discrimination,

(3) a description of the categories of data the system processes as inputs and outputs,

(4) an overview of categories of data used by the Deployer to customize the AI System,

(5) metrics used to evaluate the performance and known limitations of the system,

(6) description of transparency measures including disclosures,

(7) description of post-deployment monitoring and safeguards, and

(8) in case of a substantial modification, a statement from the Deployer about the extent to which the AI System was used in a manner which was consistent with or varied from the Developer's intended use.

[11] "The risk management policy and program must 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." It must be regularly reviewed and updated.