Westlaw Today
powered by Reuters

Colorado sets baseline for comprehensive AI regulation in U.S.

By Monique Bhargava, Esq., Reed Smith LLP*

MAY 28, 2024

Key takeaways

- The Colorado law is the first comprehensive law targeting AI in the US, borrowing several principles from the EU AI Act
- Focused on high-risk AI systems used to make consequential decisions in areas that have a significant impact on consumers
- Establishes a duty of care to prevent algorithmic discrimination when using high-risk AI systems

On May 17, 2024 Colorado Governor Polis signed into law Senate Bill 24-205 "Concerning Consumer Protections in Interactions with Artificial Intelligence Systems" (Colorado AI Act), the first comprehensive AI law in the US. Similar to the EU AI Act and the Office of Management and Budget's Memorandum M-24-10, the Colorado AI Act adopts a risk-based approach, primarily targeting the developers and deployers of high-risk AI systems.

The law incorporates now familiar practices of risk management through documentation, impact assessments, and robust governance processes. In particular, the law heavily focuses on the duty of care to mitigate the risk of algorithmic discrimination.

Developers and deployers of high-risk AI systems will have until February 1, 2026 to develop processes to comply with the law's requirements.

What to do to prepare:

For all AI systems, ensure proper disclosures are made to consumers that interact with the systems.

- For high-risk AI systems, the Colorado AI Act formalizes what has already been good industry practice, establishing robust AI governance processes. Developers should focus on establishing documentation around the purpose, intended uses, benefits, and limitations of the high-risk AI system, as well as procedures to mitigate risks, and provide instructions to deployers. Developers should also implement risk management policies and procedures incorporating standards from industry guidance such as the NIST AI Risk Management Framework.
- Develop AI impact assessments to ensure that deployers ask for the documentation and information required to be able to properly evaluate high-risk AI systems, and for developers, that they have documented, where necessary, the processes and information to provide to deployers.
- Implement the key compliance indicators to establish a rebuttable presumption that reasonable care was used to mitigate algorithmic discrimination when developing or deploying high-risk AI systems, such as processes to identify and mitigate risk, disclosing known and reasonably foreseeable risks, and proper reporting protocols.

Below we provide a high-level overview of the requirements of the Colorado AI Act.

I. What AI systems and uses are covered?

The Colorado AI Act primarily applies to high-risk AI systems, including systems that, when used, make, or are a substantial factor in making, consequential decisions in familiar areas of priority.

The Colorado AI Act adopts a risk-based approach, primarily targeting the developers and deployers of high-risk AI systems.

Specifically, consequential decisions have a material legal, or similarly significant, effect on the provision, denial, cost, or terms of (1) education services or opportunities, (2) employment or employment opportunities, (3) health care, (4) housing, (5) insurance, (6) financial and lending services, (7) legal services, and (8) essential government services.

Whether use of an AI system is a substantial factor will depend on whether the use (1) assists in making a consequential decision, (2) is capable of altering the outcome of a consequential decision, and (3) is generated by an AI system.

Other substantial factors include output generated by an AI system such as content, decisions, predictions, or recommendations about a consumer (i.e., a Colorado resident) that is used as a basis for making consequential decisions about the consumer.

Notably, high-risk AI systems expressly exclude AI systems that are intended to perform a narrow procedural task or are used to detect patterns or deviations in decision-making (provided the systems do not influence a prior assessment by a person, without sufficient additional human review).

Other impactful exclusions include anti-fraud technologies that do not use facial recognition, database and data storage technologies, and cybersecurity technologies.

Also excluded from high-risk AI systems are generative AI tools such as chatbots, which are used to provide users with information, make referrals or recommendations, or answer questions, provided they are subject to use policies that prohibit the generation of content that is discriminatory or harmful. What is considered discriminatory or harmful, however, is not defined.

II. Emphasis on algorithmic discrimination

Notably unique to the Colorado AI Act is the duty of care for both developers and deployers to protect consumers from any known or reasonably foreseeable risk of algorithmic discrimination as a result of the intended uses of a high-risk AI system.

Algorithmic discrimination is any use of an AI system that results in unlawful differential treatment or an impact that disfavors an individual or group of individuals on the basis of an actual or perceived protected classification, including age, color, disability, ethnicity, genetic information, national origin, race, religion, reproductive health, sex, veteran status, proficiency in the English language, or any other classification protected under state or federal law.

Notably unique to the Colorado AI Act is the duty of care for both developers and deployers to protect consumers from any known or reasonably foreseeable risk of algorithmic discrimination.

It should be noted that use of AI systems solely to expand a participant pool to increase diversity or redress historical discrimination is not deemed algorithmic discrimination.

- Developers are required to provide deployers with information on the known or reasonably foreseeable risks of algorithmic discrimination resulting from the intended uses of the high-risk AI system and describe what steps have been taken to evaluate and mitigate those risks prior to making the high-risk AI system available for use.
- Deployers are required to ensure their risk management policy specifies the principles, processes, and personnel used to identify and mitigate risks of algorithmic discrimination and annually review the deployment of the high-risk AI system for algorithmic discrimination.

Finally, if the developer discovers through its own testing, or is notified by a deployer, that its high-risk AI system has caused or is reasonably likely to have caused algorithmic discrimination, the developer has an obligation to notify the Colorado attorney general, as well as other developers and deployers, of the high-risk AI system within 90 days.

Similarly, a deployer that discovers that a high-risk AI system has caused algorithmic discrimination must report the same to the attorney general within 90 days. These reporting obligations have the potential to increase the risk of investigation and reinforce the importance of compliance to help avoid regulatory scrutiny.

III. What do developers of high-risk AI systems need to do?

Developers of high-risk AI systems or "intentionally and substantially modifying" a high-risk AI system have obligations that center around providing deployers with information about the high-risk AI system that is sufficiently comprehensive to reflect the developer's internal governance and risk mitigation processes. The types of information developers need to provide include:

- The purpose of the AI system, including information on benefits, intended and reasonably foreseeable uses, intended outputs, limitations, and the harms associated with inappropriate use.
- High-level summaries of the types of data used to train the AI system, and the data governance measures used to ensure the data is suitable and the developer has mitigated for biases.
- Instructions for the deployers on how the AI systems should be used and monitored.
- Other information that may be reasonably necessary for deployers to understand and monitor the AI systems, complete impact assessments, and meet their compliance requirements under the law.

Developers are also required to provide and keep updated a public statement disclosing the types of high-risk AI systems that they develop, and how they mitigate against algorithmic discrimination.

IV. What do deployers of high-risk AI systems need to do?

The Colorado law requires deployers to establish a risk management policy and program for the lifecycle of any high-risk AI system. Deployers should take into account industry standards deemed reasonable under the law, such as the NIST AI Risk Management Framework and the ISO/IEC 42001 standard.

The policy should also be shaped around the nature and scope of the high-risk AI system and the sensitivity and volume of the data that is processed.

Deployers are required to conduct impact assessments annually and within 90 days of any qualifying medication. Impact assessments should account for:

- The purposes, intended uses and benefits, and context of the system being deployed.

- Identification and mitigation of the known or reasonably foreseeable risks of algorithmic discrimination.
- Categories of data processed as input in order to customize the system, and what outputs are produced.
- Metrics used to evaluate the performance and known limitations of the system, including cataloging modifications and whether the system is used in a manner consistent with the intended uses described by the developer.
- Post-deployment monitoring, including what oversight, use, and learning processes are in place to address identified issues.
- What transparency is provided to the consumer about the use of high-risk AI systems.

Deployers will also be required to notify consumers when using a high-risk AI system to make consequential decisions about the consumer, including disclosing the purpose of the system, the nature of the decisions made, a plain language description of the system, and the deployer's contact information.

Where applicable, the notice should include the consumer's right to opt out of profiling under Colorado's privacy law. Deployers will also be required to issue a separate, general website notice outlining the types of high-risk AI system that the deployer is using, how the deployer manages risks of algorithmic discrimination, and the nature, source, and extent of information collected and used by the deployer.

Where a consequential decision is made that is adverse to a consumer's interests, the consumer is entitled to a statement outlining the reasons for the decision, how the high-risk AI system contributed to the decision, and the nature and sources of the data that was processed, as well as being given the opportunity to correct any personal data that was processed in making the decision and an opportunity to appeal.

V. How will the law be enforced?

Violations of the Colorado AI Act will be deemed an unfair trade practice. The attorney general has exclusive enforcement authority — there is no private right of action.

Importantly, the law contains incentives for compliance; namely, affirmative defenses are available to developers and deployers where:

- A violation is discovered and cured as a result of (a) feedback encouraged by the developer or deployer, (b) adversarial testing, or (c) internal review processes; **AND**
- The developer or deployer is in compliance with the NIST AI Risk Management Framework, the ISO/IEC 42001 standard, or any other risk management framework with equivalent standards or that has been approved by the attorney general.

Further, there will be a rebuttable presumption that both deployers and developers used reasonable care to avoid algorithmic discrimination where the deployer or developer demonstrates its complied with the requirements of the law.

The attorney general may issue additional rules regarding the requirements for the applicability of rebuttable presumptions and affirmative defenses.

VI. What's next?

The attorney general is authorized to promulgate rules under the Colorado AI Act regarding documentation and requirements for developers, requirements for contents of notices and disclosures, and specific requirements for risk management policies and programs as well as impact assessments. Companies should continue to watch for additional guidance from the attorney general's office on these particulars.

While the Colorado law is currently the most comprehensive state law addressing AI systems, other states and jurisdictions continue to move toward legislation. Earlier this year, Utah and Tennessee enacted targeted legislation to address, respectively, disclosure and consumer deception, and deepfakes.

Meanwhile, New York City has a longer-standing law regulating the use of automated employment decision-making tools. Finally, California continues to advance its latest round of draft regulations, which contain stringent requirements for automated decision-making technologies.

About the author

Monique Bhargava is a partner in **Reed Smith LLP**'s entertainment and media group and co-lead of the AI section within the firm's emerging technologies practice in Chicago. She helps clients navigate compliance and legal risks relating to technology, media, advertising and data initiatives, including the development, implementation and use of artificial intelligence technologies. Her work includes advising on intellectual property and privacy issues, governance, false advertising, data collection, management and protection, digital marketing and content, social media, adtech and martech. She can be reached at mbhargava@reedsmith.com. This article was originally published May 21, 2024, on the firm's website. Republished with permission.

This article was published on Westlaw Today on May 28, 2024.