June 24, 2024

If you have any questions regarding the matters discussed in this memorandum, please contact the attorneys on the last page or call your regular Skadden contact.

One Manhattan West
New York, NY 10001
212.735.3000

Colorado's Landmark AI Act: What Companies Need To Know

Colorado has become the first state to enact a comprehensive law relating to the development and deployment of certain artificial intelligence (AI) systems. The Colorado Artificial Intelligence Act (CAIA), which **will go into effect on February 1, 2026**, adopts a risk-based approach to AI regulation that shares some similarities with the EU AI Act.

The Colorado law may spur other states to adopt similar legislation, potentially creating a patchwork of state AI laws with which companies must comply absent any omnibus federal regulation.

Key Points

- The CAIA is primarily focused on **high-risk artificial intelligence systems**, which is defined as any system that, when deployed, makes — or is a substantial factor in making — a "consequential decision." As discussed further below, **consequential decisions** generally relate to those involving education, employment, financial services, housing, health care or legal services.
- The CAIA is designed to protect against **algorithmic discrimination**, namely unlawful differential treatment that disfavors an individual or group on the basis of protected characteristics.
- The law imposes various obligations relating to documentation, disclosures, risk analysis and mitigation, governance, and impact assessments for **developers** and **deployers** of high-risk AI systems.
- With respect to ***all*** AI systems that interact with consumers, deployers must ensure that consumers are aware they are interacting with an AI system.
- The state attorney general can bring an action for violations of the CAIA as an unfair or deceptive trade practice; there is no private right of action available.

Overview

The CAIA, which was enacted on May 17, 2024, focuses on the development and deployment of "high-risk" AI systems and their potential to cause "algorithmic discrimination," which is defined as any condition in which the use of an AI system results in an unlawful differential treatment or impact that disfavors an individual or group of individuals on the basis of 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 other classification protected under the laws of Colorado or federal law.

A "high-risk" AI system is defined as any system that, when deployed, makes — or is a substantial factor in making — a "consequential decision"; namely, a decision that has a material effect on the provision or cost of:

- education enrollment or an education opportunity,
- employment or an employment opportunity,
- a financial or lending service,
- an essential government service,
- health care services,
- housing,
- insurance, or
- a legal service.

The CAIA imposes a series of obligations on developers and deployers of high-risk AI systems. A "**developer**" refers to an individual or entity doing business in Colorado that develops or intentionally and substantially modifies a high-risk AI system. A "**deployer**" refers to an individual or entity doing business in Colorado that deploys a high-risk AI system.

Both developers and deployers are required to use reasonable care to protect consumers from any known or reasonably foreseeable risk of algorithmic discrimination arising from the use of high-risk AI systems.

Developer Responsibilities

Documentation Requirements

A developer is required to make available to deployers or other developer(s) of the high-risk AI system a variety of statements and documentation, including:

- A general statement describing the reasonably foreseeable uses and known harmful or inappropriate uses of the system.
- Documentation disclosing the type of data used to train the system, any known or reasonably foreseeable limitations of the system, the purpose of the system, and the intended benefits and uses of the system.
- Documentation describing:
  - How the system was evaluated for performance.
  - Measures taken to mitigate the effects of algorithmic discrimination.
  - Data governance measures (including measures used to examine the suitability of data sources, possible biases and appropriate mitigation).
  - The intended outputs of the system.
  - How the system should be used, not be used and be monitored by an individual.
- Any additional documentation reasonably necessary to assist deployers in understanding the outputs and monitor performance of the system for algorithmic discrimination.
- Documentation and information necessary for a deployer to complete an impact assessment as required under the CAIA (see below for more).

Disclosures

Developers are also required to clearly display on their website or in a public use case inventory an up-to-date disclosure of any high-risk AI systems they have developed and make available how they manage known or reasonably foreseeable risks of algorithmic discrimination.

Disclosures to the Colorado Attorney General

Within 90 days of a developer discovering, or learning from a credible source, that their high-risk AI system has caused or is reasonably likely to cause algorithmic discrimination, they must inform the Colorado attorney general and all known deployers of the system.

The attorney general may require developers to provide certain of the documentation described above. Developers can designate such documentation as proprietary so as to avoid being disclosable under the Colorado Open Records Act, and developers do not waive attorney-client privilege by sharing this information and documentation with the attorney general.

Deployer Responsibilities

Notification to Consumers

Deployers must notify consumers when they have deployed a high-risk AI system to make — or to be a substantial factor in making — a consequential decision about the consumer before the decision is made. This disclosure must include:

- A description of the high-risk AI system and its purpose.
- The nature of the consequential decision.
- Contact information for the deployer.
- Instructions on how to access the required website disclosure (see below for more).

Information regarding the consumer's right to opt out of the processing of the consumer's personal data for profiling.

Handling Adverse Decisions

Where a high-risk AI system reaches a decision that is adverse to the consumer, the deployer must provide the consumer with a statement regarding:

- The reason for the consequential decision.
- The degree to which the high-risk AI system contributed to the decision.
- The type of data that was processed by the system and the sources of that data.

The consumer must be given the opportunity to correct any incorrect personal data used as well as an opportunity to appeal the adverse decision and request human review.

Disclosures

Deployers must clearly and readily make available on their website:

- The type of high-risk AI systems that are currently deployed.
- How they manage known or reasonably foreseeable risks of algorithmic discrimination that may arise.
- The nature, source, and extent of the information collected and used by the deployer in connection with the AI system.

Disclosure of AI Systems That Interact With Consumers

Deployers that make available any AI system that interacts with consumers (even if not high-risk) must disclose to the consumer that they are interacting with an AI system, unless it would be obvious to a reasonable person.

Establishing Reasonable Care

If the attorney general brings an enforcement action against a deployer of a high-risk AI system, there is a rebuttable presumption that the deployer used reasonable care as required under the CAIA if they satisfy the following criteria:

- Have an up-to-date risk management policy and program that specifies the principles, processes and personnel the deployer uses to identify, document and mitigate risks of algorithmic discrimination.
  - The CAIA cites the National Institute of Standards and Technology's (NIST's) "Artificial Intelligence Risk Management Framework" as a benchmark for the required risk management programs, but it allows for other national or international frameworks (such as ISO/IEC 42001), or any other framework designated by the Colorado AG.
  - The CAIA also notes the policy must be reasonable considering the deployer's size and complexity, the nature and scope of the high-risk AI system, and the sensitivity and volume of data processed.
- Perform an impact assessment that is reevaluated at least annually and within 90 days after any substantial modification to the high-risk AI system.
  - The assessment must include, among other matters:
    - A statement of the purpose, intended use cases and benefits afforded by the high-risk AI system.
    - An analysis of whether the system poses any known or reasonably foreseeable risks of algorithmic discrimination and, if so, the nature of such risk and the steps taken to mitigate the risks.
    - A description of the categories of data that the system processes as inputs and produces as outputs.
    - The metric used to evaluate the performance and known limitations of the system.
- Conduct an annual review of the high-risk AI system to ensure that it is not causing algorithmic discrimination.

Deployers will therefore want to make sure they have these policies and procedures in place even prior to any attorney general action.

Disclosures to the Colorado AG

Like developers, deployers of high-risk AI systems are required to notify the attorney general of any algorithmic discrimination within 90 days of discovery.

Exemptions

Deployers that meet the following criteria are exempt from the CAIA, other than the requirement to notify consumers that a consequential decision was made about them using a high-risk AI system:

- The deployer employs less than 50 people.
- The deployer does not use its own data to train the system.
- The system deployed is used for its intended purpose (as specified by the developer and required under the CAIA).
- The deployer makes available to consumers the impact assessment provided by the developer.

There are also exemptions for systems that have been approved by a federal agency, entities subject to the Health Insurance Portability and Accountability Act and certain other entities that are subject to existing laws and regulations.

Colorado's Landmark AI Act: What Companies Need To Know

Enforcement by the Attorney General

The attorney general has exclusive authority to enforce the CAIA. Developers or deployers bear the burden of demonstrating that the requirements set forth by the CAIA have been satisfied.

Affirmative Defenses

Developers and deployers facing an enforcement action have an affirmative defense if they have:

- cured a violation as a result of their own internal reviews or by "red teaming" (*i.e.*, following an internal process to discover risks) or external feedback, and
- complied with the latest version of the NIST AI risk management framework, another nationally or internationally recognized AI risk management framework or a framework chosen by the attorney general.

Additional Regulations To Be Developed

The attorney general has the right, but is not required, to promulgate rules as necessary in order to implement and enforce the regulations set forth in the CAIA. Such rules can pertain to, among other matters:

- The detailed documentation required from developers.
- The contents of consumer notices and disclosures.
- The content of the risk management policies and impact assessments.

What Should Organizations Be Doing?

Although the CAIA does not go into effect until February 2026, Colorado businesses that have developed or deployed high-risk AI systems, or that are planning to do so, or deployers using high-risk AI systems to make consequential decisions concerning Colorado consumers, would be well served to review the requirements of the CAIA and begin to build out a governance and compliance program so they are not left scrambling to comply.

In addition, the CAIA requirements may shape how companies go about developing and deploying these systems over the next 20 months. Steps that companies might consider include:

1. Develop a statement of the purpose, intended use cases and benefits afforded by the high-risk AI system.
2. Adhere to the NIST risk management framework or another comparable nationally or internationally recognized framework.
3. Begin developing the documentation that will be required of developers and deployers, or at least establish a process for how that documentation will be developed and maintained.
4. Start creating the necessary infrastructure (*i.e.*, principles, processes and personnel) to perform impact assessments, conduct annual system reviews and report any adverse/discriminatory findings of high-risk AI systems.
5. If planning to claim an exemption, consider the documentation that will be required to establish that exemption.

Contacts

Stuart D. Levi
Partner / New York
212.735.2750
stuart.levi@skadden.com

Ken D. Kumayama
Partner / Palo Alto
650.470.4553
ken.kumayama@skadden.com

William E. Ridgway
Partner / Chicago
312.407.0449
william.ridgway@skadden.com

Mana Ghaemmaghami
Associate / New York
212.735.2594
mana.ghaemmaghami@skadden.com

MacKinzie M. Neal
Associate / New York
212.735.2856
mackinzie.neal@skadden.com