Insights: Alert

Byte-sized Justice: Colorado's New AI Bill Seeks to Address Algorithmic Discrimination

June 6, 2024

Written by **Jon Neiditz, Meghan K. Farmer**

Written by John Brigagliano, Jon Neiditz and Meghan Farmer with contributions from Summer Associate Carter Christopher

General Background

On May 17, 2024, Governor Jared Polis of Colorado enacted the Colorado Artificial Intelligence (AI) Act (CAIA), marking a significant milestone as what its sponsor calls a “chassis,” an initial framework focusing at first on bias, for the first significant piece of AI safety legislation in the United States. The CAIA defines high-risk AI systems and delineates duties for both developers and deployers of such systems. Developers must exercise due diligence to prevent algorithmic discrimination and provide comprehensive documentation to deployers (and the Colorado Attorney General). Deployers are chiefly tasked with establishing risk management policies and conducting annual impact assessments. Exclusive enforcement lies with the Attorney General, and consumers are granted rights such as pre-use notice and the ability to opt-out of some decision-making (subject to material exceptions).

Defenses and exceptions are also outlined, including carveouts for trade secret protection and exemptions for small businesses, insurers, and consumer financial institutions meeting specific criteria. Those exceptions, paired with the law's focus on high-risk AI systems and on discrimination-related risks, cabin the law's scope so that the law does not comprehensively regulate AI (despite what you might read in some legal media). Yet the focus on discrimination can be viewed as part of the “chassis” strategy, permitting expansion, and the focus on high-risk systems may be precisely the way to encourage innovation while prioritizing regulation appropriately.

What Is a High-Risk AI System?

A high-risk AI system is “any artificial intelligence system that, when deployed, makes, or is a substantial factor in making, a consequential decision.”[1] A consequential decision is a decision that has a legal or significant effect on “(a) education enrollment or an education opportunity; (b) employment or an employment opportunity; (c) a financial or lending service; (d) an essential government service; (e) health-care services; (f) housing; (g) insurance; OR (h) a legal service."[2]

What Duties Do Developers Have?

Under the CAIA, there are distinct requirements for developers and deployers. A developer, as defined by the law, is an individual or entity operating within Colorado who creates or alters an artificial intelligence system.

Developers' duties are summarized as followed:

- Exercise due diligence in safeguarding consumers against any known or foreseeable risks of algorithmic discrimination from intended or contracted uses
- Disclose information about their system's use, risks, data, biases, intentions, mitigation protocols, and other particulars
- Prepare publicly available documents that detail high-risk AI systems for deployers
- Promptly notify the AG upon discovery of algorithmic discrimination

What Duties Do Deployers Have?

As the entity interfacing with Colorado residents, deployers shoulder a greater load of responsibilities compared to developers under the CAIA. A deployer, as defined by the legislation, refers to an individual or entity conducting business within the state who puts into operation a high-risk artificial intelligence system.

Duties for deployers include the following:

- Protect consumers from known or reasonably foreseeable risks of algorithmic discrimination
- Establish a robust risk management policy that explains procedures and personnel involved in identifying and mitigating algorithmic discrimination
- Conduct an annual impact assessment that provides details on the system's purpose, intended uses, risks of discrimination, mitigation strategies, and other relevant responsibilities
- Issue a pre-deployment statement of use to consumers disclosing the use of AI, its decision-making process, and the option to opt out
- Promptly notify the AG upon discovery of algorithmic discrimination

Industry and Data-Specific Considerations

i. **Financial Institutions**
Despite regulating discrimination related to lending, insurance, and healthcare, the law also exempts many institutions that offer such services. Although the CAIA lacks the broad exemption for any data or company subject to the Gramm-Leach-Bliley Act's (GLBA) privacy rule (as typically occurs in consumer privacy laws), the CAIA nonetheless exempts banks and credit unions subject to regulatory oversight that is equivalent or stricter than the CAIA and require algorithmic discrimination mitigation.

ii. **Healthcare and Insurance Exemptions**
Opposite those material exceptions in consumer finance, the law's HIPAA “exemption” looks like it was negotiated to be virtually meaningless by exempting only healthcare decisions that are not high risk (when the law regulates only high-risk AI systems). Insurers that are subject to Colorado algorithmic and predictive discrimination regulations in Section 10-3-1104.9 are also exempt from the act.

iii. **FCRA Application**

Under the Colorado Privacy Act (CPA), data subject to the Fair Credit Reporting Act (FCRA) is excluded from the law's purview (likely to avoid dormant commerce clause challenges). The CAIA differs from the CPA in this regard. So while many companies avoided offering a right to opt out of “profiling” under the CPA on the basis that all in scope use cases were subject to the FCRA, that basis for avoiding compliance will not be available under the CAIA.

iv. **Ad Placement**

The new law creates an interesting question of whether displaying an ad for a service like housing or healthcare counts as a consequential decision within the law's scope. Could the legislation extend far enough that a targeted online offer (generated by AI systems in the ad tech ecosystem) for healthcare or lending could not be delivered in a discriminatory manner? While we do not know for sure how regulators would interpret this issue, it is worth considering when addressing potential compliance challenges with the new law. Moreover, combating discriminatory deliver of online ads could present a major tradeoff between privacy and safety (i.e., inferring a user's demographics to combat discriminatory practices).

Enforcement and Rulemaking Authority

Under the Colorado Artificial Intelligence Accountability Act (CAIA), enforcement falls squarely under the purview of the Attorney General, who wields exclusive authority. The Attorney General is also empowered to promulgate rules aimed at implementing and enforcing the act. These rules encompass various facets, including disclosures, mandates concerning risk management policies, stipulations regarding impact assessments, criteria for rebuttable presumptions, and specifications for affirmative defenses. The penalty for violating this Act is up to $20,000 per consumer relationship or transaction in violation of the law.

Defenses

The CAIA also offers defenses and safe harbors. Those who adhere to relevant bill provisions maintain a rebuttable presumption of reasonable care. Additionally, an affirmative defense is available if a developer or deployer identifies and rectifies a violation internally, following the NIST Artificial Intelligence Risk Management Framework. These measures aim to provide guidance and protection for stakeholders involved in AI development and deployment, promoting accountability and adherence to established standards.

General Exceptions

An already narrow bill that focuses almost exclusively on algorithmic discrimination, the CAIA creates exceptions that further narrows the scope and applicability of the act. One of the biggest exceptions is the disclosure of trade secrets. If a trade secret is protected from disclosure by federal law, then there is no duty to disclose for compliance with this act.[3] The CAIA also exempts banks and credit unions that are subject to a similar or more strict regulatory system applicable to AI use and require regular auditing of the systems.

Some small businesses are also exempt from deployer responsibilities. If a deployer employs fewer than fifty full-time employees, does not train a high-risk AI system with its own data, limits the use of the high-risk AI systems to those previously disclosed, and provides impact assessments to consumers from developer, then the deployer can skip risk-management programs, assessments, and notices.

Consumer Rights

While the Colorado act does not permit private legal action, the law does grant consumers specific rights, including to challenge high-risk decisions. Before AI systems are used to make consequential decisions, consumers must receive pre-use notices. They must also be informed if an AI system is meant to interact with them and, under the much narrower Colorado Privacy Act, have the option to opt-out of profiling in automated decisions.

Moreover, consumers possess enforceable rights if they face adverse consequential decisions. They are entitled to explanations detailing the primary reasons behind decisions, the AI system's contribution, the data used, and its source. Additionally, they have the right to correct any inaccuracies in their personal data utilized by the AI system and (subject to some exceptions) can appeal decisions for human review.

In Comparison with Utah

Utah recently enacted S.B. 149, also known as the AI Law, which mandates that businesses and individuals utilizing generative artificial intelligence to engage with consumers disclose their use of GenAI for commercial activities overseen by the Division of Consumer Protection. This disclosure must be prominently displayed and easily noticeable. However, it is only required if the consumer explicitly asks whether they are interacting with a human or AI. Additionally, the law stipulates that regulated professions must inform individuals of their interaction with GenAI through the same medium of communication, even without a specific request.

In contrast, Colorado's recent legislation takes a different approach, focusing less on AI's direct engagement with consumers and more on its role in making consequential decisions. Nevertheless, Colorado similarly mandates disclosure of AI systems that directly interface with consumers.

Conclusion

Colorado's latest legislation represents a significant stride in the national legislative landscape concerning artificial intelligence regulation. These emerging laws aim to govern the conduct of consumers, businesses, and industries, fostering responsible AI utilization. While only five states have enacted AI-related legislation, seven others are actively engaged in the legislative process or have made attempts to pass similar measures. Colorado's initiative-taking stance positions it as a pioneering state in this domain, reflecting the evolving dynamics of AI regulation on a national scale.

As AI legislation begins to become increasingly relevant, businesses need to be proactive in their implementation of AI systems. Business should understand the trends of legislation, staying informed of the industry best practices to ensure compliance. Businesses also need to take an inside look at their AI technology, application, and use to understand the effect the new law has on a specific practice or industry. Employers should identify any risk or gaps between their system and legal compliance. Businesses should practice strategy that ensures compliance, while maximizing the benefits of AI.

Footnotes

---

[1] 6-1-1701(9)(a)

[2] 6-1-1701 (3)(a)-(h)

[3] 6-1-1703 (8)

Related People

Jon Neiditz

t 404.815.6004

jneiditz@ktslaw.com

Meghan K. Farmer

t 404.541.6816

mfarmer@ktslaw.com