

California to consider regulating AI algorithm bias again — with even higher stakes

By Eric He

04/05/2024 08:00 AM EDT

Last year, California lawmakers attempted to establish greater regulations around companies using artificial intelligence — but they were unsuccessful amid broad industry opposition. Now that the bill is in its second iteration, the conversation around the technology has grown more expansive and a couple of industry stakeholders are working to advance the legislation.

[AB 2930](#), authored by Assemblymember [Rebecca Bauer-Kahan](#) (D-Orinda), would require companies to perform impact assessments before releasing AI tools and prevent them from deploying if the algorithm is found to result in discriminatory outcomes.

The proposal once again has the support of Workday — which [has been working](#) to push its own model AI legislation in various states — and has also gained the backing of Microsoft, but remains likely to face opposition from other business interests.

Bauer-Kahan [said in a statement](#) that the bill would help ensure AI does not “drag us backwards with biased results.”

“Algorithms are already making decisions about our lives, and there’s no check on the bias,” Bauer-Kahan said. “These hidden forces are embedding discrimination further into our system and it’s unacceptable.”

WHAT’S IN THE BILL?

This Pro Bill Analysis is based on the [text of the bill](#) as introduced on Feb. 15.

The bill would amend the [California Business and Professions Code](#) to bar companies from using artificial intelligence that results in algorithmic discrimination, which would be defined as technology that contributes to “unjustified differential treatment” or disfavoring people based on [protected classes under state law](#), such as ethnicity, sexual orientation and religion (Sec. 1).

By 2026, the measure would establish regulations for anyone who uses artificial intelligence to make a “consequential decision,” meaning something that impacts a person’s access to (Sec. 1):

- Government benefits, services or civil fines
- Employment, including pay or promotion, hiring or termination and having an algorithm allocate tasks that limit, segregate or classify employees
- Education, including testing, detecting cheating or plagiarism, accreditation, certification, admissions and financial aid or scholarships
- Housing, including the ability to rent
- Essential utilities, like electricity, heat, water, internet and transportation
- Family planning
- Adoption services, reproductive services or assessments related to child protective services

- Health care or insurance, including mental health, dental and vision
- Financial services, such as services provided by mortgage companies, brokers or creditors
- The criminal justice system, including risk assessments for pretrial hearings, sentencing and parole
- Legal services
- Private arbitration
- Mediation
- Voting

The bill would specifically regulate any AI used through an “automated decision tool” that is specifically developed and marketed, or modified, to be a “controlling factor” in making a consequential decision.

The measure would apply to “deployers” — or anyone who uses AI to make a consequential decision — as well as “developers,” defined as anyone who designs, codes or produces a tool using AI to be a “controlling factor” in making a consequential decision — including any use by a third party.

Any person, business or agency that deploys or develops an automated decision tool would have to conduct a risk-based impact assessment that includes:

- A statement of purpose outlining benefits, uses and how it would be deployed
- A description of what outputs the AI would generate and how it would be used as a controlling factor in making a consequential decision
- A summary of the data collected about people and processed using AI to make a consequential decision
- An analysis of potential adverse impacts of AI on the basis of sex, race, color, ethnicity, religion, age, national origin, limited English proficiency, disability, veteran status or genetic information
- A description of how the AI could be used or monitored to make a consequential decision, as well as any safeguards to protect against algorithmic discrimination

If a deployer or developer performs a significant update to the system, it would have to perform another impact assessment “as soon as feasible” (Sec. 1).

Additionally, deployers would have to describe how the AI will be evaluated for validity or relevance. Deployers with fewer than 25 employees would be exempt from the assessment and having to set up a governance program — unless they used a tool that impacted 1,000 people or more per year.

Before an AI user makes a consequential decision with the help of the technology, they must notify anyone who would be impacted by that decision. The notification would need to include the reason for using AI, the deployer’s contact information and a breakdown of human components of the technology and how automated components are used in decision-making.

If it is “technically feasible,” a deployer would be prohibited from using only AI to make a consequential decision if an impacted person requests not to be subject to an automated decision tool. Instead, they would have to use an alternative selection process or accommodation, with the ability to “reasonably request, collect and process information” from the person for identification purposes. The deployer could then use AI if the person does not provide that information.

AI developers would have to — without needing to divulge trade secrets — provide deployers with a statement on the intended uses of the technology, along with documentation containing:

- The known limitations of the AI, including “reasonably foreseeable risks” of discrimination from the algorithm
- The type of data used to program or train the AI

— How the AI was evaluated for validity and explainability before sale or licensing

Companies that develop or use AI would have to establish a governance program and create safeguards to protect against foreseeable risks of discrimination from an AI-run algorithm. Safeguards would be applied consistent with what the AI is being used for, the role and size of the deployer or developer, the context of the deployer or developer in connection with the AI and how “technically feasible” it is for deployers and developers to manage the risks associated with AI.

The governance program would have to be overseen by at least one employee, who would have the authority to tell the company that an AI tool is in violation of the measure and the company would be required to assess the issue. The program would also have to implement safeguards against “reasonably foreseeable risks” of discrimination from an AI algorithm, and run the impact assessments required by the measure. The results of the assessments would have to be maintained for two years (Sec. 1).

Each year, companies would need to conduct a comprehensive review of its policies, practices and procedures pursuant to the bill. If there are material changes in technology, the risks associated with the AI, the state of technical standards or changes in operations, the company would have to “evaluate and make reasonable adjustments” to its safeguards.

Companies would have to make their AI policies publicly available, including the types of AI tools being used and how the company would manage the risk of future algorithmic discrimination.

Allegations of algorithmic discrimination would be investigated by the state’s Civil Rights Department, which could direct a company to provide its impact assessment within seven days. A company would not be able to claim attorney-client privilege or work-product protections. Trade secrets included in the assessment would not be subject to public records requests and could be redacted if it is included as part of a request.

Companies that fail to provide the impact assessment to the department could be fined up to \$10,000 per day.

The measure would allow the California Attorney General and local prosecutors to bring a civil lawsuit against a company for violations. Companies could be fined up to \$25,000 per violation of algorithmic discrimination. A court could also award injunctive relief, declarative relief and attorney’s fees.

Companies would have a cure provision of 45 days to address any alleged violations after a written notice from an attorney before a lawsuit is filed. If the violation is cured, the plaintiff would not be able to sue for injunctive relief.

The bill would explicitly make it illegal for companies to retaliate against anyone who asserts their rights under this measure (Sec. 1).

The measure would not apply to cybersecurity-related technology.

Sections 2 and 3 of the bill contain technical provisions related to proprietary information and state reimbursement (Secs. 2, 3).

WHO ARE THE POWER PLAYERS?

AB 2930 is the second attempt by Assemblymember [Rebecca Bauer-Kahan](#) (D-Orinda) to clamp down on how companies use algorithms run by artificial intelligence. While it does not have any co-authors yet, Assembly Speaker [Robert Rivas](#) (D-Hollister) signed onto [last year’s bill](#).

The bill text resembles [draft language](#) pushed by the workforce management software company **Workday**, which has lobbied lawmakers in several states to pass similar bills, [according to Recorded Future News](#).

Chandler Morse, vice president of public policy at Workday, [said in a statement](#) that the company was “pleased to have contributed” to AB 2930.

“With a growing appetite to regulate AI at the state level, California can lead the nation in establishing meaningful guardrails,” Morse said. “This new bill would establish concrete requirements and strong protections for consequential decisions driven by AI.”

Microsoft, which did not have a position on last year's bill, is a full-throated supporter this year. **Robyn Hines**, senior director of state government affairs, said in a statement that the legislation will lead to "strong legal guardrails" for AI and empower "safe and transparent AI solutions."

Bauer-Kahan [told POLITICO](#) that having Microsoft on board is an admittance from a major power player in the AI space that it can comply with regulations, and it "takes some of the wind out of the sails of opposition."

"That's an exciting move, to see a responsible AI provider actually stepping up to do what's right in the anti-discrimination space," she said.

Opposition is still likely to put up a fight, though. The **California Chamber of Commerce** — which opposed last year's bill — has not taken a formal position on AB 2930 but spokesperson Denise Davis anticipates it will be against the bill again.

However, Davis said in an email that the Chamber agrees with Bauer-Kahan's underlying goal and is "looking forward to working with her to resolve our concerns."

The Chamber was among dozens of business groups last year that signed onto an [opposition letter](#) calling it "critical to reduce bias and discrimination in consequential decisions impacting people" but said discrimination exists whether the decision is being made by humans or technology.

"We believe it is critical that any regulatory efforts proceed with precision, particularly as technology is still developing and has the potential to reduce, if not one day eliminate, such undesirable outcomes," the letter stated.

The **California Retailers Association**, **California Grocers Association** and **NetChoice** were also among the groups that penned the April 2023 letter.

Bauer-Kahan said opponents this year will probably include "Other companies [and] all the people who called me last year and said, 'Why don't you just exempt my industry?'"

WHAT'S HAPPENED SO FAR?

AB 2930 is largely similar to last year's [AB 331](#), with a few changes that were suggested in the opposition coalition's letter.

For example, this year's version would not allow a private individual to sue, notwithstanding existing workplace discrimination laws. The opposition last year called a private right of action "simply not feasible for such an important and growing technology," and "highly problematic and chilling."

Additionally, companies would not have to prove that no further violations will occur to meet the right to cure provision. Industry stakeholders said in the letter that it would be unrealistic for companies to guarantee that they would not violate the law in the future.

AB 2930 would also allow the Civil Rights Department to investigate complaints of algorithmic discrimination — something that was not in last year's bill.

Similar bills have been introduced this year in [Illinois](#), [Rhode Island](#), [Connecticut](#), [New York](#) and [Washington](#), all of which have elements of Workday's reported draft language. But so far, none of the bills have advanced through a committee.

Critics of the legislation told [Recorded Future News](#) they fear companies would claim that they do not use AI as a "controlling factor" in making decisions, thus exempting them from the regulation because the definition of "automated decision tool" would not apply. Further, enforcement could be lax due to the lack of an independent auditor overseeing compliance. The California proposal would only require governance programs set up by the companies themselves.

WHAT'S NEXT?

The bill is currently awaiting hearings in the Assembly Privacy and Consumer Protection Committee and the Assembly Judiciary Committee. Bauer-Kahan chairs the privacy committee, and Assemblymember [Ash Kalra](#) (D-San Jose) — who has an [AI proposal](#) of his own regulating the technology's use in the entertainment industry — helms judiciary.

Microsoft's support may sway industry stakeholders or lessen opposition, and the ever-growing developments around AI will likely heighten the stakes of negotiations compared to last year.

Meanwhile, on the Senate side, Sen. [Scott Wiener](#) (D-San Francisco) [has a bill](#) that would place safeguards on large AI models — another sweeping proposal that is expected to face industry pushback. Wiener's bill [cleared its first test](#) in the Senate Judiciary Committee on April 2.

WHAT ARE SOME STORIES ON THE BILL?

[Read POLITICO news on AB 2930.](#)

Jeremy B. White contributed to this report.