

Date of Hearing: June 23, 2026

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

SB 947 (McNerney) – As Amended June 16, 2026

PROPOSED AMENDMENTS

SENATE VOTE: 29-9

SUBJECT: Employment: automated decision systems

SYNOPSIS

Automated Decision Systems (ADS) are typically predictive artificial intelligence (AI) systems that produce simplified outputs – such as scores, classifications, or recommendations – to assist or replace human discretionary decisionmaking. ADS can process enormous datasets, identify hidden patterns, and make decisions with efficiency and scale that vastly exceeds human capabilities. But relying on ADS to make life-impacting decisions can be hazardous if deployed without proper oversight: the datasets they are trained on are often unrepresentative or contaminated with bias, the inferences they draw from those datasets are often inscrutable, and these systems can fail to accurately account for the complexity of human behavior. Without human-centered oversight, particularly in consequential contexts such as employment, housing, healthcare, and criminal justice, these impacts can be irreparable.

One especially impactful area in which ADS are being deployed at scale is employment. Employers are increasingly using ADS to manage workforce operations, including hiring, setting schedules, assigning tasks, and tracking performances, as well as for disciplinary purposes, such as automated tracking of workplace violations. This bill would establish a comprehensive set of requirements governing the use of ADS in the workplace by requiring human oversight and independent corroboration when employers use ADS for decisions that impact workers' livelihoods. Certain types of ADS, such as those that infer sensitive information about workers or attempt to predict their behavior, are prohibited. And with respect to especially consequential decisions involving discipline, termination, or deactivation, the bill prohibits employers from relying solely on the ADS to make the decision. The bill also establishes anti-retaliation protections for workers and provides for enforcement by the Labor Commissioner and public prosecutors. The bill is a narrower version of SB 7 (McNerney, 2025), which was passed by the Legislature but ultimately vetoed by the Governor.

This bill is sponsored by the California Federation of Labor Unions AFL-CIO. It is supported by a large coalition of labor organizations and advocacy groups, including the California Nurses Association and TechEquity Action. Opponents include a broad coalition of industry groups led by California Chamber of Commerce.

This bill was previously heard by the Labor Committee, where it passed on a 5-1 vote. If passed by this committee, it will next be referred to the Judiciary Committee.

EXISTING LAW:

- 1) Establishes the California Privacy Protection Agency (Privacy Agency) and vests it with full administrative power, authority, and jurisdiction to implement and enforce the California Consumer Privacy Act of 2018. (Civ. Code § 1798.199.10.)
- 2) Requires the California Privacy Protection Agency to issue regulations governing access and opt-out rights with respect to businesses' use of automated decisionmaking technology. (Civ. Code § 1798.185.)
- 3) Establishes the Civil Rights Department, and sets forth its statutory functions, duties, and powers. (Gov. Code § 12930.)
- 4) Establishes the Fair Employment and Housing Act. (Gov. Code § 12900 *et seq.*) Makes it an unlawful employment practice, unless based upon a bona fide occupational qualification, for an employer to refuse to hire or employ the person or to refuse to select the person for a training program leading to employment, or to bar or discharge the person from employment or from a training program leading to employment, or to discriminate against the person in compensation or in terms, conditions, or privileges of employment based upon specified characteristics, including race, religious creed, color, national origin, ancestry, physical disability, mental disability, reproductive health decisionmaking, medical condition, genetic information, marital status, sex, gender, gender identity, gender expression, age, sexual orientation, or veteran or military status. (Gov. Code § 12940.)
- 5) Requires the California Department of Technology to conduct a comprehensive inventory of all high-risk ADS that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency. Defines the following terms:
 - a) "Artificial intelligence" means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.
 - b) "Automated decision system" means a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons. "Automated decision system" does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.
 - c) "High-risk automated decision system" means an ADS that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to, or approval for, housing or accommodations, education, employment, credit, health care, and criminal justice. (Gov. Code § 11546.45.5.)

THIS BILL:

- 1) Defines, among other terms:

- a) “ADS output” as any information, data, assumptions, predictions, scoring, recommendations, decisions, or conclusions generated by an ADS.
 - b) “Employer” as any person who directly or indirectly, or through an agent or any other person, employs or exercises control over the wages, benefits, other compensation, hours, working conditions, access to work or job opportunities, or other terms or conditions of employment, of any worker. This shall include all branches of state government, including all cities, counties, charter counties, municipalities, charter municipalities, cities and counties, special districts, transit districts, the University of California upon agreement by the regents, the California State University, community college districts, school districts, or any other governmental entity, as well as a labor contractor of a person otherwise defined as an employer.
 - c) “Employment-related decision” as any decision by an employer that impacts wages, wage setting, benefits, compensation, work hours, work schedule, performance evaluation, hiring, discipline, promotion, termination, job tasks, skill requirements, work responsibilities, assignment of work, access to work and training opportunities, productivity requirements, workplace health and safety, and any other terms or conditions of employment.
 - d) “Predictive behavior analysis” as any system that predicts a worker’s behavior, beliefs, or intentions.
 - e) “Worker” as any natural person who is an employee of, or an independent contractor providing service to, or through, a business or a state or local governmental entity in any workplace.
 - f) “Worker data” as any information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with, a worker, regardless of how the information is collected, inferred, or obtained.
- 6) Prohibits an employer from using an ADS that does any of the following:
- a) Prevents compliance with or results in a violation of any federal, state, or local labor, occupational and safety, employment, or civil rights laws or regulations.
 - b) Infers certain protected characteristics about a worker.
 - c) Conducts predictive behavior analysis on a worker with the intention of using the analysis to make an employment-related decision.
 - d) Predict or take adverse action against a worker for exercising their legal rights, including rights guaranteed by state and federal employment and labor law.
- 7) Prohibits an employer from relying solely on an ADS when making a disciplinary, termination, or deactivation decision. If an employer primarily relies upon an ADS output to make a disciplinary, termination, or deactivation decision, requires the employer to direct a human reviewer to conduct an independent investigation into the decision. If the employer cannot corroborate the output, or the human reviewer concludes the output is inaccurate,

incomplete, or misleading, prohibits the employer from using the output to discipline, terminate, or deactivate a worker.

- 8) Grants workers the right to request and receive a copy of the most recent 12 months of the worker's data primarily used by an ADS to make a disciplinary, termination, or deactivation decision once every 12 months. Requires employers to provide that data in a manner that anonymizes customers, workers, and individuals' personal information.
- 9) With respect to ADS used to make a discipline, termination, or deactivation decision, requires the employer to provide the affected worker with a written post-use notice containing specified information at the time the employer informs the worker of the decision.
- 10) With respect to a data access request, requires an employer to provide the worker with a plain-language written document accessible away from the workplace that contains specified information, including the specific decision for which the employer used the ADS, the specific worker data that the ADS used, the specific ADS output produced by the ADS, any additional corroborating or supporting information used in addition to the ADS output in making the decision, the name of the vendor or entity that created the ADS and the product name of the ADS, and a copy of any completed impact assessments regarding the ADS in question.
- 11) Establishes anti-retaliation protections for workers who attempt to use their rights provided under the bill, as specified.
- 12) Grants enforcement authority to the Labor Commissioner or public prosecutors. An employer is subject to \$500 per violation and claimants may seek damages, including punitive damages, injunctive relief, and fees and costs.
- 13) Provides that it does not preempt any city, county, or city and county ordinance that provides equal or greater protection to workers who are covered by the bill.
- 14) Specifies that an employer who complies with the requirements related to notice and appeal under the bill is not required to comply with any substantially similar notice and appeal provisions related to automated decision systems used in employment-related decisions required under any other state law, other than employers that are businesses subject to the California Consumer Privacy Act of 2018.
- 15) Specifies that it does not apply to parties covered by a collective bargaining agreement if the agreement explicitly waives this part in clear and unambiguous terms, expressly provides for the wages or earning, working conditions, and other terms and conditions of work, and provides protection from algorithmic management.
- 16) Provides that its provisions are severable.

COMMENTS:

- 1) **Author's statement.** According to the author:

Businesses are increasingly using AI to boost efficiency and productivity in the workplace. But there are currently no safeguards to prevent machines from unjustly or illegally impacting workers' livelihoods and working conditions. SB 947 establishes post-use notifications when ADS is used in the workplace, and ensures employers have a human in the loop when making critical employment decisions.

SB 947 does not prohibit ADS in the workplace, rather it will establish guardrails to ensure that California businesses are not operated by robo bosses, because there will be a human in the loop. AI must remain a tool controlled by humans, not the other way around.

2) Background. AI and GenAI. In brief, AI is the mimicking of human intelligence by artificial systems such as computers. AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process: numbers, text, audio, video, or movement. AI is not fundamentally different from other computer functions; its novelty lies in its application. Unlike traditional computer functions, AI can accomplish tasks that are normally performed by humans. AI that is trained on small, specific datasets to make recommendations and predictions is sometimes referred to as “predictive AI.” This differentiates it from GenAI, which is trained on massive datasets in order to produce detailed text and images. When Netflix suggests a TV show to a viewer, that recommendation is produced by predictive AI that has been trained on the viewing habits of Netflix users. When ChatGPT generates text in clear, concise paragraphs, it uses GenAI that has been trained on the written contents of the internet.

Automated decision systems. Automated decision systems (ADS) typically use predictive AI to produce simplified outputs – such as scores, classifications, or recommendations – to assist or replace human discretionary decisionmaking.¹ ADS can process enormous datasets, identify hidden patterns, and make decisions with efficiency and scale that vastly exceeds human capabilities. This has led to profoundly beneficial applications and breakthroughs.²

But relying on ADS can be hazardous if the systems are not trained carefully or tested thoroughly: the datasets they are trained on are often unrepresentative or contaminated with bias, the inferences ADS draw from those datasets are often inscrutable, and these systems can fail to accurately account for the complexity of human behavior. When deployed without proper oversight in consequential contexts such as employment, housing, healthcare, and criminal justice, the impacts of flawed ADS can be devastating.

Algorithmic discrimination. There is a well-known saying in computer science: “garbage in, garbage out.” The performance of an ADS is directly impacted by the quality, quantity, and relevance of the data used to train it.³ If the data used to train the ADS contain bias, the tool’s outputs will be similarly biased:

¹ Government Code section 11546.45.5(a)(1) defines an ADS as “a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons.”

² See e.g. Santariano & Metz, “Using A.I. to Detect Breast Cancer That Doctors Miss,” *New York Times* (Mar. 5, 2023), <https://www.nytimes.com/2023/03/05/technology/artificial-intelligence-breast-cancer-detection.html>.

³ Rohit Sehgal, “AI Needs Data More Than Data Needs AI,” *Forbes* (Oct. 5, 2023), <https://www.forbes.com/sites/forbestechcouncil/2023/10/05/ai-needs-data-more-than-data-needs-ai/>.

Algorithmic discrimination occurs when automated systems contribute to unjustified different treatment or impacts disfavoring people based on their race, color, ethnicity, sex (including pregnancy, childbirth, and related medical conditions, gender identity, intersex status, and sexual orientation), religion, age, national origin, disability, veteran status, genetic information, or any other classification protected by law.⁴

Over the past thirty years, industries and governmental entities have been forced to contend with this problem as they have attempted to introduce ADS into their workflows. In 2015, Amazon opted against automating their hiring process when they realized that their ADS-enabled system was excluding women from the pool of acceptable candidates because it had been trained to vet applicants by observing patterns in resumés submitted to the company over a 10-year period. Most came from men, a reflection of inequities across the tech industry.⁵

Unsafe or ineffective systems. In addition to discriminatory outcomes, some ADS are unsafe or ineffective regardless of who the subjects of the tool’s prediction are. As Princeton researchers Arvind Narayanan and Sayash Kapoor put it bluntly in *AI Snake Oil*: “In contrast to generative AI, predictive AI often does not work at all.”⁶ The correlations that machine-learning ADS rely on may have little to do with the attributes they purportedly measure. Some hiring tools trained on videos of successful employees are used to assess the fitness of job applicants who are required to record video responses to specific prompts. Researchers have found that such tools can easily be manipulated by making simple changes to the subject’s appearance (such as wearing glasses) or to the background of the room (such as adding more books to a bookshelf), leading to increased scores. Journalist Hilke Schellmann found she was able to obtain consistently high scores despite responding to a hiring tool’s prompt by reading an irrelevant Wikipedia entry in German.⁷

In another case, “[a] company installed AI-powered cameras in its delivery vans in order to evaluate the road safety habits of its drivers, but the system incorrectly penalized drivers when other cars cut them off or when other events beyond their control took place on the road. As a result, drivers were incorrectly ineligible to receive a bonus.”⁸

Especially questionable are ADS that purportedly forecast individual human behavior. A recently released study compiled a list of 47 applications of ADS that use machine learning to predict the future behavior or outcomes for individuals in eight domains: criminal justice, healthcare, welfare, finance, education, workplace, marketing, and recommender systems. The study concluded that such tools frequently fall well short of their purported benefits. The authors argue that developers and deployers of such systems should have the burden of demonstrating that their

⁴ White House Archives, “Algorithmic Discrimination,” <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/algorithmic-discrimination-protections/>.

⁵ Jeffrey Dastin, “Amazon scraps secret AI recruiting tool that showed bias against women,” *Reuters* (Oct. 9, 2018), <https://www.reuters.com/article/amazoncom-jobs-automation/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSL2N1VB1FQ/>.

⁶ *AI Snake Oil*, *supra*, at p. 9.

⁷ See Hilke Schellmann, *The Algorithm: How AI Decides Who Gets Hired, Monitored, Promoted, and Fired and Why We Need to Fight Back Now* (1st ed. 2024).

⁸ *Blueprint*, *supra*, p. 17.

tools are not harmful.⁹ As Narayanan and Kapoor write: “Accurately predicting people’s social behavior is not a solvable technology problem and determining people’s life chance on the basis of inherently faulty predictions will always be morally problematic.”¹⁰

Robot bosses. Employers are increasingly using ADS to manage workforce operations, including hiring, setting schedules, assigning tasks, and tracking performances. Three key functions are 1) ranking workers, 2) predicting their future behavior or characteristics, and 3) giving directions to workers.

ADS tools that rank workers do so by scoring or comparing their performance, behavior, or traits using metrics such as productivity, accuracy, manager evaluations, and customer ratings. These systems often generate alerts or recommendations for employers, such as coaching prompts and suggested interventions. According to the UC Berkeley Labor Center, examples of these tools include:

- *Central*: A platform that collects real-time data on worker task completion, accuracy, training progress, and sales, and generates worker rankings, gamified scoring systems, and performance dashboards for managers.
- *WorkTime*: A platform that monitors office workers’ computer activity and generates productivity and distraction score rankings, with team and individual performance comparisons displayed on leaderboards or manager dashboards.
- *Harver*: A hiring platform that uses behavioral and personality assessments to score and rank job candidates by predicted fit, enabling employers to set cutoff thresholds and identify top job matches.¹¹

ADS tools that make predictions about worker behavior or characteristics use personal information to identify patterns and correlations in datasets about other workers to predict things such as performance, ability to work on teams, stress tolerance, likelihood of joining a union, or becoming pregnant. For example, Perceptryx “analyzes employee surveys and HR data to predict engagement, retention risk, and vulnerability to union organizing, using tools like a Union Vulnerability Index to flag high-risk groups or locations and inform employer response strategies.”¹²

ADS tools that give workers directions about their job tasks use data from workplace surveillance devices to deliver real-time instructions on which tasks to perform and in what order. Examples, as described by the UC Berkeley Labor Center, include:

⁹ Angelina Wang et al. 2023. “Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy.” In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (Chicago, IL, USA: ACM, 2023), 626.

¹⁰ *AI Snake Oil*, *supra*, at p 15.

¹¹ Annette Bernhardt & Lisa Krege, “Electronic Monitoring and Automated Decision Systems: Frequently Asked Questions,” UC Berkeley Labor Center (May 2025), p. 2, <https://laborcenter.berkeley.edu/wp-content/uploads/2025/05/Electronic-Monitoring-and-Automated-Decision-Systems-FAQ.pdf>.

¹² *Id.* at p. 3.

- *Amadeus HoTSOS*: A system used in hotels to direct housekeeping, maintenance, and guest service tasks. The platform sends real-time instructions to workers based on guest check-ins, service requests, or sensor data. As conditions shift, the system reorders tasks and reassigns them among staff.
- *Trackforce*: A system used in the security sector to dispatch real-time instructions to security guards. It assigns patrols, alarm responses, or site-specific checks based on live inputs like incidents or visitor activity. As new issues arise, the system updates workflows, verifies task completion order, and alerts supervisors to delays or deviations.¹³

Blueprint for an AI bill of rights. In 2022, the White House Office of Science and Technology Policy released the *Blueprint for an AI Bill of Rights*, which identifies five principles that should “guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence.”¹⁴ As summarized in the *Blueprint*, the principles are as follows:

- *Safe and Effective Systems*: You should be protected from unsafe or ineffective systems. Automated systems should be developed with consultation from diverse communities, stakeholders, and domain experts to identify concerns, risks, and potential impacts of the system. Systems should undergo pre-deployment testing, risk identification and mitigation, and ongoing monitoring that demonstrate they are safe and effective based on their intended use, mitigation of unsafe outcomes including those beyond the intended use, and adherence to domain-specific standards. Outcomes of these protective measures should include the possibility of not deploying the system or removing a system from use. Automated systems should not be designed with an intent or reasonably foreseeable possibility of endangering your safety or the safety of your community. They should be designed to proactively protect you from harms stemming from unintended, yet foreseeable, uses or impacts of automated systems. You should be protected from inappropriate or irrelevant data use in the design, development, and deployment of automated systems, and from the compounded harm of its reuse. Independent evaluation and reporting that confirms that the system is safe and effective, including reporting of steps taken to mitigate potential harms, should be performed and the results made public whenever possible.
- *Algorithmic Discrimination Protections*: You should not face discrimination by algorithms and systems should be used and designed in an equitable way. . . . Designers, developers, and deployers of automated systems should take proactive and continuous measures to protect individuals and communities from algorithmic discrimination and to use and design systems in an equitable way. This protection should include proactive equity assessments as part of the system design, use of representative data and protection against proxies for demographic features, ensuring accessibility for people with disabilities in design and development, pre-deployment and ongoing disparity testing and mitigation, and clear organizational oversight. Independent evaluation and plain language

¹³ *Ibid.*

¹⁴ The White House, *Blueprint for an AI Bill of Rights*, (Oct. 2022), p. 14, <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/> (*Blueprint*). Despite the use of the term “AI” in its title, the *Blueprint* focuses on ADS.

reporting in the form of an algorithmic impact assessment, including disparity testing results and mitigation information, should be performed and made public whenever possible to confirm these protections.

- *Data Privacy*: [. . .] Designers, developers, and deployers of automated systems should seek your permission and respect your decisions regarding collection, use, access, transfer, and deletion of your data in appropriate ways and to the greatest extent possible; where not possible, alternative privacy by design safeguards should be used [. . .] Enhanced protections and restrictions for data and inferences related to sensitive domains, including health, work, education, criminal justice, and finance, and for data pertaining to youth should put you first. In sensitive domains, your data and related inferences should only be used for necessary functions, and you should be protected by ethical review and use prohibitions. [. . .]
- *Notice and Explanation*: You should know that an automated system is being used and understand how and why it contributes to outcomes that impact you. Designers, developers, and deployers of automated systems should provide generally accessible plain language documentation including clear descriptions of the overall system functioning and the role automation plays, notice that such systems are in use, the individual or organization responsible for the system, and explanations of outcomes that are clear, timely, and accessible. Such notice should be kept up-to-date and people impacted by the system should be notified of significant use case or key functionality changes. You should know how and why an outcome impacting you was determined by an automated system, including when the automated system is not the sole input determining the outcome. [. . .]
- *Human Alternatives, Consideration, and Fallback*: You should be able to opt out from automated systems in favor of a human alternative, where appropriate. Appropriateness should be determined based on reasonable expectations in a given context and with a focus on ensuring broad accessibility and protecting the public from especially harmful impacts. [. . .]

The Legislature, via SCR 17 (Dodd, 2023), adopted these principles. These principles inform efforts to regulate ADS, including in this bill and other related bills such as AB 1018 (Bauer-Kahan, 2025).

3) What this bill would do. This bill would establish a comprehensive set of requirements governing the use of ADS in the workplace by requiring human oversight and independent corroboration when employers use ADS for decisions that impact workers' livelihoods. Key aspects of the bill follow.

Prohibited uses of ADS. The bill prohibits an employer from using an ADS that does any of the following:

- Prevents compliance with or results in a violation of labor, occupational and safety, employment, or civil rights laws or regulations.

- Infers certain protected characteristics about a worker.
- Conducts predictive behavior analysis on a worker with the intention of using the analysis to make an employment-related decision.
- Predicts or takes adverse action against a worker for exercising their legal rights.

Right to access data. The bill requires an employer to allow a worker to access their own worker data collected or used by an ADS.

Provisions governing discipline, termination, or deactivation decisions. With respect to decisions involving discipline, termination, or deactivation, the bill requires employers to:

- Not rely solely on the ADS when making such decisions. Employers must use human reviewers to conduct their own investigation and compile corroborating or supporting information for the decision.
- Provide the affected worker with a written post-use notice at the time the employer informs the worker of the decision.
- When responding to a data access request from a worker, provide to the worker with a written, plain language document using a simple and easy-to-use method that is accessible away from the workplace and that contains specified information.

Anti-retaliation provisions and enforcement. The bill establishes anti-retaliation protections for workers who attempt to use their rights provided under the bill, as specified. The bill grants enforcement authority to the Labor Commissioner or public prosecutors. An employer is subject to \$500 per violation and claimants may seek damages, including punitive damages, injunctive relief, and fees and costs.

4) SB 7 comparison. SB 947 shares much of its language with SB 7 (McNerney, 2025). That bill was ultimately vetoed by the governor, who shared the following veto message:

I share the author’s concern that in certain cases unregulated use of ADS by employers can be harmful to workers. However, rather than addressing the specific ways employers misuse this technology, the bill imposes unfocused notification requirements on any business using even the most innocuous tools. This proposed solution fails to directly address incidents of misuse.

Moreover, this measure proposes overly broad restrictions on how employers may use ADS tools. For example, prohibiting an employer from using customer ratings as the primary input data for an ADS takes away a potentially valuable tool for rewarding high-performing employees. To the extent that customer reviews are unfairly or inappropriately used to make decisions about a worker, legislation should address those specific scenarios rather than ban this practice altogether.¹⁵

¹⁵ Gavin Newsom, *SB 7 Veto*, (Oct. 13, 2025), <https://www.gov.ca.gov/wp-content/uploads/2025/10/SB-7-Veto.pdf>.

Acknowledging these concerns, SB 947 is significantly narrower compared to SB 7. In contrast to its predecessor, SB 947:

- Does not include pre-use notice requirements for employers.
- Does not require employers to notify job applicants about their use of an ADS.
- Does not prohibit the use of customer ratings as the only or primary input source for making employment-related decisions, directly responding to a concern raised in the governor's veto message.
- Requires specified information be provided to a worker who makes a data access request.
- Requires that if an employer cannot corroborate an ADS output, or a human reviewer has concluded that an ADS output is inaccurate, incomplete, or misleading, the employer shall not use the ADS output to discipline, terminate, or deactivate a worker.
- Explicitly includes in its definition of "employer" the California State University, community college districts, and the University of California upon agreement by the regents.

5) Committee amendments. Both SB 947 and SB 7 contain the following exemption:

This part does not prohibit any employer from complying with regulatory or contractual requirements in the provision of products or services to the federal government.

According to the author, the goal of this language is to exempt ADS with national security implications – for example, systems that make critical employment decisions with respect to the operation of aircraft in the national airspace. As written, however, the exemption is broad enough to encompass a wide range of non-national security activities. The author has agreed to replace this exemption with language adapted from AB 1883 (Bryan, 2026), which contains a narrower exemption addressing the same issue:

(a) This part shall not apply to the use of [an automated decision system] to the extent that use of the [system] is required by, or reasonably necessary to comply with, a federal statute, federal regulation, or binding federal contract relating to any of the following:

(1) The development of aircraft for use in the national airspace.

(2) The development of products or services for national security, military, space, or defense purposes.

(b) The exemption in this section applies only to the operations covered by a federal statute, federal regulation, or binding federal contract outlined in paragraphs (1) and (2) of subdivision (a).

ARGUMENTS IN SUPPORT:

A coalition of labor, tech industry oversight, and privacy organizations including California Federation of Labor Unions, Electronic Frontier Foundation, and TechEquity Action writes the following in support:

Employers' drive to increase worker productivity and lower costs is nothing new. What is different today is the power, speed and scale, and secrecy of workplace technologies such as automated decision-making systems (ADS), often used for algorithmic or automated management. Employer use of ADS in the workplace is widespread. A 2025 ResumeBuilder survey of employers found that 60% of managers use AI systems to make critical decisions about workers, including raises, terminations, and layoffs. Of those, 20% let AI make final decisions without any human input or oversight. "Robot bosses" can do serious harm to workers. Endlessly increasing efficiency through eliminating routine tasks and increasing work speeds can lead to fatigue, burn-out, excessive injuries, and other harm, as seen in Amazon warehouses that relied on algorithms to implement speed quotas which forced workers to skip bathroom breaks and skirt safety measures.

...

SB 947 allows employers to use ADS, but requires human oversight and independent corroboration of firing, disciplinary, and deactivation decisions. Additionally, SB 947 prohibits employers from using an ADS to make employment related decisions on workers based on predictive analysis, thus protecting workers from being profiled and disciplined based on actions they have not committed. SB 947 will prevent the outsourcing of decisions that impact workers' lives to machines. It allows for the use of technology and tools to make workplaces more productive and efficient but ensures human oversight to prevent abuse and mistakes.

ARGUMENTS IN OPPOSITION:

A coalition of business trade associations including California Chamber of Commerce, TechNet, CalBroadband, and the Civil Justice Association of California writes the following in opposition:

The bill broadly targets businesses of all sizes, across every industry, and regulates even low-risk applications of automated decision systems (ADS). Significantly, we are disappointed that the bill undoes many amendments taken in SB 7, its predecessor bill from 2025 that was vetoed, marking a step backwards in progress on this issue. SB 947 will drive up costs for consumers and employers because it would impose significant compliance burdens and any misstep would lead to costly litigation for even the smallest of employers. While we appreciate concerns over employees being unfairly disciplined or terminated solely based on automated tools, SB 947 is not tailored to those scenarios and does not consider the benefits of ADS technology. Unfortunately, we believe SB 947 will have an undesired chilling effect on technology and make it that much harder to develop the very tools that can help combat bias in decision making.

REGISTERED SUPPORT / OPPOSITION:

Support

California Faculty Association
California Professional Firefighters

California School Employees Association
California Teachers Association
California Working Families Party

Opposition

American Petroleum and Convenience Store Association Apca
American Staffing Association
Associated Equipment Distributors
Associated General Contractors of California
Associated General Contractors San Diego
Associated General Contractors, California
Association of California Healthcare Districts (ACHD)
Association of California School Administrators
Bay Area Council
California Apartment Association
California Association of Recreation & Park Districts
California Association of School Business Officials (CASBO)
California Association of Winegrape Growers
California Broadband & Video Association
California Chamber of Commerce
California Coalition on Workers Compensation
California County Superintendents
California Farm Bureau
California Fuels and Convenience Alliance
California Grocers Association
California Hospital Association
California Landscape Contractor's Association
California Landscape Contractors Association
California League of Food Producers
California Manufacturers & Technology Association (CMTA)
California Manufacturers and Technology Association
California Restaurant Association
California Retailers Association
California Retailers Association
California Special Districts Association
California Staffing and Recruiting Association
California State Association of Counties (CSAC)
California Trucking Association
California's Credit Unions
Chamber of Progress
Cinema Association of California
Civil Justice Association of California (CJAC)
County of Fresno
County of Kern
Family Business Association of California
Fresno County Board of Supervisors
Leadingage California
League of California Cities
Lyft

National Association of Mutual Insurance Companies
Personal Insurance Federation of California
Public Risk Innovation, Solutions, and Management (PRISM)
Rural County Representatives of California (RCRC)
Self Storage Association
Shrm California
Society of Human Resources Management
Technet
Uber
Urban Counties of California (UCC)
Western Growers Association

Oppose Unless Amended

Business Software Alliance
Gaslamp Quarter Association
Insights Association
Long Beach Restaurant Association
Los Angeles County Business Federation (BIZFED) (UNREG)

Analysis Prepared by: Slater Sharp / P. & C.P. / (916) 319-2200