

## ASSEMBLY THIRD READING

AB 2575 (Ortega)

As Amended April 23, 2026

Majority vote

**SUMMARY**

Requires a health facility, clinic, physician's office, or office of a group practice that uses or deploys a clinical decision support system (CDSS), as defined, for patient care to disclose required information to any licensed health care professional or other person using a CDSS or viewing outputs from a CDSS. Requires the disclosure to include, among other things, a notice that a worker providing direct patient care is authorized to override the output of a CDSS if, in the judgment of the worker acting in their scope of practice, an override is appropriate for the patient, or as necessary to meet the applicable standard of care or to comply with the law. Prohibits an employer from retaliating or discriminating against a worker providing patient care based solely on the worker's override of, or reliance on, the output of a CDSS. Authorizes a worker who is subject to retaliation or discrimination in violation of this bill to file a complaint with the Labor Commissioner against an employer. Prohibits a defendant who developed, modified, selected, or deployed a CDSS that is alleged to have harmed a plaintiff from asserting a defense that the failure of a licensed health care professional or other health care worker to override an output of the CDSS is a superseding cause severing the defendant's liability for the alleged harm.

**COMMENTS**

*Artificial Intelligence (AI).* AI is the mimicking of human intelligence by artificial systems. AI uses algorithms, or 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; unlike other computer functions, however, AI is able to accomplish tasks that are normally performed by humans. Most modern AI tools are created through a process known as "machine learning." Machine learning involves techniques that enable AI tools to learn the relationship between inputs and outputs without being explicitly programmed. The next step is "training," the process of exposing a naïve AI to data. The algorithm that an AI develops during training is known as its "model." Models trained on small, specific datasets in order to make recommendations and predictions are referred to as "predictive AI." This differentiates them from "generative AI" (GenAI), which are trained on massive datasets in order to produce detailed text, images, audio, and video. Because it can process a range of data sources and create novel outputs, and because it can convincingly mimic human capabilities and convincingly generate perfectly worded nonsense, GenAI poses unique opportunities and challenges.

*Administrative and clinical-adjacent uses of AI applications in health care.* Hospitals, clinics, physician groups, and health plans are leveraging GenAI to automate a wide range of routine back-office tasks as well as those tasks that provide administrative support for clinical work.

For instance, electronic health record (EHR) systems are being equipped with GenAI functionality that allows health care providers to automatically generate billing codes, improving accuracy and completeness by checking for errors, omissions, and compliance with current requirements. At the same time, health plans and insurers are using AI on the other end of the

transaction to automate and streamline multiple functions, including processing claims and evaluating prior authorization requests. According to the *Wall Street Journal*, UnitedHealth Group said it now has a thousand AI applications in production, even as a class action lawsuit is advancing through the courts accusing the insurer of using AI algorithms instead of medical professionals to illegally deny Medicare Advantage claims. Other types of administrative tasks, such as appointment scheduling and other routine, non-clinical communication have significant potential to be automated. Researchers also believe AI can assist in generation of quality metrics, which are important for measuring health system performance but often rely on data that must be manually extracted from EHRs. Similar to the Layer Health project mentioned above, a pilot study found that large language models (LLMs) could perform accurate extractions of specific data from these patient records for use in calculating complex quality measures.

Finally, although they state they are not seeking to diagnose or prescribe, a company called Hippocratic AI seeks to usher in a world of "healthcare abundance" through the development and deployment of "health care AI agents" who interact with patients on behalf of health care providers. The company describes these voicebot agents as being designed to live within a liminal space in the health care system: accomplishing a number of common tasks that are often performed by medical assistants or clinical staff like nurses, such as case management, appointment preparation, follow-up from procedures, but that may not require a license.

*Diagnostics, Treatment and CDSS.* In addition to the clinical-adjacent applications discussed above, clinical applications of GenAI technology are advancing rapidly and substantially. Google's MedPaLM-2 LLM achieved expert-level scores on the United States Medical Licensing Examination, with physicians preferring AI answers to those from other physicians on eight of nine clinical axes. GenAI can now synthesize, augment, and interpret heterogeneous complex images across various modalities, such as X-rays, MRI, and CT scans. AI algorithms can also assist in diagnosing dental health conditions through image analysis and data interpretation.

In recent years, minimally invasive surgical techniques such as laparoscopic surgery and robotic surgery have become increasingly prevalent. The ultimate goal of robotic surgery development is the creation of fully autonomous AI-powered surgical instruments. Data from the 2023 American Hospital Association Annual Survey Information Technology Supplement reflect 65% of United States hospitals used predictive models, and 79% of those used models from their EHR developer. Hospitals use AI and predictive models to predict health trajectories or risks for inpatients, identify high-risk outpatients to inform follow-up care, monitor health, and recommend treatments. Clinical decision support systems are also being deployed and designed to aid physicians in diagnosing, managing, and treating patients in outpatient settings.

*Racial, ethnic, and gender bias.* The performance of an AI is directly impacted by the quality, quantity, and relevance of the data used to train it. If the data used to train the AI is biased, the tool's outputs will be similarly biased and the results can be inaccurate when applied to populations not reflected in the training data.

When automated decision systems are deployed in healthcare, biased historical data can lead to patients being recommended substandard care on the basis of their race or ethnicity. In 2007, an automated decision system was developed to help doctors estimate whether it was safe for people who had delivered previous children through cesarean section to deliver subsequent children vaginally. The system considered relevant factors as it made its decision, such as the woman's age, her reason for the previous cesarean, and how long ago the cesarean had been performed.

However, a 2017 study found that the system was biased; it predicted Black and Latino people were less likely to have a successful vaginal birth after a cesarean than similar non-Hispanic white women. As a result, doctors performed more cesareans on Black and Latino people than on white people. Such discrepancies can potentially perpetuate historical biases and lead to worse health outcomes. The University of California (UC), San Francisco also reported bias in an algorithm used to identify potential appointment no-shows to facilitate double-booking for that appointment. The program was confirmed to result in low-resourced and marginalized populations being double-booked more often than others, reflecting underlying structural inequalities and highlighting how these tools, if not studied and corrected for bias, that can create feedback loops that worsen discrimination.

*Cognitive biases and cognitive burden.* Bias exhibited by an AI model based on underlying training data is not the only bias that may influence how an AI system works when deployed. Its effectiveness can also be impacted by predictable patterns of human error called cognitive biases. Reviewing an AI system's output for errors or omissions is a substantively different cognitive task than generating a clinical note or medical advice, and the use of AI systems raise questions about how cognitive bias evoked by AI assistance with clinical tasks might affect clinical judgement or practice in ways that are difficult to understand, predict, and measure.

Research shows automation bias, or placing undue confidence in and over-relying on automated outputs, is a problem in many fields. Automation-induced complacency, or insufficient monitoring of automation output, is also a concern. Over time, these biases can lead to people being less likely to catch errors or to disagree with what was written. There are many factors that can exacerbate the potential danger posed by automation bias in CDSS, including, for instance, if an AI model's process to arrive a given output lacks transparency or is not explainable, if the model is implemented with inadequate training of end users, or if a clinician is under significant time pressure or cognitive burden that limits their practical ability to systematically assess and effectively integrate the additional information provided by an AI system with their clinical knowledge and experience.

Although there are many examples of promising AI applications for improving clinical decision-making, a 2023 experimental study demonstrated some concerning results. It tested the efficacy of AI models designed to assist clinicians in diagnosing chronic obstructive pulmonary disease, pneumonia, or heart failure from a radiograph. Although assistance from a carefully designed AI model slightly improved clinicians' accuracy in diagnosis as compared the clinicians who received no assistance (76-78% versus 73%), in cases where clinicians were provided AI support using a systematically biased model, diagnostic accuracy dropped substantially to 62%. In other words, receiving support from a bad AI system actually made clinicians significantly worse at diagnosing conditions than simply relying on their own clinical judgement.

This study showed that having a "clinician-in-the-loop" overseeing the AI does not overcome the challenges of poor-performing AI systems, regardless of whether the clinicians are given information explaining how the AI arrived at its output. A commentary on the study, "Automation Bias and Assistive AI Risk of Harm From AI-Driven Clinical Decision Support," points to automation bias as the culprit for these troubling outcomes.

*Safety and effectiveness.* In some cases, an AI model's accurate predictions may nevertheless lead to bad decisions. In one example, a hospital trained AI models on a dataset of 15,000 pneumonia patients in order to develop a model that could identify which pneumonia patients were at the

greatest risk, in order to triage new patients. During testing, it was discovered that one of the most accurate models recommended outpatient status for asthmatics. This is a life-threateningly dangerous error based on an accurate statistical correlation, namely, asthmatics are less likely to die from pneumonia than the general population precisely because asthma is such a serious risk factor that asthmatics automatically get elevated care.

The "generative" aspect of GenAI models mean they may produce incorrect outputs, including "confabulations" and "hallucinations"—confidently stated but erroneous content that may mislead or deceive users. GenAI's well-reported challenges with factual correctness are particularly problematic in health care, where inaccuracies can cause serious harm. Recent problems include incorrect differential diagnosis and invalid scientific citations.

*Health AI nutrition labels.* The Coalition for Health AI (CHAI) has developed applied model cards, like "nutrition labels," to be used by health systems to create a more standard way to present foundational information about AI solutions. CHAI says that applied model cards are one example of many independent quality assurance resources that can empower health systems to assess and analyze innovative health AI solutions. Their use may simplify the process of sharing validation and testing results, as well as accelerate the development of models through trusted, independent entities. Health AI Partnership notes that model cards serve as standardized documentation capturing essential details about AI systems, including their purpose, performance metrics, ethical considerations, and limitations. By providing a structured inventory with comprehensive information, the initiative aims to foster trust among stakeholders, including developers, users, and regulators. An article titled "Could transparent model cards with layered accessible information drive trust and safety in health AI?" published in *NPJ Digital Medicine*, says that, with respect to avoiding duplication and ensuring integration with regulatory approaches, it is particularly relevant to consider the international requirements for labelling and for the 'Instructions for Use' (IFU) of medical devices, as internationally AI-models performing medical purposes are regulated as medical devices. Existing regulations require that devices are prepared with both a label and an IFU (with a small number of exceptions for software devices if it can be demonstrated that they have use explanations through their interfaces).

### **According to the Author**

Health care workers are facing new challenges as AI is integrated into their workplaces. They are pressured by employers to defer to AI systems that may be opaque, erroneous, or systemically biased. They face an added risk of professional and legal blame when they follow algorithmic recommendations that fail. The author argues this bill preserves healthcare workers' ability to follow their professional judgment by prohibiting employer retaliation when a worker overrides or follows a recommendation, and requires transparency for AI tools so that patients and providers understand how they work and what the risks are. The author concludes that, overall this bill requires that AI tools are used to support clinical judgement, not replace it, ensuring that human expertise and patient safety remain the focus of California's health care system.

### **Arguments in Support**

The California Nurses Association (CNA) is a co-sponsor of this bill and states in support that as hospitals rapidly adopt AI-driven technologies that influence clinical decisions and working conditions, California must ensure that these tools support safe patient care rather than undermine professional standards or patient protections. California currently lacks clear guardrails governing the use of AI in health care settings. CNA states that hospitals and clinics now use AI tools in electronic health records, CDSS, remote monitoring platforms, staffing

management software, and administrative workflows. These systems generate patient acuity scores, treatment recommendations, discharge planning prompts, insurance determinations, and nurse workload assignments. As hospitals and other health care entities expand their adoption of these tools, they increasingly shape both clinical decision-making and working conditions. CNA argues that state law provides few standards to ensure transparency, protect professional judgment, or establish accountability when these technologies contribute to harmful outcomes. CNA contends that evidence increasingly shows that many AI tools used in health care raise serious safety and accuracy concerns, and lack of transparency prevents clinicians and patients from understanding how AI systems influence care decisions. CNA concludes that this bill establishes clear guardrails by requiring transparency regarding when health care entities use AI in patient care, protecting clinicians' professional judgement and ability to override AI-driven decisions and ensuring accountability for developers and deployers when AI systems cause harm.

The California Federation of Labor Unions (Labor Fed) is a co-sponsor of this bill and states in support that an estimated 65% of U.S. hospitals are already using AI tools, most commonly to predict inpatient health trajectories. In addition, hospitals and clinics use AI for electronic health records, staffing systems, CDSS, remote monitoring platforms, and administrative decision-making. These tools can influence patient acuity scores, treatment recommendations, insurance determinations, discharge planning, and nurse workloads. Labor Fed argues that despite their widespread use, patients and health care workers often receive little information about when these tools are used, how they function, what data they rely on, or what risks and limitations they carry. This lack of transparency, combined with the expanding role of AI in clinical and workplace decisions, has significant implications for patient safety, professional practice, and accountability in health care. Labor Fed concludes that this bill directly addresses the lack of transparency, worker protections, and accountability for the use of AI in health care by requiring pre-notification of the use of AI tools in patient care, protects clinician's professional judgement and ability to override AI-driven decisions, and ensures accountability for developers and deployers when AI systems cause harm.

Health Access California supports this bill stating there are very few laws regulating the use of AI, especially in healthcare. As more industries have incorporated AI into their operations, growing data, litigation and research findings have highlighted the risks of deploying these tools without appropriate guardrails. While AI models can support the delivery of health care administratively, there are numerous ways that algorithms, often improperly trained, relying upon biased datasets and poorly implemented, may negatively impact consumers and their health care. Health Access California argues that significant concerns remain on how AI may impact healthcare equity, access, and quality. In a recent study, data showed that while 65% of health facilities have incorporated AI models into their systems, only 9% evaluated their models for biases by testing them in real world scenarios. The lack of evaluation is troubling for consumers, as the lack of proper implementation and biased datasets put Californians at risk of further health disparities. Additionally, research has found that AI models further perpetuate racial biases in healthcare.

### **Arguments in Opposition**

A coalition including ADVAMED, America's Physician Groups, ATA Action, Biocom, CalChamber, the California Hospital Association, California Association of Health Facilities, California Association of Health Plans, California Life Sciences, the California Medical Association, California Radiological Society, California Society of Pathologists, CPCA

Advocates, Kaiser Permanente, and Ochin, states in opposition that AI has the potential to improve nearly every aspect of health care, including quality, patient experience, and affordability. At the same time, the health care field does face unique considerations when using AI. Health care leaders and policymakers must understand and balance the potential benefits and risks to ensure that AI is used safely, effectively, and equitably. They argue that the framework created by this bill is overly broad, impossible to implement, and likely to hinder beneficial patient outcomes. It would affect existing AI tools and systems that have been used successfully in health care for many years — from basic medication safety alerts to well-established clinical scoring tools — by subjecting them to onerous requirements that negate their tested and proven benefits. Along with overly restrictive disclosure, liability, and labor provisions, these requirements would create enormous waste in the system and reduce the time clinicians have to spend with patients — without any clear corresponding benefits. The opponents argue that it would also hinder technological advancement and, troublingly, exacerbate existing health disparities by impeding the ability of health care providers, particularly those serving vulnerable communities, to leverage AI tools to improve patient outcomes and the health of the populations they serve.

TechNet, the Civil Justice Association of California, and Connected Health Initiative also oppose this bill stating that the definition of a covered tool could encompass a wide range of technologies beyond traditional CDSS tools, including notetaking software, administrative triage tools, scheduling systems, and other routine technologies used in modern health care delivery. They argue that by applying extensive disclosure and compliance requirements across this broad category, this bill risks sweeping in low-risk operational tools that do not meaningfully implicate clinical decision-making. A more targeted, risk-based approach would better align regulatory obligations with actual patient safety considerations. They also argue that requiring plain language descriptions of complex validation processes or training methodologies in medical records would create an administrative burden without improving patient understanding or outcomes.

## FISCAL COMMENTS

According to the Assembly Appropriations Committee, costs of an unknown but potentially significant amount to the California Department of Public Health to adopt regulations and add enforcement of this bill's requirements into its facility surveys, among other activities (Licensing and Certification Fund).

The Department of Industrial Relations expects costs to the Labor Commissioner to exceed \$300,000 annually for staff workload, including enforcement deputies, support, and supervisory staff, to enforce the retaliation prohibition (Labor Enforcement and Compliance Fund).

The University of California (UC) estimates costs of approximately \$30 million per year, ongoing (General Fund). UC states compliance would require a cross-departmental effort from enterprise information technology, legal, human resources, and clinical operations at each UC Health location and across the system. Costs would principally arise from the annual requirements to audit, track, and extract highly specific proprietary data from third-party vendors for every clinical decision system.

Costs of an unknown but potentially significant amount to the Department of Justice (DOJ) to bring enforcement actions as authorized by this bill. Actual costs will depend on the number of enforcement actions pursued by DOJ and the amount of additional work created by each action,

but costs may be in the hundreds of thousands of dollars annually (Unfair Competition Law Fund).

Cost pressures of an unknown but potentially significant amount to the courts to adjudicate any additional filings (Trial Court Trust Fund, General Fund). It generally costs approximately \$1,000 to operate a courtroom for one hour. Although courts are not funded based on workload, increased pressure on the Trial Court Trust Fund may create a demand for increased funding for courts from the General Fund. The state budget provides annual General Fund backfills to the Trial Court Trust Fund to offset revenue reductions, totaling approximately \$117.3 million in 2025-26.

The Legislative Analyst's Office recently warned of General Fund structural deficits of around \$35 billion per year in the 2027-28 fiscal year and ongoing.

## VOTES

### **ASM HEALTH: 11-1-4**

**YES:** Bonta, Addis, Aguiar-Curry, Ahrens, Caloza, Mark González, Patel, Rogers, Schiavo, Sharp-Collins, Stefani

**NO:** Johnson

**ABS, ABST OR NV:** Chen, Carrillo, Patterson, Sanchez

### **ASM LABOR AND EMPLOYMENT: 5-2-0**

**YES:** Ortega, Elhawary, Kalra, Lee, Ward

**NO:** Lackey, Chen

### **ASM PRIVACY AND CONSUMER PROTECTION: 9-4-2**

**YES:** Bauer-Kahan, Bryan, Lowenthal, McKinnor, Ortega, Pellerin, Ward, Wicks, Wilson

**NO:** Macedo, DeMaio, Hoover, Irwin

**ABS, ABST OR NV:** Patterson, Petrie-Norris

### **ASM APPROPRIATIONS: 10-4-1**

**YES:** Wicks, Aguiar-Curry, Calderon, Caloza, Fong, Mark González, Krell, Pellerin, Sharp-Collins, Solache

**NO:** Hoover, Dixon, Ta, Tangipa

**ABS, ABST OR NV:** Pacheco

## UPDATED

VERSION: April 23, 2026

CONSULTANT: Logan Hess / HEALTH / (916) 319-2097

FN: 0002980