

## ASSEMBLY THIRD READING

AB 1979 (Bonta)

As Amended April 23, 2026

Majority vote

**SUMMARY**

Requires a health facility, clinic, physician's office, or office of a group practice to ensure that no clinical decision, as specified, is based solely on the output of a clinical decision support system (CDSS), as defined, and that a licensed health care professional exercises independent professional judgment when reviewing and approving a clinical decision that is based on the output of a CDSS. Specifies that this bill does not prohibit the use of artificial intelligence (AI) for documentation and communication that does not involve the application of professional judgment, including automated messages to inform patients of updates to their health records. Clarifies, for the purposes of the Confidentiality of Medical Information Act (CMIA) that a business that offers a healthcare chatbot, as defined, to a consumer is a provider of health care subject to the requirements of the CMIA.

**COMMENTS**

*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. CDSS 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.

*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.

*Direct to consumer AI tools.* Direct to consumer (DTC) health-related AI tools are rolling out at a rapid pace. This year, major AI developers at OpenAI, Anthropic, and Microsoft have announced ChatGPT Health, Claude for Health Care, and Copilot Health, respectively. The CMIA prohibits a provider of health care, among others, from intentionally sharing, selling, using for marketing, or otherwise using any medical information, for any purpose not necessary to provide health care services to a patient, except as provided. The CMIA makes a business that offers software or hardware to consumers, including a mobile application designed to maintain medical information to make the information available to an individual or a provider of health care, for purposes of allowing the individual to manage their information or for the diagnosis, treatment, or management of a medical condition of the individual, a provider of health care subject to the requirements of the CMIA. In announcements of these DTC health AI apps, developers have made clear that their tools do not diagnose or treat patients, or claim that consumer facing elements are intended to assist with preparing for doctor's appointments. Given this focus, there may be room for clarity for what it means to manage an individual's information.

In January 2025, the California Department of Justice released legal advisories regarding the use of AI in health care, which cite these provisions of the CMIA:

"California state medical privacy laws provide protections that are, in some cases, more stringent than federal health privacy laws like HIPAA. The CMIA and the Information Practices Act govern use and disclosure of Californians' medical information. Covered entities must preserve confidentiality of patients' medical information and ensure that patients have access to that information. Sensitive information, including mental and behavioral healthcare and reproductive and sexual healthcare (e.g., abortion and gender affirming care), receive heightened protections. Medical privacy laws apply to governmental healthcare agencies, medical providers, and insurance plans, as well as businesses that offer software or hardware to consumers for the purposes of managing medical information, diagnosis or treatment, or management of medical conditions, via mobile applications or other related devices."

### **According to the Author**

AI is rapidly integrating into our health care system and reshaping our personal experience with health care. While this technology holds a lot of promise, there is no question that without careful consideration of the potential perpetuation of biases, risks to patient safety, and challenges of clinical workers knowing what to question and what to trust, the deployment of AI in health care can do more harm than good. The author notes that a 2023 study found that, while carefully crafted AI could slightly improve diagnostic accuracy for certain disorders, in cases where clinicians were provided AI support using a systematically biased model, diagnostic accuracy dropped substantially to 62% (from 73%). This also demonstrates that having a human-in-the-loop is not a panacea for all the challenges that AI can present. Providing health care requires compassion, empathy, and real-world judgment that cannot be captured in patterns and algorithms. Technology should assist human clinicians, not replace them. As AI deploys into health care settings, it is also reaching consumers directly through applications like Copilot and ChatGPT offering to connect directly to personal medical records. The author concludes that voluntary commitments to protect this sensitive information are not enough, we must ensure any entity accessing medical records for managing health is abiding by the law.

### **Arguments in Support**

The California Nurses Association (CNA) is the sponsor of this bill and states in support that it would protect against the threats presented by the rapid expansion of AI in health care to standards of care in clinical practice and patient privacy. CNA argues that caring for patients is deeply human work, requiring professional judgment, clear accountability, and strong protection for patients' sensitive medical information. CNA contends this bill rejects the false promise of a human-in-the-loop model where AI performs the substantive clinical decision making and the clinician is left to supervise or sign off. CNA argues that this bill preserves a human in command standard, where the licensed health care professional remains the actual decision maker in health care clinical practice. In plain terms, if California law requires a licensed clinician to perform a health care activity, an AI tool cannot be used in that clinician's place. CNA concludes this bill also ensures that AI health applications remain subject to California's medical confidentiality laws, so patients do not lose privacy protections simply because their information passes through a digital platform.

Consumer Watchdog supports this bill and states that by ensuring that health care entities and AI developers cannot use AI to replace licensed clinical judgment or bypass health privacy protections it safeguards patient safety, medical privacy, and the integrity of health care practice. AI is rapidly entering the places patients receive care and consumer-facing health applications, often unbeknownst to patients. Too often the existing guardrails that are meant to protect patients are being blurred. Patient safety demands clear guardrails on the use of AI technologies. Consumer Watchdog argues this bill would protect patients by prohibiting health care entities from using AI to perform functions that require the clinical judgment of licensed health care professionals. Allowing untested and unaccountable AI technologies to perform or guide these functions risks patient harm and undermines public trust in the health care system.

The Board of Behavioral Sciences (BBS) also supports this bill stating that the BBS believes it is essential to provide guardrails to ensure that emerging technologies enhance, rather than undermine, the delivery of safe and effective care. By reinforcing the requirement for independent professional judgment and prohibiting the inappropriate delegation of licensed functions, this bill supports patient safety and upholds the integrity of clinical practice.

### **Arguments in Opposition**

A coalition including ataAction, America's Physician Groups, the California Association of Health Plans, California Radiological Society, Civil Justice Association of California, Epic, and Kaiser Permanente opposes this bill stating that it will significantly hinder the ability of the entire health care system to provide safe and effective care for Californians and still imposes rigid statutory constraints on rapidly evolving tools that are already deeply embedded across clinical operations. The coalition argues that this bill would require inventorying and legally re-adjudicating a wide range of existing tools across EHR workflows, patient education, care navigation, imaging, medical devices, and enterprise analytics. They state this bill would cause them to redesign or disable many current capabilities, create new compliance, audit, and documentation processes to demonstrate "independent professional judgment" and absorb downstream productivity losses.

The California Chamber of Commerce (CalChamber) also opposes this bill citing concerns with the ambiguity of which outputs would be considered a "clinical decision." CalChamber states that it strongly believes that AI should support — not replace — licensed health care

professionals, who retain full decision-making authority at all times. Unfortunately, this bill would impose sweeping oversight requirements on not just AI systems, but also a wide range of patient-centered technologies. CalChamber argues that this bill could disrupt or create unintended challenges to existing, widely used, and beneficial technologies embedded in electronic health record platforms, where AI supports functions such as predictive alerts, clinical summarization, and documentation. These capabilities are deeply integrated into clinical workflows and, without clarifying a "clinical decision," could compromise workflow efficiency and interoperability and cannot be removed without compromising data quality, workflow efficiency, and interoperability.

## **FISCAL COMMENTS**

According to the Assembly Appropriations Committee, costs of an unknown amount to the California Department of Public Health (DPH) to ensure health facilities, clinics, and physician offices comply with the provisions of this bill. If DPH adopts regulations, one-time costs could be in the hundreds of thousands of dollars (Licensing and Certification Fund).

The Department of Consumer Affairs reports most healing arts boards anticipate an increase in workload due to a higher volume of complaints resulting from this bill. However, the boards are unable to estimate the increase in complaint volume or related enforcement costs due to the lack of data on how frequently AI violations occur (various special funds).

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). Actual costs will depend on the number of cases filed and the amount of court time needed to resolve each case. 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.

## **VOTES**

### **ASM HEALTH: 12-3-1**

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

**NO:** Johnson, Patterson, Sanchez

**ABS, ABST OR NV:** Chen

### **ASM PRIVACY AND CONSUMER PROTECTION: 11-4-0**

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

**NO:** Macedo, DeMaio, Hoover, Patterson

**ASM APPROPRIATIONS: 11-4-0**

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

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

**UPDATED**

VERSION: April 23, 2026

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

FN: 0002898