

Date of Hearing: April 21, 2026

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

AB 2575 (Ortega) – As Amended April 9, 2026

PROPOSED AMENDMENTS

SUBJECT: Health care services: artificial intelligence

SYNOPSIS

The rapid advancement of generative artificial intelligence (GenAI) over the past few years has proven transformative for a wide variety of industries, including health care. These tools hold enormous promise: they can help doctors and nurses sort through vast medical literatures, support earlier and more accurate diagnoses, lessen administrative burdens, and expand access to care in remote or underserved areas through digital interfaces. At the same time, the known limitations of AI can prove especially consequential in high-risk clinical settings. GenAI systems can “hallucinate,” producing confidently-presented but incorrect information – if a health care AI tool hallucinates, it can lead to misdiagnosis, inappropriate treatment, or delayed care. AI systems trained on incomplete or unrepresentative data may produce biased outputs, perpetuating existing disparities in health outcomes. Many health professionals argue – including in opposition letters submitted in response to this measure – that AI is best utilized as a tool to augment and support licensed health care professionals.

This bill would require specified information about AI healthcare tools be disclosed to health workers, would prohibit the replacement of professional judgement by AI tools in healthcare, and would prohibit retaliation or discrimination against a worker who overrides, or complies with “technology,” as defined. The bill would also prohibit a defendant who developed, modified, or used an artificial intelligence tool from asserting as a defense that a failure of a health care worker to override an output of the tool is a superseding cause severing the defendant’s liability for the alleged harm. Committee amendments would narrow disclosure requirements and recast several of the bill’s declarations of public policy.

The bill is co-sponsored by the California Nurses Association and the Labor Federation. It is supported by a variety of labor groups, along with Health Access California and Oakland Privacy. The bill is opposed by a coalition of associations representing physicians, hospitals and health systems, health plans, life sciences, and other health care stakeholders, including California Medical Association, California Hospital Association, and California Chamber of Commerce. The bill was previously heard in the Health Committee, where it passed in an 11-1 vote, and the Labor and Employment Committee, where it passed in a 5-2 vote.

EXISTING LAW:

1) 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. (Health & Saf. Code § 1339.75.)

- 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. (Gov. Code § 11546.45.5.)
 - c. “Generative artificial intelligence” means artificial intelligence that can generate derived synthetic content, including images, videos, audio, text, and other digital content. (Health & Saf. Code § 1339.75.)
 - d. “Health care provider” means a person licensed or certified pursuant to Division 2 (commencing with Section 500) of the Business and Professions Code. (Health & Saf. Code § 1339.75.)
 - e. “Physician’s office” means an office of a physician in solo practice. (Health & Saf. Code § 1339.75.)
 - f. “Office of a group practice” means an office or offices in which two or more physicians are legally organized as a partnership, professional corporation, or licensed not-for-profit corporation. (Health & Saf. Code § 1339.75.)
- 2) Provides for the licensure and regulation of clinics and health facilities by the State Department of Public Health. (Health & Saf. Code § 1200 *et seq.*, Health & Saf. Code § 1250 *et seq.*)
 - 3) Establishes “healing arts” boards for the licensure and regulation of various healing arts professions and vocations by boards established within the Department of Consumer Affairs, and establishes the standards for licensure or certification of health professionals. Existing law makes certain violations of specified provisions relating to healing arts by a licensee or registrant a crime. (Bus. & Prof. Code § 500 *et seq.*)
 - 4) Provides various types of facilities, places, or buildings that fall under the definition of “health facility.” (Health & Saf. Code § 1250.)
 - 5) Requires a health facility, clinic, physician’s office, or office of a group practice that uses generative artificial intelligence (GenAI) to generate written or verbal patient communications pertaining to patient clinical information, as defined, to ensure that those communications include both a disclaimer that indicates to the patient that a communication was generated by generative artificial intelligence, as specified, and clear instructions describing how a patient may contact a human health care provider, employee, or other appropriate person. (Health & Saf. Code § 1339.75.)
 - 6) Establishes the Division of Labor Standards Enforcement, under the direction of the Labor Commissioner, within the Department of Industrial Relations and charges the Labor Commissioner with enforcement of various labor laws, including investigation of employee complaints. (Lab. Code § 79 *et seq.*)

- 7) Authorizes the Labor Commissioner to investigate employee complaints. The Labor Commissioner may provide for a hearing in any action to recover wages, penalties, and other demands for compensation, including liquidated damages, as specified, properly before the Division of Labor Standards Enforcement or the Labor Commissioner, including orders of the Industrial Welfare Commission, and shall determine all matters arising under their jurisdiction. (Lab. Code § 98(a).)
- 8) Provides that it shall not be a defense, and a defendant may not assert, that an artificial intelligence system autonomously caused harm to a plaintiff in an action against a defendant who developed, modified, or used an AI system that is alleged to have caused a harm to the plaintiff. (Civ. Code § 1714.46.)

THIS BILL:

- 1) 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. “Clinic” has the same meaning as defined in Section 1200 of the Health and Safety Code.
 - c. “Clinical decision support system” means a computerized system or tool that does both of the following:
 - i. Supports decisionmaking related to patient care based on algorithms, or models, based in clinical practice guidelines or that derive relationships from training data, including algorithms or models that are developed using unsupervised learning models.
 - ii. Produces an output that results in a prediction, classification, recommendation, evaluation, or analysis.
 - d. “Covered tool” means a tool, system, or device that includes artificial intelligence or a clinical decision support system.
 - e. “Health facility” has the same meaning as defined in Section 1250 of the Health and Safety Code.
 - f. “Office of a group practice” has the same meaning as defined in Section 1339.75 of the Health and Safety Code.
 - g. “Physician’s office” has the same meaning as defined in Section 1339.75 of the Health and Safety Code.
 - h. “Technology” means scientific hardware or software, including artificial intelligence and clinical decision support systems, used to achieve a medical or nursing care objective at a health facility.

- 2) Prohibits a defendant who developed, modified, or used artificial intelligence or a clinical decision support system alleged to have harmed the 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 artificial intelligence or clinical decision support system is a superseding cause severing the defendant's liability for the alleged harm.
- 3) Requires a health facility, clinic, physician's office, or office of a group practice that uses or deploys an AI tool for patient care to provide the following information about the tool to any licensed health care professional or other person using a covered tool or viewing outputs from a covered tool, or to a patient who may be affected by the covered tool:
 - a. Details on the covered tool, including developer, funding source, any foundation model used, and description of output.
 - b. Intended use of the covered tool, including intended patient population, intended users, and intended decisionmaking role.
 - c. Cautioned out-of-scope use of the covered tool, including known risks and limitations.
 - d. List of the inputs into the covered tool.
 - e. Description of how the covered tool generates outputs.
 - f. Development details of the covered tool, including, but not limited to, all of the following:
 - i. Description of the training set or clinical research underlying recommendations, including demographic representativeness and known biases based on protected characteristics.
 - ii. Description of the relevance of training data to deployed setting.
 - iii. Process used to ensure fairness in development of the intervention.
 - g. Description of the validation process.
 - h. Qualitative measures of performance.
 - i. Description of ongoing maintenance of intervention implementation and use.
 - j. Description of updates and continued validation or fairness assessment process.
 - k. Notice that health care entities and developers may be liable for harm that results from the use of artificial intelligence in patient care.
 - l. Notice that a worker providing direct patient care is permitted to override the output of a covered tool if, in the judgment of the worker acting in their scope of practice, such an override is appropriate for the patient, or as necessary to comply with applicable law,

including civil rights law.

- 4) Requires the disclosure provided pursuant to (3) be provided consistent with the following:
 - a. To a new licensed health care professional or other person upon hire, onboarding, or credentialing, if that individual will likely use the covered tool or view outputs from the covered tool.
 - b. At least 90 days before a new covered tool is first deployed for patient care.
 - c. At least 90 days before a material change in the use, function, intended users, intended patient population, or decisionmaking role of an existing covered tool.
 - d. On or before February 1, 2028, and annually thereafter, by providing an updated inventory of all covered tools currently in use or deployed for patient care.
- 5) Declares it is the policy of the state that:
 - a. A worker providing direct patient care be free to use their professional judgment to make assessments and decisions within their scope of practice as appropriate for their patients.
 - b. A worker should not be penalized for relying in good faith on technology that the licensed health care professional's employer has selected or approved for their use in patient care.
- 6) Prohibits an employer from using or deploying technology to replace or eliminate a worker's use of professional judgment in patient care.
- 7) Prohibits an employer from retaliating or discriminating against a worker providing patient care based on their override of or compliance with "technology."

COMMENTS:

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

Healthcare 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. AB 2575 preserves healthcare workers' ability to follow their professional judgment by prohibiting employer retaliation when a worker overrides or follows a recommendation. AB 2575 requires transparency for AI tools so that patients and providers understand how they work and what the risks are. Overall, AB 2575 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 healthcare system.

- 2) **Background.** This bill relates to an AI tool's ability to contribute to the clinical decision-making process. Given the central role of AI in this discussion, it is appropriate to briefly examine how these systems work, and the role they play in health care spaces.

AI and GenAI. “Artificial intelligence” refers to 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, including numbers, text, audio, video, or other data.¹ “Generative artificial intelligence” (GenAI) is a subset of AI that produces outputs closely resembling human-created content.²

Compared to conventional computer programs, which act according to pre-programmed rules, GenAI models “learn” from examples such as books, articles, photos, film, or music. This learning occurs within “neural networks” – massive systems of nodes linked by adjustable connections – that encode statistical patterns gleaned from data. During training, data is broken into fundamental units known as “tokens” – groups of syllables, pixels, or musical notes, for example – that can be represented numerically. A naïve neural network is exposed to an incomplete sequence of tokens and prompted to predict the next token in the sequence. If the prediction is incorrect, the network adjusts the strengths of its connections in order to minimize error and improve its next prediction. This process continues iteratively until the neural network can reliably emulate the human-created content it was trained on. A trained neural network embedded in a GenAI system is known as a “model,” and the strengths of its connections are known as its “model weights.”³

LLMs do not fundamentally understand the text they produce. They calculate one token at a time – if they predict that the next word or symbol in an outputted sentence should be a period, then the sentence ends. Otherwise, the sentence continues. It is a testament to the ingenious architecture of the deep neural nets powering these systems that their outputs are remotely coherent. But while the text these systems produce is cogent, it is not always correct. According to Melanie Mitchell, an AI researcher at the Santa Fe Institute, “These systems live in a world of language. . . . That world gives them some clues about what is true and what is not true, but the language they learn from is not grounded in reality. They do not necessarily know if what they are generating is true or false.”⁴

There is a famous saying in computer science: “garbage in, garbage out.” The performance of an AI is directly impacted by the quality, quantity, and relevance of the data used to train it.⁵ If the data used to train the AI is biased, the tool’s outputs will be similarly biased. In their work on mitigating bias in artificial intelligence, the Berkeley Haas Center for Equity, Gender and Leadership (Center) tracks publicly available instances of bias in AI systems using machine learning. In their analysis of around 133 biased systems across industries from 1988 to the present day, the Center found that 44% (59 systems) demonstrate gender bias, with 26% (34 systems)

¹ AB 2885 (Bauer-Kahan & Umberg; Ch. 843, Stats. 2024) defined AI as “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.”

² AB 2013 (Irwin, Ch. 817, Stats. 2024) defined GenAI as “artificial intelligence that can generate derived synthetic content, such as text, images, video, and audio, that emulates the structure and characteristics of the artificial intelligence’s training data.”

³ IBM, What is generative AI?, <https://www.ibm.com/think/topics/generative-ai>; IBM, What is machine learning?, www.ibm.com/topics/machine-learning.

⁴ Cade Metz, “What Makes A.I. Chatbots Go Wrong?,” *New York Times*, March 29, 2023, www.nytimes.com/2023/03/29/technology/ai-chatbots-hallucinations.html.

⁵ 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/>.

exhibiting both gender and racial bias.⁶

AI in health care. AI in health care is not new; AI algorithms, machine learning, and predictive AI models of varying degrees of sophistication have been developed and deployed for decades. INTERNIST-1, which used a search algorithm to arrive at clinical diagnoses based on patients' symptoms, was created in 1971.⁷ ELIZA, a rules-based mental health therapy chatbot program, was developed even earlier.⁸ With the recent advancement of GenAI, particularly in natural language processing, interest in, use of, and hype over AI has grown rapidly and health care applications have proliferated. According to the National Academy of Medicine (NAM), GenAI has the potential to transform health and medicine as we know it: improving health care delivery, advancing medical research, and augmenting the capacity of clinicians to provide personalized care at an unprecedented scale. However, NAM also notes that the potential for both breakthrough innovation and unintended consequences demands careful consideration.⁹

When automated decision systems are deployed in health care, biased historical data can lead to patients being recommended substandard care on the basis of 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 – a risky procedure under any circumstances. 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.¹⁰

Similarly, in 2019, a study discovered harmful racial bias in an AI tool developed by the health care company Optum – a subsidiary of UnitedHealth Group – and used by providers across the country to offer care management services. The tool assigned Black patients lower likelihoods of adverse health outcomes than equally at-risk white patients. The authors found that this happened because the tool was designed to predict health care costs instead of needs. Because the health care system has historically spent less on care for Black patients than white patients for the same health conditions, the tool was, in essence, issuing a prediction that mirrored and perpetuated past discrimination.¹¹

The University of California San Francisco reported bias in an algorithm used to identify potential appointment no-shows to facilitate double-booking for that appointment. The program

⁶ Smith, G., & Rustagi, I. (2021). When Good Algorithms Go Sexist: Why and How to Advance AI Gender Equity. *Stanford Social Innovation Review*. <https://doi.org/10.48558/A179-B138>.

⁷ Hirani, R., Noruzi, K., et al. (2024). Artificial Intelligence and Healthcare: A Journey through History, Present Innovations, and Future Possibilities. *Life* (Basel, Switzerland), 14(5), 557. <https://doi.org/10.3390/life14050557>

⁸ Weizenbaum, J. (1966, January 1). ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>

⁹ National Academy of Medicine. (2020, April 30). Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril. <https://nam.edu/publications/artificial-intelligence-in-health-care-the-hope-the-hype-the-promise-the-peril/>.

¹⁰ Caleb J Colón-Rodríguez, “Shedding Light on Healthcare Algorithmic and Artificial Intelligence Bias,” US Department of Health & Human Services Office of Minority Health, (Jul. 12, 2023). <https://pmc.ncbi.nlm.nih.gov/articles/PMC6875681/>.

¹¹ Obermeyer, Z., Powers, B., et.al. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science* (New York, N.Y.), 366(6464), 447–453. <https://doi.org/10.1126/science.aax2342>

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, can create feedback loops that worsen discrimination.¹²

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 to 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 life-threateningly dangerous error was based on an accurate statistical correlation: asthmatics are, in fact, less likely to die from pneumonia than the general population precisely *because* asthma is such a serious risk factor that asthmatics automatically receive elevated care.¹³

3) **What this bill would do:** This bill would require a health facility, clinic, physician's office, or office of a group practice that uses or deploys an AI tool for patient care to provide specified information about the tool to any licensed health care professional or other person using or viewing outputs generated by the tool, or to any patient whose care may be affected by the tool. This bill would additionally prohibit an employer from using technology to replace or limit a worker's use of professional judgement in patient care, and would prohibit an employer from retaliating or discriminating against a worker based on their override of, or compliance with, technology whose use was approved by the employer for patient care. The bill would declare it to be the policy of the state that a worker providing direct patient care be free to use their professional judgement to make assessments and decisions within their scope of practice as appropriate for patients, and that they should not be penalized for relying in good faith on technology that the licensed health care professional's employer approved for use in patient care. Finally, the bill would prohibit a defendant who developed, modified, or used an AI tool or clinical decision support system that is alleged to have harmed the plaintiff from asserting, as a defense, that the failure of a licensed health care professional or other health care worker to override an output of the tool or system is a superseding cause severing the defendant's liability.

A coalition representing physicians, hospitals and health systems, health plans, life sciences, and other health care stakeholders opposes this bill on the grounds that it is overly broad and "likely to hinder beneficial patient outcomes." They write:

[AB 2575] 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. It would also hinder technological advancement and, troublingly, exacerbate existing health disparities by impeding the ability of health care providers, particularly those serving vulnerable

¹² University of California Presidential Working Group on Artificial Intelligence. (2021, October). Responsible AI: Recommendations to Guide the University of California's Artificial Intelligence Strategy.

<https://www.ucop.edu/ethics-compliance-audit-services/compliance/uc-ai-working-group-final-report.pdf>

¹³ Christian, B. (2020). The Alignment Problem: Machine Learning and Human Values. Norton, pp. 82-84.

communities, to leverage AI tools to improve patient outcomes and the health of the populations they serve.

Oakland Privacy pushes back on the suggestion that the bill's disclosure requirements will prove to be overly expansive in their letter of support:

In response to objections that the disclosure requirements are too expansive, we will note that health care is a life and death situation, making it one of the areas where one would want to err on the side of more information rather than less. We also are open to reasonable streamlining of the disclosure agreements as long as they remain informative enough for patients and workers to decide when a tool is helpful and when it is not.

TechNet, Civil Justice Association of California, and Connected Health Initiative describe a need for a tighter definition of "covered tool" and related terms:

While the bill is framed as addressing clinical decision support systems, its requirements apply broadly to any "covered tool" used for "patient care". This definition could encompass a wide range of technologies beyond traditional CDS tools, including notetaking software, administrative triage tools, scheduling systems, and other routine technologies used in modern health care delivery.

By applying extensive disclosure and compliance requirements across this broad category, the 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 go on to describe potential pitfalls in prohibiting tools that could "replace" a worker's use of professional judgement:

Lastly, we have significant concerns regarding the language in Section 3 subsection (c), which states that AI technology cannot be deployed if it will "replace or limit a worker's use of professional judgment in clinical care." We are concerned that this language lacks clarity and may result in organizations not deploying systems that could make providers more effective and efficient in their practice. For instance, AI models in radiology that can screen images and detect cancerous tumors is technology that has long been in use, which has "replaced" a portion of the radiologists' "professional judgment." Radiologists have embraced this technology for well over a decade, and it has resulted in huge benefits to society.

4) **Committee amendments.** To address these various concerns, the author has agreed to amend the bill in the following ways:

- a. Throughout the bill, instead of using the terms "*artificial intelligence*," "*covered tool*," "*technology*," and "*clinical decision support system*," recast bill exclusively in terms of "*clinical decision support systems*," defined to mean "*an automated decision system or generative artificial intelligence system whose outputs are used to inform clinical decisionmaking with respect to the provision, timing, or course of patient care.*" **This definition will be shared with AB 1979 (Bonta)**, which will be heard in the same hearing as this bill. Insert definitions for "*automated decision system*" and "*generative*

artificial intelligence system” accordingly.

- b. In Section 2, remove “funding source” and “foundation model used” from information required to be provided to licensed health care professional or other person using a clinical decision support system.
- c. In Section 2, remove the requirement that a licensed health care professional using a clinical decision support system receive notice that “health care entities and developers may be liable for harm that results from the use of artificial intelligence in patient care.”
- d. In Section 2, remove patient disclosure requirements.
- e. In Section 2, remove “ample time” timing requirement.
- f. In Section 3, remove the provision containing a public policy statement related to workers relying in good faith on technology.
- g. In Section 3, remove the prohibition on use or deployment of technology to replace or limit a worker’s use of professional judgement in healthcare.
- h. In Section 3, replace existing language related to retaliation or discrimination with a prohibition on retaliation or discrimination against a worker providing direct patient care “based solely on the worker’s override of, or reliance on, the output of a clinical decision support system,” and clarify that the provision does not affect a worker’s duty to meet the applicable standard of care, act within their scope of practice, or exercise independent professional judgement in providing direct patient care.
- i. Various minor technical and conforming changes throughout the bill.

The full text of the bill as proposed to be amended follows:

SECTION 1. Section 1714.48 is added to the Civil Code, to read:

1714.48. (a) For purposes of this section, the following definitions shall apply:

(1) “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.

(x) (A) “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.

(B) “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.

(2) “Clinical decision support system” means *an automated decision system or generative artificial intelligence system whose outputs are used to inform clinical decisionmaking with respect to the provision, timing, or course of patient care.* ~~a computerized system or tool that does both of the following:~~

~~(A) Supports decisionmaking related to patient care based on algorithms, or models, based in clinical practice guidelines or that derive relationships from training data, including algorithms or models that are developed using unsupervised learning models.~~

~~(B) Produces an output that results in a prediction, classification, recommendation, evaluation, or analysis.~~

(x) “Generative artificial intelligence” has the same meaning as defined in Section 1339.75.

(b) In an action against a defendant who developed, modified, selected, or deployed ~~artificial intelligence or a~~ clinical decision support system that is alleged to have caused harm to the plaintiff, it shall not be a defense, and the defendant may not assert, that the failure of a licensed health care professional or other health care worker to override an output of the ~~artificial intelligence or~~ clinical decision support system is a superseding cause severing the defendant’s liability for the alleged harm.

(c) This section does not limit or preclude a defendant from presenting either of the following:

(1) Any other affirmative defense, including evidence relevant to causation or foreseeability.

(2) Other evidence relevant to the comparative fault of any other person or entity.

SEC. 2. Section 1339.76 is added to the Health and Safety Code, to read:

1339.76. (a) A health facility, clinic, physician’s office, or office of a group practice that uses or deploys a ~~covered tool~~ **clinical decision support system** for patient care shall provide written notice of required information, described in subdivision (b), to any licensed health care professional or other person using a **clinical decision support system** ~~covered tool~~ or viewing outputs from a **clinical decision support system** ~~covered tool~~.

(b) Required information under subdivision (a) shall include all of the following:

(1) Details on the **clinical decision support system** ~~covered tool~~, including developer, ~~funding source, any foundation model used,~~ and description of output.

(2) Intended use of the **clinical decision support system** ~~covered tool~~, including intended patient population, intended users, and intended decisionmaking role.

(3) Cautioned out-of-scope use of the **clinical decision support system** ~~covered tool~~, including known risks and limitations.

- (4) List of the inputs into the *clinical decision support system*~~covered tool~~.
- (5) Description of how the *clinical decision support system* ~~covered tool~~ generates outputs.
- (6) Development details of the *clinical decision support system*~~covered tool~~, including, but not limited to, all of the following:
 - (A) Description of the training set or clinical research underlying recommendations, including demographic representativeness and known biases based on protected characteristics.
 - (B) Description of the relevance of training data to deployed setting.
 - (C) Process used to ensure fairness in development of the intervention.
- (7) Description of the validation process.
- (8) Qualitative measures of performance.
- (9) Description of ongoing maintenance of intervention implementation and use.
- (10) Description of updates and continued validation or fairness assessment process.
- ~~(11) Notice that health care entities and developers may be liable for harm that results from the use of artificial intelligence in patient care.~~
- (12) Notice that a worker providing direct patient care is ~~permitted to~~**may** override the output of a *clinical decision support system* ~~covered tool~~ if, in the *independent professional* judgment of the worker acting ~~in~~**within** their scope of practice, ~~such an~~**the** override is ~~appropriate for the patient~~**necessary to meet the applicable standard of care or, or as necessary to comply with applicable law, including civil rights law.**
- (c) (1) A disclosure made pursuant to this section shall be provided consistent with all of the following:
 - (A) To a new licensed health care professional or other person upon hire, onboarding, or credentialing, if that individual will likely use the *clinical decision support system* ~~covered tool~~ or view outputs from the *clinical decision support system* ~~covered tool~~.
 - (B) At least 90 days before a new *clinical decision support system* ~~covered tool~~ is first deployed for patient care.
 - (C) At least 90 days before a material change in the use, function, intended users, intended patient population, or decisionmaking role of an existing *clinical decision support system*~~covered tool~~.
 - (D) On or before February 1, 2028, and annually thereafter, by providing an updated inventory of all *clinical decision support systems* ~~covered tools~~ currently in use or deployed for patient care.

~~(2) The disclosure shall be provided in plain language to, and linked in the health record of, any patient whose care may be affected by the output of the covered tool or whose health information may be used as an input to the covered tool.~~

~~(3) The disclosure shall be provided with ample time for the licensed health care professional or other person to review and make reasoned decisions based on their professional judgment on whether and how to use the covered tool.~~

(d) (1) A violation of this section by a licensed health facility is subject to the enforcement mechanisms described in Article 4 (commencing with Section 1290) of Chapter 2.

(2) A violation of this section by a licensed clinic is subject to the enforcement mechanisms described in Article 4 (commencing with Section 1235) of Chapter 1.

(3) A violation of this section by a physician is subject to the jurisdiction of the Medical Board of California or the Osteopathic Medical Board of California, as appropriate.

(4) A violation of this section constitutes “unfair competition” as defined in Section 17200 of the Business and Professions Code and is punishable as prescribed in Chapter 5 (commencing with Section 17200) of Part 2 of Division 7 of the Business and Professions Code.

(e) For purposes of this section, the following definitions shall apply:

(1) “Artificial intelligence” has the same meaning as in Section 1339.75.

(x) (A) “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.

(B) “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.

(2) “Clinic” has the same meaning as defined in Section 1200.

(3) “Clinical decision support system” means ***an automated decision system or generative artificial intelligence system whose outputs are used to inform clinical decisionmaking with respect to the provision, timing, or course of patient care.*** ~~a computerized system or tool that does both of the following:~~

~~(A) Supports decisionmaking related to patient care based on algorithms, or models, based in clinical practice guidelines or that derive relationships from training data, including such algorithms or models that are developed using unsupervised learning models.~~

~~(B) Produces an output that results in a prediction, classification, recommendation, evaluation, or analysis.~~

~~(4) “Covered tool” means a tool, system, or device that includes artificial intelligence or a clinical decision support system.~~

(x) “Generative artificial intelligence” has the same meaning as defined in Section 1339.75.

(5) “Health facility” has the same meaning as defined in Section 1250.

(6) “Office of a group practice” has the same meaning as defined in Section 1339.75.

~~(7) “Patient clinical information” has the same meaning as defined in Section 1339.75.~~

(8) “Physician’s office” has the same meaning as defined in Section 1339.75.

SEC. 3. Article 2.7 (commencing with Section 2820) is added to Chapter 2 of Division 3 of the Labor Code, to read:

Article 2.7. Health Information Technology: Worker Rights

2820. For the purposes of this article, the following definitions shall apply:

(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.

(x) (1) “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.

(2) “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.

(b) “Clinical decision support system” means ***an automated decision system or generative artificial intelligence system whose outputs are used to inform clinical decisionmaking with respect to the provision, timing, or course of patient care.*** ~~a computerized system or tool that does both of the following:~~

~~(1) Supports decisionmaking related to patient care based on algorithms, or models, based in clinical practice guidelines or that derive relationships from training data, including such algorithms or models that are developed using unsupervised learning models.~~

~~(2) Produces an output that results in a prediction, classification, recommendation, evaluation, or analysis.~~

(x) “Generative artificial intelligence” has the same meaning defined in Section 1339.75 of the Health and Safety Code.

~~(c) “Technology” means scientific hardware or software, including artificial intelligence and clinical decision support systems, used to achieve a medical or nursing care objective at a health facility.~~

2821. (a) It is the public policy of the State of California that a worker providing direct patient care be free to use their professional judgment to make assessments and decisions within their scope of practice as appropriate for their patients.

~~(b) It is the public policy of the State of California that a worker providing direct patient care should not be penalized for relying in good faith on technology that the licensed health care professional’s employer has selected or approved for their use in patient care.~~

~~(c) An employer shall not use or deploy technology to replace or limit a worker’s use of professional judgment in patient care.~~

~~(db) (1)~~ An employer shall not retaliate or discriminate against a worker providing direct patient care based *solely* on ~~both of the following~~:

~~(1) The~~ *the* worker’s override of, or request to override, *reliance on*, the output of technology if, in the judgment of the worker acting in their scope of practice, such an override is appropriate for the patient, or as necessary to comply with applicable law, including civil rights law.

~~(2) The worker’s compliance with the output of technology if the technology was provided or approved by the worker’s employer for patient care.~~ *a clinical decision support system.*

(2) This subdivision does not affect a worker’s duty to meet the applicable standard of care, act within their scope of practice, or exercise independent professional judgement in providing direct patient care.

~~(ec)~~ A worker who is subject to retaliation or discrimination in violation of this article has the right under this article to file a complaint with the Labor Commissioner against an employer who retaliates or discriminates against the worker.

ARGUMENTS IN SUPPORT: California Nurses Association, a co-sponsor of this bill, writes in support:

California currently lacks clear guardrails governing the use of artificial intelligence in health care settings. Hospitals and clinics now use AI tools in electronic health records, clinical decision-support systems, 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. Yet state law provides few standards to ensure transparency, protect professional judgment, or establish accountability when these technologies contribute to harmful outcomes.

California Labor Federation, the other co-sponsor of this bill, writes in support:

Currently, 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, clinical decision supports systems, remote monitoring platforms, and administrative decision-making. These tools can influence patient acuity scores, treatment recommendations, insurance determinations, discharge planning, and nurse workloads.

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. For example, algorithmic systems may generate patient acuity scores, treatment prompts, or discharge recommendations without disclosing how those outputs were produced, what data they relied on, or their known limitations. Without this information, clinicians cannot fully evaluate the reliability of these tools, and patients have no visibility into technologies affecting their care.

ARGUMENTS IN OPPOSITION: A coalition of physicians, hospitals and health systems, health plans, life sciences, and other health care stakeholders, including California Medical Association, California Hospital Association, and California Chamber of Commerce writes in opposition:

Artificial intelligence is not an aspiration in health care. Rather, it is simply a reality that is saving lives in California today. We have a shared obligation and commitment to ensure that these tools are developed and deployed responsibly, equitably, and transparently. AB 2575, as drafted, would not achieve these goals. Instead, it would bury clinicians in unworkable disclosure requirements, create perverse liability incentives, undermine patient safety systems, impair clinical quality oversight, and ultimately reduce patient access to beneficial technology, with the greatest harm falling on the communities that can least afford it.

REGISTERED SUPPORT / OPPOSITION:

Support

California Federation of Labor Unions, Afl-cio (Co-Sponsor)
 California Nurses Association (Co-Sponsor)
 American Federation of State, County and Municipal Employees, Afl-cio
 California Alliance for Retired Americans
 California Democratic Party Rural Caucus
 California Faculty Association
 California Pan - Ethnic Health Network
 California School Employees Association
 Cft– a Union of Educators & Classified Professionals, Aft, Afl-cio
 Consumer Watchdog
 Health Access California
 Oakland Privacy
 TechEquity Action

Opposition

Advanced Medical Technology Association (ADVAMED)

Adventist Health

America's Physician Groups

Biocom

California Chamber of Commerce

California Hospital Association

California Life Sciences

California Medical Association (CMA)

California Radiological Society

California Society of Pathologists

Civil Justice Association of California (CJAC)

Connected Health Initiative

Lake Elsinor Chamber of Commerce

Menefee Valley Chamber of Commerce

Ochin, INC.

Southwest California Legislative Council

TechNet

Temecula Chamber of Commerce

Wildomar Chamber of Commerce

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