Office of Senate Floor Analyses
(916) 651-1520 Fax: (916) 327-4478

UNFINISHED BUSINESS

Bill No: SB 503
Author: Weber Pierson (D), et al.
Amended: 8/19/26
Vote: 21

SENATE HEALTH COMMITTEE: 10-0, 4/9/25
AYES: Menjivar, Durazo, Gonzalez, Grove, Limón, Padilla, Richardson, Rubio, Weber Pierson, Wiener
NO VOTE RECORDED: Valladares

SENATE JUDICIARY COMMITTEE: 13-0, 4/29/25
AYES: Umberg, Niello, Allen, Arreguín, Ashby, Caballero, Durazo, Laird, Stern, Valladares, Wahab, Weber Pierson, Wiener

SENATE APPROPRIATIONS COMMITTEE: 6-0, 5/23/25
AYES: Caballero, Seyarto, Cabaldon, Grayson, Richardson, Wahab
NO VOTE RECORDED: Dahle

SENATE FLOOR: 38-0, 5/29/25
AYES: Allen, Alvarado-Gil, Archuleta, Arreguín, Ashby, Becker, Blakespear, Cabaldon, Caballero, Cervantes, Choi, Cortese, Dahle, Durazo, Gonzalez, Grayson, Grove, Hurtado, Jones, Laird, McGuire, McNerney, Menjivar, Niello, Ochoa Bogh, Padilla, Pérez, Richardson, Rubio, Seyarto, Smallwood-Cuevas, Stern, Strickland, Umberg, Valladares, Wahab, Weber Pierson, Wiener
NO VOTE RECORDED: Limón, Reyes

ASSEMBLY FLOOR: Passed, 8/24/26

SUBJECT: Health care services: artificial intelligence

SOURCE: Author

DIGEST: This bill (1) requires developers of clinical decision support systems, as defined, to make reasonable efforts to mitigate the known or reasonably

foreseeable risk of biased impacts resulting from the use of these systems; and, (2) requires developers to provide deployers of systems identified as presenting a risk of biased impacts with specified information, including documentation of how the system was evaluated for performance, limitations, and mitigation of biased impacts.

Assembly Amendments revised and recast the provisions of this bill, including by changing terminology from patient care decision support tools to clinical decision support systems to conform with other health care AI bills, deleted a provision requiring the tools to be tested for bias every three years, added a requirement that developers make certain information available to users of clinical decision support systems, and made other technical, clarifying, and conforming changes.

ANALYSIS:

Existing law:

1) Establishes the Department of Health Care Access and Information (HCAI), and designates HCAI as the state agency designated to collect health facility data for use by all state agencies, including various financial data reports. [Health and Safety Code (HSC) §128730, et seq.]

2) Requires a health facility, clinic, physician's office, or office of a group practice that uses generative AI to generate written or verbal patient communications pertaining to patient clinical information to ensure that those communications include a disclaimer that indicates to the patient that the communication was generated by generative AI, and clear instructions on how to contact a human person. Defines "artificial intelligence," for these purposes, as an engineered or machine-based system that varies in 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. [HSC §1339.75]

3) Requires the Department of Technology to conduct, in coordination with other interagency bodies as it deems appropriate, a comprehensive inventory of all high-risk automated decision systems that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, and state agency. Defines "high-risk automated decision system" as an automated decision system that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to, or approval for, housing or accommodations, education, employment, credit, health care, and criminal justice. [Government Code (GOV) §11546.45.5]

4) Establishes the California Privacy Protection Agency (CPPA) to implement and enforce the California Privacy Rights Act of 2020, and provides the CPPA with the full administrative power, authority and jurisdiction to implement and enforce the California Consumer Privacy Act of 2018, including responsibilities to update existing regulations and adopt new regulations. [Civil Code (CIV) §1798.100, et seq.]

5) Establishes the Unruh Civil Rights Act that entitles all persons within the jurisdiction of this state to full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever no matter what their sex, race, color, religion, ancestry, national origin, disability, medical condition, genetic information, marital status, sexual orientation, citizenship, primary language, or immigration status. [CIV §51(b)]

This bill:

1) Requires developers and deployers, as each of these are defined, to make reasonable efforts to identify clinical decision support systems developed for use by deployers that are known or have a reasonably foreseeable risk of biased impacts resulting from the deployment of the system in health programs or activities.

2) Requires developers to make reasonable efforts to mitigate the known or reasonably foreseeable risk of biased impacts resulting from the use of the clinical decision support systems identified pursuant to 1) above.

3) Requires developers to provide deployers of clinical decision support systems identified pursuant to 1) above a statement describing the intended uses and known or reasonably foreseeable risks associated with use of the clinical decision support system, and documentation disclosing all of the following:
   a) High-level summaries of the type or types of data used to train the clinical decision support system, including a description of its demographic representativeness when demographic data is available;
   b) How the clinical decision support system was evaluated for performance, limitations, and mitigation of biased impacts;
   c) High-level summaries of data governance measures used to evaluate training datasets and measures used to determine the suitability of data sources and possible biases;
   d) Intended benefits and expected outputs of the clinical decision support system;
   e) Any known or reasonably foreseeable risk of biased impacts and efforts to mitigate those risks;

   f) Recommendations regarding how the clinical decision support system should be used and monitored, and how risks should be mitigated; and,
   g) Any other documentation that is reasonably necessary to help deployers understand the outputs and monitor the performance of the clinical decision support system for biased impacts.

4) Permits developers, in complying with the disclosures in 3) above, to adhere to nationally recognized or widely adopted industry standards that are developed through multistakeholder consensus, relevant to bias testing and artificial intelligence systems in health care, or provide the results from regularly conducted algorithmic impact assessments using industry-accepted methodologies, including information that is relevant and appropriate for the clinical decision support system's intended use, to the extent that those standards or assessments exist.

5) Requires the documentation required under 3) above to be made available to deployers upon request or at the time of initial sale, whichever is earlier, and, as applicable, upon release of material updates to the clinical decision support system.

6) Specifies that a deployer that also serves as a developer of a clinical decision support system is not required to generate documentation for 3) above if the system is developed and deployed for internal use and information consistent with 3) above is readily available to the deployer.

7) Requires deployers to regularly monitor the clinical decision support systems identified pursuant to 1) above and to take reasonable and proportionate steps to mitigate the known or reasonably foreseeable risk of biased impacts.

8) Defines various terms for purposes of this bill, including the following:
   a) "Biased impact," means an adverse impact, including diminished access to health care, quality of care, or outcomes, on an individual based on their protected characteristics;
   b) "Clinical decision support system" means an artificial intelligence system that produces a prediction, classification, recommendation, evaluation, or analysis that aids clinical decisionmaking related to timing of care, diagnosis, or treatment. Excludes systems that provide appointment management, such as booking, canceling, and rescheduling appointments, appointment reminders, patient education and previsit materials and preparation, and payment processing, to the extent that the independent performance of these activities by the system does not require a professional license.

c) "Developer" means a person, partnership, state or local governmental agency, corporation, or deployer that designs, codes, substantially modifies, or otherwise produces a clinical decision support system for commercial or public use.
d) "Deployer" means a health facility, clinic, physician's office, or office of a group practice that uses a clinical decision support system.
e) "Protected characteristic" means any characteristic listed in the Unruh Civil Rights Act, as specified.

9) Specifies that a person, partnership, state or local governmental agency, or corporation can be both a developer and a deployer.

Background

According to the author of this bill:

> This bill is a crucial step toward ensuring fairness in healthcare by addressing the racial biases embedded in AI models and systems. These technologies are becoming more prevalent in healthcare, yet research has shown that these systems can produce biased outputs that disproportionately affect communities of color. Without proper oversight, these biases can go unchecked, deepening existing disparities in our healthcare system. This bill establishes a framework to identify and address biased outputs, ensuring that AI tools in healthcare are fair, accurate, and equitable for all patients. By taking a proactive approach, this bill helps protect patients from the effects of biased algorithms and ensures that technological advancements serve all communities equitably.

What is AI? In a draft regulation, the CPPA defines AI as follows: AI means a machine-based system that infers, from the input it receives, how to generate outputs that can influence physical or virtual environments. The AI may do this to achieve explicit or implicit objectives. Outputs can include predictions, content, recommendations, or decisions. Different AI varies in its levels of autonomy and adaptiveness after deployment. For example, AI includes generative AI models, such as large language models (LLMs), that can learn from inputs and create new outputs, such as text, images, audio, or video; and facial- or speech-recognition or -detection technology.

According to an article in MIT News, from the Massachusetts Institute of Technology, before the generative AI boom of the past few years, when people talked about AI, typically they were talking about machine-learning models that can learn to make a prediction based on data. For instance, such models are trained,

using millions of examples, to predict whether a certain X-ray shows signs of a tumor or if a particular borrower is likely to default on a loan. Generative AI can be thought of as a machine-learning model that is trained to create new data, rather than making a prediction about a specific dataset. A generative AI system is one that learns to generate more objects that look like the data it was trained on. Generative AI systems such as ChatGPT are powerful because they are trained on extremely large datasets, which could potentially take advantage of nearly all the information on the internet.

According to an article published by CalTech Science Exchange, today's generative AI models produce content that often is indistinguishable from that created by humans. Generative AI has successfully written news articles, created realistic artwork, and composed aesthetically pleasing music. While these applications sometimes make glaring mistakes (sometimes referred to as hallucinations), they are being used for many purposes, including in health care.

Generative AI in health care. According to a July 25, 2023 viewpoint article in the *Journal of the American Medical Association*, "Generative AI in Health Care and Liability Risks for Physicians and Safety Concerns for Patients," generative AI is being heralded in the medical field for its potential to ease the long-lamented burden of medical documentation by generating visit notes, treatment codes, and medical summaries. This article stated that to date, no court in the U.S. has considered the question of liability for medical injuries caused by relying on AI-generated information. The article warned that to minimize risks, even medical professionals must exercise caution when relying on AI-generated information, due to the "black-box" nature of how these systems generate output, because it may be impossible for the physician to independently evaluate the accuracy of the AI's output. A March 2024 article in Implement Science stated that the utility and impact of generative AI in health care remain poorly understood, with concerns about ethical and legal implications, integration into healthcare service delivery, and workforce utilization. The *American Medical Association* published "Principles for Augmented Intelligence Development, Deployment, and Use," which was approved by the American Medical Association (AMA) Board of Trustees on November 14, 2023. According to this AMA document, as the number of AI-enabled health care tools and systems continue to grow, these technologies must be designed, developed, and deployed in a manner that is ethical, equitable, responsible, and transparent. The AMA notes that while the FDA regulates AI-enabled medical devices, many types of AI-enabled technologies fall outside the scope of FDA oversight, including AI that may have clinical applications, such as some clinical decision support functions. The AMA report specifically looked at transparency in use of AI-enabled systems, and stated that it is essential that use of

AI in health care be transparent to both physicians and patients, and disclosure should contribute to physician and patient knowledge and not create unnecessary administrative burden. The AMA report states that when AI is utilized in health care decision-making, that use should be disclosed and documented in order to limit risks to, and mitigate inequities for, both physicians and patients. The AMA noted that while transparency does not necessarily ensure AI-enabled tools are accurate, secure, or fair, it is difficult to establish trust if certain characteristics are hidden. According to the AMA, when AI is used in a manner, which directly impacts patient care, access to care, or medical decision-making, that use of AI should be disclosed and documented to both physicians and patients. The opportunity for a patient to request additional review from a licensed clinician should be made available upon request. The AMA policy also states that AI tools or systems cannot augment, create, or otherwise generate records, communications, or other content on behalf of a physician without that physician's consent and final review.

Racial bias in algorithms and AI. Several studies have identified risks of racial bias in algorithms and AI. A October 2019 article in the journal *Science*, "Dissecting racial bias in an algorithm used to manage the health of populations," found that Black patients assigned the same level of risk by the algorithm are sicker than White patients. The study estimated that this racial bias reduces the number of Black patients identified for extra care by more than half. According to the study, bias occurs because the algorithm uses health costs as a proxy for health needs, and because less money is spent on Black patients who have the same level of need, the algorithm falsely concludes that Black patients are healthier than equally sick White patients. The journal *NPJ Digital Medicine* published a study in 2023, "Large language models propagate race-based medicine," that assessed four large language models used in healthcare systems. The study found that all models had examples of perpetuating race-based medicine in their response to questions. Models were not always consistent in their response when asked the same question repeatedly, and that based on their findings, these large language models could potentially cause harm by perpetuating debunked, racist ideas.

Trustworthy AI. The federal Health and Human Services Agency (HHS) has developed a "playbook" for AI development with the goal of maintaining public trust by using AI solutions that are "ethical, effective, and secure," according to the HHS Chief AI Officer. With regard to HHS agencies, principles have been established for trustworthy AI that "fosters public trust and confidence while protecting privacy, civil rights, civil liberties, and American values, consistent with applicable laws." The playbook describes AI as whether the solution or system:

a) Performs tasks under varying and unpredictable circumstances without significant human oversight, or can learn from experience and improvement performance when exposed to data sets;
b) Uses computer software, physical hardware, or other technology to solve tasks that require human-like perception, thinking, planning, learning, communication, or physical action;
c) Thinks or acts like a human, including the use of cognitive architecture or neural networks (e.g., developed to mimic the underlying mechanisms of the human mind);
d) Relies on a set of techniques, including machine learning, to approximate a cognitive task; and,
e) Is designed to act rationally by utilizing intelligent software or an embodied robot to achieve goals using perception, planning, reasoning, learning, communicating, decision-making, and action.

The HHS principles for use of trustworthy AI in government cover six areas: fair and impartial application; transparent and explainable data use; responsible and accountable governance and policies; robust and reliable outputs; safe and secure from risks; and, that privacy should be respected, data not used beyond its intended and stated use, and that it's use is approved by the data owner or steward.

FISCAL EFFECT: Appropriation: No Fiscal Com.: Yes Local: No

According to the Assembly Appropriations Committee analysis, costs to the University of California (UC) of an unknown but potentially significant amount. According to UC Office of the President (UCOP), costs for vendor auditing will likely be passed on to UC health systems. UCOP also states the third-party audit requirement will disincentivize AI development within academic medical centers by adding substantial new cost for implementing internally-developed AI solutions. UCOP indicates these costs would reduce hospital revenue and create cost pressures for UC and the General Fund.

SUPPORT: (Verified 8/24/26)

America's Physician Groups
California Association of Health Plans
California Medical Association
California Radiological Society
Kaiser Permanente
OCHIN Inc.
University of California

OPPOSITION: (Verified 8/24/26)

Advanced Medical Technology Association

ARGUMENTS IN SUPPORT: A coalition representing physicians, hospitals, physician groups, health plans, and other stakeholders, including the California Medical Association, Kaiser Permanente, and the California Association of Health Plans, submitted a letter of support to the most recent version of this bill. These supporters state that this bill would provide valuable information to health care deployers such as clinics, rural hospitals, independent physician practices, and others, before making significant technology/AI investments. Under this bill, developers would be required to provide information that will significantly help small health care providers to make informed decisions and better understand the strengths and limitations of tools they are considering using. This bill would equip deployers with necessary information to mitigate bias and risks of such tools and systems. These supporters state that AI holds extraordinary potential to transform and improve every aspect of health care, and this bill helps to level the playing field by ensuring all purchasers, regardless of size, have access to basic information necessary to determine whether a tool is appropriate for their patients and clinical workflows.

ARGUMENTS IN OPPOSITION: The Advanced Medical Technology Association (AdvaMed) submitted a letter of opposition to the most recent version of this bill. According to AdvaMed, unlike many other industries, the use of AI in medical technology is already subject to robust oversight and regulation by the FDA throughout the device's lifecycle. FDA's premarket review of AI-enabled devices includes an assessment of clinical relevant information, potential risk, the mitigation of unwanted bias, and performance validation on the device's intended use and patient populations. Additionally, medical device manufacturers must follow strict federal requirements and regulations once devices are on the market, including monitoring and reporting of device malfunctions and serious adverse events. AdvaMed states that while they value the goal of this bill, they are concerned with the potential unintended consequences. AdvaMed's concerns include the lack of clear definitions; that it overlooks the existing regulatory framework; and that it would establish duplicative California-specific requirements, and risks to confidentiality and intellectual property.

Prepared by: Vincent Marchand / Health / (916) 651-4111
8/24/26 17:09:33

****** END ******