

Date of Hearing: July 16, 2025

Fiscal: Yes

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

SB 503 (Weber Pierson) – As Amended July 10, 2025

SENATE VOTE: 38-0

PROPOSED AMENDMENTS

SUBJECT: Health care services: artificial intelligence

SYNOPSIS

In 2020, the Legislature enacted AB 3121 (Weber, Ch. 319, Stats. 2020), which established the Task Force to Study and Develop Reparation Proposals for African Americans, with Special Consideration for African Americans Who Are Descendants of Persons Enslaved in the United States. In 2023, the Task Force released a report detailing both the historical injustices of slavery and their continuing influence in contemporary society. The report also offered a series of recommendations for addressing those enduring harms.

With respect to healthcare, the Task Force observed: “Racial disparities in African American health outcomes occur today as a culmination of historical racial inequality, discriminatory health policy, and persistent racial discrimination in many sectors of life in the United States.” As technology advances, these disparities can manifest in new ways, including the underrepresentation of people of color, particularly Black Californians, in the data used to train algorithms and artificial intelligence (AI) systems. This underrepresentation may contribute to disparate outcomes, such as misdiagnoses, denial or degradation of care, and neglect by systems intended to promote health and wellbeing.

This bill, sponsored by the California Black Health Network, seeks to address these historical and structural inequities by imposing requirements on developers and deployers of AI systems used in healthcare settings. Specifically, it requires developers and deployers to identify AI systems they use to support clinical decisionmaking or resource allocation that are known to, or may reasonably be foreseen to, pose a risk of adverse impact on a protected class. Additionally, the bill imposes a duty on both developers and deployers to take reasonable measures to mitigate risks of biased outputs. Deployers must also monitor the AI systems under their control and take reasonable corrective action if biased impacts are detected.

This measure is supported by Kaiser Permanente, the California Medical Association, and the California Hospital Association. The bill has no registered opposition. It passed the Health Committee on a 16-0 vote.

Committee amendments, outlined in Comment #6, would enhance accountability for developers of these technologies by requiring independent third-party audits to assess compliance with this bill beginning in 2030.

THIS BILL:

- 1) Requires developers of artificial intelligence systems, and health facilities, clinics, physician's offices, or offices of a group practice, to have an ongoing duty to make reasonable efforts to identify artificial intelligence systems used to support clinical decisionmaking or health care resource allocation that are known or have a reasonably foreseeable risk of biased impacts in the system's outputs resulting from use of the system in health programs or activities.
- 2) Requires developers and deployers shall make reasonable efforts to mitigate the risk for biased impacts in the system's outputs resulting from use of the system in health programs or activities.
- 3) Requires deployers shall regularly monitor these artificial intelligence systems and take reasonable and proportionate steps to mitigate any bias that may occur.
- 4) Defines the terms as follows:
 - a. "Biased impact" to mean an unintended adverse impact, including diminished access to health care, quality of care, or outcomes, on an individual based on their protected characteristics.
 - b. "Deployer" to mean a person, partnership, state or local governmental agency, corporation, or developer that uses an artificial intelligence system to support clinical decisionmaking or health care resource allocation
 - c. "Developer" to mean a person, partnership, state or local governmental agency, corporation, or deployer that designs, codes, substantially modifies, or otherwise produces an artificial intelligence system for commercial or public use to support clinical decisionmaking and or health care resource allocation.

EXISTING LAW:

- 1) Prohibits, pursuant to the ACA, an individual, on the ground prohibited under title VI of the Civil Rights Act of 1964, title IX of the Education Amendments of 1972, the Age Discrimination Act of 1975, or section 504 of the Rehabilitation Act of 1973, be excluded from participation in, be denied the benefits of, or be subjected to discrimination under, any health program or activity, any part of which is receiving Federal financial assistance, including credits, subsidies, or contracts of insurance, or under any program or activity that is administered by an Executive Agency or other entity, as provided. (42 U.S.C. § 18116 ("Section 1557").)
- 2) Provides that the enforcement mechanisms provided for and available under such title VI, title IX, section 504, and Age Discrimination Act shall apply for purposes of violations of the above. Authorizes the Secretary to promulgate regulations relevant thereto. (42 U.S.C. § 18116.)
- 3) Provides the following regulatory guidelines with regard to the above:
 - a) A covered entity must not discriminate on the basis of race, color, national origin, sex, age, or disability in its health programs or activities through the use of patient care decision support tools.

- b) A covered entity has an ongoing duty to make reasonable efforts to identify uses of patient care decision support tools in its health programs or activities that employ input variables or factors that measure race, color, national origin, sex, age, or disability.
 - c) For each patient care decision support tool identified in (b), a covered entity must make reasonable efforts to mitigate the risk of discrimination resulting from the tool's use in its health programs or activities. (45 C.F.R. § 92.210.)
- 4) Establishes the Unruh Civil Rights Act (“Unruh”), which provides that all persons within the jurisdiction of this state are free and equal, and 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, are entitled to the full and equal accommodations, advantages, facilities, privileges, or services in all business establishments of every kind whatsoever. (Civ. Code § 51.)
- 5) Requires the California Department of Technology (CDT) to conduct a comprehensive inventory of all high-risk automated decision systems (ADS) that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency. It defines the relevant terms:
- 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. “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) “High-risk automated decision system” means an ADS that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to, or approval for, housing or accommodations, education, employment, credit, health care, and criminal justice. (Gov. Code § 11546.45.5.)

COMMENTS:

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

SB 503 is a crucial step towards ensuring fairness in healthcare by addressing the racial biases embedded in AI models and systems. This technology is 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 will require collaboration between developers and healthcare facilities to identify AI tools used in the delivery of patient care and proactively work towards meaningfully reducing bias. By requiring identification, mitigation, and oversight, SB 503 will help promote safety, equity, and exceptional performance while protecting patients against avoidable harm.

2) **California Reparation Report and Healthcare.** In 2020, the Legislature enacted AB 3121 (Weber, Ch. 319, Stats. 2020), which established the Task Force to Study and Develop Reparation Proposals for African Americans, with special consideration for those who are descendants of individuals enslaved in the United States. The Task Force’s work culminated in the publication of The California Reparations Report in 2023, which documents the ongoing and compounding harms experienced by African Americans as a result of slavery and its enduring legacy in contemporary society. With respect to healthcare, the report states:

Centuries of exposure to racism has contributed to a serious decline in African American physical and mental health. African Americans die at disproportionately higher rates from preventable health problems. Doctors are more likely to misdiagnose African Americans, leading to disparate outcomes in mental health. African American women face high rates of maternal death and adverse birth outcomes—even Black women with the highest education attainment have the worst birth outcomes across all women in America. African American children face poverty, malnutrition, and worse health than that of white American children. The mismanagement of public health crises by county, state, and federal governments has resulted in an undue burden of disease and death in African American communities—particularly during the COVID-19 pandemic. Despite this, in the face of overwhelming oppression, African American healthcare providers, patients, and community members, nonetheless, have worked to build healthy communities and fight for a more equitable healthcare system.¹

The report further documents the historical injustices within California’s healthcare system, which has systematically and consistently disadvantaged Black Californians. It includes a section highlighting how advances in technology have often failed to benefit Black communities, resulting in continued disparities in health outcomes and access to healthcare services:

Algorithms are widely used in U.S. hospitals to refer people to health programs that improve a patient’s care— however, at least one widely-used algorithm was found to systematically discriminate against Black patients. This algorithm led to African American patients receiving less referrals for programs that provided personalized care—despite being just as sick as white patients.

African Americans are less likely to be treated for skin diseases due to the lack of medical research and training for diagnosing skin conditions for those with darker skin. Most medical textbooks and journals that assist dermatologists in diagnosing skin disorders do not include images of skin conditions as they appear on African Americans. Images of darker skin with skin conditions caused by COVID-19, skin cancer, psoriasis, rosacea, and melanoma often do not appear in medical textbooks and journals. Doctors routinely miss these diagnoses for African American patients because they are not trained to identify or treat skin conditions for African American patients. Consequently, discriminatory medical research and technology has resulted in worsening health disparities that harm African Americans.²

¹ Task Force to Study and Develop Reparation Proposals for African Americans, “Chapter 12 - Mental and Physical Harm and Neglect”, *Department of Justice* (June 29, 2023), <https://oag.ca.gov/ab3121/report>.

² *Ibid.*

3) **AI and GenAI.** The development of GenAI is creating exciting opportunities to grow California's economy and improve the lives of its residents. GenAI can generate compelling text, images and audio in an instant – but with novel technologies come novel safety concerns.

In brief, AI is the mimicking of human intelligence by artificial systems such as computers. AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process: numbers, text, audio, video, or movement. AI is not fundamentally different from other computer functions; its novelty lies in its application. Unlike normal computer functions, AI is able to accomplish tasks that are normally performed by humans.

AI that are trained on small, specific datasets in order to make recommendations and predictions are sometimes referred to as “predictive AI.” This differentiates them from GenAI, which are trained on massive datasets in order to produce detailed text and images. When Netflix suggests a TV show to a viewer, the recommendation is produced by predictive AI that has been trained on the viewing habits of Netflix users. When ChatGPT generates text in clear, concise paragraphs, it uses GenAI that has been trained on the written contents of the internet.

4) **AI in Healthcare.** As noted in the background paper for the Joint hearing between this Committee and the Assembly Health Committee, AI in healthcare is not a new phenomenon:

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 years. Some of the first applications were developed in the 1970's and 1980's. INTERNIST-1, which used a search algorithm to arrive at clinical diagnoses based on patients' symptoms, was created in 1971. ELIZA, a rules-based mental health therapy chatbot program, was developed even earlier. In 2007, IBM created the open-domain question-answering system, “Watson.” In 2011, Watson won first place on Jeopardy and, in 2017, neurologists used it to identify key proteins that are altered in patients with Amyotrophic lateral sclerosis (ALS). Later, scientists at GoogleDeepMind shared a 2024 Nobel Prize in Chemistry for developing an AI model called AlphaFold2 to predict a protein's 3D structure from its amino-acid sequence, which is reportedly accelerating breakthroughs in biology and drug development.

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 and large language models (models designed for natural language processing tasks, or LLMs) have 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.³

Nevertheless, the maturity of a technology does not guarantee that it functions as intended. AI systems are often trained on historical datasets, which may underrepresent certain demographic groups and contribute to disparate health outcomes. These datasets frequently overrepresent

³ Joint Informational Hearing Assembly Health and Privacy & Consumer Protection Committees Generative Artificial Intelligence in Health Care background paper can be found at <https://ahea.assembly.ca.gov/hearings/2025-2026-informational-hearings>.

white individuals, who have historically accessed healthcare at higher rates than people of color. As a result, the outputs of such systems may perpetuate or exacerbate existing health disparities. The Assembly Health Committee Analysis outlines several examples of how these disparities have manifested in practice:

In their work on mitigating bias in AI, 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) exhibiting both gender and racial bias.

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— a procedure that carries some risk. 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, perpetuating historical racial and ethnic biases.

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 healthcare costs instead of needs. Because the healthcare system has historically spent less on care for Black patients than white patients for the same health conditions, the tool was issuing a prediction that mirrored and perpetuated past discrimination.

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

5) Federal Section 1557 Regulations and what this bill would do. Section 1557, the civil rights provision of the Affordable Care Act (ACA), prohibits discrimination on the basis of race, color, national origin, sex, age, or disability in certain health programs and activities.⁴ It authorizes the Secretary of Health and Human Services (HHS) to issue regulations to implement these nondiscrimination requirements, and in 2024, HHS promulgated the following regulations:

⁴ Details regarding Section 1557 of the ACA can be found at <https://www.hhs.gov/civil-rights/for-individuals/section-1557/fs-sex-discrimination/index.html>.

- A covered entity must not discriminate on the basis of race, color, national origin, sex, age, or disability in its health programs or activities through the use of patient care decision support tools.
- A covered entity has an ongoing duty to make reasonable efforts to identify uses of patient care decision support tools in its health programs or activities that employ input variables or factors that measure race, color, national origin, sex, age, or disability.
- For each patient care decision support tool identified in ii) above, a covered entity must make reasonable efforts to mitigate the risk of discrimination resulting from the tool's use in its health programs or activities.

Section 1557 applies only to “covered entities,” meaning health programs and activities that receive federal financial assistance from HHS. Examples of covered entities include hospitals, health clinics, physicians’ practices, community health centers, nursing homes, rehabilitation centers, health insurance issuers, and state Medicaid agencies. The Federal Register outlines the scope of the term “patient care decision support tools,” which includes: automated decision systems (ADS); artificial intelligence (AI); flowcharts; formulas; equations; calculators; algorithms; utilization management applications; software as medical devices (SaMDs); software in medical devices (SiMDs); screening, risk assessment, and eligibility tools; and diagnostic and treatment guidance tools.⁵

This bill codifies similar principles to those articulated in Section 1557 by imposing a duty to mitigate harms associated with AI systems, many of which are likely to fall within the categories described by the Federal Register and used in the healthcare sector. Notably, the bill extends obligations to developers, who are not covered under Section 1557, as well as to health facilities, clinics, physician offices, and group practices.

Under the bill, developers and deployers must exercise due diligence to identify AI systems used to support clinical decision-making or healthcare resource allocation that present a foreseeable risk of biased outcomes. The bill defines “biased impact” as an unintended adverse effect on healthcare access, quality, or outcomes based on a protected characteristic, as defined under the Unruh Civil Rights Act, a broader set of characteristics than those covered under Section 1557. In addition, the bill requires both developers and deployers to take reasonable steps to mitigate the risk of biased impacts. Deployers are also obligated to regularly monitor these AI systems and implement reasonable measures to address any bias that arises.

6) Amendments. A major challenge in imposing requirements on developers and deployers of AI systems is ensuring compliance and determining whether a standard of care is being upheld, particularly when many users of these technologies are not well-versed in relevant processes or standards. To address this concern, the Committee has adopted the use of third-party auditors as a guiding principle this session (as reflected in Bauer-Kahan AB 1018 and Wiener SB 53, also being heard today), with auditing requirements set to take effect in 2030. This delayed onset provides a four-year runway for the development of a market. Accordingly, the author has agreed to amendments that would adopt an auditing requirement for developers of these technologies

⁵ 89 Fed. Reg. 37522

that are used to support clinical decision-making or health care resource allocation which will also begin in 2030. The inclusion of this language would encompass deployers who also serve as developers. The intent is to capture only the largest and most well-resourced entities that are capable of complying with the auditing requirement.

However, due to timeline constraints, the current amended language will broadly apply to developers. In upcoming discussions with stakeholders, if it becomes apparent that the provision may inadvertently capture safety net providers or other resource-constrained health organizations, the Committee is committed to working with the author to adopt further amendments to ensure the intent is properly effectuated.

To be clear, the auditing requirement will not apply to entities that only deploy the technology. The author has also agreed to adopt various clarifying and technical amendments. The amendments are as follows:

SECTION 1. Section 1339.76 is added to the Health and Safety Code, to read:

1339.76. (a) (1) Developers ~~and deployers of artificial intelligence systems, and health facilities, clinics, physician's offices, or offices of a group practice,~~ shall have an ongoing duty to make reasonable efforts to identify artificial intelligence systems used to support clinical decisionmaking or health care resource allocation that are known or have a reasonably foreseeable risk of biased impacts in the system's outputs resulting from use of the system in health programs or activities.

(2) Developers and deployers shall make reasonable efforts to mitigate the risk for biased impacts in the system's outputs resulting from use of the systems ***identified pursuant to paragraph (1)*** in health programs or activities.

(3) Deployers shall regularly monitor these artificial intelligence systems ***identified pursuant to paragraph (1)*** and take reasonable and proportionate steps to mitigate any bias that may occur.

(b) For purposes of this section, a person, partnership, state or local governmental agency, or corporation may be both a developer and a deployer.

(c)(1) Except as provided in subparagraph (B), beginning January 1, 2030, and at least annually thereafter, a developer shall submit their artificial intelligence systems to an independent third-party auditor to assess whether the developer has complied with their duties pursuant to subdivision (a).

(2) An entity subject to paragraph (1) shall make a high-level summary of the results of an audit required by that paragraph publicly available at no cost to a person who accesses the developer's internet website.

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

(1) "Biased impact" means an unintended adverse impact, including diminished access to health care, quality of care, or outcomes, on an individual based on their protected characteristics.

(2) "Clinic" has the same meaning as defined in Section 1200 or 1200.1.

(3) “Deployer” means a *health facility, clinic, physician’s office, or office of a group practice* ~~person, partnership, state or local governmental agency, corporation, or developer~~ that uses an artificial intelligence system to support clinical decisionmaking or health care resource allocation.

(4) “Developer” means a person, partnership, state or local governmental agency, corporation, or deployer that designs, codes, substantially modifies, or otherwise produces an artificial intelligence system for commercial or public use to support clinical decisionmaking or health care resource allocation.

(5) “Health care provider” means a person licensed or certified pursuant to Division 2 (commencing with Section 500) of the Business and Professions Code.

(6) “Health facility” has the same meaning as Section 1250.

(7) “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 nonprofit corporation licensed according to subdivision (a) of Section 1204.

(8) “Artificial intelligence” has the same meaning as in Section 11546.45.5 of the Government Code.

(9) “Physician’s office” means an office of a physician in solo practice.

(10) “Protected characteristic” means a characteristic listed in subdivision (b) of Section 51 of the Civil Code.

(d) This section is in addition to and does not supplant or replace any other applicable provision of state law regulating the use of artificial intelligence or automated decision systems. Compliance with this section shall not be used as a defense to a claim of unlawful discrimination.

ARGUMENTS IN SUPPORT: Kaiser Permanente, writes in support:

Kaiser Permanente is pleased to support SB 503 (Weber-Pierson), which as amended June 30th, would require artificial intelligence (AI) developers and deployers to make reasonable efforts to mitigate biased impacts when AI tools are used in health programs and activities. As amended, the bill would also require deployers to regularly monitor such AI systems and tools to take steps in mitigating potential bias outcomes.

As one of the nation’s largest integrated health care organizations with a long history of adopting new technologies and applying them broadly, we are leading the responsible use of artificial intelligence across the health care sector. We use AI-based tools to improve health outcomes and the member care experience. And we use them without compromising patient safety, privacy, or the quality of our care. AI never makes medical decisions — our physicians and care teams do.

Artificial intelligence (AI) holds extraordinary potential to transform and improve every aspect of health care by enhancing care delivery and high quality, patient experiences, health outcomes, clinical safety, and affordability. AI has played a role in health care for decades

and has improved overall health outcomes. Artificial intelligence has also played a significant role in supporting a health care field that adheres to high standards of patient safety, privacy, responsible use, and ethical decision-making. These same principles must guide AI's continued development and deployment to ensure it improves patient care while maintaining trust.

Kaiser Permanente believes SB 503 is a measured and practical first approach in ensuring AI systems and tools are safe, secure, and reliable. AI must be trustworthy and for patients and the larger community to trust AI, we must do our best to prevent negative outcomes and maximize positive outcomes that benefit people. SB 503 is crucial for building trust in AI.

REGISTERED SUPPORT / OPPOSITION:

Support

California Hospital Association
California Medical Association (CMA)
Kaiser Permanente
Oakland Privacy

Opposition

None on file.

Analysis Prepared by: John Bennett / P. & C.P. / (916) 319-2200