

SENATE THIRD READING
SB 503 (Weber Pierson)
As Amended September 4, 2025
Majority vote

SUMMARY

Requires developers and deployers of artificial intelligence (AI) systems in specified health care applications to take steps to identify, mitigate, and monitor biased impacts. *Beginning in 2027, requires developers and deployers to provide the California Department of Public Health (DPH) with a report identifying their efforts to comply with the bill's requirements and maintain documentation of their compliance efforts. Specifies DPH is not required to independently inspect, test, or evaluate the functionality of an AI system.*

Major Provisions

COMMENTS

AI in Health Care. AI is the mimicking of human intelligence by artificial systems. 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 in health care is not new; AI models of varying degrees of sophistication have been developed and deployed in the health care setting for decades, and the use of AI is growing. According to the National Academies of Medicine, advancements in AI have the potential to transform health and medicine as we know it, but the potential for both breakthrough innovation and unintended consequences demands careful consideration.

Informational Hearing. On May 28, 2025, the Assembly Committees on Health and Privacy & Consumer Protection held an informational hearing titled, "*Generative Artificial Intelligence in Health Care: Opportunities, Challenges, and Policy Implications*." The hearing examined the history of AI in health care; current applications; a range of challenges that pose barriers to responsible, effective adoption of AI, including bias and other challenges; other considerations, such as workforce, cost, reimbursement, liability, and data privacy; and existing regulatory frameworks and best practices, including federal and state law and regulation and private efforts.

Current Use Cases. Researchers, health care providers and facilities, health plans and others are deploying AI for a range of tasks in biomedical and health research, as well as various administrative and clinical use cases.

Racial, Ethnic and Gender Bias in AI. 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. If the data used to train the AI is biased, the tool's outputs will be similarly biased and the results can be inaccurate when applied to populations not reflected in the training data. When automated decision systems are deployed in healthcare, biased historical data can lead to patients being recommended substandard care on the basis of their race or ethnicity. Numerous examples of bias have been described in the literature.

Affordable Care Act (ACA) Anti-Discrimination Regulations. Section 1557, the civil rights provision of the ACA, prohibits discrimination on the grounds of race, color, national origin, sex, age, or disability in certain health programs and activities. Section 1557(c) of the ACA

authorizes the Health and Human Services (HHS) Secretary to promulgate regulations to implement the nondiscrimination requirements of Section 1557. Section 1557 only applies to "covered entities," that is, health programs and activities that receive federal financial assistance from HHS. Examples of types of covered entities under Section 1557 include hospitals, health clinics, physicians' practices, community health centers, nursing homes, rehabilitation centers, health insurance issuers, and state Medicaid agencies.

2024 HHS Office of Civil Rights (OCR) regulations issued pursuant to Section 1557 prohibit a covered entity from discriminating 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 (automated or AI tools used to support clinical decision-making). The regulation requires a covered entity to make reasonable efforts to identify uses of patient care decision support tools that may promote bias and make reasonable efforts to mitigate the risk of discrimination resulting from the tool's use in its health programs or activities.

Alignment with Federal Section 1557 Regulations. This bill is similar to federal Section 1557 regulations in that deployers of these tools are required to make "reasonable efforts to mitigate the risk of discrimination on the basis of a protected characteristic resulting from the tool's use in its health programs or activities," similar to the federal regulation. However, the bill is different than federal regulations in key ways. It is narrower than the federal regulation in the types of tools it covers, it applies to fewer *health care entities* than the federal regulation applies to but extends its requirements to *developers*, which are not generally covered under Section 1557 regulations, and it specifically requires monitoring, which is not explicitly addressed in regulations, and it defines "protected characteristic" as those laid out in state civil rights law, which is broader than the federal regulation, which prohibits discrimination based on "race, color, national origin, sex, age, or disability." It also requires developers to submit to independent third-party audits.

How is Bias in AI Systems Identified, Mitigated, and Monitored? There is widespread awareness that bias is a problem that needs attention from developers and deployers of AI, and there is ongoing work to develop ways to measure and address it. Identifying, monitoring, and mitigating for bias may require a deployer to be aware of the literature on what types of systems may be prone to bias, understand how a model's training data compares to their patient population, conduct sensitivity analyses to see how calibrating a model in different ways effects the outputs of the model, and make technical adjustments to a model. For instance, a clinic might have to calibrate or adjust the model in a specific way to ensure it works effectively for their particular patient population. As industry standards continue to develop to support such monitoring, it is possible that more of these this work to mitigate bias will be done on the front end in the development of AI solutions, and thus the work of addressing bias may become less burdensome on individual providers over time.

Private and Industry Efforts on Transparency, Disclosure, and Evaluation. Transparency about AI models is critical to identifying the potential for bias, and deployers often do not receive standardized information on model development that is needed to meaningfully identify and address bias. An August 2022 survey by the Office of California Attorney General (AG) Rob Bonta found hospitals are using decision-making tools to make judgments about patients across many contexts, they rely on the vendor's assessment that the tools they use are ethical and unbiased, and they lack insight into vendors' data modeling. To address this lack of transparency and improve the effective and appropriate deployment of AI technology in health care, the

Coalition for Health AI (CHAI), a large national collaborative effort of health systems, public and private organizations, academia, patient advocacy groups, and AI experts, has released a draft template for an "applied model card" that is intended to be published by a developer and describe key information about health AI models, in a manner somewhat similar to a "Nutrition Facts" label. If the model card were widely adopted, a number of elements included therein would assist a deployer of an AI model in identifying, mitigating, and monitoring for bias.

Enforcement. DPH licenses and oversees clinics and health facilities. Other deployers to whom this bill applies, including physician's offices and offices of a group practice, are operated under the auspices of physician licensure and are not specifically regulated in California. Developers, similarly, are not subject to any specific state oversight regime that would monitor compliance. *This bill, however, requires developers to report compliance to DPH.*

According to the Author

This bill is a crucial step towards ensuring fairness in health care 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. The author states that 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, the author notes this bill will help promote safety, equity, and exceptional performance while protecting patients against avoidable harm. This bill was inspired by the 2023 California Reparations Task Force Report.

Arguments in Support

Oakland Privacy, a citizen's coalition that works regionally to defend the right to privacy and enhance public transparency, writes in support that it is sensible to place into California statute a corresponding version of federal regulations that they believe support California's goals and priorities. Oakland Privacy notes, given the interesting definitions of "DEI" and "discrimination" being promulgated by some federal agencies, it makes even more sense to capture the intent of federal regulations before they are twisted into shapes quite unlike their original intent. The California Medical Association also supports this bill.

Arguments in Opposition

TechNet argues in opposition that this bill defines a "developer" so broadly that it could encompass any company whose AI system might be used in a healthcare setting—regardless of whether the company designed or intended its system to support clinical decision-making or health-resource allocation. AdvaMed, Biocom, and California Life Sciences express similar concerns to several provisions of this bill, citing inadequate attention to existing regulatory frameworks. Several organization opposed a prior version of the bill, which required developers to submit their systems to third-party audits to demonstrate compliance. This audit requirement has been removed from the bill.

FISCAL COMMENTS

According to the Assembly Committee on Appropriations:

- 1) The California Department of Public Health (CDPH) estimates costs of approximately \$9.3 million in fiscal year (FY) 2026-27 and FY 2027-28 and \$9.1 million in FY 2028-29 to develop regulations and ensure compliance (Licensing and Certification Fund).
- 2) 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.
- 3) CDPH states it must (1) develop regulations defining "reasonable efforts" to mitigate risk for biased impacts in the AI system's outputs and (2) survey health facilities to ensure the facilities are making reasonable efforts to identify AI systems used to support clinical decision-making or health care resource allocation and to mitigate the risk for biased impacts in the system's outputs. CDPH estimates it will require 47 to 47.5 full-time equivalent (FTE) positions in FY 2026-27 through FY 2028-29, and 46.5 ongoing FTE positions starting FY 2029-30 to survey health facilities to ensure they are making reasonable efforts to identify AI systems used to support clinical decision-making or health care resource allocation and making reasonable efforts to mitigate the risk for biased impacts in the AI system's outputs.

VOTES

SENATE FLOOR: 38-0-2

YES: 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

ABS, ABST OR NV: Limón, Reyes

ASM HEALTH: 16-0-0

YES: Bonta, Chen, Addis, Aguiar-Curry, Caloza, Rogers, Flora, Mark González, Elhawary, Patel, Ellis, Celeste Rodriguez, Sanchez, Schiavo, Sharp-Collins, Stefani

ASM PRIVACY AND CONSUMER PROTECTION: 13-0-2

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

ABS, ABST OR NV: DeMaio, Macedo

ASM APPROPRIATIONS: 15-0-0

YES: Wicks, Sanchez, Arambula, Calderon, Caloza, Dixon, Elhawary, Fong, Mark González, Ahrens, Pacheco, Pellerin, Solache, Ta, Tangipa

UPDATED

VERSION: September 4, 2025

CONSULTANT: Lisa Murawski / HEALTH / (916) 319-2097

FN: 0001675