

Date of Hearing: April 23, 2024

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

AB 3030 (Calderon) – As Amended April 11, 2024

AS PROPOSED TO BE AMENDED

SUBJECT: Health care services: artificial intelligence

SYNOPSIS

Unlike artificial intelligence (AI), which has existed for more than 50 years, generative artificial intelligence (GenAI) is more or less a fundamentally new technology. As such, its quirks and limitations are still being actively explored. One such limitation – the capacity of GenAI systems to hallucinate, or create convincing but ultimately untrue “facts” out of whole cloth – can be amusing in certain situations and deadly serious in others. If a chatbot misspells “Berenstain Bears,” the hallucination is amusing. If a GenAI medical assistant directs an elderly patient to open their medicine cabinet and swallow thirty tablets of Tylenol, the hallucination becomes deadly serious. The latter situation has not yet occurred, but combining the known limitations of GenAI technology with its rapid adoption across industries will inevitably have unforeseen consequences.

This bill seeks to combat such high-risk scenarios. It would require health care providers to disclose their use of GenAI when communicating with patients, so that patients are able to apply an appropriate degree of skepticism to any instructions provided or information requested. Committee amendments would expand the types of communications affected to include all written or verbal communications, and would specify how disclosure should be provided in a variety of common communication scenarios. Committee amendments also adjust the definition of “GenAI,” insert definitions for “AI” and “health care provider,” and remove liability exemptions from the language of the bill.

This bill is author-sponsored and has no support or opposition.

SUMMARY: Requires health practitioners to disclose their use of GenAI to generate patient communications. Specifically, **this bill:**

- 1) Requires a health facility, clinic, physician’s office, or office of a group practice that uses GenAI to generate patient communications through an online interface, or via telephone, to ensure those communications include:
 - a) A disclaimer that indicates the communication was generated by AI. If the communication was generated by GenAI and reviewed by a human health care provider, the disclaimer may indicate this.
 - b) Clear instructions enabling a patient to navigate to the health entity’s internet website, or other platform, in order to communicate directly with a health care provider.

- 2) Prevents an entity or health care provider that fails to comply with the requirements of the bill from being subjected to any disciplinary actions related to licensure or certification solely because of their failure to comply.
- 3) States that existing Health and Safety Code Section 1290, which punishes as an infraction or misdemeanors specified violations, does not apply to a violation of this bill's provisions.
- 4) Defines "generative artificial intelligence tool" to mean a tool that uses machine learning systems capable of generating text, images, code or other types of content, often in response to a prompt entered by a user.

EXISTING LAW:

- 1) Establishes the California Department of Public Health which, among other functions, licenses and regulates health facilities. Defines "health facility" to mean a facility, place, or building that is organized, maintained, and operated for the diagnosis, care, prevention, and treatment of human illness, physical or mental, including convalescence and rehabilitation and including care during and after pregnancy, or for any one or more of these purposes, for one or more persons, to which the persons are admitted for a 24-hour stay or longer, including, but not limited to, hospitals, nursing facilities, and hospice facilities. (Health & Saf. Code § 1250.)
- 2) Establishes the Department of Consumer Affairs, which licenses and oversees health care professionals through various healing arts boards. (Bus. & Prof. Code § 100, Bus. & Prof. Code § 500 *et seq.*)
- 3) Outlines penalties resulting from various violations of the Health and Safety Code. (Health & Saf. Code § 1290.)
- 4) Defines "clinic" to mean an organized outpatient health facility that provides direct medical, surgical, dental, optometric, or podiatric advice, services, or treatment to patients who remain less than 24 hours, and that may also provide diagnostic or therapeutic services to patients in the home as an incident to care provided at the clinic facility. (Health & Saf. Code § 1200.)

FISCAL EFFECT: As currently in print, this bill is keyed fiscal.

COMMENTS:

1) **"But doctor, I am Pagliacci!"** The development of GenAI is creating exciting opportunities to grow California's economy and improve the lives of its residents. GenAI in particular can generate compelling text, videos, and audio in an instant – but with novel technologies come novel safety concerns. As GenAI-enabled products and systems become embedded into California's most private, sensitive spaces – healthcare facilities, for example – the Legislature may wish to carefully consider the various benefits and risks associated with these technologies. There are three primary risks related to the use of AI in healthcare:

1. **Bias:** broadly affects AI, including GenAI.
2. **Memory:** broadly affects AI, especially GenAI.
3. **Hallucination:** predominantly affects GenAI.

2) **What is GenAI?** AI 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: numbers, text, audio, video, or movement. 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. A hypothetical product that a healthcare provider might use to generate patient communications would be an example of GenAI.

Large language models. Large language models (LLMs) are a type of GenAI that have been specifically designed to understand, generate, and work with human language. These models are trained on vast quantities of text sourced from the internet and historical literature, including medical textbooks and other healthcare-related information. Chatbots that produce clear, coherent text or speech in response to user queries are considered to be LLMs.

2) **Risk 1: Bias.** There is a common saying in computer science: “garbage in, garbage out.” The performance of an AI tool is directly impacted by the quality, quantity, and relevance of the data used to train it. Deployed with care, AI should be able to reduce the bias introduced by humans into decision-making processes; if deployed thoughtlessly or maliciously, however, these same tools have the potential to amplify historical inequities and harm vulnerable populations.

AI and healthcare. When AI tools 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 algorithm 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. The algorithm considered various health relevant factors as it made its recommendation, 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 original algorithm was biased; it predicted Black and Latino people were less likely to have a successful vaginal birth after a cesarean than non-Hispanic white women. As a result, doctors performed more cesareans on Black and Latino people than on white people.¹ Such discrepancies tend to perpetuate historical biases – Black Americans, for example, have historically received a lower standard of healthcare than their white counterparts.²

Voice recognition systems. Beyond health outcomes, the very interactions between humans and GenAI can themselves be subject to bias. If a voice-recognition system is trained using audio data that excludes certain accents, it can fail to understand accented speech once deployed:

“Clow-dia,” I say once. Twice. A third time. Defeated, I say the Americanized version of my name: “Claw-dee-ah.” Finally, Siri recognizes it. Having to adapt our way of speaking to

¹ Caleb J. Colón-Rodríguez, “Shedding Light on Healthcare Algorithmic and Artificial Intelligence Bias,” *U.S. Department of Health & Human Services Office of Minority Health*, Jul. 12, 2023, minorityhealth.hhs.gov/news/shedding-light-healthcare-algorithmic-and-artificial-intelligence-bias.

² California Task Force to Study and Develop Reparation Proposals for African American, “Final Report,” California.

interact with speech recognition technologies is a familiar experience for people whose first language is not English or who do not have conventionally American-sounding names. I have even stopped using Siri because of it.³

A 2020 University of Michigan study examined the ability of several common automated speech recognition (ASR) systems to understand accents. It tested five state-of-the-art ASR systems – developed by Amazon, Apple, Google, IBM, and Microsoft – to transcribe structured interviews with 42 white speakers and 73 Black speakers. The study found that all five ASR systems exhibited substantial racial disparities.⁴

3) **Risk 2: Memory.** Just as humans cannot intentionally forget information they have learned, it is not currently possible to remove data from trained AI.⁵ Unlike an Excel spreadsheet, which stores data in neat columns, AI stores data in unpredictable and irreversible ways. The only foolproof way to remove data from an AI is to retrain it from scratch.

For a healthcare-related GenAI system to be useful, it will likely need to access patient information. Medical information is some of the most private, sensitive information to exist; coupled with a tool that possesses an elephantine memory, the widespread adoption of GenAI in healthcare threatens to be a privacy nightmare.

4) **Risk 3: Hallucination.** When an AI produces a result that is not grounded in reality, it is said to be “hallucinating.” Text-generators such as ChatGPT do not fundamentally understand the text they are producing. They calculate one word or symbol at a time – if they estimate that the next word/symbol to output should be a period, then the sentence ends. Otherwise, the sentence continues. It is a testament to the ingenious architecture of these systems that their outputs are remotely coherent. But while the text these systems produce is cogent, it is not always correct.

“These systems live in a world of language,” said Melanie Mitchell, an A.I. researcher at the Santa Fe Institute. “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.”⁶

This capacity to hallucinate is seriously concerning when an AI tool is embedded in a healthcare system. Doctor-patient trust is a critical component of the medical process; patients need to know that the precise, complicated sets of instructions they are often given by their doctors can be believed. Patients should know when a trusted healthcare provider is on the other end of a communication, and when a communication has instead been produced by GenAI.

Can a GenAI communication tool explain its reasoning? Exactly how GenAI tools arrive at conclusions and generate outputs is not always obvious. These systems are sometimes criticized

³ Claudia Lopez Lloreda, “Speech Recognition Tech Is Yet Another Example of Bias,” Jul. 5, 2020, www.scientificamerican.com/article/speech-recognition-tech-is-yet-another-example-of-bias/.

⁴ Allison Koenecke et al., “Racial disparities in automated speech recognition,” *PNAS*, vol. 117, no. 14, p. 7684-7689, Mar. 23, 2020.

⁵ Stephen Patis, “A.I.’s un-learning problem: Researchers say it’s virtually impossible to make an A.I. model ‘forget’ the things it learns from private user data,” *Yahoo! Finance*, Aug. 30, 2023, finance.yahoo.com/news/un-learning-problem-researchers-virtually-164342971.html.

⁶ Cade Metz, “What Makes A.I. Chatbots Go Wrong?,” *New York Times*, Mar. 29, 2023, www.nytimes.com/2023/03/29/technology/ai-chatbots-hallucinations.html

for being “black boxes” whose internal logic cannot be clearly explained. If a decision-making process cannot be clearly understood, can the decision reached be fully trusted? Experts are currently researching ways to solve this problem. In the meantime, policymakers should carefully consider how and when to limit the use of “black box” AI.

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

Across the state, pilot programs are testing the use of generative Artificial Intelligence (GenAI) as a tool to assist clinicians with patient communications. As AI becomes increasingly integrated in our healthcare systems, it is important to maintain the trust between a patient and their provider, while ensuring the accuracy of information being communicated to patients. AB 3030 would require healthcare providers who use this technology to provide a disclaimer that the communication was AI-generated, along with clear instructions for how a patient can directly communicate with a healthcare provider.

6) **What this bill would do.** This bill would require health care providers to disclose their use of GenAI when communicating with patients through online interfaces or over the phone, so that patients are able to apply an appropriate degree of skepticism to the instructions provided or information requested. Healthcare providers would be exempted from liability related to licensure solely on the basis of failing to disclose the use of GenAI in a patient communication.

7) **Analysis.** This bill adopts an efficient approach to regulating GenAI: it identifies a high-risk use case for the technology, and it outlines how the technology’s use should be adjusted in order to minimize risk. In collaboration with the author, this committee has proposed a series of amendments aimed at clarifying and strengthening the provisions of this bill.

Communication types. This bill addresses two types of communications between healthcare providers and patients: communications through online interfaces and telephone calls. These two examples are representative of broader categories, namely verbal and written communications. A proposed amendment would expand the types of communication this bill affects, as well as describe exactly how disclaimers should be provided in several common communication situations: occasional physical or digital communications (mail, email, messages in patient portals), continuous digital communications (chatbots), audio communications (telephone calls and voice messages), and visual communications (video telehealth).

Liability exemptions. As written, this bill exempts healthcare providers from consequences related to licensure or certification solely based on a failure to provide appropriate disclaimers. The bill also broadly exempts healthcare providers from penalties outlined in Health and Safety Code section 1290. It is not clear why these exemptions should exist for GenAI-related communication tools specifically. As outlined above, communications provided by GenAI systems should, in fact, be treated with more caution and care than communications written or dictated by human healthcare providers. Patients should be aware of the fact that a healthcare provider was not directly involved in crafting the directions they receive, or in requesting their personal medical information. Such exemptions, if enacted into law, would likely incentivize the adoption of GenAI technology at the cost of patient safety. Furthermore, these exemptions arguably undermine the bill’s goal of ensuring patients are aware when GenAI, rather than a medical professional, is the source of the information or instructions they are receiving. A proposed amendment would remove this exemption from liability.

Defining GenAI. Few serious attempts have been made to define GenAI for legislative purposes. The Biden White House’s AI executive order defines it thus:

“Generative AI” means the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content.

For the purposes of this bill, this Committee proposes a related definition:

“Generative artificial intelligence” means artificial intelligence that can generate derived synthetic content, including images, videos, audio, text, and other digital content.

This definition declines to mention that GenAI “emulates the structure and characteristics of input data.” It is true that GenAI emulates the structure and characteristics of *training data* – but with respect to input data, as the term is usually used, GenAI is able to produce fundamentally different outputs. For example, a GenAI system may accept text data as an input and produce visual data as an output.

Other definitions. Committee amendments propose to add definitions for “artificial intelligence” and “health care provider” into the bill. Various other minor clarifying changes have also been made, as reflected below.

- 8) **Proposed amendments.** The author has agreed to amendments, discussed above, that would revise and recast the bill in its entirety. The full draft text of the amended bill is reproduced below:

SECTION 1. Section 1316.9 is added to the Health and Safety Code, to read:

1316.9. (a) A health facility, clinic, physician’s office, or office of a group practice that uses generative artificial intelligence to generate written or verbal patient communications shall ensure that those communications include both of the following:

(1) A disclaimer that indicates to the patient that the communication was generated by generative artificial intelligence.

(i) For written communications involving physical and digital media, such as letters, emails, and other occasional messages, the disclaimer shall appear prominently at the start of each communication.

(ii) For written communications involving continuous online interactions, such as chat-based telehealth, the disclaimer shall be prominently displayed throughout the interaction.

(iii) For audio communications, the disclaimer shall be provided verbally at the start and end of the interaction.

(iv) For video communications, the disclaimer shall be prominently displayed throughout the interaction.

(2) Clear instructions permitting a patient to communicate with a human health care provider.

(b) If a written communication is generated by generative artificial intelligence and reviewed by a human health care provider, the disclaimer shall indicate that the communication was generated by generative artificial intelligence and reviewed by a human.

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

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

(2) “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.

(3) “Generative artificial intelligence” means artificial intelligence that can generate derived synthetic content, including images, videos, audio, text, and other digital content.

(4) “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 not-for-profit corporation licensed according to subdivision (a) of Section 1204.

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

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

9) **Related legislation.** AB 2811 (Lowenthal, 2024) would require attorneys to disclose their use of GenAI in the composition of court documents. This bill is pending in Assembly Judiciary Committee.

AB 2930 (Bauer-Kahan, 2024) would require deployers of AI-powered automated decision tools to disclose the use of these tools when making a consequential decision. This bill is pending in Assembly Judiciary Committee.

AB 3211 (Wicks, 2024) would require conversational AI systems to disclose their creation of synthetic content to users. This bill is pending in Assembly Judiciary Committee.

SB 942 (Becker, 2024) would require covered providers of generative AI systems to include visible disclosures in their outputs. This bill is pending in Senate Governmental Organization Committee.

SB 1154 (Hurtado, 2024) would require a person that uses a pricing algorithm to recommend or set a price or commercial term to clearly disclose that fact to the customer, current or prospective employee, or independent contractor. This bill is pending in Senate Judiciary Committee.

SB 1229 (Nguyen, 2024) would require a property and casualty insurer to disclose to an applicant or insured entity if it uses AI to make decisions related to applications for residential property insurance, automobile insurance, claims review, or other coverage of losses. This bill contains an urgency clause and is pending in Senate Insurance Committee.

REGISTERED SUPPORT / OPPOSITION:

Support

None on file.

Opposition

None on file.

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