

TRANSPARENCY

Telling patients a machine wrote it: the evidence on disclosure, and the laws that now require it

ABSTRACT

Practices that use language models to draft messages and notes, or to record visits, face two questions: whether to tell patients, and whether to ask them. This paper reviews the evidence and the statutes as of September 2026. Patients say they want notice and cannot reliably detect AI authorship themselves. A disclosure label lowers rated trust by a small, consistent margin, including when a physician is said to supervise. Fuller explanation lowers consent to ambient recording. The statutes diverge: California exempts clinician-reviewed communications, Texas requires disclosure regardless of review, Illinois and Maine require written consent in mental health, and Louisiana requires spoken notice before AI transcription. This review found no Tennessee provider disclosure statute.

Type Evidence and regulatory review **References** 26 **Reading time** 14 min

Last reviewed September 2026

1 Notice and consent are different questions

A practice that lets a language model draft a portal reply, summarize a visit or write a note faces two separate questions about the patient. The first is notice: whether the patient is told that software produced some of what they read, or of what was written about them. The second is consent: whether the patient is asked before software listens to the visit. Different laws and different studies govern each.

The short answer, as of September 2026, has four parts. Patients say they want to be told: in a national probability sample, 62.7% called it “very true” that being notified mattered to them.¹ Telling them has a cost that is real, small and consistent (Cohen’s *d* of 0.26 to 0.29 in one experiment), and the cost does not disappear when the label says a physician supervised the machine.^{2,3} Consent to recording falls when patients are given a fuller explanation.⁴ And the statutes that now require disclosure disagree about what triggers it; California’s exempts any AI-generated communication a licensed clinician has read and reviewed.⁵

For a Tennessee practice the statutory layer is thin: within the limits of the search described in section 7, this review found no Tennessee statute requiring a clinician to tell a patient that AI was used. Tennessee’s wiretap statute permits a party to a conversation to record it,⁶ and a 2026 statute bars representing an AI

system as a qualified mental health professional.⁷ What follows reports what the laws say; it is not legal advice.

2 What patients say they want, and what they were actually asked

A widely repeated figure in this area is Pew Research Center's: 60% of US adults said they would be uncomfortable if their own health care provider relied on artificial intelligence, against 39% comfortable.⁸ It says nothing about disclosure. The survey of 11,004 adults was fielded December 12–18, 2022, and asked about a provider relying on AI “to do things like diagnose disease and recommend treatments.” The report includes no question about drafted messages or notes, or about whether respondents wanted to be told.

The same report shows how much the headline averages over. For their own care, 65% would want AI used in skin cancer screening, while 31% would want it guiding pain management after surgery. Among those who had heard a lot about AI, comfort split 50–50; among those who had heard a little or nothing, 63% and 70% were uncomfortable.⁸ The 60% is an attitude toward delegated clinical judgment that varied with task and familiarity, measured once, in late 2022.

The question this paper needs was asked directly by Platt and colleagues in NORC's probability-based AmeriSpeak panel: 2,021 adults, June 27 to July 17, 2023.¹ Asked how true it was that being notified about the use of AI in their health care was important to them, 62.7% answered “very true” and 4.8% “not at all true”; the mean was 3.39 on a four-point scale (95% CI 3.33 to 3.44), rising with age from 3.14 at 18 to 29 years to 3.57 at 60 and older, and with education. The authors conclude that the question is not whether to notify patients but when and how.

Nor can notice be left to inference. Nov and colleagues showed 392 respondents ten patient questions answered by a provider or a chatbot, telling them half came from each.⁹ Chatbot responses were identified correctly 65.5% of the time (1,284 of 1,960) and provider responses 65.1% (1,276 of 1,960): wrong about one time in three.

3 Telling patients costs a little trust, and supervision does not buy it back

62.7%

of US adults: being notified of AI use is “very true” for them

0.13 pts

lower satisfaction, 5-point scale, with an AI rather than human disclosure

81.6% → 55.3%

ambient-recording consent, basic versus fuller explanation

Three experiments vary what the reader is told while holding the rest fixed. Reis and colleagues ran two preregistered online experiments with 2,280 participants, attributing identical medical advice to a human physician, to AI, or to a physician working with AI.² Both AI labels were rated less reliable and less empathetic, with Cohen's d of 0.26 to 0.29 in the first study and at least 0.21 in the second, where willingness to follow the advice also fell. Comprehensibility did not differ. The "human + AI" label, which describes supervised drafting, was penalized about as much as AI alone. The one behavioral measure, saving a link in the second study, did not differ between the human and AI labels (19.3% against 18.5%; $p = .789$; 22.9% for human + AI).

The same group then showed 1,276 US adults, sampled to census quotas in January 2025, a family physician's advertisement that did or did not mention AI for administrative, diagnostic or therapeutic purposes.³ Every AI mention lowered ratings of competence, trustworthiness and empathy. Willingness to book, on a five-point scale, was 3.61 without AI and 3.32, 3.16 and 3.15 with administrative, diagnostic and therapeutic use; effect sizes reached $d = 0.41$. Even administrative use, the category nearest message drafting, was penalized.

Cavalier and colleagues tested drafted messages directly.¹⁰ Of 2,511 members of Duke University Health System's patient advisory committee, 1,455 (57.9%) responded in late 2023, older and more educated than non-respondents. The AI drafts came from GPT-3.5 and were reviewed by study physicians, who made minimal changes. Participants slightly preferred them to human-written replies: by 0.30 points for satisfaction, 0.28 for usefulness and 0.43 for feeling cared for, on five-point scales. Disclosure moved the other way. Satisfaction was 0.13 points lower with an AI disclosure than a human one (95% CI 0.05 to 0.22) and 0.09 points lower than with none (95% CI 0.01 to 0.17). Whatever the author or disclosure, more than 75% were satisfied. In a follow-up answered by 893 participants, the most-preferred AI disclosure, chosen by 33%, was "This message was written by Dr. T. with the support of automated tools." The authors concluded that disclosure should be maintained anyway.

In 40 interviews, Owens and colleagues found that patients valued clinician oversight of AI-drafted portal messages and wanted assurance that clinicians remained accountable; most endorsed disclosure, and they differed on its timing and format.¹¹ The base is narrow: of 23 studies in a 2025 systematic review of drafted patient messages, only two focused on patients or laypeople rather than clinicians, the Duke survey and the human-versus-AI identification test.¹²

One result matters more than the others: a supervised-AI label was discounted like an unsupervised one.² Patients do not appear to treat a clinician's review as making the machine's involvement irrelevant, which is the premise of the California exemption in section 6.

WHAT THESE STUDIES DO NOT ESTABLISH

Every result here is a rating from a survey, a vignette or a mock advertisement, not behavior inside a care relationship. No study located for this review measured what disclosure does in a live deployment to portal use, complaints, adherence or outcomes, and the single behavioral measure did not move. The samples were an advisory panel and online panels, not clinic populations. The direction of the effect is consistent; its consequence for care is unmeasured.

4 Consent to be recorded falls as patients learn more

The evidence on the second question is thinner still. Lawrence and colleagues evaluated consent among 121 pilot users of ambient documentation at a New York academic center between March and December 2024: 18 clinicians and 103 patients, recruited through an online research platform from among those who reported experiencing the tool.⁴ Consent was most often a verbal conversation before the encounter, varying with time, knowledge and the relationship. Of the patients, 74.8% were comfortable or very comfortable with the technology. Given basic information, 84 (81.6%) consented. Given details of its AI features, data storage and corporate involvement, 57 (55.3%) did.

That fall is the finding a consent process has to face: consent obtained with less explanation is higher and less informed, so a consent rate means little without the script that produced it. Topic mattered too: patients said they would be more likely to self-censor about mental health (35.0%), sexual health (40.8%) and illicit activity (51.5%). For errors in the note, 64.1% held the physician accountable; for data breaches, 76.7% held the vendor responsible.

5 Recording law answers a different question

Recording statutes answer a narrower question than ambient AI raises: whether a recording is an unlawful interception. Federal law permits a person not acting under color of law to intercept a communication to which that person is a party, or where one party has given prior consent, unless the purpose is criminal or tortious.¹³ Tennessee's statute is written the same way.⁶ California makes it an offense to record a confidential communication "without the consent of all parties."¹⁴ A one-party rule settles whether the recording is a crime, not whether the patient was informed or a vendor may process what was recorded.

The first lawsuits sit in that gap. In December 2025 KPBS reported a proposed class action in San Diego Superior Court alleging that Sharp HealthCare recorded visits with an ambient AI tool, that more than 100,000 patients may have been recorded, and that the tool inserted chart statements that patients had been advised of the recording and had consented when, the plaintiff says, they had not.¹⁵ These are pleading allegations reported by journalists, not findings; Sharp declined to comment.

Statutes aimed at AI recording

Louisiana's Act 649, effective August 1, 2026, is summarized on the legislature's own bill page as requiring a provider "to obtain a patient's consent prior to recording a medical visit."¹⁶ The enacted text requires something else: a licensed professional "shall verbally disclose the use of any recording device, software, or service to a patient before recording any part of an appointment or treatment to be transcribed by artificial intelligence," subject to possible board discipline, with immunity from civil liability absent gross negligence or willful misconduct. It is a notice statute, not a consent statute.

Illinois and Maine went further, for mental health only. Illinois's Wellness and Oversight for Psychological Resources Act, effective on signature in August 2025, bars AI from therapeutic decision-making and limits licensed professionals to administrative and supplementary uses, with civil penalties of up to \$10,000 per violation; where a session is recorded or transcribed, the client must be told in writing that AI will be used and for what purpose, and must consent.¹⁷ Maine's Public Law chapter 687, approved April 13, 2026, follows Illinois's structure and its definition of consent as an express, specific, written and revocable agreement,¹⁷ and adds three things.¹⁸ The notice must explain how session data "will be stored, retained, used for training and deleted upon termination" of services. Consent may also be given by initialing a specific section of a general consent-to-treatment agreement. And a licensee may not refuse services solely because a client declined AI.

6 The disclosure statutes disagree about the trigger

California: disclosure unless reviewed

AB 3030, operative January 1, 2025 at Health and Safety Code § 1339.75, applies to a health facility, clinic, physician's office or group practice office that uses generative AI to produce written or verbal communications to patients about their clinical information.⁵ These must carry a disclaimer that AI generated them, placed by medium (at the beginning of a letter or email, throughout a chat or video, spoken at the start and end of audio), with instructions for reaching a human. Scheduling, billing and other administrative matters are excluded; physicians answer to the Medical Board of California or the Osteopathic Medical Board. None of it applies where the communication is "read and reviewed by a human licensed or certified health care provider."

The statute therefore reaches unreviewed output, such as an automated chat or an unattended voice line, and not a draft a clinician reviews before it is sent. It treats review as a substitute for disclosure. The evidence does not support that as a statement about patients: a physician-plus-AI label was discounted like an AI label.² The exemption is a judgment that review is the safeguard, not a finding about what patients want.

Texas: disclosure regardless of review

Texas has two overlapping duties and no review exemption. Senate Bill 1188, effective September 1, 2025, permits a practitioner to use AI “for diagnostic purposes,” including recommendations on a diagnosis or course of treatment, on conditions that include acting within the scope of licensure and reviewing all records created with AI consistently with Texas Medical Board standards. Such a practitioner “must disclose the practitioner’s use of that technology to the practitioner’s patients.”¹⁹ Review is a condition of use, not a substitute for telling, and the section sets no form or timing.

House Bill 149, effective January 1, 2026, supplies both, with a broader trigger. If an AI system “is used in relation to health care service or treatment,” the provider must disclose it to the recipient or a personal representative no later than the date the service is first provided, or as soon as reasonably possible in an emergency.²⁰ It must be clear, conspicuous, in plain language and free of dark patterns, and it is required even where AI involvement would be obvious to a reasonable consumer. The attorney general enforces the chapter and may not sue until 60 days after giving notice.

Utah: disclosure when the patient interacts with the machine

Utah’s 2024 act required anyone providing the services of a regulated occupation to “prominently disclose when a person is interacting with a generative artificial intelligence in the provision of regulated services,” spoken at the start of an oral exchange and sent before a written one.²¹ In 2025 that section was replaced by § 13-75-103, which keeps the duty and timing but confines them to a “high-risk artificial intelligence interaction,” defined to include an interaction that collects health data or provides medical or mental health advice.²² In health care that narrowing removes little. The framing is the patient interacting with generative AI, which describes a chatbot or voice agent more plainly than a clinician’s drafting tool.

Rules against impersonation are not disclosure rules

California’s AB 489 (Stats. 2025, ch. 615, operative January 1, 2026) makes it unlawful to use AI to provide health information, advice or services while representing that a licensed professional is providing them.²³ Tennessee’s Public Chapter 647, effective July 1, 2026, bars developers and deployers from representing an AI system as being, or able to act as, a qualified mental health professional, with a \$5,000 civil penalty per violation under the Consumer Protection Act.⁷ Both forbid a false claim; neither requires a practice to volunteer that it uses AI.

TABLE 1

State duties to disclose AI use or AI recording to patients, and adjacent rules, as of September 2026.

Instrument	Trigger	What is required	Reviewed output exempt?	In effect
California H&S Code § 1339.75 ⁵	Generative AI produces patient communications about clinical information	Disclaimer placed by medium; route to a human	Yes	Jan. 1, 2025
Texas H&S Code § 183.005 ¹⁹	AI used for diagnostic purposes	Disclose to patients; form and timing unspecified	No; review is a condition of use	Sept. 1, 2025
Texas Bus. & Com. Code § 552.051(f) ²⁰	AI used in relation to a health care service or treatment	Clear, plain-language notice by first date of service	No	Jan. 1, 2026
Utah Code § 13-75-103 ²²	Patient interacting with generative AI in a high-risk interaction	Prominent notice, spoken at start or written before	Not addressed	May 7, 2025
Illinois 225 ILCS 155 ¹⁷	AI supplementary support where a therapy session is recorded or transcribed	Written notice of use and purpose; consent	No	Aug. 2025
Maine P.L. ch. 687 ¹⁸	The same, for mental health licensees	Notice including data retention and training use; written, revocable consent	No	Approved Apr. 13, 2026; effective date not confirmed
Louisiana R.S. 37:22.1 ¹⁶	Recording to be transcribed by AI	Verbal disclosure before recording	No	Aug. 1, 2026
Tennessee § 39-13-601(b)(5); Pub. Ch. 647 ^{6,7}	Recording; AI presented as a mental health professional	One party's consent suffices; impersonation barred; no disclosure duty found	Not applicable	In force; July 1, 2026

7 Professional guidance asks for more than the statutes do

The American Medical Association's principles, dated November 2024, scale disclosure to the risk of harm.²⁴ AI that affects access to care or point-of-care decisions should be disclosed and documented, with an opportunity to request review by a licensed clinician; AI that directly affects care or the record should be documented in the record; patients engaging directly with AI should be told at the start. And where patient-facing content is generated by AI, the document says, that use should be disclosed or otherwise noted within the content itself. That principle carries no review exemption, and the same document says AI may not generate records or communications on a physician's behalf without that physician's consent and final review. On the AMA's account, review and disclosure are both expected, not alternatives.

The Federation of State Medical Boards adopted policy in April 2024 that physicians should "disclose to patients when and how AI is used in their care," and that they retain the duty to review records created with AI and remain accountable for resulting harms.²⁵ Neither document is law. Both describe a

standard a board could apply to a complaint, and both ask for more than California's statute.

The federal posture points the other way. Executive Order 14365, signed December 11, 2025, directed the Attorney General to form an AI Litigation Task Force to challenge state AI laws, told the Secretary of Commerce to identify state laws that "may compel AI developers or deployers to disclose or report information in a manner that would violate the First Amendment," and asked the Federal Communications Commission to consider a preemptive federal disclosure standard.²⁶ An executive order directs agencies; it does not itself repeal a state statute.

STATUS, AND WHAT THIS REVIEW DID NOT CHECK

Enactment, effective dates and text were checked against legislative and agency sources where retrievable; otherwise the reference list says what was read instead. This review did not survey litigation. The Tennessee negative rests on a bounded search: the provisions cited here and a law-firm survey of 2026 state enactments, not the whole Tennessee Code or board rules. Because the December 2025 order directs federal challenges to state AI laws, the status of any statute here should be re-verified before reliance.

8 Properties a system needs to satisfy all of them

1. **Provenance recorded per artifact.** California's duty turns on review, Texas's on the purpose of AI use, Utah's on whether the patient interacted with generative AI.^{5,19,20,22} A system that cannot say, for a given message or note, whether generative AI produced text in it, who reviewed it and when, cannot show which rule applied.
2. **Disclosure as a property of the channel.** One statute places the disclaimer differently for a letter, a chat and a call,⁵ and the jurisdiction decides whether any is needed.
3. **Wording that is true.** The disclosure Duke respondents most often preferred (33%) attributes authorship to the clinician "with the support of automated tools."¹⁰ It is accurate only when the clinician reviewed the draft; attached to unreviewed output it becomes a false statement.
4. **Consent to record as its own event, entered by a person.** It needs a time, an actor, the information given and a revocation path; in mental health, Maine requires consent to be written and revocable.¹⁸ The Sharp allegations describe the failure mode: consent text produced by the tool it authorizes.¹⁵ Because consent falls as information rises, the script belongs with the consent.⁴

5. **Refusal that leaves care intact.** Maine bars refusing services to a client who declines AI.¹⁸ Elsewhere the evidence points the same way: more than four in ten fully informed patients did not consent,⁴ so working without the tool has to be ordinary.
6. **AI involvement in the record.** The AMA asks for documentation in the medical record where AI directly affects care or the record itself.²⁴

9 What survives, and what would settle it

Five statements are supported. Patients say they want to be told: 62.7% in a national probability sample.¹ They cannot reliably tell for themselves.⁹ Telling them lowers rated trust by a small, consistent margin, and a supervised-AI label does not avoid it.^{2,3,10} Fuller explanation lowers consent to ambient recording.⁴ And the law is a patchwork: California exempts reviewed communications, Texas requires disclosure regardless of review, Utah reaches patient-facing interaction, Illinois and Maine require written consent in mental health, Louisiana requires spoken notice before AI transcription, and this review found no Tennessee statute requiring any of them.^{5,16,17,18,19,20,22}

What is not known is whether disclosure changes anything patients do. Every effect here is a rating, and the one behavioral measure did not move.² The study that would settle it is a randomized comparison of disclosure formats inside a live deployment, measuring portal engagement, complaints, opt-outs and adherence rather than satisfaction. Until then, the case for telling patients rests on what they say they want and what the professional bodies ask; the case against rests on small rating penalties: 0.09 to 0.13 points on a five-point scale for drafted messages, and standardized differences of about 0.2 to 0.4 in vignettes.^{2,3,10}

REFERENCES

Entries 1 to 4 and 8 to 12 are the empirical studies: entries 2, 3 and 10 are the experiments that carry section 3, entry 4 is the only study of ambient-recording consent this review located, and entry 12 is a systematic review used to show how narrow the base is. All are surveys, experiments on online panels or qualitative studies; none measures behavior in care. Entries 5 to 7, 13, 14 and 16 to 23 are statutes, cited for what they say rather than as evidence. Entries 5, 6, 7, 14, 17 and 23 were not read on an official legislature site: entry 5 was confirmed against the Medical Board of California's summary and a reproduction of the code, entries 6, 14 and 17 were read from reproductions of the code or enrolled text, entry 7 from a copy of the chaptered act, and entry 23 through a state licensing board's advisory; entry 17's public act number comes from a legislative tracker. Entry 15 is journalism about pleading allegations, entries 24 and 25 are professional policy without legal force, and entry 26 is an executive order.

1. Platt J, Nong P, Carmona G, et al. *Public Attitudes Toward Notification of Use of Artificial Intelligence in Health Care*. JAMA Network Open. 2024;7(12):e2450102. doi:10.1001/jamanetworkopen.2024.50102 NATIONAL SURVEY
2. Reis M, Reis F, Kunde W. *Influence of believed AI involvement on the perception of digital medical advice*. Nature Medicine. 2024;30(11):3098–3100. doi:10.1038/s41591-024-03180-7 RANDOMIZED EXPERIMENT
3. Reis M, Reis F, Kunde W. *Public Perception of Physicians Who Use Artificial Intelligence*. JAMA Network Open. 2025;8(7):e2521643. doi:10.1001/jamanetworkopen.2025.21643 RANDOMIZED EXPERIMENT
4. Lawrence K, Kuram VS, Levine DL, et al. *Informed Consent for Ambient Documentation Using Generative AI in Ambulatory Care*. JAMA Network Open. 2025;8(7):e2522400. doi:10.1001/jamanetworkopen.2025.22400 QUALITATIVE STUDY
5. California Legislature. *Assembly Bill No. 3030, Health care services: artificial intelligence*. Stats. 2024, ch. 848; operative Jan. 1, 2025; codified at Cal. Health & Safety Code § 1339.75. [leginfo.legislature.ca.gov](https://leginfo.ca.gov); summarized in Medical Board of California, *GenAI Notification Requirements*, mbc.ca.gov STATUTE
6. Tennessee General Assembly. *Interception of wire, oral or electronic communications*. Tenn. Code Ann. § 39-13-601(b)(5); text read from the 2024 code as reproduced at law.justia.com. STATUTE
7. Tennessee General Assembly. *Public Chapter No. 647* (S.B. 1580, 114th General Assembly), relative to mental health; approved Apr. 1, 2026; effective July 1, 2026; adding T.C.A. § 33-1-205 and amending § 47-18-104(b). publications.tnsosfiles.com; text read from the chaptered copy at legiscan.com because the official file refused automated access. STATUTE
8. Tyson A, Pasquini G, Spencer A, Funk C. *60% of Americans Would Be Uncomfortable With Provider Relying on AI in Their Own Health Care*. Pew Research Center. Feb. 22, 2023. pewresearch.org NATIONAL SURVEY
9. Nov O, Singh N, Mann D. *Putting ChatGPT's Medical Advice to the (Turing) Test: Survey Study*. JMIR Medical Education. 2023;9:e46939. doi:10.2196/46939 SURVEY
10. Cavalier JS, Goldstein BA, Ravitsky V, et al. *Ethics in Patient Preferences for Artificial Intelligence–Drafted Responses to Electronic Messages*. JAMA Network Open. 2025;8(3):e250449. doi:10.1001/jamanetworkopen.2025.0449 SURVEY
11. Owens K, Jayaram A, Chowdhury A, et al. *Patient Perspectives on AI-Drafted Electronic Portal Messages*. JAMA Network Open. 2026;9(7):e2622463. doi:10.1001/jamanetworkopen.2026.22463 QUALITATIVE STUDY

12. Hu D, Guo Y, Zhou Y, et al. *A systematic review of early evidence on generative AI for drafting responses to patient messages*. npj Health Systems. 2025;2:27. doi:10.1038/s44401-025-00032-5 SYSTEMATIC REVIEW
13. United States Congress. *Interception and disclosure of wire, oral, or electronic communications prohibited*. 18 U.S.C. § 2511(2)(d). uscode.house.gov STATUTE
14. California Legislature. *Eavesdropping on or recording confidential communications*. Cal. Penal Code § 632; text read from the 2025 code as reproduced at law.justia.com because the official site refused automated access. STATUTE
15. KPBS. *Lawsuit claims Sharp HealthCare secretly recorded exam room conversations without patient consent*. Dec. 11, 2025. Reporting on a proposed class action in San Diego Superior Court; allegations, not findings. kpbs.org JOURNALISM
16. Louisiana Legislature. *HB 475, 2026 Regular Session*. Act No. 649; signed June 2, 2026; effective Aug. 1, 2026; enacting La. R.S. 37:22.1, "Recordings; artificial intelligence; disclosure." legis.la.gov STATUTE
17. Illinois General Assembly. *Wellness and Oversight for Psychological Resources Act*. HB 1806, 104th General Assembly; Public Act 104-0054; 225 ILCS 155; effective on signature, August 2025. Illinois Department of Financial and Professional Regulation, *Gov. Pritzker Signs Legislation Prohibiting AI Therapy in Illinois*, Aug. 4, 2025, idfpr.illinois.gov; statute text read from reproductions at law.justia.com because ilga.gov could not be reached. STATUTE
18. Maine Legislature. *An Act to Regulate the Use of Artificial Intelligence in Providing Certain Mental Health Services*. LD 2082 (HP 1397), 132nd Legislature; Public Law chapter 687; approved Apr. 13, 2026. legislature.maine.gov STATUTE
19. Texas Legislature. *Senate Bill 1188*. 89th Leg., R.S. (2025); adding Tex. Health & Safety Code ch. 183, including § 183.005; effective Sept. 1, 2025. capitol.texas.gov STATUTE
20. Texas Legislature. *Texas Responsible Artificial Intelligence Governance Act*. H.B. 149, 89th Leg., R.S. (2025); Tex. Bus. & Com. Code § 552.051; effective Jan. 1, 2026. capitol.texas.gov STATUTE
21. Utah Legislature. *Artificial Intelligence Amendments*. S.B. 149 (2024 Gen. Sess.); enacting Utah Code § 13-2-12; effective May 1, 2024. le.utah.gov STATUTE
22. Utah Legislature. *Artificial Intelligence Consumer Protection Amendments*. S.B. 226 (2025 Gen. Sess.); enacting Utah Code §§ 13-75-101 to 13-75-106 and repealing § 13-2-12; effective May 7, 2025. le.utah.gov STATUTE

- 23.** California Legislature. *Assembly Bill No. 489*. Stats. 2025, ch. 615; operative Jan. 1, 2026; as summarized in California Board of Psychology, *AB 489 legislative advisory*. psychology.ca.gov
STATUTE
- 24.** American Medical Association. *Augmented Intelligence Development, Deployment, and Use in Health Care*. November 2024. ama-assn.org **POSITION PAPER**
- 25.** Federation of State Medical Boards. *Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice*. Adopted by the FSMB House of Delegates, April 2024. fsmb.org **GUIDANCE**
- 26.** Executive Office of the President. *Ensuring a National Policy Framework for Artificial Intelligence*. Executive Order 14365; signed Dec. 11, 2025; 90 Fed. Reg. 58499 (Dec. 16, 2025); Doc. No. 2025-23092. federalregister.gov **EXECUTIVE ORDER**

RELATED PAPERS

Regulating AI in certified EHR technology

The federal regimes, and the state layer that binds vendors rather than clinicians

AI-drafted replies to patient messages

What deployments of drafted portal replies measured about time, burden and safety

Ambient documentation and the clinician's day

The measured effect of ambient scribes, at the size the trials support