

A WYOMING TELEHEALTH NETWORK WEBINAR

AI-Powered Note-Taking in Clinical Practice

Consent, privacy, and what a safe workflow actually looks like

Dr. Gigi Magan, MD MPH

Family Medicine · September 21, 2026



Disclosures



Financial disclosures

- No financial relationships with any ambient AI scribe vendor.
- No conflicts of interest
- I speak today in a personal capacity - my views, not those of my employer or any health center.



About this talk

- Product names appear for educational comparison only; no endorsement is implied.
- Legal content is educational, not legal advice - implementation decisions belong with your compliance and counsel.
- Evidence current as of September 2026 - plus de-identified field notes from my own daily ambient-scribe use in bilingual primary care.
- Clinical field notes are de-identified composites from my own practice; details are changed and no individual patient is described.

What Wyoming telehealth providers asked for

Wyoming Telehealth Network survey - 55 responses from telehealth providers across Wyoming, June–July 2026 [43]

60%

already use AI in their work or practice
(31 of the 52 who answered)

[43]

4 in 10

written responses raised privacy,
security, ethics, legal, or consent
questions

19 of 47 [43]

1 in 4

asked about AI for notes or
documentation

11 of 47 [43]

In your words: “what is allowed and what is not” · “help, or just create note bloat?” · “patients that want to opt out” · “I don’t know even the basics... I would like to start there.”

If you remember nothing else...

Real

The benefits are real: burnout drops and attention returns to the patient - but measured time savings are modest.

Part 2 · the evidence [12][13][14]

Uneven

The risks are unevenly distributed: errors concentrate in accented speech, multilingual visits, and poor audio - telehealth included.

Part 3 · where it breaks [11][20]

Fixable


Consent, pause workflows, and contract terms are where clinicians still hold the pen.

Parts 4–6 · consent and your workflow [22][24]

Where we're going

- 1 How AI note-taking works** - the basics, and why it is not dictation ~5 min
- 2 Does it actually help?** - time, burnout, and the note-bloat question ~5 min
- 3 Where it breaks** - errors, review, and the telehealth visit ~8 min
- 4 Consent & privacy** - Wyoming law, HIPAA, and where the audio goes ~8 min
- 5 Behavioral health** - why the most sensitive visits need different rules ~5 min
- 6 Monday morning** - vendor questions, one checklist, then your questions ~5 min

The pipeline: from conversation to chart



The draft is not the record.

- Steps 1–3 usually run on a third-party vendor's cloud - audio and transcript leave your walls [6][7].



The signed note is yours.

- Accountability concentrates at step 5.

This is not dictation



Dictation / transcription

- Maps your words to text, one-to-one
- Errors are mishearings - they look wrong and are easy to spot
- The output is a record of what you said



Generative ambient AI

- Writes the note word-by-word by predicting what most plausibly comes next, from training on vast clinical text [39]
- Errors are plausible inventions - they look right and hide in plain sight
- The output is a reconstruction of what typically happens in visits like yours



Why this explains “I never said that, but it added it”

- The model fills gaps with what usually follows - a “normal exam,” a standard counseling line - because statistically, it usually does [11][39]. Hallucination is a property of the method, not a glitch.

Two very different things get called “AI note-taking”



A general-purpose chatbot

- ChatGPT, Gemini, Copilot: built for everything, not for the chart
- Consumer accounts come with no BAA - patient details typed or pasted in can be an impermissible disclosure of PHI [23]
- Removing the name is not de-identification: HIPAA lists 18 identifiers [47]
- Enterprise versions may offer a BAA - ask, and get it in writing



A clinical AI scribe

- Built for one job: visit audio in, draft note out
- The vendor signs a BAA and works inside or alongside your record
- Retention and model-training terms live in the contract
- Still needs the patient’s consent - and your review

The test is not the brand. It is two questions: is there a BAA, and what does the contract say happens to the data?

The evidence base at a glance (September 2026)

Randomized trials 3 published, each in one academic system [12][44][45]

Large multisite observational 8,581 clinicians, 5 academic centers, difference-in-differences [13]

Cohort / QI studies Kaiser NorCal, Penn, 6-system burnout study; 13-mo ED cohort [7][14][15][38]

Qualitative & note-quality Interview studies; blinded note-quality comparisons [10][16][18][19]

Gaps

- No patient-outcome trials
- Follow-up mostly 30–90 days
- Little from rural or safety-net settings

Most studies measure process (minutes, burnout scores) - not diagnostic accuracy, treatment adherence, or patient outcomes.

The biggest study to date: modest and uneven

Rotenstein et al., JAMA 2026 - 8,581 clinicians (1,809 adopters), 5 academic systems, difference-in-differences

-13.4

min total EHR time / 8
scheduled pt-hrs

95% CI 9.1–17.7

-16.0

min documentation time /
8 pt-hrs

95% CI 13.7–18.3

+0.49

visits per week

95% CI 0.17–0.81

≈0

change in after-hours EHR
time

not significant

Largest gains: primary care, advanced practice clinicians, female clinicians, and clinicians using the scribe in ≥50% of visits. Time saved was partly reinvested in more visits - efficiency for the system is not automatically relief for the clinician.

Help, or note bloat? AI drafts vs human notes



Palm et al. 2025 - 97 visits; authors work for a scribe vendor

- Overall quality close: 4.20 vs 4.25 / 5 (physician comparison notes slightly higher, $P=.04$) [18]
- **Hallucinations: 31% of AI notes vs 20% of physician notes written from the same recordings ($P=.01$) [18]**
- AI notes: more thorough & organized, less succinct; reviewers still preferred them, 47% vs 39%



Reddy et al. 2026 - blinded, VHA

- 11 AI tools vs 18 human note-writers, 30 blinded raters, 5 standardized cases [19]
- **Humans scored higher in all 10 quality domains [19]**
- Biggest AI gaps: thoroughness, organization, usefulness. Humans wrote in simulated conditions, which may favor them [19]

Four ways an AI note goes wrong



Hallucination

- Content that was never said: exam findings not performed, medications never discussed, instructions never given.



Omission

- Clinically relevant content silently dropped - arguably more dangerous, because absence is invisible on review.



Misattribution

- Patient's words recorded as clinician's (or vice versa): "I think it's my heart" becomes an assessment.



Contextual misinterpretation

- Right words, wrong meaning - sarcasm, rule-outs, hypotheticals, family history mistaken for the patient's.

Taxonomy: Topaz et al. 2025 [11] · hallucinations in 31% of ambient notes [18] · 5-platform bench test: 26.3% of key elements missed or wrong, ~3 potentially harmful errors per case [40]

Four field notes from my own exam room



The invented medication

- I said “Macrobid”; the draft said “migrabit” - a medication that does not exist. Plausible-looking errors sail straight through a tired skim.



The phantom tortillas

- Dietary counseling about tortillas appeared in a visit where no one said the word - the model added culturally “expected” content.



The invisible exam

- The patient points to where it hurts; the microphone hears nothing. Unverbalized findings vanish - or come back as a “normal exam.”



The too-tidy note

- Drafts read more “resolved” than the visit felt - hedges become clean plans. And a deskilling worry: when the note writes itself, some of the thinking that writing forced is lost.

Personal, de-identified composites from my own bilingual primary-care practice - illustrations of the published taxonomy [11], details changed; not study data.

The signature is the safety system



Why clinicians edit AI drafts

- Accuracy & specialty-specific precision
- Billing, coding & documentation standards
- Medico-legal protection
- Personal style & voice
- 30-clinician qualitative study, UC Irvine [16]




The automation-bias risk

- Human-factors research: trust in automation grows with exposure; verification effort decays
- Editing is unpaid, invisible work - the incentive is to skim
- Not yet directly measured in scribes - treat as a predictable risk to design against, not a proven effect

Design implication: build review time into templates and audit signed notes - don't rely on vigilance.

The microphone is not colorblind



Average word error rate across 5 commercial ASR systems (Apple, Amazon, Google, IBM, Microsoft) [20]



Koenecke et al., PNAS 2020

- Error rates nearly 2× higher for Black speakers (0.35 vs 0.19)
- >20% of Black speakers' audio effectively unusable (WER>0.5) vs <2% for white speakers [20]



In the clinic, today

- Safety-net preprint, 54,160 visits: non-English visits had 21–25% lower odds of a draft kept mostly as written; a bilingual clinician did not close the gap [50]
- Not yet peer reviewed; vendor-defined metric that does not measure accuracy. Independent, stratified testing is still rare [10][21]

The telehealth visit: a six-move playbook



1 • Say who is on the call

“Is anyone else in the room with you?” The AI cannot tell a patient from a spouse off camera - and everyone it hears deserves to know it is listening.



2 • Narrate what you see

The AI hears the visit; it does not watch the video. Say your observations aloud - affect, gait, the rash on camera - or they never existed.



3 • Anchor every number

Doses, units, and intervals are most fragile on compressed audio - “15” and “50” sound closer on a weak connection than they look on paper.



4 • Summarize once, clearly

After a choppy connection, an interpreter, or a long history, one brief spoken summary gives the note a clean anchor.



5 • Verify who said what

On review, confirm that a family member’s comments or an interpreter’s questions did not become the patient’s history.



6 • Budget the editing time

Phone-only, low-bandwidth, and interpreter visits routinely need more correction - plan review time for them.

Rural corollary: phone-only and low-bandwidth visits give the weakest audio. Decide whether the scribe belongs in that visit at all.

What patients actually think - when you ask

Shah et al., JAMIA Open 2026 - 2,202 patient surveys after AI-scribed outpatient visits, Stanford, May–Jul 2025

70.1%

found the ambient AI scribe helpful

[37]

73.6%

wanted it used in their future visits

[37]

≠

small but significant differences by
gender, age, and race

[37]

Patients are largely receptive - but detail changes the answer: the share who felt comfortable fell from 81.6% to 55.3% once patients heard detail on AI features, data storage, and corporate involvement (NYU, n=103) [41]. Transparency costs some comfort - and it is the only consent that counts. [37][41]

Recording consent: where Wyoming sits

34

states: one-party consent (the clinician's own consent can suffice)

[22]

16

states: all-party or mixed rules by this count; other sources list 11 to 12

[22]

WY

is one-party - legally, your own consent can suffice. Ask anyway.

Wyo. Stat. §7-3-702 [46]

All-party states (in-person clinical conversations)

- CA · FL · IL · MD · MA · MT · NH · PA · WA - plus OR (all-party in person). Next door: Montana requires that all parties know of the recording (a clear warning satisfies it); ID, UT, CO, NE, SD are one-party.
- **Telehealth rule of thumb: the stricter state's rule governs - know where your patient is sitting.**

“We have a BAA” is necessary - and not enough



What HIPAA does here

- Scribe vendors handling PHI for you are business associates - a BAA is mandatory (45 CFR §160.103, §164.502(e)) [23]
- Sets security & permitted-use rules for the vendor
- Treatment-operations uses generally don't require patient authorization under HIPAA itself



What HIPAA does not do

- It is a floor, not a ceiling - it does not preempt stricter state law (45 CFR §160.203) [23]
- HIPAA compliance ≠ consent: state recording law and your professional ethics are separate questions
- It does not answer retention, training-use, or recording-consent questions - your contract and workflow must
- California class actions (2025–26) allege missing consent even though the vendor signs BAAs; pending, nothing decided [24][25]

A consent workflow that actually works

1 • Notify early Scheduling & check-in mention it; posted notice; in the patient's language

2 • Clinician asks Verbal ask at visit start - every visit, not just the first

3 • Document Consent or decline recorded in the EHR (discrete field, auditable)

4 • Honor the no Tool off, no friction, zero care consequences - and staff believe it


5 • Re-ask when it matters Sensitive visits: behavioral health, adolescents, DV, substance use

A script that works

"I use a secure AI tool that listens during our visit and drafts my note, so I can focus on you instead of the keyboard. The audio is processed by [vendor] and deleted after [X days]. I review and correct everything before it goes in your chart. Is that OK with you? You can say no - it won't change your care at all."

Fill the brackets from your actual contract. If you can't fill them, that's a finding.

The “shadow record” problem



Only the last box has settled legal status. Everything upstream is the shadow record.



Open questions with no federal standard

- Is the audio/transcript part of the designated record set (patient access rights)? Is it discoverable in litigation?
- It may contain sensitive disclosures the clinician deliberately left out of the note (e.g., a DV disclosure)
- Retention varies by vendor & contract - e.g., one large Canadian pilot kept audio 30 days [31]. Set YOUR policy in the contract.

Behavioral health: why the rules are different



Why therapy content is different in kind

- The tool records the whole session - not just a summary of it
- HIPAA gives psychotherapy notes extra protection: process detail is too sensitive for the general record [48]
- A session transcript recreates that detail outside the protected container; SUD records add 42 CFR Part 2 [49]
- It can resurface in custody, employment, immigration, or security-clearance matters [42]



What a careful workflow builds in

- Default off; opt in session by session, never through a blanket form
- Not in crisis, IPV, or trauma-processing sessions - this is where the pause matters
- In writing from the vendor: retention, who can access, no training use
- One signal to watch: AI-scribed primary care notes documented more psychiatric symptoms, yet documented interventions were less likely [17]

“A record of someone’s lowest moments can be used against them...” - Nicole Alvarez, Center for American Progress, quoted by CalMatters, June 2026 [42]

More documentation, less management

Castro et al., JAMA Psychiatry 2026 - 20,302 primary-care annual-visit notes, Mass General Brigham, matched case-control



Documentation went UP

- AI-scribed notes showed more neuropsychiatric symptoms in all 6 RDoC domains, as estimated by a language model reading the notes [17]



Management went DOWN

- Lower odds of a documented psychiatric intervention (referral, new diagnosis, or antidepressant): aOR 0.83 (95% CI 0.72–0.95) [17]
- Human scribes: no difference (aOR 0.97, 95% CI 0.85–1.11)

Hypothesis-generating, not causal - but it challenges the assumption that richer notes produce better care.

The deliberate pause



When I pause the recording

- Immigration status - the patient's or a family member's
- Intimate partner violence and abuse disclosures
- Substance use and stigmatized mental-health content
- Anything a patient would not say if a stranger might read it



Why a pause beats an edit

- Deleting a line from your draft does not delete it from the vendor's transcript
- What is never recorded can never be produced - the pause is data minimization
- You can still chart what matters: pause, listen, dictate the clinically relevant piece afterward
- The vendor ask: one-tap pause, early deletion, on-device processing [11][31]

"I'm going to pause this tool so we can talk privately." In my clinic: "Voy a pausar esta herramienta para que hablemos en privado." One sentence turns the pause into trust.

Six questions to ask any AI note-taking vendor



1 • Will you sign a BAA?

No BAA, no patient information. That one answer sorts clinical tools from consumer chatbots.



2 • What do you keep, and how long?

Audio, transcript, and draft are three separate answers. Ask where each one lives and when it is deleted.



3 • Do you train on our data?

The answer you want: never without explicit opt-in - written into the contract, not the brochure.



4 • Can we delete on request?

When a patient withdraws mid-visit, can you purge what was captured - and confirm it in writing?



5 • How does it handle our visits?

Telehealth audio, accents, interpreters, three or more voices. Ask for performance by group, not one average.



6 • Is there a one-tap pause?

If pausing takes three menus, nobody pauses. Try it yourself during the demo.

If you cannot get these answers in writing, that is your finding. Small practices: questions 1 to 3 cost nothing to ask.

A safe workflow on one slide

Before

Know your BAA, retention, and training terms. Tell patients at scheduling or check-in.

CHOOSE · is this visit right for the scribe?

During

Ask out loud, every visit. Say who is on the call. Narrate what you see. Pause for anything sensitive.

CAPTURE · say what you want documented

After

Read for what looks right, not only what looks wrong. Check meds, numbers, who said what.

CHECK · the signed note is yours

Screenshot this one. Three habits, in order: choose, capture, check - the “3 Cs,” a framework from a 2026 narrative review [39].

If you remember five things

- ✦ **The benefit is real but modest and early.** Minutes per day, a genuine burnout signal, mostly short follow-up. The tool you pick changes the result. [12][13][14]
- ✦ **The draft is not the record.** Hallucination and omission are expected behaviors; your review is the safety system. [11][16][18]
- ✦ **Consent is a conversation, not a form.** Wyoming is a one-party state, and asking is still the standard: verbal, documented, easy to decline. [22][27][46]
- ✦ **Know where the data goes.** BAA, retention, and training use, in writing. A consumer chatbot is not a clinical scribe. [23][47]
- ✦ **Some visits should not be recorded.** Crisis, IPV, substance use, therapy process: pause, or default off. [42][48][49]



Questions

Ambient AI can give us back our attention - whether it also protects our patients' trust is up to how we deploy it.

Thank you!

Dr. Gigi Magan, MD
drgigimagan.com · drgigimagan.substack.com

STAY IN TOUCH

Let's connect

LinkedIn



linkedin.com/in/yimdriuskamagan

Substack



drgigimagan.substack.com

Website



drgigimagan.com

Scan to follow my work. @DrGigiMagan · MedTech & Heart: medicine and technology, with heart. Happy to share the consent script and the checklist.

References (1 of 5)

- [1] Sinsky C, et al. Allocation of Physician Time in Ambulatory Practice: A Time and Motion Study in 4 Specialties. *Ann Intern Med*. 2016;165(11):753-760. doi:10.7326/M16-0961
- [2] Arndt BG, et al. Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations. *Ann Fam Med*. 2017;15(5):419-426. doi:10.1370/afm.2121
- [3] Saag HS, et al. Pajama Time: Working After Work in the Electronic Health Record. *J Gen Intern Med*. 2019;34(9):1695-1696. doi:10.1007/s11606-019-05055-x
- [4] Muhiyaddin R, et al. Electronic Health Records and Physician Burnout: A Scoping Review. *Stud Health Technol Inform*. 2022;289:481-484. doi:10.3233/SHTI210962
- [5] American Medical Association. Augmented Intelligence Research: physician sentiment surveys - 2023, 2024, and 2026 waves (2026 Physician Survey on Augmented Intelligence, n≈1,700, fielded fall 2025, published Mar 2026). ama-assn.org
- [6] Leung TI, Coristine AJ, Benis A. AI Scribes in Health Care: Balancing Transformative Potential With Responsible Integration. *JMIR Med Inform*. 2025;13:e80898. doi:10.2196/80898
- [7] Tierney AA, et al. Ambient Artificial Intelligence Scribes to Alleviate the Burden of Clinical Documentation. *NEJM Catalyst Innov Care Deliv*. 2024;5(3). doi:10.1056/CAT.23.0404
- [8] Tierney AA, et al. Ambient AI Scribes: One-Year Findings [The Permanente Medical Group]. *NEJM Catalyst Innov Care Deliv*. 2025. doi:10.1056/CAT.25.0040
- [9] Sweller J. Cognitive load during problem solving: Effects on learning. *Cognitive Science*. 1988;12(2):257-285. doi:10.1207/s15516709cog1202_4
- [10] Shah SJ, et al. Physician Perspectives on Ambient AI Scribes. *JAMA Netw Open*. 2025;8(3):e251904. doi:10.1001/jamanetworkopen.2025.1904
- [11] Topaz M, Peltonen LM, Zhang Z. Beyond human ears: navigating the uncharted risks of AI scribes in clinical practice. *NPJ Digit Med*. 2025;8:569. doi:10.1038/s41746-025-01895-6
- [12] Lukac PJ, et al. Ambient AI Scribes in Clinical Practice: A Randomized Trial. *NEJM AI*. 2025. doi:10.1056/Aloa2501000

References (2 of 5)

- [13] Rotenstein LS, et al. Changes in Clinician Time Expenditure and Visit Quantity With Adoption of Artificial Intelligence-Powered Scribes: A Multisite Study. JAMA. 2026. doi:10.1001/jama.2026.2253
- [14] Olson KD, et al. Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout. JAMA Netw Open. 2025;8(10):e2534976. doi:10.1001/jamanetworkopen.2025.34976
- [15] Duggan MJ, et al. Clinician Experiences With Ambient Scribe Technology to Assist With Documentation Burden and Efficiency. JAMA Netw Open. 2025;8(2):e2460637. doi:10.1001/jamanetworkopen.2024.60637
- [16] Guo Y, et al. Clinicians' rationale for editing ambient AI-drafted clinical notes: persistent challenges and implications for improvement. J Am Med Inform Assoc. 2026;33(7):1345-1353. doi:10.1093/jamia/ocag059
- [17] Castro VM, et al. Psychiatric Documentation and Management in Primary Care With Artificial Intelligence Scribe Use. JAMA Psychiatry. 2026;83(3):281-286. doi:10.1001/jamapsychiatry.2025.4303
- [18] Palm E, et al. Assessing the quality of AI-generated clinical notes: validated evaluation of a large language model ambient scribe. Front Artif Intell. 2025;8:1691499. doi:10.3389/frai.2025.1691499 (authors affiliated with Suki AI, a scribe vendor)
- [19] Reddy A, et al. Rapid Evaluation of Artificial Intelligence Technology Used for Ambient Dictation in Primary Care. Ann Intern Med. 2026;179(6):765-772. doi:10.7326/ANNALS-25-02772
- [20] Koenecke A, et al. Racial disparities in automated speech recognition. Proc Natl Acad Sci USA. 2020;117(14):7684-7689. doi:10.1073/pnas.1915768117
- [21] Blaseg E, Huffstetler A. Artificial Intelligence Scribes Shape Health Care Delivery [Graham Center Policy One-Pager]. Am Fam Physician. 2025;111(4):304-305
- [22] Anderson TN, Sinha MS, Cohen IG, Jones RL. Ambient AI in Clinical Practice - The Legal Landscape of Recording Consent Requirements. NEJM AI. 2026;3(6). doi:10.1056/Alp2600203

References (3 of 5)

- [23] HIPAA: 45 CFR §160.103 (business associate); §164.502(e) (BAA requirement); §160.203 (state-law preemption floor)
- [24] Washington et al. v. Sutter Health; Memorial Health Services, Inc.; MemorialCare Medical Foundation. No. 4:26-cv-03012 (N.D. Cal., filed Apr 7, 2026); analysis: Alston & Bird / JD Supra (2026)
- [25] Saucedo v. Sharp HealthCare (San Diego Super. Ct., filed Nov 26, 2025) - first major ambient-AI recording-consent class action; pending
- [26] Title VI, Civil Rights Act of 1964 (42 USC §2000d); ACA §1557 and 45 CFR Part 92 (2024 final rule) - language access / meaningful access requirements
- [27] American Medical Association. Augmented Intelligence Development, Deployment, and Use in Health Care: AMA Principles. November 2024. [ama-assn.org](https://www.ama-assn.org)
- [28] Daneshvar N, et al. Artificial Intelligence in the Provision of Health Care: An American College of Physicians Policy Position Paper. Ann Intern Med. 2024;177(7):964-967. doi:10.7326/M24-0146
- [29] 21st Century Cures Act software exclusions, FDCA §520(o) (21 USC §360j(o)); Bipartisan Policy Center. FDA Oversight: Understanding the Regulation of Health AI Tools
- [30] NHS England. Guidance on the use of AI-enabled ambient scribing products in health and care settings. April 2025 (v1); updated 2026. [england.nhs.uk](https://www.england.nhs.uk)
- [31] Beny H. AI scribes in health care raise risks for patients and privacy. Policy Options (IRPP). April 2026
- [32] Therapeutic Goods Administration (Australia). Digital scribes guidance (2025); ACSQHC AI Safety Scenario: Ambient Scribe (2025); Safer Care Victoria Ambient AI Scribes Advisory (2025)
- [33] Regulation (EU) 2024/1689 (EU AI Act); Council of the EU, Digital Omnibus provisional agreement, May 2026
- [34] Buller R. AI Tools Promise Better Care but Challenge Safety-Net Providers. California Health Care Foundation. March 2025. [chcf.org](https://www.chcf.org)

References (4 of 5)

- [35] Kanaparth NS, et al. Real-World Evidence Synthesis of Digital Scribes Using Ambient Listening and Generative AI for Clinician Documentation Workflows: Rapid Review. JMIR AI. 2025;4:e76743. doi:10.2196/76743
- [36] HRSA Health Center Program: national program data (2024). bphc.hrsa.gov
- [37] Shah SJ, et al. Patient perspectives on clinicians' use of ambient AI scribes. JAMIA Open. 2026. doi:10.1093/jamiaopen/ooag104
- [38] Preiksaitis C, et al. Ambient AI Scribes and Emergency Department Documentation Burden: Retrospective Cohort Study. JMIR AI. 2026;5:e92193. doi:10.2196/92193
- [39] Kuhn TW. From Clinical Encounter to Draft Documentation: A Mechanistic Narrative Review of Ambient Scribe Technology. Cureus. 2026;18(6):e110201. doi:10.7759/cureus.110201
- [40] Anderson TN, et al. Evaluating the Quality and Safety of Ambient Digital Scribe Platforms (5-platform simulation study; unedited notes). Mayo Clin Proc Digit Health. 2025;3(4):100292. doi:10.1016/j.mcpdig.2025.100292
- [41] Lawrence K, et al. Informed Consent for Ambient Documentation Using Generative AI in Ambulatory Care. JAMA Netw Open. 2025;8(7):e2522400. doi:10.1001/jamanetworkopen.2025.22400
- [42] Your medical provider might be recording your mental health care visits. CalMatters / The Markup. June 16, 2026. calmatters.org (journalism, not peer reviewed)

References (5 of 5)

- [43] Wyoming Telehealth Network. AI survey of Wyoming telehealth providers, June–July 2026 (n=55). Unpublished; shared with permission. Open-text themes coded by the speaker.
- [44] Afshar M, et al. A Pragmatic Randomized Controlled Trial of Ambient Artificial Intelligence to Improve Health Practitioner Well-Being. NEJM AI. 2025;2(12). doi:10.1056/Aloa2500945
- [45] Chowdhury A, et al. Comparing ambient scribes: a randomized crossover clinical trial addressing ambient scribe technologies' impact on physician burnout. J Am Med Inform Assoc. 2026;33(5):990-999. doi:10.1093/jamia/ocag018
- [46] Wyo. Stat. §7-3-702 (one-party consent to recording); Mont. Code Ann. §45-8-213 (knowledge of all parties required; a warning satisfies it)
- [47] HIPAA de-identification standard: 45 CFR §164.514(b) (Safe Harbor, 18 identifiers; Expert Determination)
- [48] HIPAA psychotherapy notes: 45 CFR §164.501 (definition); §164.508(a)(2) (authorization required)
- [49] Confidentiality of Substance Use Disorder Patient Records: 42 CFR Part 2 (2024 final rule)
- [50] Aldis R, Wang S, Sage M, Metzmaker M, Galvin H. Performance of an Ambient Generative AI Documentation Tool in a Linguistically Diverse Clinical Setting. medRxiv. Posted August 17, 2026. doi:10.64898/2026.08.14.26360467 (preprint, not peer reviewed)