



WELCOME TO THE 2026 AI COHORT SERIES

RAISING AWARENESS, DRIVING CHANGE, AND COMBATING CHRONIC CONDITIONS



AI Scribes in Action: Lessons from a Practicing Physician

Yimdriuska Magan, MD, MPH

August 3rd, 2026 10:00AM PT


DISCLAIMER

This session is purely for informational purposes.

The California Telehealth Resource Center has no relevant financial interest, arrangement, or affiliation with any organizations related to commercial products or services discussed during today's session.

The California Telehealth Resource Center (CTRC) and all resources and activities produced or supported by the CTRC are made possible by grant number U1UTH42520-01-01 from the Office for the Advancement of Telehealth, Health Resources and Services Administration, DHHS. This information or content and conclusions are those of the CTRC and should not be construed as the official position or policy of, nor should any endorsements be inferred by HRSA, HHS or the U.S. Government.

MICROSOFT TEAMS TIPS



Open the Chat

Please open the chat and use it liberally; we want to hear from you!

Raise your Hand or React

Attendees will have an opportunity to ask questions. Raise your hand or send reactions!

Muted on Entry

You are muted on entry.

Session Recording

Today's session is being recorded! Keep an eye out for attendee announcements when recordings are available.

THANK YOU TO OUR FUNDERS



Established in 2006, the California Telehealth Resource Center (CTRC) exists to share **unbiased, no-cost digital health resources and consultative support services** with providers and patients located across all 58 California counties and beyond. As part of a collaboration of 12 regional and 2 national Telehealth Resource Centers, CTRC is **committed to implementing digital health programs for rural and underserved communities**.

The work of the CTRC would not be possible without funding and support from the Health Resources and Services Administration's **Office for the Advancement of Telehealth** and the **CA State Office of Rural Health**, part of the Department of Health Care Access and Information.

MEET YOUR SPEAKER

AI Scribes in Action: Lessons from a Practicing Physician



Yimdriuska (Gigi) Magan, MD, MPH
Co-founder Alma First

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 July 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.

If you remember nothing else...

Real

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

The evidence - sections 3–4 [12][13][14]

Uneven

The risks are unevenly distributed: errors concentrate on accented, multilingual, and safety-net patients.

Safety & equity - sections 4, 7 [11][20]

Fixable

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

Consent & governance - sections 5–6 [22][24]

By the end of this hour, you can...

- 1 Explain how ambient AI scribes work - and name the failure mode at each stage of the pipeline.
- 2 Appraise the 2024–2026 evidence on documentation time, burnout, and note quality, including its limits.
- 3 Build a legally sound, in-language consent workflow - and explain why a line in the intake packet is not one.
- 4 Apply an equity lens to any ambient AI deployment, especially in safety-net and telehealth settings.

Where we're going

- 1 The documentation crisis** - why this technology exploded ~3 min
- 2 How ambient scribes work** - the pipeline and its failure points ~5 min
- 3 What the evidence shows** - RCT, multisite, burnout, note quality ~10 min
- 4 Patient safety** - hallucination, omission, automation bias, algorithmic bias ~5 min
- 5 Consent & the legal landscape** - CIPA, the lawsuits, and a workflow that works ~8 min
- 6 Guidance & governance** - AMA/AAFP/ACP · US patchwork vs NHS model ~5 min
- 7 Equity & the safety net** - who the evidence leaves out ~5 min
- 8 Monday morning** - practical takeaways + your questions ~10 min



1

The documentation crisis

Why a technology this young is already everywhere.

The arithmetic of the clinic day

≈2 hrs

of EHR + desk work for every 1 hour of
direct patient care

Researchers timed every task, 57 physicians, 4
specialties [1]

+1–2 hrs

of after-hours work nightly, mostly EHR
tasks

Same cohort, nightly diaries [1]


~1.4 hrs

after clinic per day for family physicians -
“pajama time”

EHR event logs [2]; term quantified in [3]

Where this number comes from? Researchers timed every task for 57 physicians across four specialties over 430 hours in the clinic; 21 of them also kept nightly diaries. The 2-to-1 figure combines EHR work *and* all other desk work, not just note-writing. It comes from a small set of high-performing practices and is now nearly a decade old (2016). So the exact ratio varies by setting, but every study since has pointed the same direction: documentation is a consistent driver of burnout. [4].

Adoption more than doubled in three years



AMA physician surveys

- 38% (2023) → 66% (2024) → 81% (2026, n≈1,700) - use “more than **doubled**” since 2023 [5]
- 76% now believe AI improves patient care (65% in 2023) [5]



Same survey, same physicians:

- ~40% “equally excited and concerned”; 88% want robust safety validation; top ask = clear liability frameworks [5]




2

How ambient AI scribes work

Two technologies stacked - and a signature at the end.

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.

Every stage has a failure mode



ASR stage - mishearing

- Error rates nearly 2× higher for Black speakers in landmark testing of commercial ASR [20]
- Degrades with accents and limited-English-proficiency speech [11]



LLM stage - miswriting

- Hallucination: content never said in the room
- Omission · misattribution (who said it) · misinterpretation [11]



Review stage - missing it

- Editing is real work: accuracy, billing, liability, style [16]
- Review fatigue invites rubber-stamping (automation bias)



Storage stage - what remains

- Audio + intermediate transcript may persist on vendor servers
- Retention, training use, discoverability = contract questions [11][24]

Why clinicians love them anyway



The cognitive-load argument

- Documenting while interviewing = two concurrent high-demand tasks; performance degrades on both (cognitive load theory) [9]
- Decoupling the note from the conversation returns attention to the patient
- Matters most for nuance: psychosocial cues, hesitation, what isn't said



I feel like I can be more face-to-face with them and have more eye contact... instead of staring at the computer.

Physician interviewee, Stanford ambient-scribe pilot - where 68% of patient-engagement comments were positive [10]



3

What the evidence shows

One RCT, one big multisite study, and a lot of 30-day pilots.

The evidence base at a glance (mid-2026)

Randomized trial 1 published RCT - single academic center, 2 months [12]

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
- No FQHC studies

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

The one RCT: platform choice matters

Lukac et al., NEJM AI 2025 - UCLA, pragmatic 3-arm RCT • 238 physicians, 14 specialties • Nov 2024–Jan 2025

–9.5%

time-in-note, Nabla arm vs control

95% CI –17.2 to –1.8 • P=.02 [12]

–1.7%

DAX Copilot arm — not significant

95% CI –9.4 to +5.9 • P=.66 [12]

Both ↓

burnout & task-load improved in both arms (secondary)

Mini-Z / NASA-TLX measures [12]

Takeaway: “ambient AI” is not one intervention. Two platforms, same trial, different results.

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.

The burnout signal - real, early, and partly vendor-adjacent



Olson et al. - 263 clinicians, 6 health systems, Abridge, 30 days [14]



Also in that study

- Self-reported after-hours documentation: -0.90 h/week (4.95 $\rightarrow$ 4.05)
- ≈ 10.8 minutes saved per workday [14]



Read with caveats

- 30-day follow-up; self-report; funded by participating institutions; two co-authors were vendor employees [14]
- Penn QI study: -20.4% note time/appt; -30.0% after-hours work [15]

What deployment at scale looks like

The Permanente Medical Group (Kaiser Permanente Northern California) - the largest published deployment

303,266

encounters in the first 10 weeks (3,442 physicians)

Oct–Dec 2023 [7]

2.58 M

encounters by year one (7,260 physicians)

1-year follow-up [8]

15,791

estimated hours saved in year one

uncontrolled estimate [8]

Quality-improvement design - no control group, so causal claims are limited. Internal note-quality review scored ~48/50; the paper does not name the vendor.

Head-to-head: AI drafts vs human notes



Palm et al. 2025 - PDQI-9 audit

- Overall quality close: 4.20 vs 4.25 / 5 (physician notes slightly higher, $P=.04$) [18]
- **Hallucinations: 31% of ambient notes vs 20% of physician notes ($P=.01$) [18]**
- AI notes: more thorough & organized, less succinct



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 deficits: thoroughness, organization, usefulness; worst case (low back pain): 20.3 vs 43.8 / 50

What we still don't know



Patient outcomes

- No trial has measured diagnostic accuracy, treatment adherence, or downstream harm - only process measures.



Durability

- Follow-up is mostly 30–90 days. Does benefit persist? Does review vigilance decay?



Head-to-head platforms

- One 2-platform RCT [12]; no standardized multi-platform comparison exists.



Safety-net settings

- Every major study: academic medical centers on Epic. FQHC evidence: none.

These gaps don't say “refuse” - they define what your own deployment must measure.

4

Patient safety

The note that lies, the note that forgets, and the reviewer who stops looking.

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.

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 recorded significantly more neuropsychiatric symptoms - across all 6 RDoC domains [17]




Management went DOWN

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

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

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

- Stanford pilot physicians: limited functionality with non-English-speaking patients [10]
- Vendors report improvement - independent, stratified validation is still missing [21]

The multilingual visit: a six-move playbook



1 • Announce the roles aloud

“I’m the doctor; this is the interpreter.” Spoken role labels give the model structure - misattribution is its default failure with three voices.



2 • Narrate the exam

The AI cannot see pointing, wincing, or gestures. Verbalize findings in the language of the visit - or they never existed.



3 • Anchor every number

Doses, units, intervals are most fragile at language switches - “0.5” and “5” sound closer in a bilingual room than they look on paper.



4 • Summarize once, clearly

After a complex history in Spanish - or any language - one brief summary in the documentation language gives the note a clean anchor.



5 • Verify who said what

On review, confirm the interpreter’s clarifying questions did not become the patient’s medical history.



6 • Budget the editing time

Code-switched and non-English drafts routinely need more correction - plan review time for them, and say so in rollout plans.

Telehealth corollary: phone interpreters + compressed audio make attribution worse - flag interpreter visits before the encounter starts.



5

Consent

The question your patients - and now the courts - are asking.

Why AI scribes are not just “a scribe in the room”



A human scribe

- Present, visible, introduceable
- Bound by your policies and HIPAA training
- Creates no copy outside your walls
- Stops writing when asked



An ambient AI scribe

- Recording may not be obvious to the patient
- Audio + transcript processed on third-party servers
- Retention & model-training use vary by contract
- Triggers recording statutes a person taking notes never does

The legal difference in one word: interception. Recording + transmission is what wiretap statutes regulate [22][24].

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: willingness fell 81.6% → 55.3% when patients heard full detail on storage and corporate access (NYU, n=103) [41]. Transparency costs uptake - and it is the only consent that counts. [37][41]

Meaningful consent: seven elements

1 • Recording - An AI tool will listen to and record this visit

2 • Processing - Audio is processed by a named third-party vendor, off-site

3 • Storage - What is kept (audio? transcript?), where, and for how long

4 • Secondary use - Whether data can train the vendor's models

5 • The note - A draft enters your record after clinician review

6 • The right to say no - Declining changes nothing about your care

7 • Revocation - You can change your mind - including mid-visit

Buried boilerplate in admission paperwork is exactly what the lawsuits attack as insufficient [24].

Recording consent: a 50-state patchwork

34

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

[22]

16

states: all-party or mixed regimes - every participant must consent

[22]

CA

is all-party - and adds statutory damages + a private right of action

CIPA §632 / §637.2 [24]

All-party core (in-person clinical conversations)

- CA · FL · IL · MD · MA · MT · NH · PA · WA - plus OR (all-party for in-person; one-party for phone). MI is contested case law; CT is all-party for phone/electronic - directly relevant to telehealth.
- **Telehealth rule of thumb: the stricter state's rule governs - know where your patient is sitting.**

The California statute stack



CIPA - Invasion of Privacy Act

- Penal Code §632: all-party consent for confidential communications. §637.2: statutory damages - greater of \$5,000 per violation or 3× actual damages + private right of action. No proof of harm required.



CMIA - Confidentiality of Medical Information Act

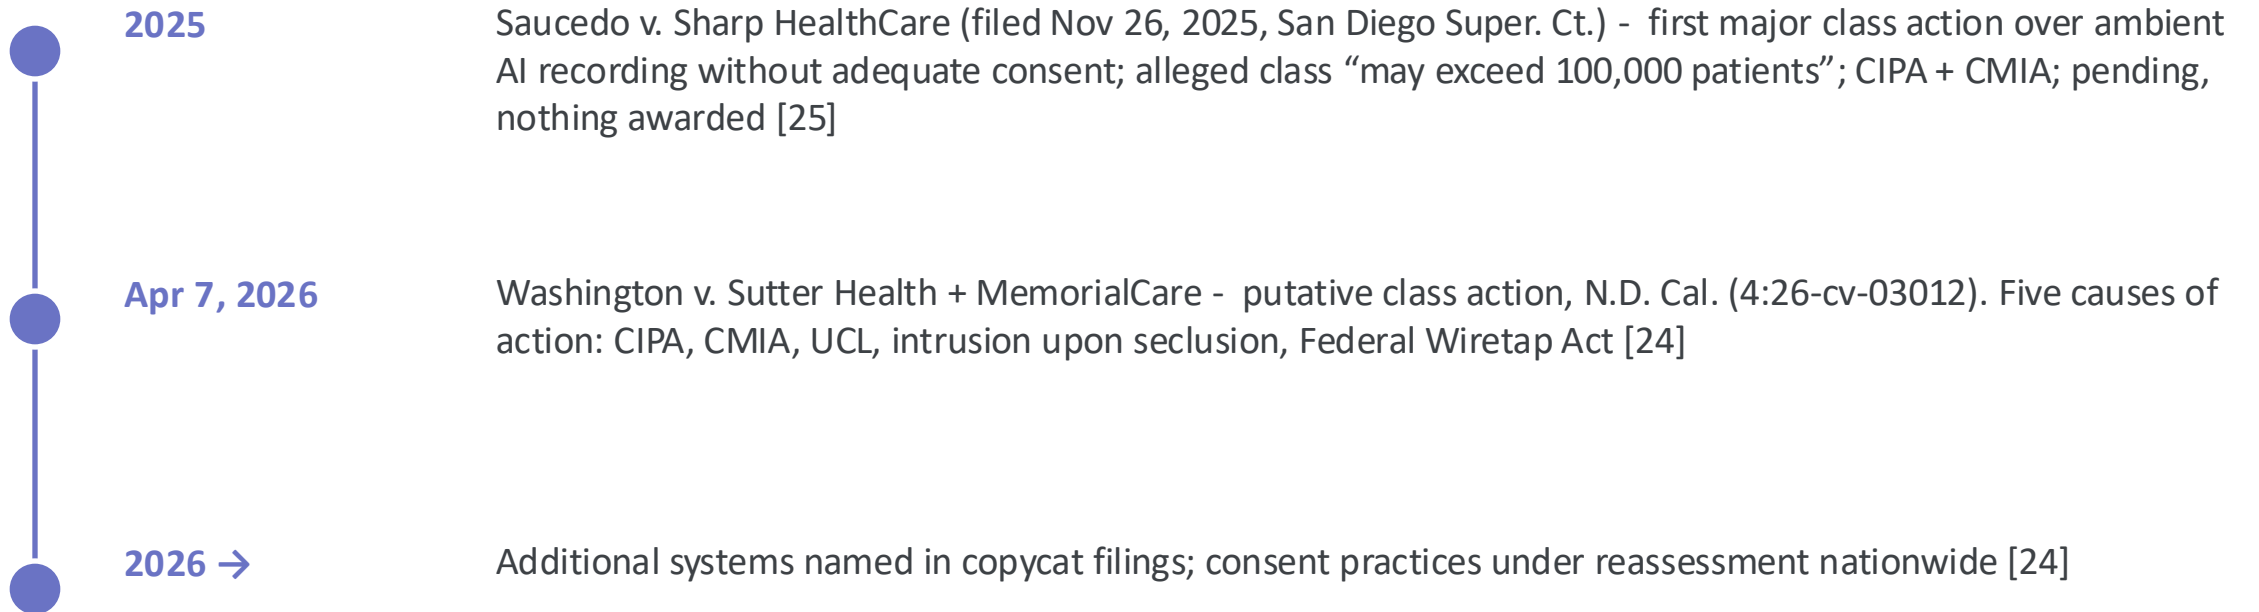
- Civil Code §56 et seq.: limits disclosure of medical information by providers and their contractors; pleaded alongside CIPA in the AI-scribe cases.



The multipliers

- UCL §17200 (unfair competition), common-law intrusion upon seclusion, and the Federal Wiretap Act (18 U.S.C. §2510) - all appear in the 2026 complaints. Damages can scale per patient, per encounter.

The litigation wave is here



Note who is NOT a defendant: the vendor. The health systems hold the consent risk.

What the complaints actually argue

1 · Harm at the moment of interception

- The violation is the recording itself without proper consent - not any later misuse of the data.

2 · Third-party transmission is the core

- Audio and transcripts flowed to an outside company's servers, processed by a non-party to the visit - potentially for product improvement.

3 · General notices ≠ consent

- Patients were not given “clear notice” that conversations would be recorded by an AI platform and sent outside the clinic. Boilerplate privacy language doesn't cover it.

4 · Damages scale per encounter

- Statutory damages assessed per violation make exposure proportional to deployment size - the more successful the rollout, the larger the class.

Unadjudicated allegations - but each theory maps to a fixable workflow gap. That's the good news.

“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 ≠ CIPA consent: the wiretap question is separate
- It does not answer retention, training-use, or recording-consent questions - your contract and workflow must

Consent that patients can actually understand



The legal frame

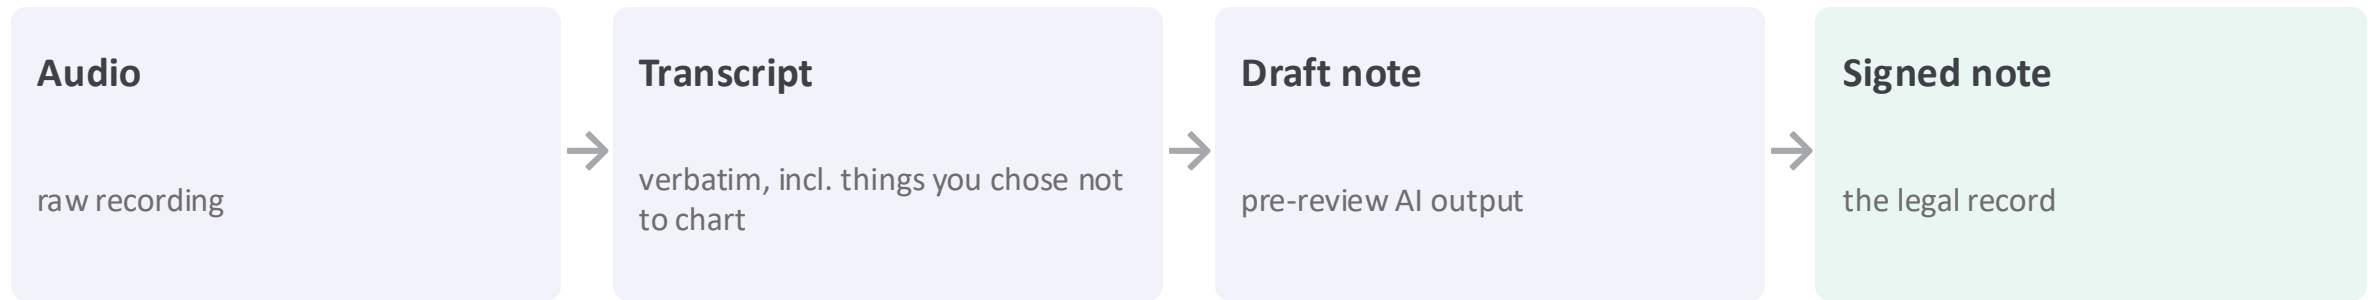
- Title VI + ACA §1557: federally funded providers owe LEP patients meaningful access - qualified interpreters, translated vital documents (45 CFR Part 92) [26]
- Applying that standard to AI-scribe consent conversations is a straightforward inference: consent you can't understand isn't meaningful
- Label it honestly: an application of the standard - no AI-scribe-specific regulation exists yet



The operational reality

- In-language consent adds a real workflow step - the step most likely to be skipped on a busy Tuesday
- Compounding problem: the same LEP patients face the tool's weakest transcription performance [10][20]
- Fix: translated scripts + interpreter briefed on the AI-scribe ask + consent status flagged before the visit

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.

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, immigration-adjacent topics

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.

When consent gets thin in the hardest visits



What clinicians reported

- Entire therapy sessions were recorded - not just a note
- Clinicians said consent felt coercive - for them and for patients
- No clear answers on how recordings were stored or who could access them
- Therapy content is different in kind - it can resurface in custody, immigration, employment



Kaiser's response - and the open gap

- Kaiser: consent is obtained; HIPAA-compliant; retention capped at 14 days
- Kaiser: the data is not used to train AI models
- Independent verification of these assurances is limited [42]
- In the most sensitive visits, this is why the pause matters

“A record of someone’s lowest moments can be used against them.” - policy analyst quoted in the CalMatters investigation

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]

“Voy a pausar esta herramienta para que hablemos en privado.” - one sentence, in the patient's language, turns the pause into a trust-building move.



6

Guidance & governance

What the societies say - and who is actually in charge.

The societies converge on three themes

AMA (2024 AI principles)

- AI use affecting care “should be disclosed and documented”
- Liability should be shared / allocated to whoever is best positioned to mitigate risk - pushing back on signer-owns-everything [27]

AAFP (Graham Center, 2025)

- Promise acknowledged - “appropriate AI oversight is clear[ly]” needed
- Warns of “chart integrity degradation, biases, and model collapse”
- Vendor accuracy claims (90–99%) vs “users have reported a different experience” [21]

ACP (2024 position paper)

- AI should reduce - not exacerbate - disparities
- Prioritize bias mitigation & diverse training data
- Physician judgment stays central [28]

Meanwhile, today's reality: malpractice doctrine and vendor terms put responsibility on the signing clinician - AMA policy argues that burden should be shared. Know the difference.

Who regulates ambient AI scribes? Mostly nobody.

 **FDA** Documentation-only scribes generally fall OUTSIDE the device definition (21st Century Cures §520(o)). No premarket review. None are FDA-cleared. [29]

 **FTC** Deceptive trade practices authority - could reach overstated accuracy claims; no scribe-specific action yet.

 **States** Wiretap/privacy statutes + private litigation = the active regulatory front (see: California). [22][24]

 **You** Institutional policy, vendor contracts, consent workflows, and QI audits carry the real load.



7

The equity lens

These tools were not trained on the patients I serve - and the evidence base still isn't.

The safety-net calculus



Why the upside is biggest here

- 31M+ patients rely on health centers; visits are complex, multilingual, SDOH-heavy [36]
- Documentation burden per visit is highest exactly here



Why the risk concentrates here too

- ASR error disparities by race, accent, language [20]
- Consent complexity: LEP, trust, interpreter-mediated visits [10][26]



CHCF, 2025 - barriers reported by California safety-net leaders

- Cost: “the pricing models don't work for the safety net” - enterprise pricing assumes enterprise margins [34]
- Language performance concerns for patients who speak languages other than English [34]
- Thin IT infrastructure & staffing; liability worries with little legal support [34]

The evidence gap is itself an equity problem

Every major study - AMCs on Epic: UCLA RCT; MGB/Emory/UCSF/Yale/UC Davis multisite; Kaiser; Penn; Stanford [7][10][12][13][15]

FQHC / CHC trials - Zero. No RCT or large observational study has been conducted in a safety-net setting

Stratified performance - No peer-reviewed study reports scribe accuracy by patient race, language, or accent in real clinics

Vendor equity data - Not published. Ask anyway - procurement is leverage

The populations most likely to benefit are the least represented in the evidence - and the most exposed to its failure modes.



8

Monday morning

Six moves that de-risk an ambient AI deployment.

TAKEAWAYS

Six moves that de-risk deployment



1 • Build the consent workflow first

Verbal ask + EHR documentation + easy decline + in-language. Not a line in the intake packet.



2 • Name an AI safety owner

One accountable person, a hazard log, incident reviews - the NHS CSO model, adopted voluntarily.



3 • Audit notes routinely

Random signed-note audits vs encounter recall; stratify by patient language & demographics.



4 • Demand stratified vendor data

Accuracy by language, accent, demographic group - before contracting. Reject unaudited claims.



5 • Contract the data lifecycle

Audio/transcript retention, deletion on signature, no training use without opt-in - in writing.



6 • Pilot with equity metrics

Decline rates by language; note quality by population; access impacts - from day one.

Each answers a failure: lawsuits (1) • unowned risk (2) • automation bias (3) • ASR bias (4) • shadow record (5) • evidence gap (6)

Three asks to leave with

Consent

Explicit, per-visit, in-language consent - verbal, documented, revocable. Not intake-packet boilerplate.

Clinicians · compliance · policy [22][24]

Pause

Build pause and opt-out moments for sensitive disclosures - and treat using them as good practice.

Clinicians · care teams · trainers [11]

Ask

Demand early deletion, no training use without opt-in, stratified accuracy - and on-device processing.

CMIOs · informaticians · advocates [30][31]

The window is now: consent norms, contract terms, and the first wave of case law are being set in the next two years - with the people in this room, or without them.

If you remember five things

- ✦ **The benefit is real but modest and early.** Minutes per day, a genuine burnout signal, ~30–90-day evidence. Platform choice 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 workflow, not a form.** Named vendor, retention, decline path, in-language - verbal and documented. [22][24][26]
- ✦ **California is the strictest consent environment in the country.** CIPA + CMIA + active class actions. Treat it as a design constraint, not a footnote. [24][25]
- ✦ **Equity must be measured, not assumed.** ASR disparities are documented; safety-net evidence is absent. Your data has to fill the gap. [20][34]



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.substack.com

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, Cristine 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
- [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/AIp2600203

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)
- [42] Kaiser mental health staff raise concerns about AI recording tool. CalMatters. June 2026. [calmatters.org](https://www.calmatters.org)

References (4 of 5)

- [35] Kanaparthi 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] Kaiser mental health staff raise concerns about AI recording tool. CalMatters. June 2026. calmatters.org

SESSION SURVEY

We value your opinion!
Please take a quick minute to complete our exit poll.



ONGOING SUPPORT




TECHNOLOGY ADVANCES

The technology continues to advance and the evidence base grows.



CONNECT WITH US

Shoot us an email, join us for a webinar, give us a call, visit the website, or even better, sign up for our newsletter to ensure you are informed on updates, events and opportunities.



OUR RESOURCES

DIY kind of person? We have numerous resources and a reliable network to get your answer. We're federally funded so our information and resources are at your disposal.



THE CTRC

CTRC is committed to advancing equitable and sustainable access to our communities via the adoption of digital health services.



AI COHORT SERIES

STAY CONNECTED



Scan the QR Code or visit HealthcareAIbootcamp.org



THANK YOU!

Next Up: OFFICE HOURS



2026 AI Cohort Series – Virtual Office Hours

National Telehealth Technology Assessment Resource Center



THANK YOU FOR ATTENDING

