Southeastern Family Medicine Forum

Top 10 Uses of Artificial Intelligence in Healthcare

What it can actually do for you — and how to tell the real from the sold

Jennifer D. Gholson, MD

Board-Certified Family Physician · For Family Physicians

SESSION OVERVIEW

Learning Objectives

1

Describe the leading applications

Identify the leading clinical and administrative uses of AI in practice today — from ambient documentation and inbox management to autonomous diagnostics and prior authorization.

2

Evaluate benefits and limitations

Weigh the benefits and limitations of these tools, including accuracy, hallucination, algorithmic bias, model drift, deskilling, and patient privacy.

3

Apply a framework for safe adoption

Apply a structured framework — PROVEN — to decide whether a given AI product is appropriate, safe, and worth adopting in your own setting.

WHERE WE ARE

This Is Not a Talk About the Future

81%

of physicians now use AI professionally — up from 38% in 2023

76%

say AI improves their ability to care for patients

70%

say AI can automate the tasks that drive burnout

88%

have some level of concern about skill loss

Physician AI use more than doubled in three years. The tools arrived before the training did.

85% of physicians want a say in AI adoption decisions at their institution — and most are not being asked. **The purpose of today is to make you the person in the room who knows what to ask.**

Source: AMA Center for Digital Health and AI. 2026 Physician Survey on Augmented Intelligence (n=1,692; fielded January–February 2026).

WHERE WE ARE

What Physicians Are Actually Using AI For

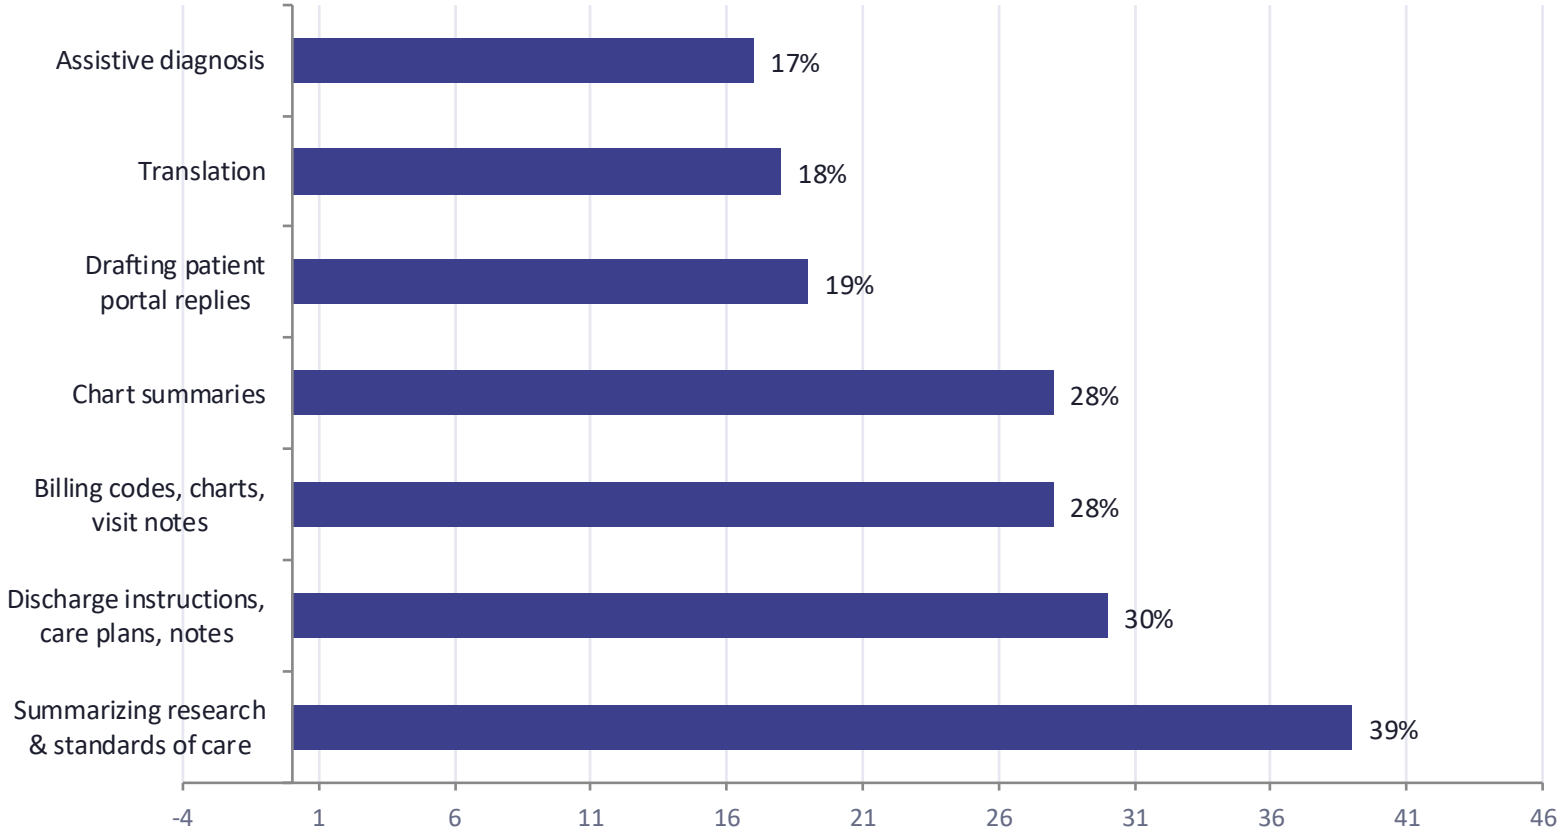

Read the shape of this

The top uses are administrative, not diagnostic.
Physicians reached first for the thing that was stealing their evenings.

Assistive diagnosis sits last, at 17%.
The clinical use cases are the most impressive — and the least adopted.

That gap is the opportunity in this room.

Source: AMA 2026 Physician Survey on Augmented Intelligence. Respondents used a mean of 2.3 use cases, up from 1.1 in 2023.

1

The Top 10

What AI can do for a family medicine practice right now

OBJECTIVE 1

The Top 10 at a Glance

1. **Ambient documentation**
   The AI scribe writes your note
2. **Inbox and chart summarization**
   Drafted replies; 400 pages in 30 seconds
3. **Diagnostic decision support**
   A second reader for the patient you're stuck on
4. **AI-ECG and the digital stethoscope**
   A weak heart on a normal-looking tracing
5. **Autonomous retinopathy screening**
   A specialist diagnosis in your own exam room
6. **Imaging and computer-aided detection**
   1,451 FDA-authorized devices and counting
7. **Risk stratification and panel management**
   Who on my panel is about to get sick?
8. **Prior authorization**
   The same weapon, pointed both ways
9. **Coding and revenue cycle**
   Getting paid for what you already did
10. **Drug discovery and precision medicine**
    A Nobel Prize, and your future formulary

TOP 10 · #1

Ambient Documentation — The AI Scribe

What it does

- **The visit is recorded; a draft note appears in the EHR before you leave the room.**
- Abridge, Microsoft DAX Copilot, Nabla, Suki, Ambience. Most integrate directly into Epic.
- Replicated findings: lower burnout scores, lower cognitive task load, less after-hours “pajama time,” more eye contact with the patient.
- UW Health pragmatic RCT: clinically meaningful burnout reduction and ~30 fewer minutes of documentation per day.
- UChicago: 8.5% less total EHR time; >15% less time writing notes.

29% and climbing

of physicians now use an ambient scribe. It is the fastest-adopted generative AI tool in the history of medicine.

The catch

A randomized trial at UCLA comparing DAX and Nabla to usual care found time-in-note fell by only 23–41 seconds, and 15% of physicians given a tool never used it.

The replicated benefit is cognitive, not chronometric — and you still sign the note.

For your practice: Start here. It is the lowest-risk, highest-satisfaction entry point — but pilot it and measure your own baseline first.

Sources: Olson KD et al. JAMA Netw Open. 2025;8(10):e2534976; UW Health pragmatic RCT (2025); UCLA 3-arm RCT (NEJM AI, 2025); Doximity 2026 State of AI in Medicine.

TOP 10 · #2

The Inbox — Drafted Replies and Chart Summarization

What it does

- **AI drafts a reply to the patient portal message; you edit and send.**
- **Chart summarization: a 400-page record from the hospital, condensed before you walk in.**
- Pre-visit summaries, hospital discharge reconciliation, specialist letter digests.
- 19% of physicians already use AI to draft portal replies; 28% use it for chart summaries.
- Studies find the drafts often save little clock time — but reduce emotional exhaustion and raise the empathy rating of the reply.

400 pages → 30 seconds

Chart summarization may be the most underrated use case in family medicine — and the most clinically useful.

The catch

The draft is a draft. It can be confidently wrong, it can be too agreeable, and it does not know what you know about this patient.

You are the author, not the approver. Read every word before it goes out under your name.

For your practice: The highest-value target may not be the reply — it is the summarization of a record you would otherwise skim.

Sources: AMA 2026 Physician Survey on Augmented Intelligence; AMA practice-time data (2024).

TOP 10 · #3

Diagnostic Decision Support — A Second Reader

What it does

- **Generates or broadens a differential for the patient you cannot figure out.**
- Large language models now perform strongly on diagnostic vignettes and board-style questions.
- Purpose-built tools (Isabel) and general-purpose models both used in practice.
- Best use is not the common case — it is the zebra, the rare disease, the patient who has been to four specialists.
- Also: patient-facing explanation, plain-language handouts, translation.

Only 17%

of physicians report using AI for assistive diagnosis — the least-adopted use case on the AMA list.

The catch

Vignette performance is not bedside performance. A model given a clean paragraph is not a model given your messy, incomplete, contradictory chart.

Automation bias is the real danger: the machine sounds certain, so we stop thinking.

For your practice: Use it to widen the differential, never to close it. Ask it what you might be missing — not what the answer is.

Source: AMA 2026 Physician Survey on Augmented Intelligence.

TOP 10 · #4

AI-ECG — Seeing a Weak Heart on a Normal Tracing

What it does

- **An AI reads a standard 12-lead ECG and flags patients with low ejection fraction — a tracing that looks unremarkable to the human eye.**
- EAGLE trial: pragmatic cluster RCT, 120 primary care teams across 45 clinics including rural sites, 22,641 adults.
- AI flag delivered through the EHR prompted an echo.
- New low-EF diagnoses rose from 1.6% to 2.1% (OR 1.32) — real patients who would otherwise have been missed.
- The Mayo algorithm is now licensed into a digital stethoscope (Eko); murmur and low-EF detection are FDA cleared.

1.6% → 2.1%

New low-EF diagnoses in a randomized trial — asymptomatic heart failure found from a test you already order.

The catch

It is a screening flag, not a diagnosis — every positive needs an echo, and that means downstream cost and some false alarms.

And it was validated in Minnesota and Wisconsin. Ask what it does in a Mississippi population.

For your practice: This is specialist-grade screening running on equipment you already own. Ask your health system whether it is enabled.

Source: Yao X et al. Artificial intelligence–enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med. 2021;27:815–819.

TOP 10 · #5

Autonomous Retinopathy Screening — In Your Own Exam Room

What it does

- **Your staff photographs the retina. The AI renders the diagnosis. No ophthalmologist reads it.**
- **The first FDA-authorized autonomous AI diagnostic in any field of medicine (2018, De Novo pathway).**
- Pivotal trial ran in primary care offices, not eye clinics: sensitivity 87.2%, specificity 90.7%, imageability 96.1%.
- Output is actionable: “more than mild diabetic retinopathy — refer,” or “negative — rescreen in 12 months.”
- LumineticsCore (formerly IDx-DR) and Eyenuk EyeArt are both cleared.

Autonomous

The AI makes the diagnosis without a physician overread — the first time that has ever been permitted in medicine.

The catch

It detects diabetic retinopathy. It does not detect glaucoma, macular degeneration, or the other things an eye exam finds.

It is a screen for one disease, not a substitute for an eye examination — and the patient needs to know that.

For your practice: In a county with no ophthalmologist, this closes a care gap at the point of care — and it is a billable service.

Source: Abràmoff MD et al. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. NPJ Digit Med. 2018;1:39.

TOP 10 · #6

Imaging and Computer-Aided Detection

What it does

- **Radiology is where regulated AI actually lives: 76% of all FDA-authorized AI devices.**
- Triage: intracranial hemorrhage, pulmonary embolism, and pneumothorax pushed to the top of the worklist.
- Detection: lung nodules, fractures, mammography, bone age.
- Colonoscopy CADe: a meta-analysis of 44 randomized trials found an 8% absolute increase in adenoma detection rate.
- You are already consuming this technology — usually without being told.

1,500 devices and climbing

FDA-authorized AI-enabled medical devices through the end of 2025 — including a record 331 in 2025 alone.

The catch

Most were cleared through the 510(k) pathway — meaning "substantially equivalent to a predicate device," not "proven to improve patient outcomes in a trial."

Clearance is a floor, not an endorsement.

For your practice: Ask your radiology group and your hospital which AI tools are reading your patients' films, and what their performance data shows.

Sources: FDA AI-Enabled Medical Device List (data through December 2025); Soleymanjahi S et al. Ann Intern Med. 2024;177:1652–1663.

TOP 10 · #7

Risk Stratification and Panel Management

What it does

- **Which of my 2,000 patients is about to get sick — and which are quietly falling through the cracks?**
- Predicting readmission, clinical deterioration, no-shows, and rising risk.
- Finding undiagnosed disease already in your data: CKD, heart failure, undocumented HCC conditions.
- Closing care gaps at scale — the overdue colonoscopy, the uncontrolled A1c, the missing mammogram.
- Directly relevant to any ACO, Medicare Advantage, or value-based contract you participate in.

2,000 patients

Your panel. Most of what a risk model finds was sitting in your EHR the whole time — unread.

The catch

Widely deployed does not mean validated. The Epic sepsis model, in use at hundreds of hospitals, achieved an AUC of just 0.63 on external validation — far below the 0.76–0.83 advertised.

Ask for the external validation, every time.

For your practice: Highest-yield in a value-based contract, where finding the patient and closing the gap is also how you get paid.

Source: Wong A et al. External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern Med. 2021;181(8):1065–1070.

TOP 10 · #8

Prior Authorization — The Same Weapon, Pointed Both Ways

What it does

- **On your side: auto-assembling the request, pulling the clinical documentation, drafting the appeal letter.**
- Growing use — physicians increasingly report using AI for insurance correspondence.
- **On the payer side: AI is being used to review — and to deny — at a scale and speed no human panel could match.**
- Denial-algorithm practices have drawn litigation, state legislation, and congressional scrutiny.
- AMA policy and multiple state medical societies are actively engaged on AI in utilization management.
- The AMA adopted additional policy in June 2026 specifically addressing AI transparency and accountability in coverage/prior-authorization decisions

13 hours a week

Physicians average 13 hours weekly on indirect patient care. Prior authorization is a large and growing share of it.

The catch

This is the one use case where the technology is being deployed against you as fast as it is being deployed for you.

A denial generated in two seconds still takes you twenty minutes to appeal. Automation without accountability is not efficiency — it is a cost shift.

For your practice: Use it to fight back. And engage your medical society — this is where policy, not procurement, decides the outcome.

Sources: AMA 2026 Physician Survey on Augmented Intelligence; AMA practice-time data.

TOP 10 · #9

Coding, Documentation Integrity, and Revenue Cycle

What it does

- **Suggests the E/M level supported by the note you actually wrote.**
- Surfaces HCC conditions documented in the chart but never coded — which drives risk adjustment and shared savings.
- Flags documentation that will not survive an audit before the claim goes out.
- Automates claim scrubbing, denial prediction, and appeal generation.
- 28% of physicians already use AI for documentation of billing codes and visit notes.

The money is real

In a risk-bearing contract, unrecorded diagnoses are unfunded care — often the fastest measurable ROI of any AI tool.

The catch

The note must reflect the visit — not the algorithm's revenue optimization.

A tool that nudges you toward a higher level of service every single time is not a coding assistant. It is a compliance liability with a friendly interface.

For your practice: Demand to see how often it recommends UP versus DOWN. If it never recommends down, that tells you what it was built to do.

Source: AMA 2026 Physician Survey on Augmented Intelligence.

TOP 10 · #10

Drug Discovery and Precision Medicine

What it does

- **AlphaFold predicted the three-dimensional structure of essentially every protein known to science — roughly 200 million of them.**
- It solved, in a few years, a problem biology had worked on for fifty.
- **The 2024 Nobel Prize in Chemistry went to this work — Demis Hassabis and John Jumper, alongside David Baker for computational protein design.**
- AI-designed drug candidates are now in human clinical trials.
- Pharmacogenomics and tumor molecular profiling are moving the same way.

200 million

Protein structures predicted — essentially every protein known to science, released openly to the world.

The catch

None of this is on your desktop on Monday.

The honest answer to “what does this do for my practice today” is: nothing. What it does is determine what is in your formulary in ten years.

For your practice: You will not use this tool. You will prescribe its output for the rest of your career.

Source: Jumper J et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596:583–589. Nobel Prize in Chemistry, 2024.

2

Where It Breaks

Hallucination · Bias · Drift · Automation bias · Deskilling

OBJECTIVE 2 · LIMITATIONS

Five Failure Modes You Must Be Able to Name

1	**Hallucination**	The model produces fluent, confident, entirely fabricated content — a citation that does not exist, a finding that was never documented. It has no native way to say "I don't know."
2	**Algorithmic bias**	A widely used commercial risk algorithm used healthcare cost as a proxy for health need. Because less is spent on Black patients, the model systematically judged them healthier. Correcting it would have raised the share of Black patients flagged for extra help from 17.7% to 46.5%.
3	**Model drift**	Performance decays as populations, coding practices, and clinical care change. A model validated in 2022 is not automatically valid in 2026 — and almost nobody is monitoring.
4	**Automation bias**	We defer to the machine, especially when tired. The confident output suppresses the doubt that would have saved the patient. Alert fatigue is the same disease in a different costume.
5	**Poor external validity**	The Epic sepsis model: advertised AUC 0.76–0.83, external validation AUC 0.63. Deployed at hundreds of hospitals for years before anyone checked.

Sources: Obermeyer Z et al. Science. 2019;366:447–453; Wong A et al. JAMA Intern Med. 2021;181:1065–1070.

OBJECTIVE 2 · LIMITATIONS

The Deskilling Signal — The Study That Should Worry Us

28.4% → 22.4%

Adenoma detection rate on colonoscopies performed WITHOUT AI, before versus after these same endoscopists began routinely using AI.

A 20% relative drop. 19 experienced endoscopists, each with 2,000+ prior colonoscopies.

The paradox

Across 44 randomized trials, AI-assisted colonoscopy RAISES adenoma detection by about 8% absolute. The tool works. And using the tool appears to erode the skill underneath it.

Why this matters to you

This is the first real-world evidence of automation-induced deskilling in medicine tied to a patient-relevant outcome.

88% of physicians in the AMA survey report some concern about skill loss — and the concern is highest among the earliest-career doctors, who have the most to lose.

The question is not whether to use AI. It is what happens to the resident who has never built a differential unaided, or written a note from a blank page.

This is not an argument against AI. It is an argument for staying a reader — not becoming a rubber stamp.

Source: Budzyń K et al. Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy. Lancet Gastroenterol Hepatol. 2025;10(10):896–903.

OBJECTIVE 2 · LIMITATIONS

Privacy, HIPAA, and Who Is Actually Liable

Consumer AI tools are not HIPAA-covered

The free or personal tier of a general chatbot is not covered by a Business Associate Agreement. Pasting patient informationinto it is a disclosure. Enterprise and health-system deployments under a signed BAA are a different matter — know which one you are using.

Read what the vendor does with your data

Does the contract permit the company to train its models on your patients' records? Who owns the data? What happens to it ifyou cancel? These answers live in the contract, not the sales deck.

Recording consent for ambient scribes

The scribe records the visit. Consent requirements vary by state and by institution. Have a scripted, documented consent — and an easy way for the patient to decline.

Liability needs clear guardrails

You are responsible for the note you sign and the decision you make. "The AI recommended it" is not a defense, and the standard of care remains the physician's. Clear liability frameworks are physicians' top regulatory ask. Physician responsibility remains, but liability allocation involving vendors and institutions is evolving and can depend on facts/contracts/jurisdiction.

3

How to Evaluate an AI Product

Because nobody else is going to do it for you

OBJECTIVE 3 · EVALUATION

The Regulatory Reality: You Are the Last Line

1,451

FDA-authorized AI-enabled medical devices through the end of 2025 — 331 in 2025 alone, a record. Roughly 76% are radiology.

Most

Most generative-AI tools physicians use—scribes, inbox assistants and general-purpose LLMs—are not FDA-authorized medical devices

What that actually means

"FDA-cleared" is not "proven."
Most AI devices clear through 510(k) — substantially equivalent to an existing product. That is a similarity test, not an outcomes trial.

Most tools you'll be sold aren't devices at all.
Scribes, summarizers, and coding assistants fall outside device regulation — as administrative support or exempt decision support.

So no regulator is checking these for you.
The evaluation is yours to do. Here is how.

August 18, 2026-**FDA released its new discussion paper, "Considerations for the Regulation of Generative AI-Enabled Medical Devices."** FDA is actively soliciting input on how these products should be evaluated before market entry and monitored afterward, including foundation models and agentic AI. Comments are open through October 19

Sources: FDA AI-Enabled Medical Device List (through December 2025); Congressional Research Service, FDA Regulation of AI-Enabled Devices, June 2026.

OBJECTIVE 3 · EVALUATION

PROVEN — Six Questions Before You Sign Anything

P	**Problem**	What exact problem does this solve, and do I actually have it? Beware the solution in search of a pain point.
R	**Regulated**	Cleared, exempt, or unregulated? Which pathway — 510(k), De Novo, PMA? Ask for the clearance number and look it up.
O	**Outcomes**	Peer-reviewed and prospective, or a white paper and a testimonial? Was it validated anywhere the vendor doesn't own?
V	**Validated in patients like mine**	What population trained it? What is the positive predictive value at MY prevalence — not the sensitivity in their cohort?
E	**Embedded**	Does it live inside my EHR and my workflow — or is it one more tab, one more login, one more click?
N	**Negotiated**	BAA. Data ownership and training rights. Indemnification. Total cost. Drift monitoring. Exit terms and data portability.

If the vendor cannot answer V, the rest does not matter. Sensitivity in their population is not performance in yours.

OBJECTIVE 3 · EVALUATION

The Vendor Meeting: What to Ask, What to Listen For

Ask the vendor this	A good answer sounds like	A bad answer sounds like
“Show me the peer-reviewed evidence.”	A published, prospective study at a site they don’t own.	“We have a white paper.” “It’s in press.”
“What population was it trained on?”	Specific demographics, sites, and years — and the known gaps.	“A large, diverse national dataset.”
“What is the PPV in a panel like mine?”	They ask about your prevalence before answering.	They repeat the sensitivity.
“How do you monitor for drift?”	Ongoing performance monitoring with a reporting path to me.	“The model is continuously improving.”
“Do you train on my patients’ data?”	A clear no, in writing, in the contract.	“All data is de-identified and secure.”
“Who is liable if it’s wrong?”	A specific indemnification clause they will show me.	“The clinician is always in the loop.”

Bring a colleague who will actually use it. And insist on a pilot with a baseline you measured yourself — before the vendor measures it for you.

OBJECTIVE 3 · EVALUATION

Red Flags — Walk Away

“It’s proprietary.”

Performance data you cannot see is performance data you cannot trust.

“FDA registered.”

‘FDA registered’ ≠ FDA authorized. Registration/listing does not mean FDA reviewed the device for safety and effectiveness.”

A single accuracy number.

“95% accurate” with no population, no denominator, and no PPV is marketing, not evidence.

Evidence is the vendor’s own study.

No independent external validation means nobody has checked their homework.

No drift monitoring.

The model will degrade. If nobody is watching, you will find out from a bad outcome.

No indemnification clause.

They keep the upside; you keep the liability. Read the contract.

You do not have to be a data scientist to do this. You have to be a skeptic with a checklist — which is what medicine trained you to be.

KEY TAKEAWAYS

Five Things to Take Back to the Office

1. AI is already in your practice. 81% of your colleagues use it — and the tools arrived before the training did.
2. The wins that pay off fastest are administrative: the note, the inbox, the prior auth, the code. Start there.
3. The wins that will impress you are the ones that see what you can't — a failing heart on a normal ECG, retinopathy in your own exam room.
4. Every benefit has a shadow: hallucination, bias, drift, automation bias, and now real evidence of deskilling.
5. Most generative AI tools used in everyday practice are not FDA-authorized medical devices. You are still the quality check. Use PROVEN — and never sign a note you didn't read.

The physician who knows how to interrogate the machine will always be more valuable than the machine.

DISCUSSION

Questions

What are you already using? What is your practice about to buy — and has anyone asked the PROVEN questions yet?

EVIDENCE BASE

References and Resources

- American Medical Association, Center for Digital Health and AI. 2026 Physician Survey on Augmented Intelligence. March 2026.
- Olson KD, et al. Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout. JAMA Netw Open. 2025;8(10):e2534976.
- Ambient AI Scribes in Clinical Practice: A Randomized Trial. NEJM AI. 2025;2(12).
- Yao X, et al. Artificial intelligence–enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med. 2021;27:815–819.
- Abràmoff MD, et al. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. NPJ Digit Med. 2018;1:39.
- Budzyń K, et al. Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy. Lancet Gastroenterol Hepatol. 2025;10(10):896–903.
- Obermeyer Z, et al. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366:447–453.
- Wong A, et al. External validation of a widely implemented proprietary sepsis prediction model. JAMA Intern Med. 2021;181(8):1065–1070.
- Jumper J, et al. Highly accurate protein structure prediction with AlphaFold. Nature. 2021;596:583–589.
- US FDA. Artificial Intelligence-Enabled Medical Devices list. www.fda.gov · Congressional Research Service. FDA Regulation of AI-Enabled Devices. June 2026.
- Coalition for Health AI (CHAI) — model cards and assurance standards. · AMA Augmented Intelligence Policy and the AMA AI toolkit for physicians.