

AI in Healthcare: Where We Are Now and Ongoing Risk Areas to Guard Against



Important Disclaimers for this Presentation



The information provided in this presentation (including these slides) does not, and is not intended to, constitute legal, medical, or other professional advice; instead, it is for informational purposes only. Information provided in this presentation should not be relied upon for personal, medical, legal, or financial decisions and you should consult an appropriate professional for specific advice that pertains to your situation. You should not act or refrain from acting based on information in this presentation without first seeking legal advice from counsel in the relevant jurisdiction. Only your attorney can provide assurances that the information contained herein—and the interpretation of it—is applicable or appropriate to your particular situation.



Health care providers should exercise their professional judgment in connection with the provision of healthcare services. The information contained in this presentation is not intended to be, nor is it, a substitute for medical diagnosis, treatment, advice, or judgment relative to a patient's specific condition.



The information in this presentation may undergo periodic changes. You should consult with your legal counsel or other professional to ensure that the information contained in this presentation has not changed.

Financial Disclosure

“None of the planners, reviewers, and faculty for this educational activity, in the past 24 months, have relevant financial relationships to disclose with ineligible companies whose primary business is producing, marketing, selling, reselling, or distributing healthcare products used by or on patients.”

Planner/Faculty: David Whitling, MD

Reviewer: Alan Lembitz, MD

There is no commercial support for COPIC CME/CE activities. In accordance with the ACCME Standards for Integrity and Independence in Accredited Continuing Education, COPIC has implemented mechanisms, prior to the planning and implementation of this LIVE CME/CE activity, to identify and mitigate conflicts of interest for all individuals in a position to control the content of the CME/CE activity.

Goals & Purpose

Clinical, social and administrative applications empowered by Artificial Intelligence, particularly those using “Generative Language Models” (like ChatGPT), are being increasingly marketed to and adopted in medical practices and health facilities. These augment other AI-empowered applications for clinical decision support, communications, image processing, document handling, etc. The benefits and risks, particularly patient safety and professional liability implications of these systems are only beginning to be discovered. COPIC, as a professional liability insurer and patient safety organization is intensely interested in understanding how AI-enabled software and devices will impact best practices in the near and intermediate future. One overarching concern for developers and users of all machine-learning technology is “the alignment problem.” This refers to disparities (or even conflicts) between goals sought by AI systems and traditional human values like privacy, equity, autonomy and human welfare. The emerging “AI Safety” movement studies ways to regulate AI systems, to monitor bias and risk, set standards for development, and devise programming and training practices to protect human rights and interests. These can be approached, as an educational agenda.

Target Audience

This **LIVE** activity is designed to meet the educational needs of healthcare professionals who diagnose and treat patients including nurses, residents, student nurses, and physicians' assistants.

Content Validity

Accredited providers will ensure that activities give a fair and balanced view of diagnostic and therapeutic options, and that they will be free of marketing or sales of products or services, in keeping with Standards [1](#), [2](#), and [5](#).

Accreditation & Designation Statements

For MDs & DOs the number of credits designated is the number of credits awarded.

Copic is accredited by the Accreditation Council for Continuing Medical Education (ACCME) to provide continuing medical education for physicians.

Copic designates this LIVE activity for a maximum of 1.0 hour AMA PRA Category 1 Credit(s)[™].

Physicians should claim only the credit commensurate with the extent of their participation in the activity.

Copic is accredited as a provider of Continuing nursing education (CNE) by the American Nurses Credentialing Center's Commission on Accreditation (ANCC).

Copic designates this activity was designated for 1.0 nursing contact hours.

Process for Claiming Credit

In order to earn CME/CE credit learners should complete the evaluation questions that will assess if practitioners have learned the most important recommendations and conclusions from this course. Each LIVE CME activity consists of the full participation of the learner, and a course evaluation. The assessment/evaluation remain locked until the learning activity is completed.

Process for Completing the Activity:

1. Read the Financial Disclosures.
2. Read the target audience, learning objectives, and financial disclosures.
3. Complete the LIVE educational activity.
4. Complete the activity evaluation/assessment on COPIC's LMS platform.

It is estimated that this activity will take approximately 1.0 hours to complete.

First, about me...

Things I am:

- Practicing Emergency Medicine Physician
- Chief Medical Informatics Officer, Boulder Community Health
- COPIC Consultant

Things I am not:

- An AI Data Scientist
- An attorney

Agenda

- How Did We Get Here?
- Truth
- Medico-Legal Pitfalls
- Agentic AI
- ChatGPTHealth
- Regulatory Environment



Where We Were

2024: Proof of Concept for AI

2025: Early Adoption and Scale

2026: Normalization

Compare with:

- 2003: First full-hospital implementation of Epic
- 2019: Implementation of organization-wide EMR at my health care organization



Where We Were



- Many healthcare organizations implemented AI tools for the first time.
- Organizational level:
 - Had to create governance and oversight largely from scratch
- Provider level:
 - Is this part of my workflow?
 - Do I trust AI?
 - Do my patients trust AI? Do I need to get their consent?

Where We Were



AI pioneers fear extinction

Our creations are as great a threat to humanity as nuclear war or pandemics, say hundreds of experts in call to regulate tech

Tom Whipple Science Editor

More than 200 of the world's most distinguished experts in artificial intelligence, including the creator of ChatGPT, have warned of the possibility that the technology could lead to the extinction of humanity.

In a joint statement, backed by the chief executives of the leading AI companies, they said that mitigating this risk "should be a global priority alongside other societal risks such as pandemics and nuclear war".

The technology would be implemented "safely and securely with guard rails". Many experts have expressed concern about the risk of models such as ChatGPT being used to propagate misinformation and cybercrime, as well as causing society-wide disruption to jobs.

The statement, co-ordinated by the US Center for AI Safety (CAS), acknowledged these worries, but said that we should also discuss more serious but less likely threats, including the potential that accelerating AI could lead to the collapse of civilisation.

plan, or unwittingly destroy us. Others worry that overreliance on systems we do not understand leaves us in catastrophic danger if they go wrong.

Don Hendryckx, director of the CAS, said that the statement was a way of "coming out" for many researchers.

"People were much too afraid to speak up earlier," he said. "This establishes it as an intellectually credible concern."

Among the signatories are Sam Altman, chief executive of OpenAI, which made ChatGPT; Geoffrey Hinton, the University of Toronto academic, often

programs to defeat the best players at Go and chess. It was also signed by the heads of Anthropic AI, Inflection AI and Stability AI, and computer science professors from Cambridge, Oxford, Harvard, Yale and Stanford.

Other academics dismissed the statement as unhelpful. Dr Mihail Juba, ethics research fellow at the Alan Turing Institute, called it a "distraction" from more pressing threats from AI.

"The narrative of super-intelligent AI is a familiar storyline from countless Hollywood blockbuster movies, and that

ethics in AI at Oxford University, was suspicious of the motives of some signatories. "I worry that the emphasis over-identical threat is distracting away from more pressing issues like the erosion or defence of democracy, that CEOs of certain companies do not want to face," she said.

"AI can create huge destruction short of existential risk."

Hendryckx said that while those backing the statement differed in their view of the severity of the risk, all believed that there was a potential for catastrophe if AI was handled badly.

SEE PAGES 18-19

A.I. 'COULD WIPE OUT HUMANITY'

Threat 'as bad as nuclear war' MUST be tackled, say tech bosses



AI creators fear the extinction of humanity

- Regulatory environment was a total mess.
 - Some states passed laws
 - Laws confused everyone
 - Adoption vastly outpaced regulation
- Pundits made wild predictions about the future.
 - AI will lead to demise of humanity.
 - Massive job losses

Where We Are Now


- Estimates vary widely, but per one survey, 58% of organizations that provide healthcare in the US are using AI tools.
- Awareness of AI in healthcare is now universal.
- Patients have rapidly adopted use of free AI tools, which are displacing traditional search engine tools like Google.



Where We Are Now



- The regulatory environment is somehow even messier!
 - The Executive Branch has officially entered the chat.
- Pundits continue to make wild predictions about the future.
 - Elon Musk: “Don’t worry about squirreling money away for retirement in 10 or 20 years...It won’t matter.
 - AI and robotics will bring about a world of “zero scarcity.”



What is AI? And What is Truth, Anyway?

It's just math.

Really, really, really complex math that is used to make predictions.

Head 1:

$$Q_1 = K_1 = V_1 = \text{Input Matrix}$$

$$A_1 = Q_1 \cdot K_1^T$$

$$A_1 = \begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix}$$

$$A_1 = \begin{bmatrix} 1 & 0 & 0 \\ 4 & 5 & 0 \\ 7 & 8 & 9 \end{bmatrix} \quad (\text{Masked})$$

$$W_1 = \text{softmax}(A_1)$$

$$O_1 = W_1 \cdot V_1$$

Head 2:

$$Q_2 = K_2 = V_2 = \text{Input Matrix}$$

$$A_2 = Q_2 \cdot K_2^T$$

$$A_2 = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$$

$$A_2 = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 5 & 6 \\ 0 & 0 & 9 \end{bmatrix} \quad (\text{Masked})$$

$$W_2 = \text{softmax}(A_2)$$

$$O_2 = W_2 \cdot V_2$$

Concatenate and Linear Transformation:

$$\text{Concatenate}([O_1, O_2])$$

(Apply Learnable Linear Transformation)

Consider your mobile phone keyboard

- Output: THE
- How does it do this?
 - There are approximately 1000 words in the dictionary that start with “th”
 - 247 words that start with “the”
 - 14 words that start with “thw”
 - ~17x more likely you meant “the”
- (This is what made the iPhone possible)
- Now scale that by trillions of calculations



Which is more likely to be malignant?



AI: Isn't looking for Truth, it is looking for what it has determined to be the most likely correct answer!

What is a Hallucination?

- When AI confabulates something
- Problem is that AI is very good at sounding convincing, and there are no cues to know when it is “confident” or not.

“The mental model I sometimes have [of AI] is as a very smart assistant who has a dozen Ph.D.s but is also high on ketamine like 30 percent of the time.”

K. Roose, NY Times 2025

<https://www.nytimes.com/2025/06/16/magazine/using-ai-hard-fork.html?smid=url-share>





Temperature

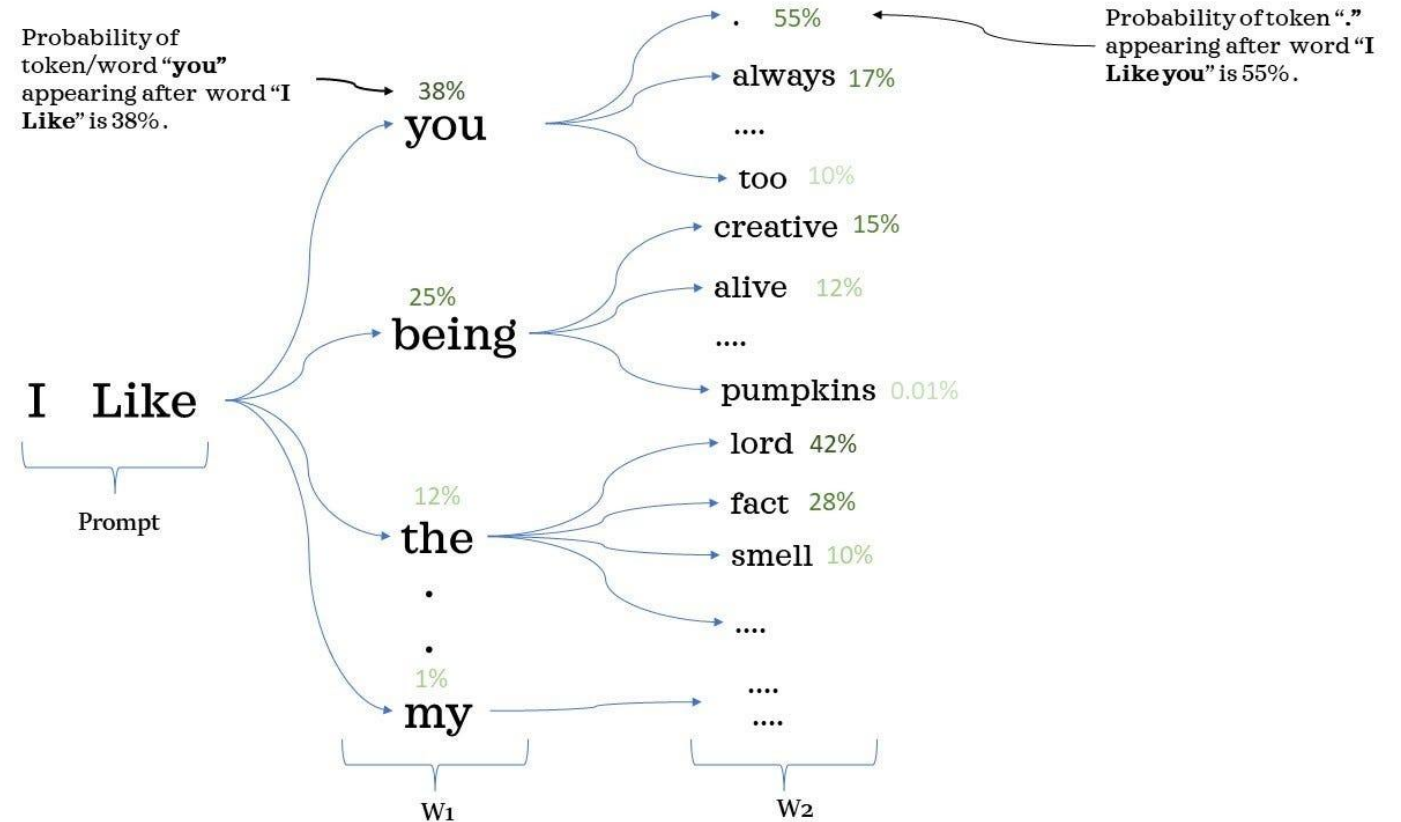
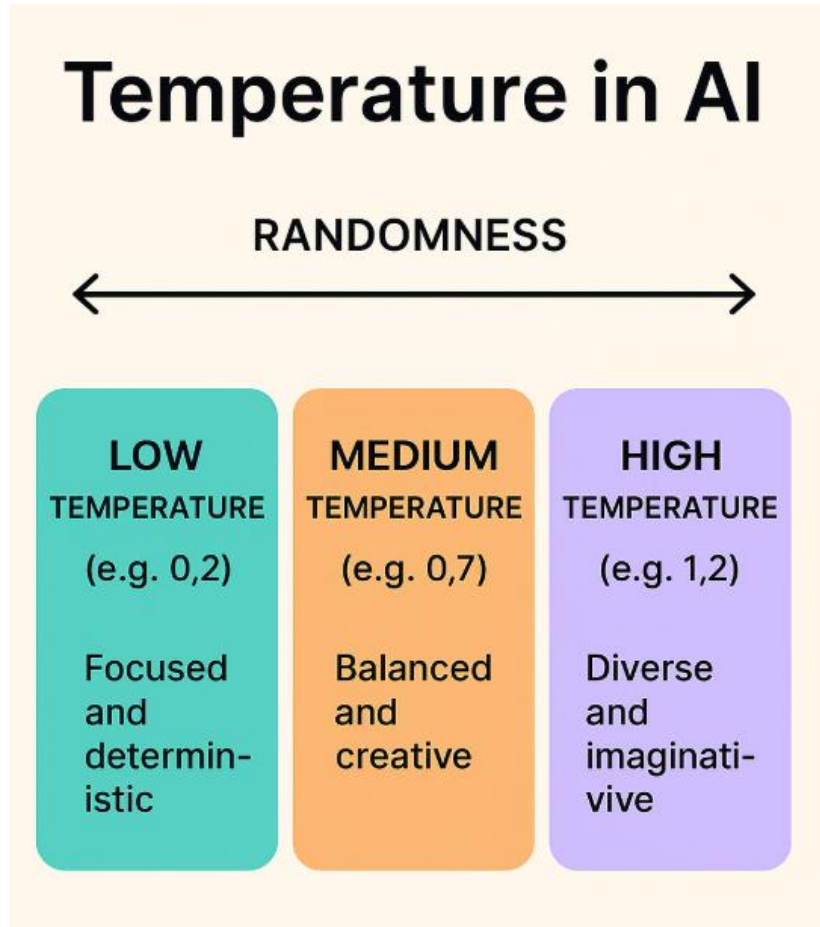
Temperature

What is **Temperature** in AI?

- A parameter that controls the randomness and creativity of a model's output.


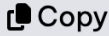
$$p_i = \frac{\exp\left(\frac{z_i}{T}\right)}{\sum_j \exp\left(\frac{z_j}{T}\right)}$$

Understanding “Temperature”


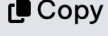


Understanding “Temperature”

Ask AI: What are 10 weird, unique, and fun things to do at the beach? Make a list without descriptions.

 ChatGPT Output (Temperature=0, Top-P = 1) 

1. Build a sandcastle
2. Collect seashells
3. Play beach volleyball
4. Fly a kite
5. Have a picnic
6. Try paddleboarding
7. Play frisbee
8. Go snorkeling
9. Take a long walk along the shore
10. Watch the sunset

 ChatGPT Output (Temperature=1.5, Top-P = 1) 

1. Wade in the shallow waters blowing bubblegum bubbles
2. Create intricate sandcastle sculptures
3. Join an improvised beach volleyball match
4. Take a scenic hike along the cliffs nearby
5. People-watch and guess each person's beach-related story
6. Collect seashells to create your own piece of beach art
7. Learn to surf or try other water sports
8. Kickstart a spontaneous sand fight
9. Attempt shoreline fishing like the locals
10. Engage in friendly competition by organizing a sandcastle building contest

Temperature

- If AI is your “Source of Truth,” recognize that the amount of truth is determined by a single variable!
- Imagine asking UpToDate or Harrison’s Principles of Internal Medicine to be...a little more creative.
- It is estimated that most healthcare-associated AI applications use a temperature setting between 0.1 to 0.3.





De-skilling

Are A.I. Tools Making Doctors Worse at Their Jobs?

Physicians are using the technology for diagnoses and more — but may be losing skills in the process.



THE LANCET
Gastroenterology & Hepatology

ARTICLES · Volume 10, Issue 10, P896-903, October 2025

[Download Full Issue](#)

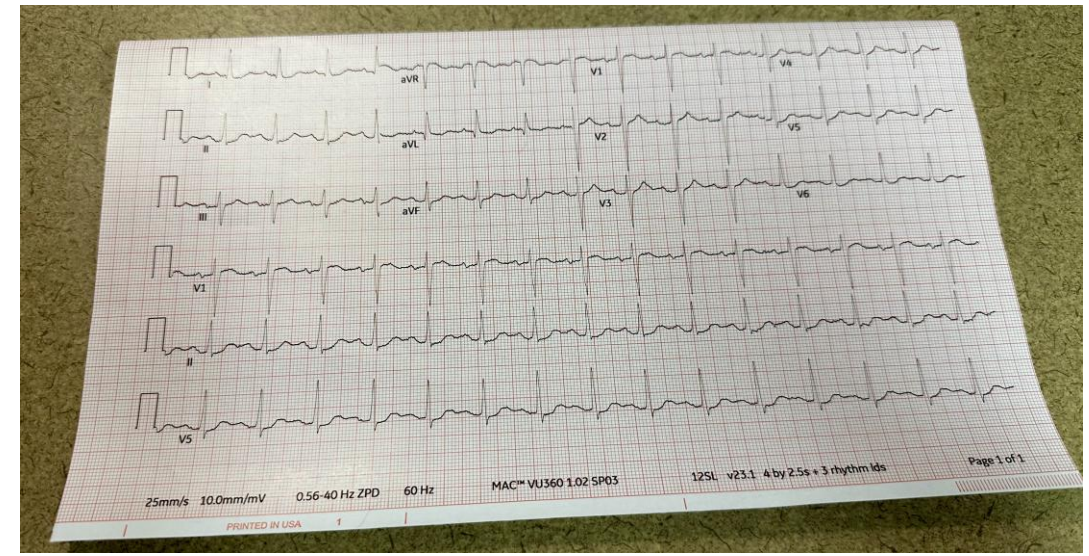
Endoscopist deskilling risk after exposure to artificial intelligence in colonoscopy: a multicentre, observational study

- 1443 colonoscopy patients in Poland
- Looked at rate of detection of adenomas before and after use of AI assistance
- Rate of detection on standard (non-AI) studies fell from **28.4%** to **22.4%** after only three months of use

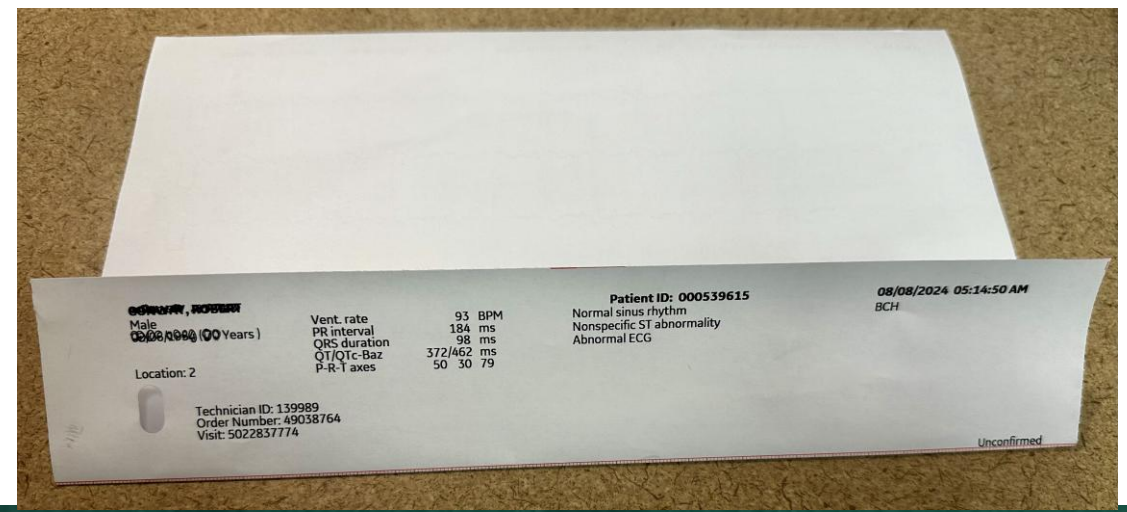
So, is AI making providers worse at their jobs?

A Joke:

How does a cardiologist read an ECG?



How does a
[insert other specialist here]
read an ECG?



AI in Diagnostics

- This is not new. Computer aided diagnostic assistance has been around for a long time.
- With AI, we are just doing it better and faster.
- Same pros, same cons, same legal issues
- New generation of physicians will be much better at using AI, probably worse at other stuff
- Do you ever use a microscope?



When will AI-enabled diagnostic tools become Standard of Care?

- Soon!
- Especially for Radiology
 - In active use for diagnosis of LVO on CTA for stroke
 - Tumor screening
- How long till AI becomes the SOC for ECGs, CXRs?

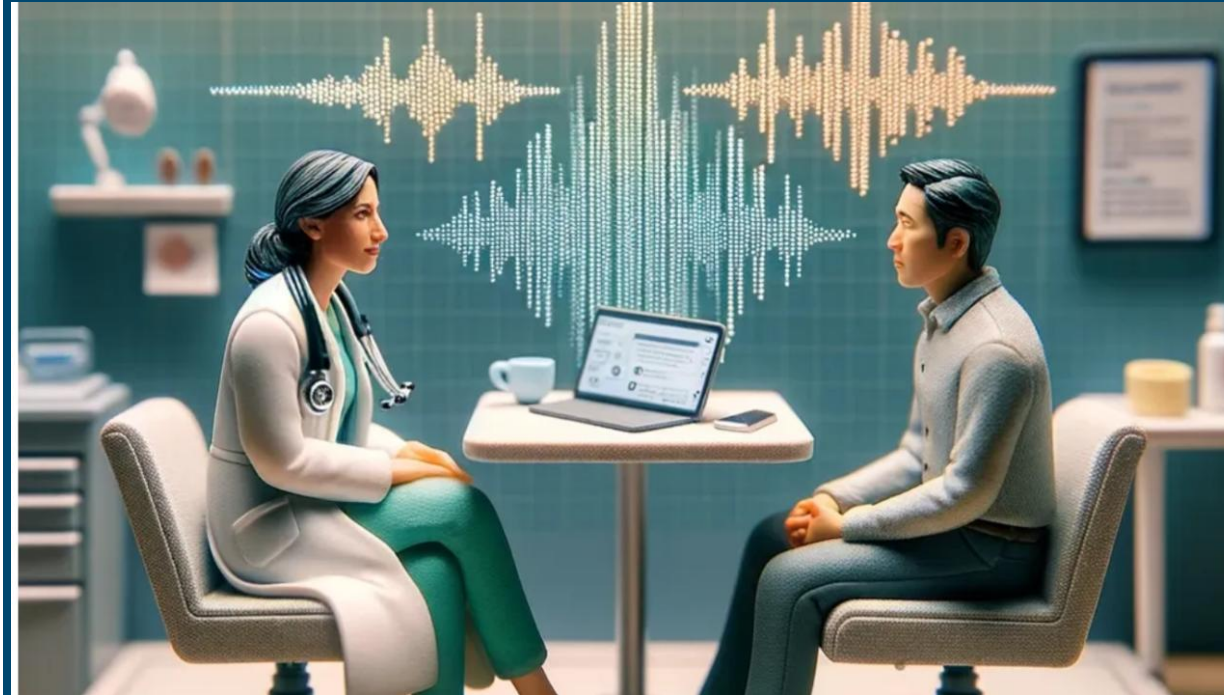




Medico-Legal Pitfalls

Ambient - Legal

- Primary use of AI by providers in healthcare to date is Ambient Virtual Scribe technology.
- AI/Ambient Scribed notes are starting to show up in court!
- To date, there has not been a case I am aware of where AI charting has been a primary issue of case. (But it's coming!)



Ambient - Legal

- Is the audio recording of my AI virtual scribe encounter discoverable in the event of a lawsuit?
- Yes!
- Plaintiffs attorneys want this.
- They want to play the file for the jury.
 - Did you interrupt the patient?
 - Did you seem rushed or busy?
 - Did you incorrectly chart some of the details?
- Find out how long your vendor holds on to this file for, and why.



Audio Recording and Consent

- Class Action lawsuits being filed in California against three large healthcare systems.
- Allege Abridge AI scribe used without consent
- “[Plaintiff] later discovered documentation in the patient portal indicating that he had been “advised” that the visit audio was being recorded and that he had “consented.” The complaint calls that language false.”

Lawsuit Alleges AI Platform Illegally Recorded Patient-Clinician Conversations

Patients Sue Two More Health Systems Over AI Scribe Use, Lack of Consent

Primary risk: Providers not proofreading their note!

And...Voice Transcription Has Already Revealed How Often Some Providers Proofread Their Notes



Fun examples:

- “the patient has a 4cm renal mass oh yeah that was insane last night I can’t believe she picked Johnny in the rose ceremony.”
- “patient was driving erotically prior to accident”
- "Nipples equal round reactive”

Easy to tell that something has gone wrong here. But not with an AI generated note!

STOP Using Disclaimers!

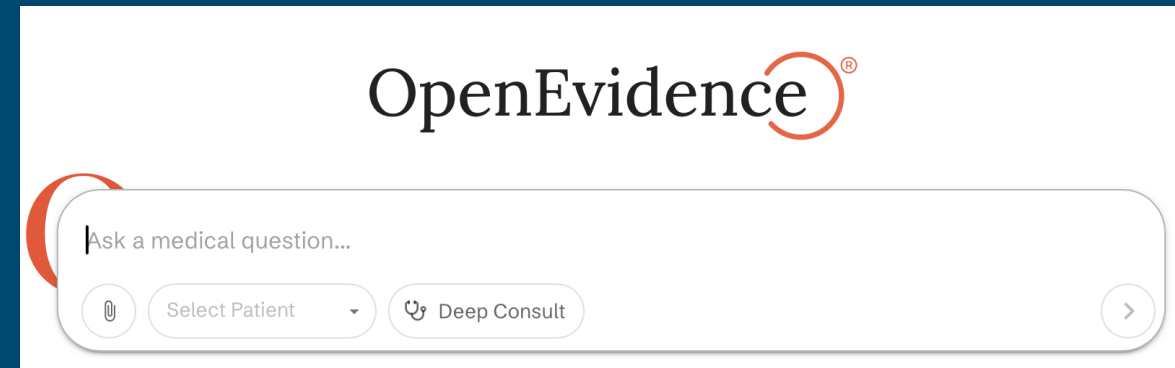
“This note was drafted using an AI-enabled ambient listening tool. Patient consent was obtained for its use. Every attempt has been made to review the chart but some charting errors may be present.”

- They don't help you.
- They give the illusion of a Get Out of Jail Free card.
- If you sign it, you own it!
- Opens door to perception of negligence in your practice.



What About Open Evidence?

- Extremely powerful
- Allows you to ask a question, or enter data
- Competitor to UpToDate
- Some organizations starting to block it.
- Free. For now...
- Can also be used by a plaintiff's attorney!

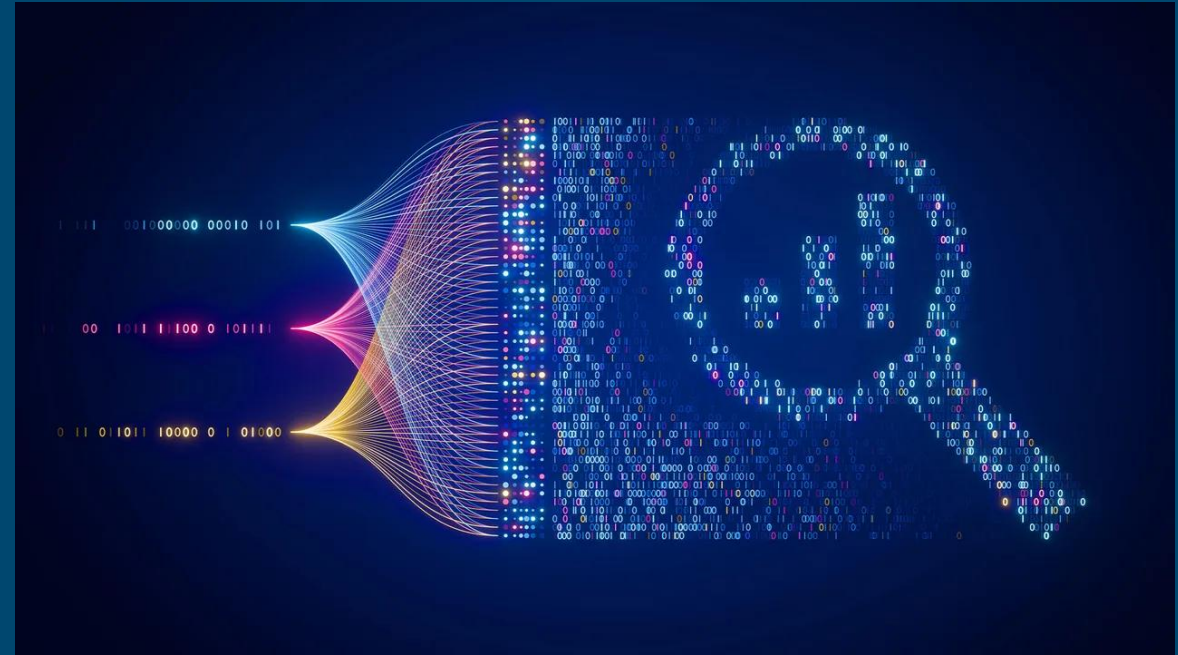




The Future – Agentic AI

What is “Agentic” AI?

- Systems that can **act autonomously to achieve goals**, rather than just responding to individual prompts or executing single tasks.
- 1st Generation: AI used to draft responses, summarize data when asked.
- Common refrain: “Don’t worry, a human provider will always look at this information before finalizing/sending it.”
- What happens when that’s no longer efficient?



Agentic ED Triage?

- Thought Experiment:
- If we fully turned over ED triage to an autonomous AI agent, what accuracy would be necessary for you to feel like it was acceptable to no longer have a human in the mix?
- 90%?
- 98%?
- 100%?



Agentic ED Triage?



(Mistry B, et al. Accuracy and Reliability of Emergency Department Triage Using the Emergency Severity Index: An International Multicenter Assessment. Ann Emerg Med. 2018 May;71(5):581-587.e3.)

- What is accuracy of ED triage for a human RN using the Emergency Severity Index (standard of care in US)?
- 59%!
- Example:
 - 70 y/o M with PAD presents to triage with painful leg.
 - Waits for 4 hours in waiting room.
 - Eventually leg pain determined to be critical ischemia, gets amputated.
 - Multiple physicians sued. Hospital sued because of **triage delay**.

Agentic ED Triage?

- My general understanding of US Legal System: **When something goes wrong, you have to be able to sue someone because of it.**
- So who do you sue in this case?
- AI tool can't explain logic.
- AI Vendor almost certainly indemnified themselves via contract.
- Hospital relying on tool, better than current standard of care.
- Will we start to insure Agentic AI as a separate entity?



ChatGPT Health



- Per OpenAI:
 - **One in four** of its users submits a healthcare-related prompt to its AI Chatbot (ChatGPT) every week!
 - **40 million people** ask ChatGPT healthcare-related questions **every day**.
 - In the US, **3 out of 5 adults** said they used AI tools for their health or healthcare in the past three months.
- This technology is here.
- Compare to urgent care experience or wait to see specialist.

ChatGPT Health

- Several vendors releasing similar products
- They promise:
 - Will not use your data to train model.
 - “Additional privacy and security protections designed for sensitive health information.”
- Not considered a “Healthcare Entity” under current law, and therefore not subject to HIPAA regulations.
- They are free to sell your data or use it for other purposes.
- They rely on “consent” frameworks rather than HIPAA regulations.



ChatGPT Health

naturemedicine

[Explore content](#) ▾ [About the journal](#) ▾ [Publish with us](#) ▾ [Subscribe](#)

[nature](#) > [nature medicine](#) > [brief communications](#) > article

Brief Communication | Published: 23 February 2026

ChatGPT Health performance in a structured test of triage recommendations

[Ashwin Ramaswamy](#) ✉, [Alvira Tyagi](#), [Hannah Hugo](#), [Joy Jiang](#), [Pushkala Jayaraman](#), [Mateen Jangda](#), [Alexis E. Te](#), [Steven A. Kaplan](#), [Joshua Lampert](#), [Robert Freeman](#), [Nicholas Gavin](#), [Ashutosh K. Tewari](#), [Ankit Sakhuja](#), [Bilal Naved](#), [Alexander W. Charney](#), [Mahmud Omar](#), [Michael A. Gorin](#), [Eyal Klang](#) & [Girish N. Nadkarni](#) ✉

- Stress test of 60 clinician-authored vignettes.
- “Among gold-standard emergencies, the system under-triaged 52% of cases, directing patients with DKA and impending respiratory failure to 24-48 hour evaluation rather than the ED.”
- If prompt included text that a friend thought it was nothing serious, 12x more likely to downplay symptoms.

The Dark Side of AI...still to come?



- Free... For Now.
 - Buy tokens in order to ask ChatGPT a question.
 - What if you have to watch an advertisement video in order to use OpenEvidence
 - What if ad was for a specific drug, that only pops up when you search for a related topic?
- Are AI companies intentionally planting doomsday scenarios in the press?

How do I make sense of all this?

AI is clearly changing health care, but is it for the better or for the worse, or some combination of both?

My advice: AI is useful in healthcare when it increases our humanity and harmful when it does the opposite.

Examples:

- Ambient scribe
- Coding
- Mandatory chatbots to schedule





Regulatory Environment

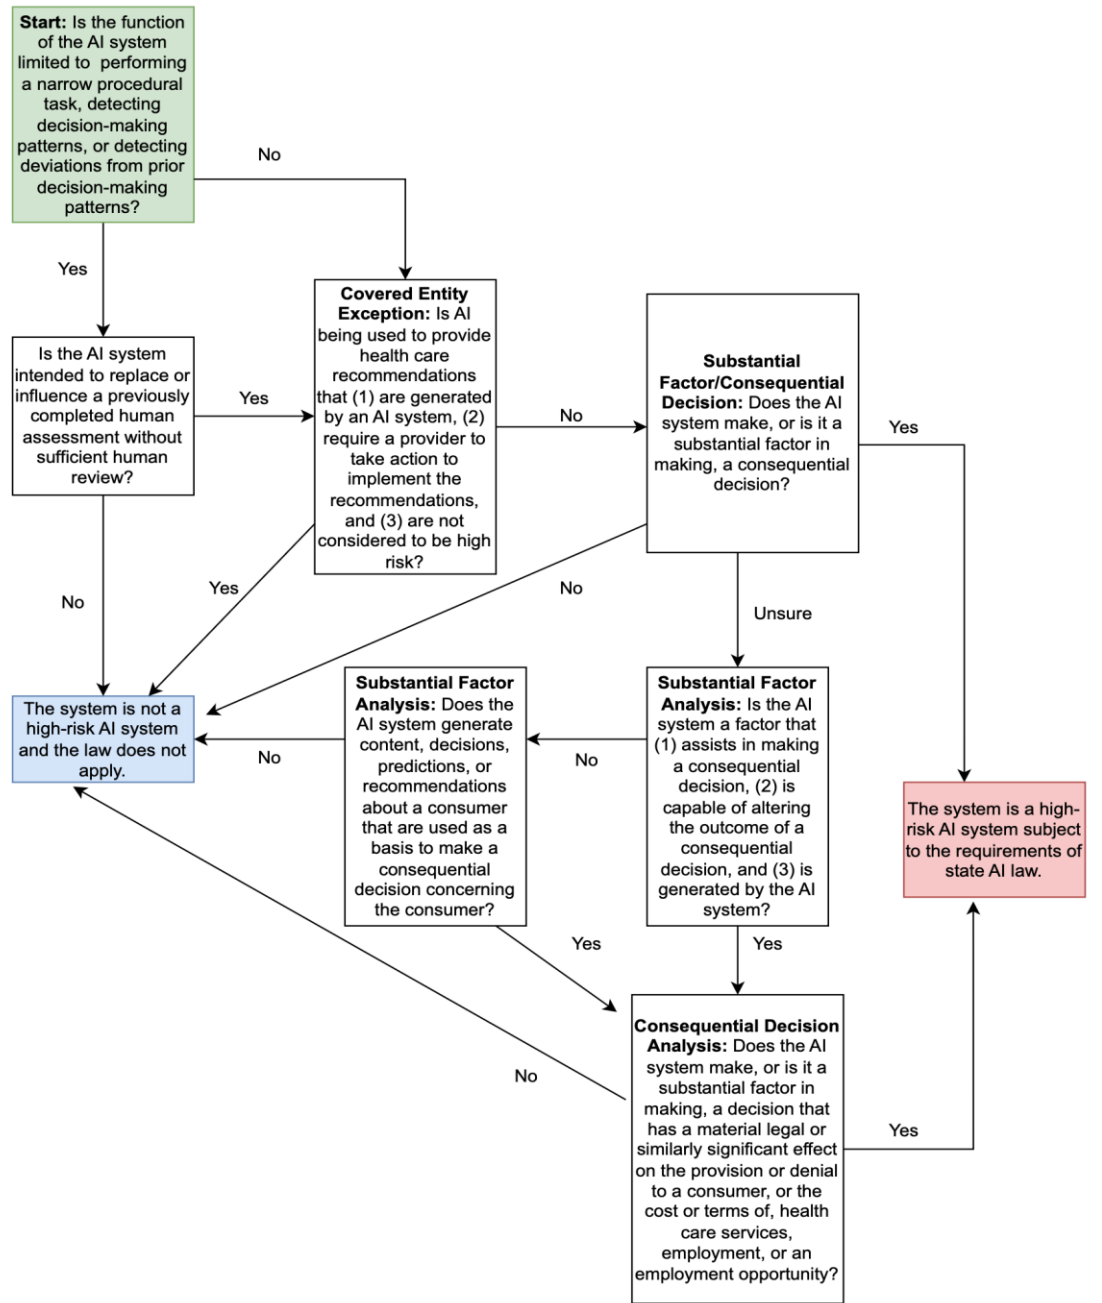
Regulatory Environment

- Is in total chaos.
- Multiple states have passed AI regulations
- Colorado: passed SB 24-205 “Concerning Consumer Protections in Interactions with AI Systems” May 2024.
- Enforcement delayed until 2/1/26, then delayed until 6/30/26
- Status in limbo.



AI Use Flowchart

Colorado AI Law Decision Matrix



Regulatory Environment Federal Level

- Biden 2023 executive order directing HHS to develop first departmental AI strategy.
- Trump rescinded order within first few weeks of taking office.
- Also suspended all meetings of the Health Information Technology Advisory Committee (HITAC) “until further notice.”
- July: “Big Beautiful Bill” initially proposed to centralize all AI regulation and forbid any state regulation of AI for ten years. That portion was removed before passage.



Executive Order

- Issued December 11, 2025
- “Ensuring a National Policy Framework for Artificial Intelligence”
- AI companies “must be free to innovate without cumbersome regulation imposed by states.”
- Creates task force to evaluate state laws and use state eligibility for Broadband Equity Access and Deployment (BEAD) funding as stick



Executive Order

- State laws are “increasingly responsible for requiring entities to embed ideological bias within models.”
- What are they talking about?
- Real world example in Healthcare:
 - AI tool used to determine estimated LOS for patients admitted to hospital.
 - Observation (< 2 midnights) vs. Inpatient (greater than 2 midnights) has BIG financial implications for Medicare patients.
 - What if Race X found to have shorter LOS than Race Y?



So, in summary...

- AI is here, has already changed Healthcare
- AI is not looking for Truth
- The skills of providers will change over time because of AI
- Stop writing disclaimers in documentation.
- Patients have already heavily adopted AI
- Regulatory environment remains in chaos.
- AI should increase our humanity, not the opposite



David Whitling, MD
dwhitling@bch.org

