AI: The Double-Edged Scalpel

COPIC Insights
Sioux Falls, SD
September 16, 2025

Faculty:
Daniel Rosenquist MD FAAFP

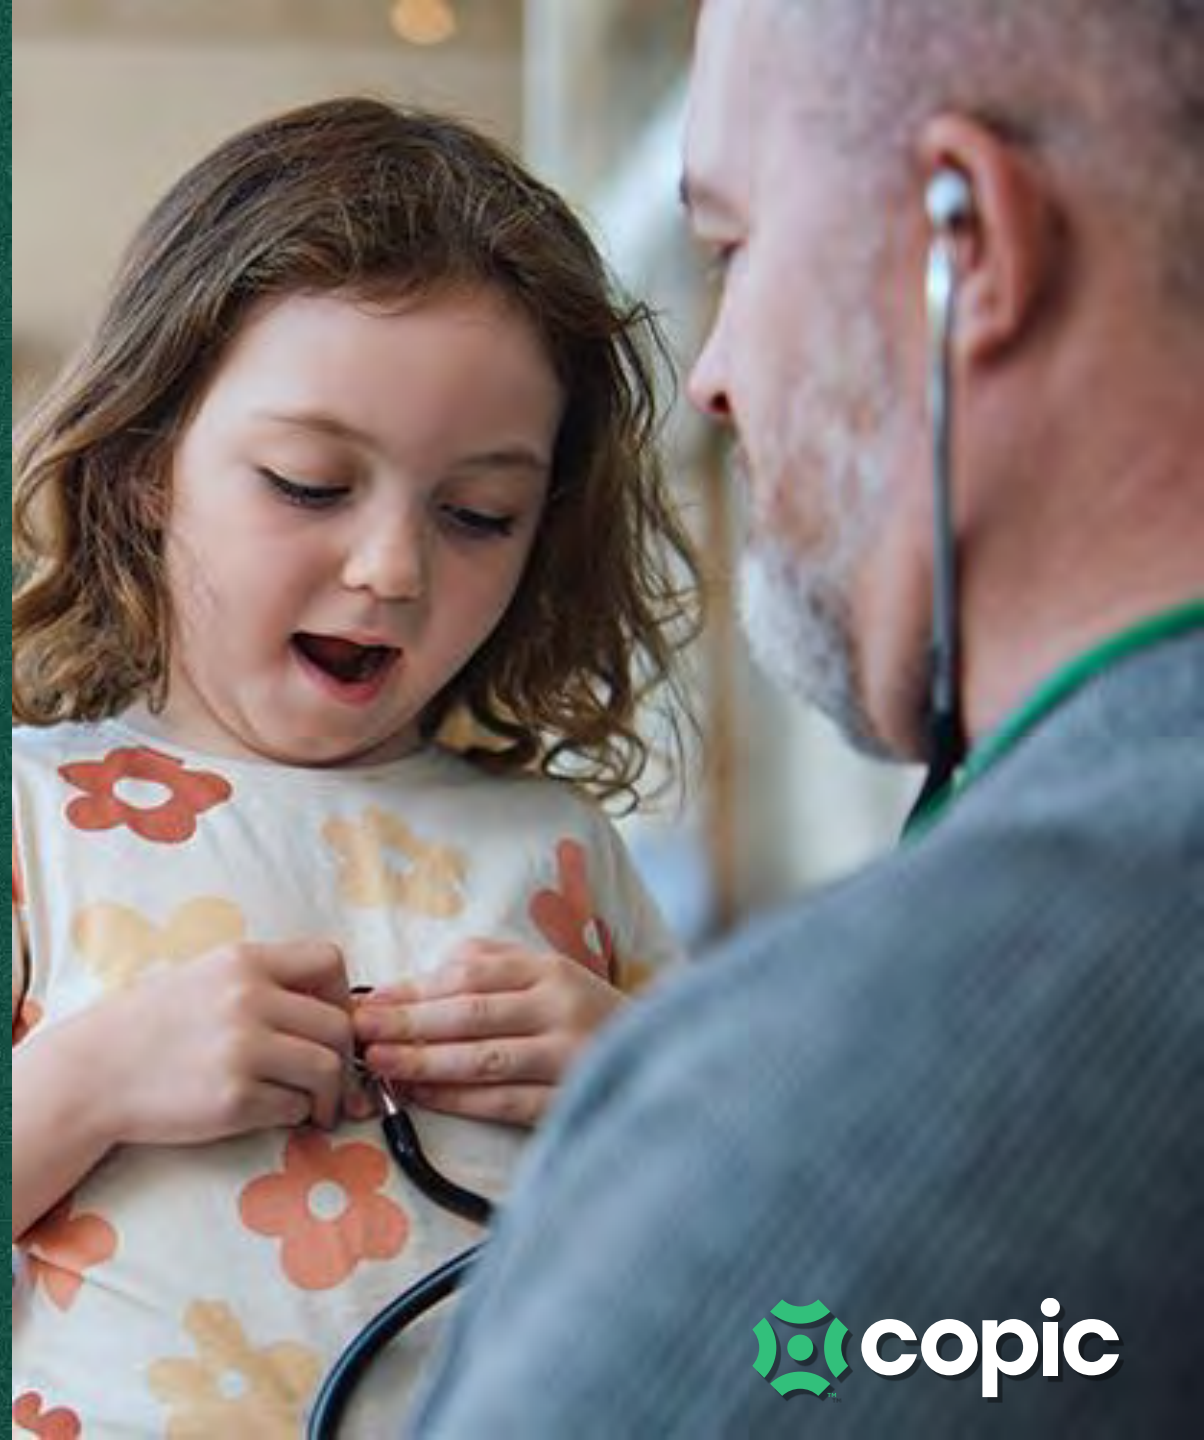

Disclaimer

The material presented here is based on formal education and training, evidence-based research, COPIC recommendations, guidelines and best practices, and personal experience. This content of this lecture is by no means legal advice. Each provider assumes responsibility of the potential risks associated with each educational management procedure, and to remain updated on best practices for any patient related procedures.

Providers participating in any COPIC related workshop, do not hold COPIC or the physician instructor liable for any consequences of individual choices, and personal practice. All patient related outcomes are ultimately the responsibility of the provider.

Financial Disclosures

“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.”

Accreditation & Designation Statements

For nursing the number of credits designated is the number of credits awarded

COPIC is accredited as a provider of Continuing nursing education (CNE) by the American Nurses Credentialing Center's Commission on Accreditation (ANCC). This activity was designated for __ nursing contact hours.

Process for Claiming Credit

To earn CNE credit learners should complete the evaluation questions that will assess if nurses have learned the most important recommendations and conclusions from this course. Each LIVE activity consists of the full participation of the learner, and a course evaluation. The evaluation will open after the learning activity is completed.

Process for Completing the Activity:

1. Read the target audience, learning objectives, and financial disclosures.
2. Complete the LIVE educational activity.
3. Complete the activity evaluation form

It is estimated that this activity will take approximately 1 hour to complete.

Course Learning Objectives

1. **Describe** selected, current and emerging applications of AI in healthcare, and some wider implications for society.

2. **Identify** key risks and limitations of using AI for medical tasks, such as communication, documentation and decision-support.

3. **Evaluate** trends and implications of AI for patient safety, professional liability and regulatory policy.

Agenda

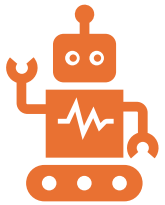

What Is AI and Why Is This All Happening Now?

Current Uses of AI in Medicine

Diagnostic

Generative

Legal Issues Related to AI

First, about me…
Things I am:

- Practicing Family Medicine Physician
- COPIC Consultant

Things I am not:

- An AI Data Scientist
- A math expert
- An attorney
- A cleverly designed cyborg, sent back in time from the future to pacify this audience*

- (*or am I?)

What is AI and Why Is This All Happening Now?

copic

- Bill Gates: Within 10 years, AI will replace many doctors and teachers—humans won’t be needed ‘for most things’
- Published Wed, Mar 26 20259:05 AM EDT

It’s just math.

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)

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

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

Remember: AI isn't looking for the truth, it is looking for what it has determined is the most likely correct answer.

- A journalist used ChatGPT and asked it who he was.
  - ***GPT-3:*** *Harry Guinness is a freelance writer and journalist based in Ireland. He has written for a variety of publications, including The New York Times, The Guardian, The Huffington Post, and Popular Mechanics.*
- Problem: He has never written for the last three.
- The New York Times is grouped far more often with The Guardian and The Huffington Post than it is with the magazines he did write for: Wired, Outside, The Irish Times
- So AI chose the option it thought was most likely.
- Who's fact-checking AI?

copic

But Why Now?

- You need three key things:
  - The Data
  - The Math
  - The Computing Power

copic

1. The Data

- Chat GPT= Generative Pre-trained Transformer

1. The Data

- Chat GPT= Generative Pre-trained Transformer
- Supervised Learning
  - AI models were initially trained with manually labeled data, like a database with pictures of animals, labeled with the names of the animals
  - This is expensive, time-consuming

1. The Data

- Chat GPT= Generative Pre-trained Transformer
- Supervised Learning
  - AI models were initially trained with manually labeled data, like a database with pictures of animals, labeled with the names of the animals
  - This is expensive, time-consuming
- Unsupervised Learning
  - Advances let AI models scour data on their own and form predictions and models.
  - Need a huge repository of digital data, with words, pictures, news etc.
  - Turn it loose on the whole internet

2. The Math

- Chat GPT= Generative Pre-trained Transformer
- Transformer Architecture proposed in research paper in 2017
- Vastly increases speed, efficiency of AI engine by allowing computer to “read” whole sentence at once.

2. The Math

- Chat GPT= Generative Pre-trained Transformer
- Transformer Architecture proposed in research paper in 2017
- Vastly increases speed, efficiency of AI engine by allowing computer to “read” whole sentence at once.
- What does the word “bank” mean in this sentence?
- *“I arrived at the bank after crossing the…”*

2. The Math

- Chat GPT= Generative Pre-trained Transformer
- Transformer Architecture proposed in research paper in 2017
- Vastly increases speed, efficiency of AI engine by allowing computer to “read” whole sentence at once.
- What does the word “bank” mean in this sentence?
- *“I arrived at the bank after crossing the…”*
- You need to know if the sentence ends in “river” or “road.”
- Older math was like reading left to right. Newer reads all the words at once. Much faster!

3. The Computing Power

AI in Diagnostics

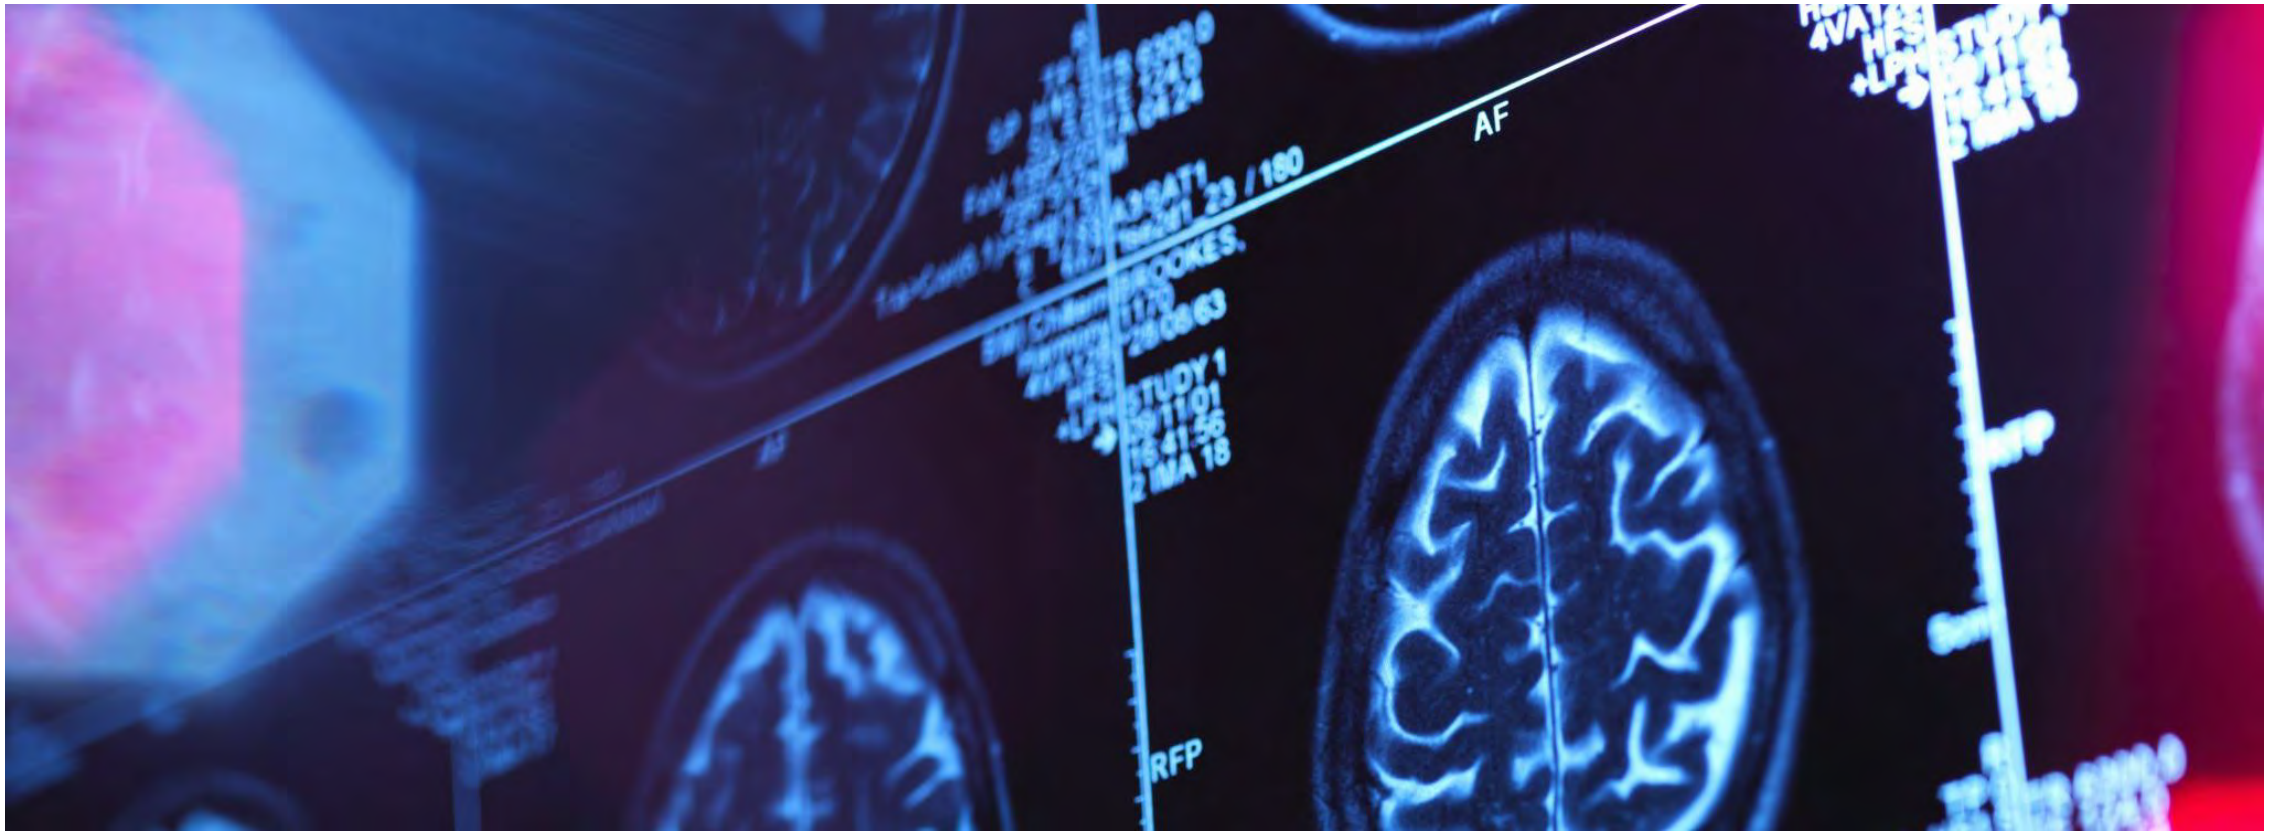

AI in Diagnostics

- This is not new. We are just doing it better and faster.
- Same pros, same cons, same legal issues

A Question?

How does a cardiologist read an ECG?

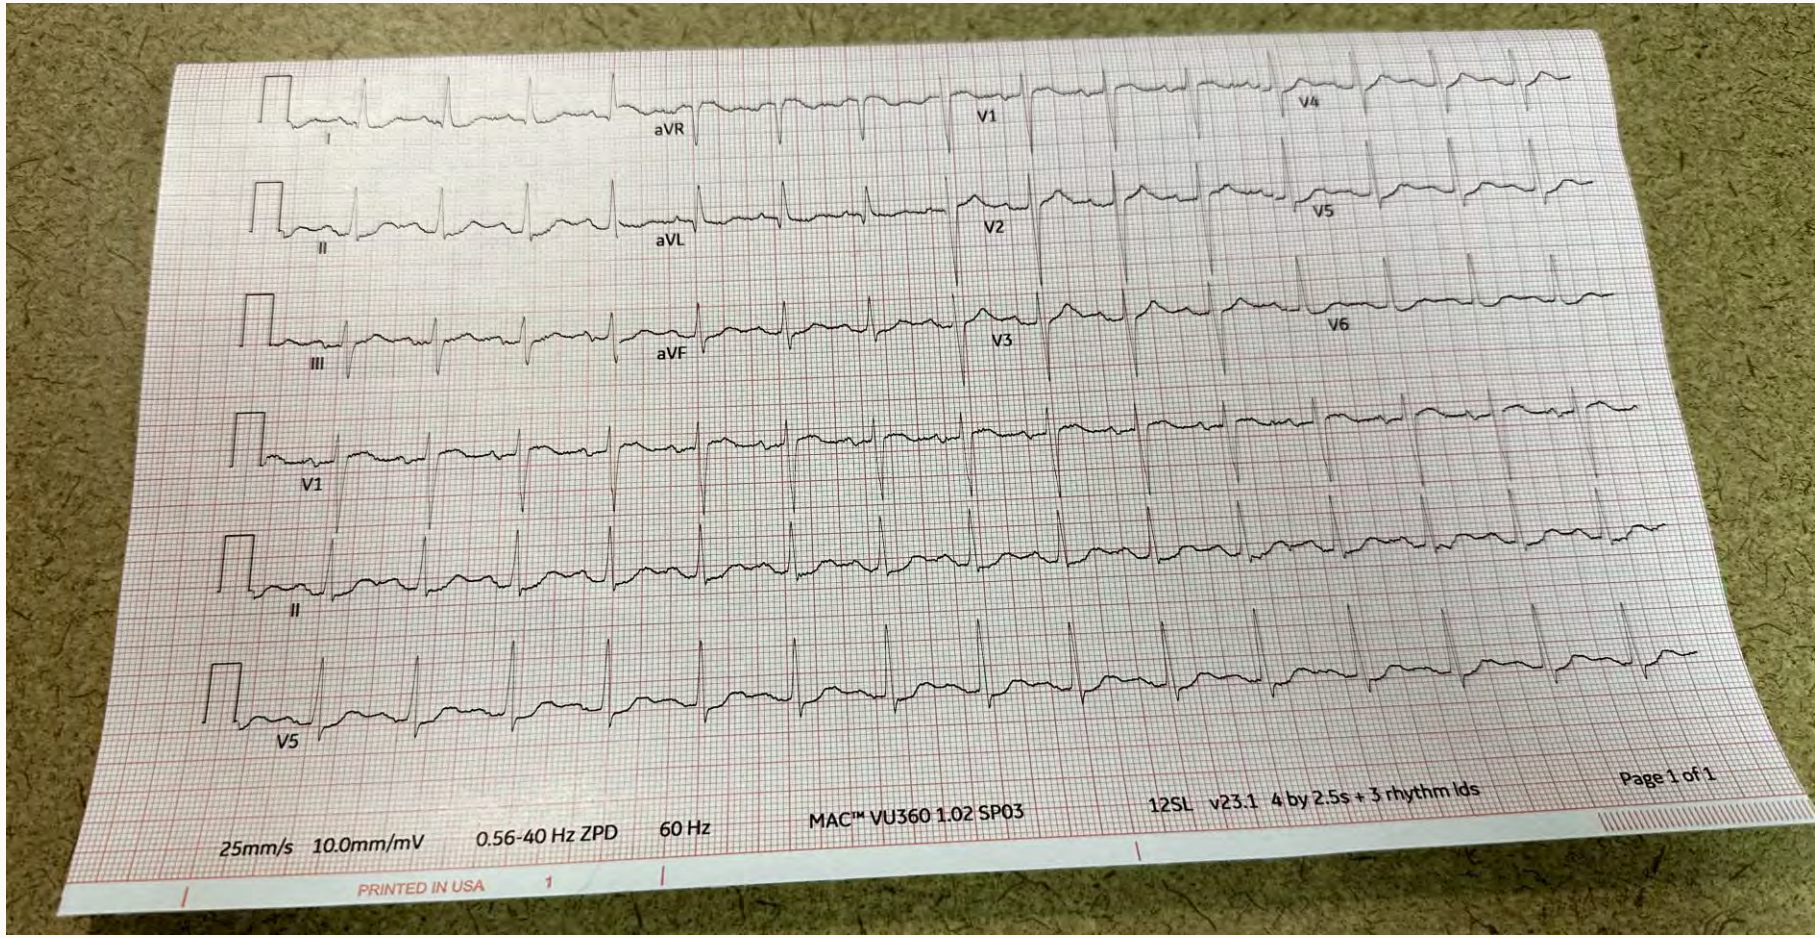

Real World:

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

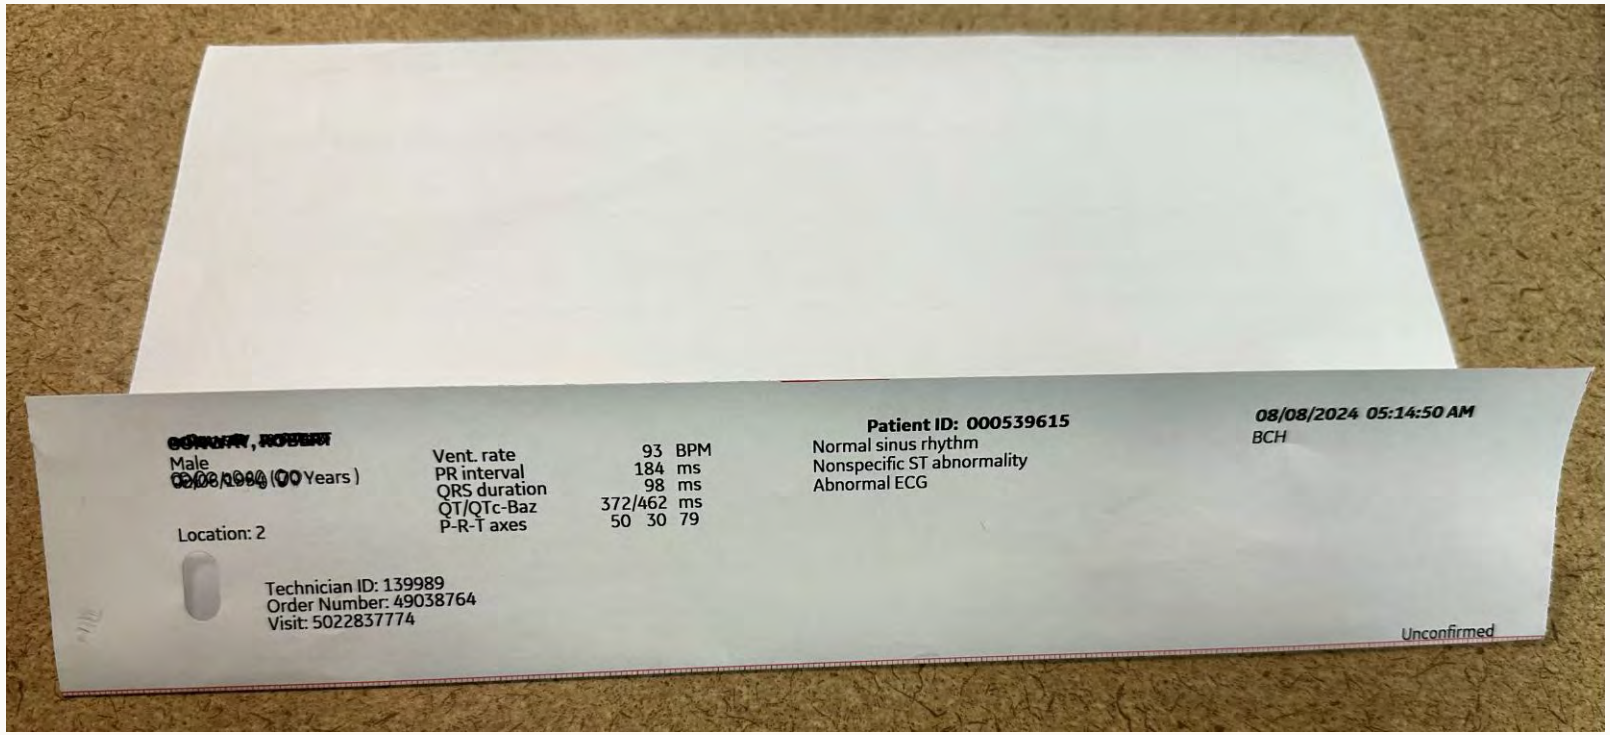

How to detect atrial fibrillation if you are an EKG machine, (since the 1980-90s):

1. Measure fiducial points
2. Compute the average variance in the R-R interval. Does it pass a certain threshold? If yes, proceed.
3. Detect atrial activity.
4. Are p-waves present? If no, proceed.
5. Are high-frequency fibrillatory waves present? If yes, result “atrial fibrillation.”

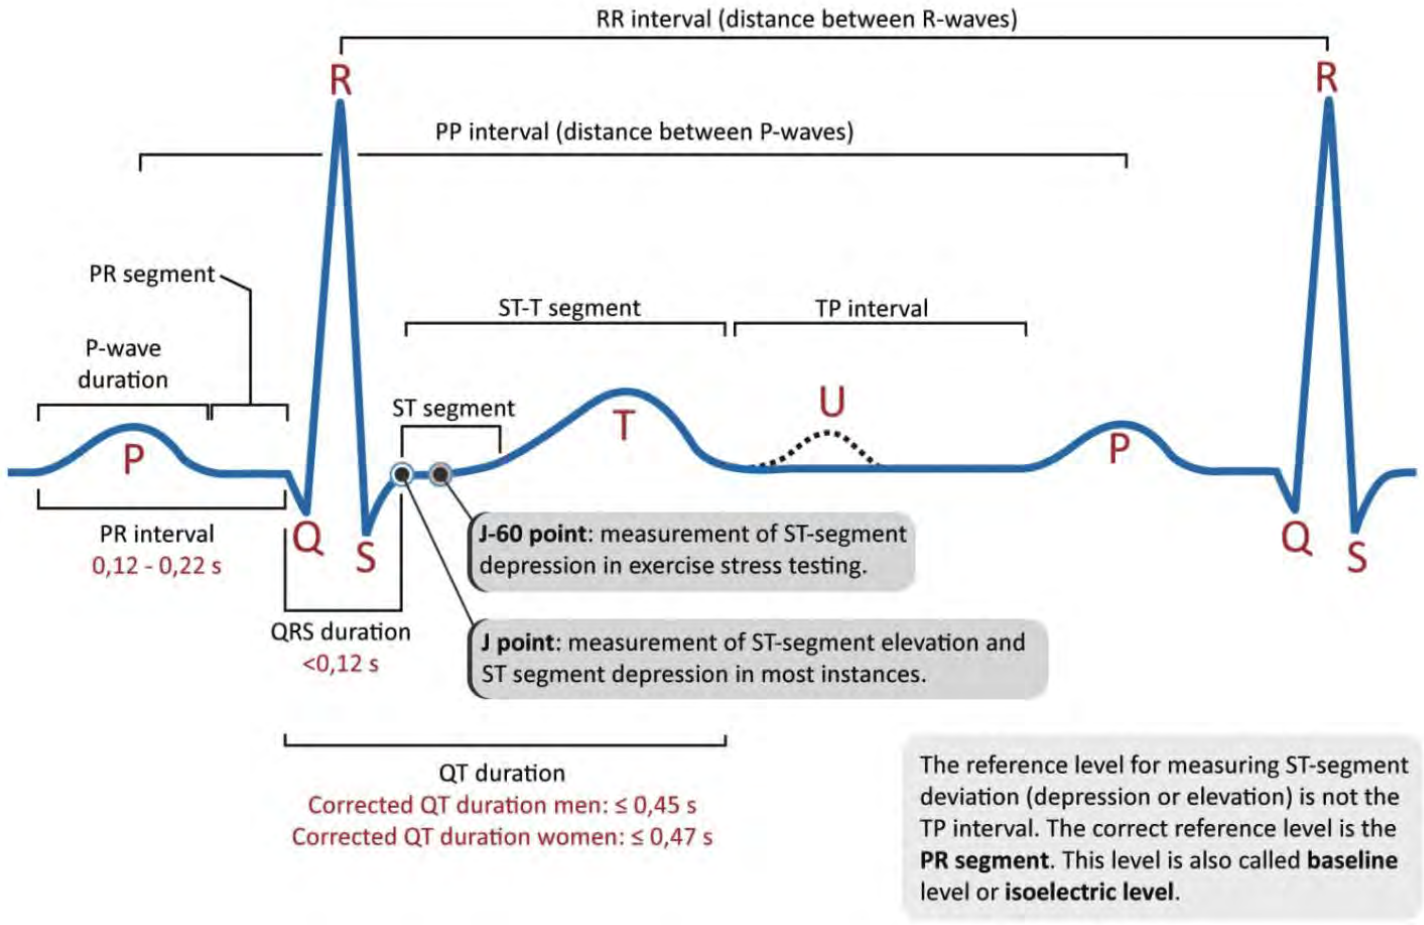

FDA has approved over 1000 AI algorithms

- Growing at about 125 per year
- 758 in Radiology
- 161 in Cardiology
- 35 in Neurology
- Others include GI, Anes, Heme, and others

As of Early 2025 per Copilot

AI for Dx Assistance

scientific reports

nature > scientific reports > articles > article

Article | Open access | Published: 13 October 2023

Differential diagnosis of neurodegenerative dementias with the explainable MRI based machine learning algorithm MUQUBIA

Silvia De Francesco, Claudio Crema, Damiano Archetti, Cristina Muscio, Robert I. Reid, Anna Nigri, Maria Grazia Bruzzone, Fabrizio Tagliavini, Raffaele Lodi, Egidio D'Angelo, Brad Boeve, Kejal Kantarci, Michael Firbank, John-Paul Taylor, Pietro Tiraboschi & Alberto Redolfi for the ADNI, Frontotemporal Lobar Degeneration Neuroimaging; NIA Alzheimer's Disease Centers; and the RIN – Neuroimaging Network

Scientific Reports **13**, Article number: 17355 (2023) | Cite this article

-A machine learning algorithm (MUQUBIA) accurately distinguished *Alzheimer's disease, frontotemporal dementia, dementia with Lewy bodies, and cognitively normal controls* using routine MRI scans and basic clinical information.

-Overall accuracy was 87.5%. Features like white matter integrity, cortical volumes and thickness were most informative.

***This system detects subtle brain patterns beyond human visual assessment.**

copic

AI Potentially Enhancing Procedures

EClinicalMedicine
Published by THE LANCET

Artificial intelligence for colorectal neoplasia detection during colonoscopy: a systematic review and meta-analysis of randomized clinical trials

—33 randomized trials, 27,404 patients

—>50% reduction of missed polyps and adenomas

—Inspection time increase only 20 seconds

Lou S, 29 November 2023

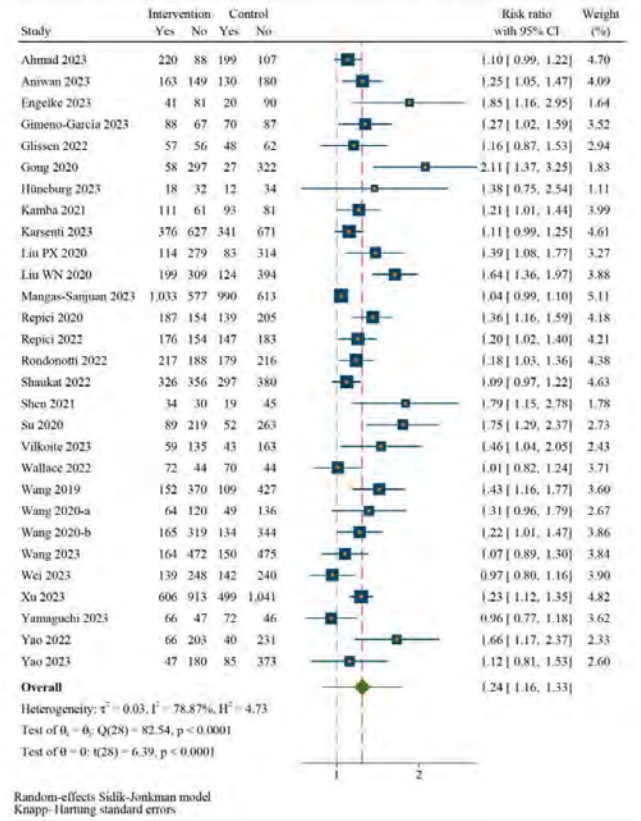

AI has a Salience Problem (something is more noticeable, but does it “matter” to the issue)

Which is more likely to be malignant?

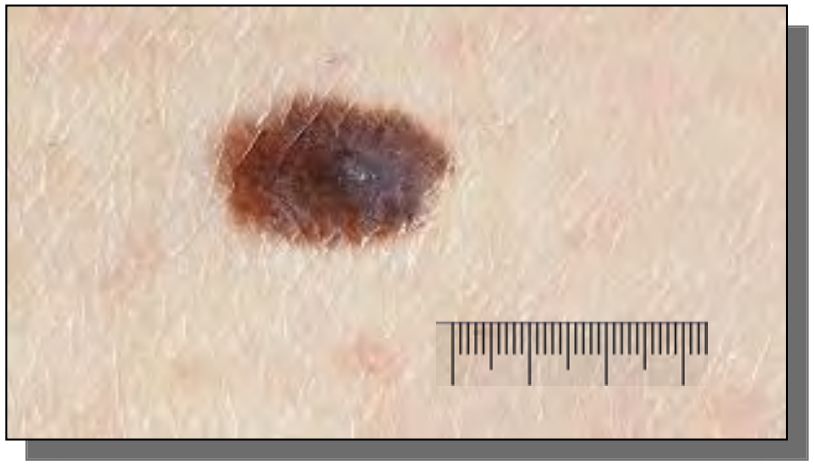

AI has Bias from Salience

- Age
- Race
- Gender
- Culture
- Sample
- Disability
- Historical
- Language
- Economic
- Geographic
- Confirmation
- Interpretability

In the algorithm—or in the world?

Algorithmic Bias

These images were all generated by AI. Notice anything wrong?

Image credits: New York Times (https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html

copic

Algorithmic Bias:

These images were generated using a model that was tweaked to avoid visual glitches.

A.I. images generated by Sina Alemohammad and others.

Image credits: New York Times (https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html

copic

Algorithmic Bias:
Second generation - AI “trained” on first set of faces, then outputs this:

A.I. images generated by Sina Alemohammad and others.

Image credits: New York Times (https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html

copic

Algorithmic Bias:
Third generation:

A.I. images generated by Sina Alemohammad and others.

copic

Algorithmic Bias
Fourth generation:

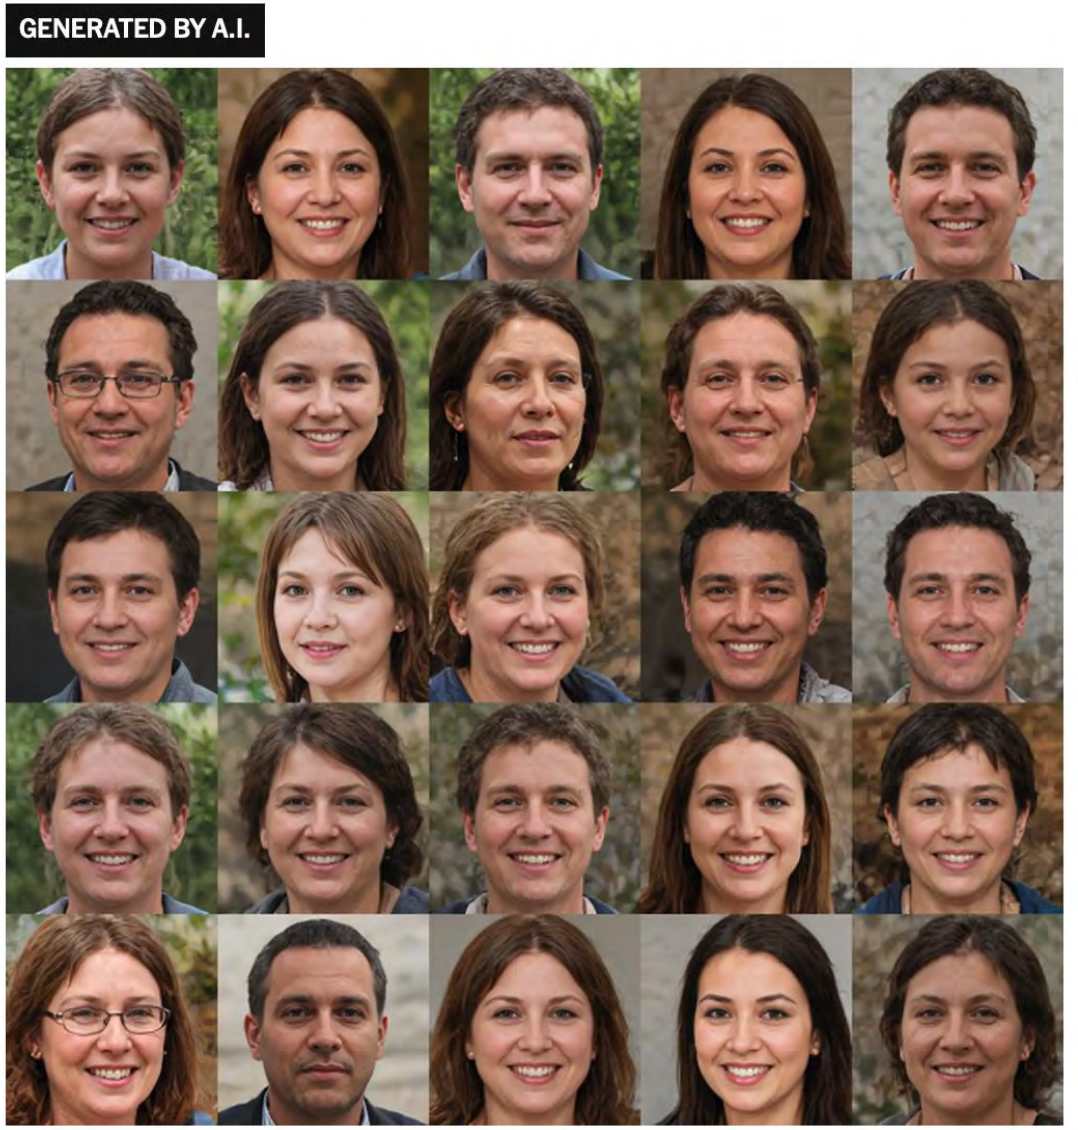

A.I. images generated by Sina Alemohammad and others.

Image credits: New York Times (https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html

copic

Algorithmic Bias:

Fifth generation. Uh oh.

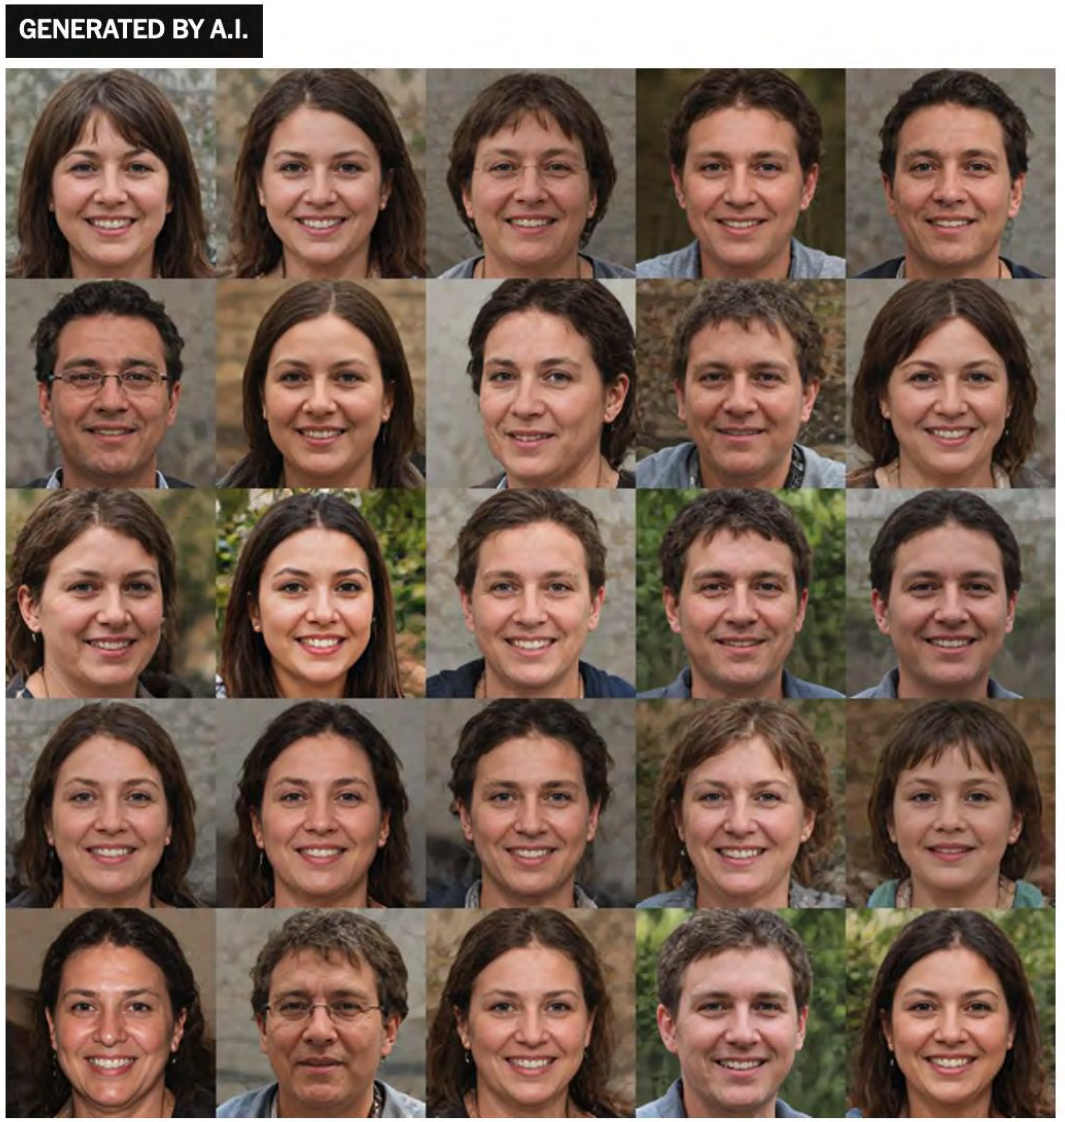

A.I. images generated by Sina Alemohammad and others.

Image credits: New York Times (https://www.nytimes.com/interactive/2024/08/26/upshot/ai-synthetic-data.html

Generative AI

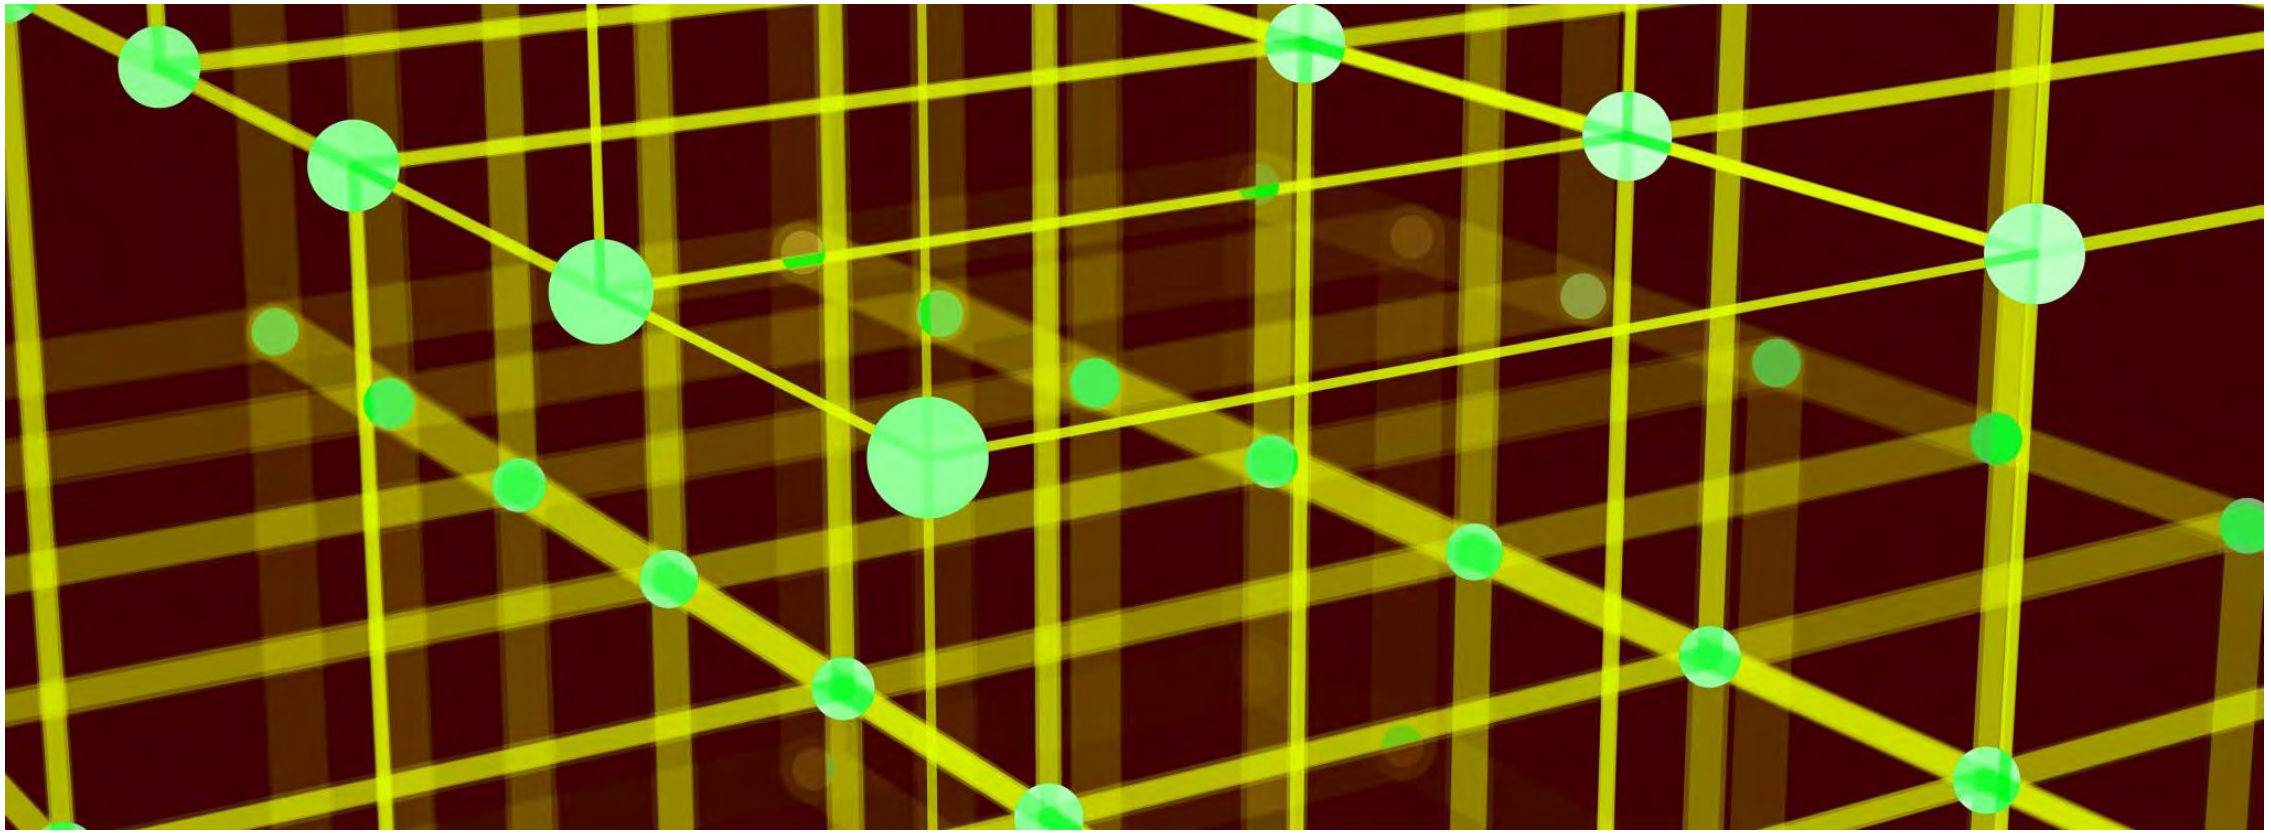

AI as Virtual Physician Scribe

- Current focus of much of the attention in healthcare.
- Uses “ambient listening” technology similar to household Alexa/Google products
- AI engine then produces progress or visit note
- Benefits:
  - Huge potential time savings for clinicians.
  - Allows physician to focus on patient, eye-to-eye contact

2-7

Minutes less per visit

50%

Less time spent on documentation

13

hours per week saved

https://www.aamc.org/news/can-artificial-intelligence-improve-doctor-patient-visits-and-reduce-burnout

AI Scribes: The Risks

- Consent:
  - State regulations vary. One Party vs Two Party State?
  - Best Practice: Always obtain consent from patients

AI Scribes: The Risks

- Consent:
  - State regulations vary. One Party vs Two Party State?
  - Best Practice: Always obtain consent from patients
- Privacy:
  - Vendors vary
  - They own the audio recording and the data
  - How long do they store the information? Is it discoverable in the case of a lawsuit?
  - Some vendors use information to help train the AI engine

AI Scribes: The Risks (continued)

Hallucinations!

- Anecdotal example: AI scribe added information in progress note that patient was scheduled for open heart surgery, which was totally false.
- Looks pretty bad if patient reviews the note, potentially very dangerous if error propagated by other readers of note.

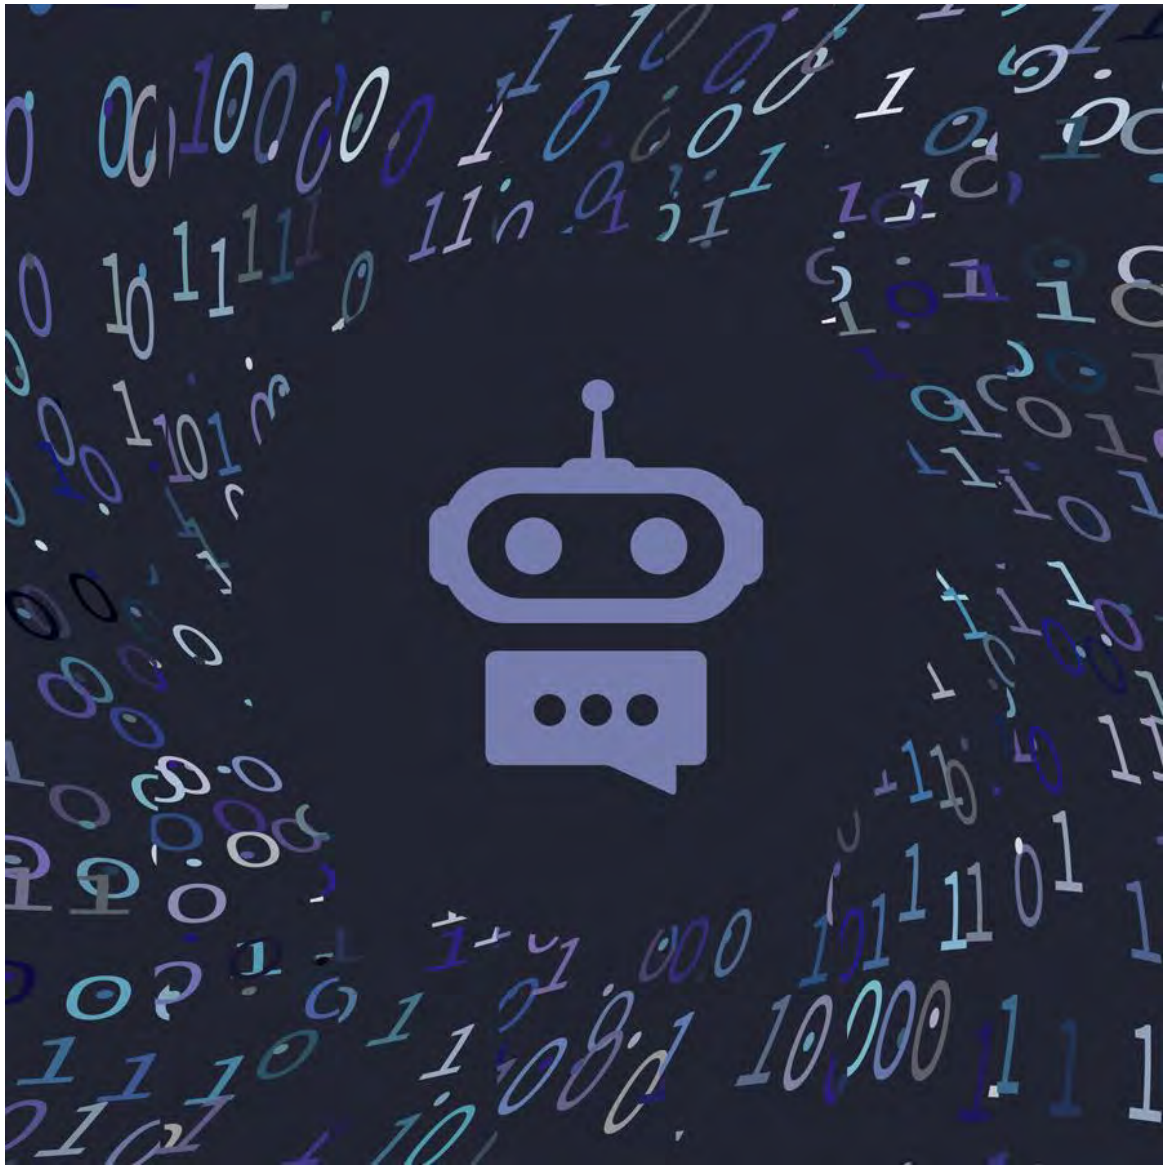

AI Scribes: The Risks (continued)

Hallucinations!

- Anecdotal example: AI scribe added information in progress note that patient was scheduled for open heart surgery, which was totally false.
- Looks pretty bad if patient reviews the note, potentially very dangerous if error propagated by other readers of note.

- The biggest risk that AI Scribe software currently poses comes from physicians not reading and correcting their note!

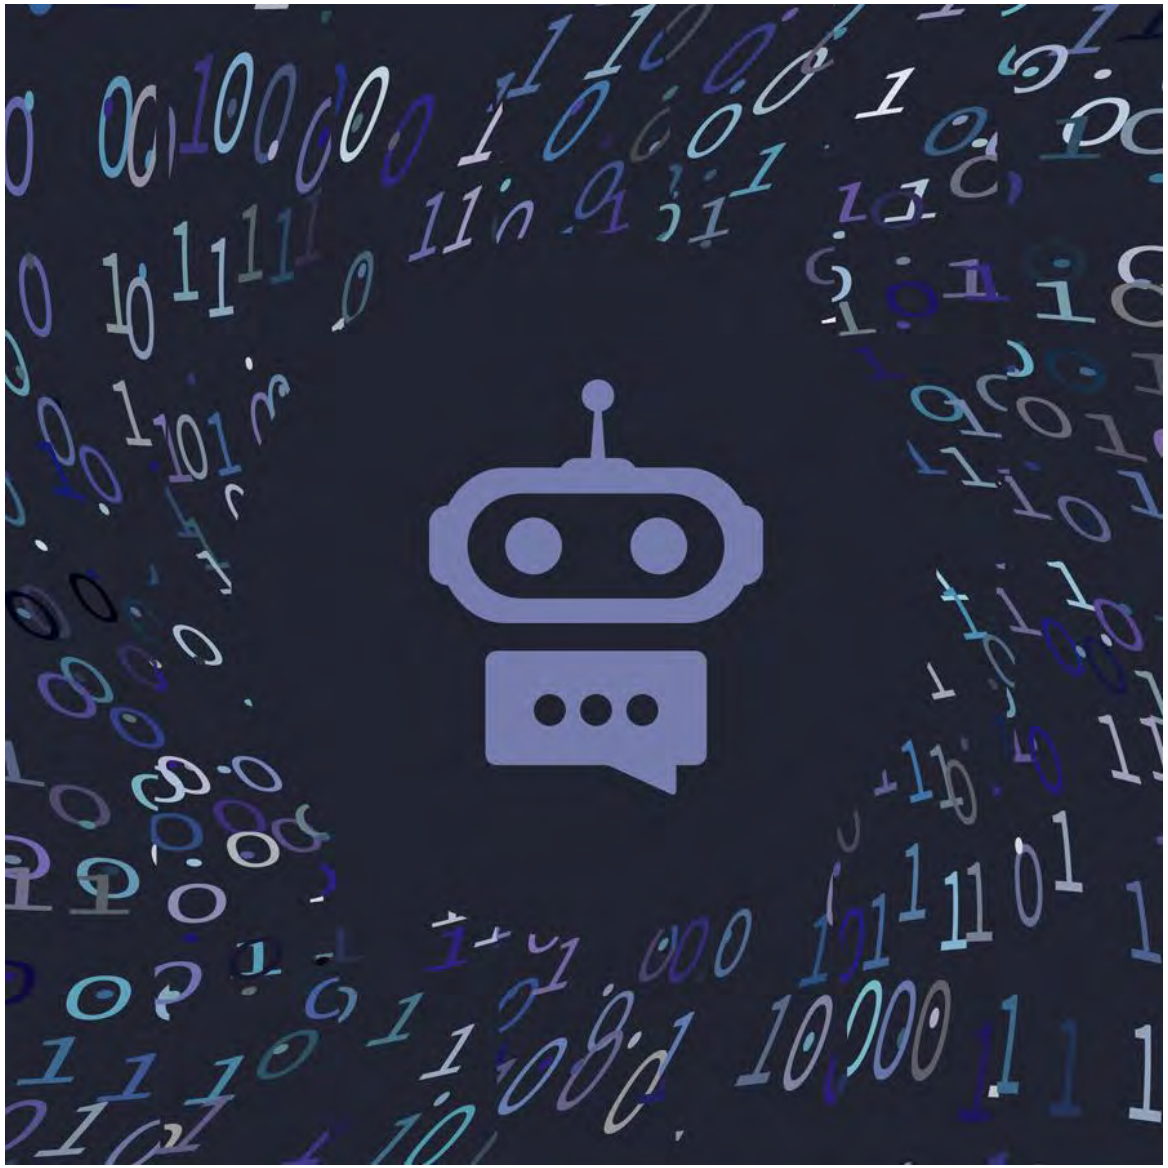

Data Drift vs. Hallucination?

-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.

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.”
  - “this is a white lady sedating medicine” (Instead of “lightly sedating medication”)
  - "Nipples equal round reactive”
- Easy to tell that something has gone wrong here.
- But what about when AI says that a patient has a history of CABG?

AI Scribes: What about the InBasket?

- Patient messages are an important source of burnout for physicians
- What if we let AI handle most of these?

That Message From Your Doctor? It May Have Been Drafted by A.I.

Overwhelmed by queries, physicians are turning to artificial intelligence to correspond with patients. Many have no clue that the replies are software-generated.

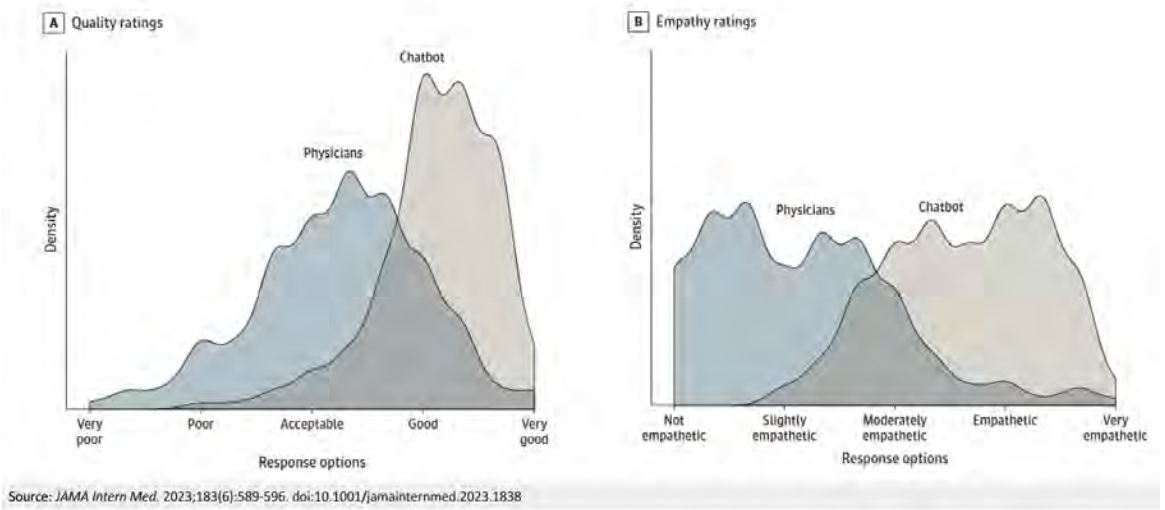

Source: *JAMA Intern Med.* 2023;183(6):589-596. doi:10.1001/jamainternmed.2023.1838

AI Scribes: What about the InBasket?

- Patient messages are an important source of burnout for physicians
- What if we let AI handle most of these?
- AI has been shown by multiple studies to produce longer messages that patients feel express more empathy than those written by actual physicians.

That Message From Your Doctor? It May Have Been Drafted by A.I.

Overwhelmed by queries, physicians are turning to artificial intelligence to correspond with patients. Many have no clue that the replies are software-generated.

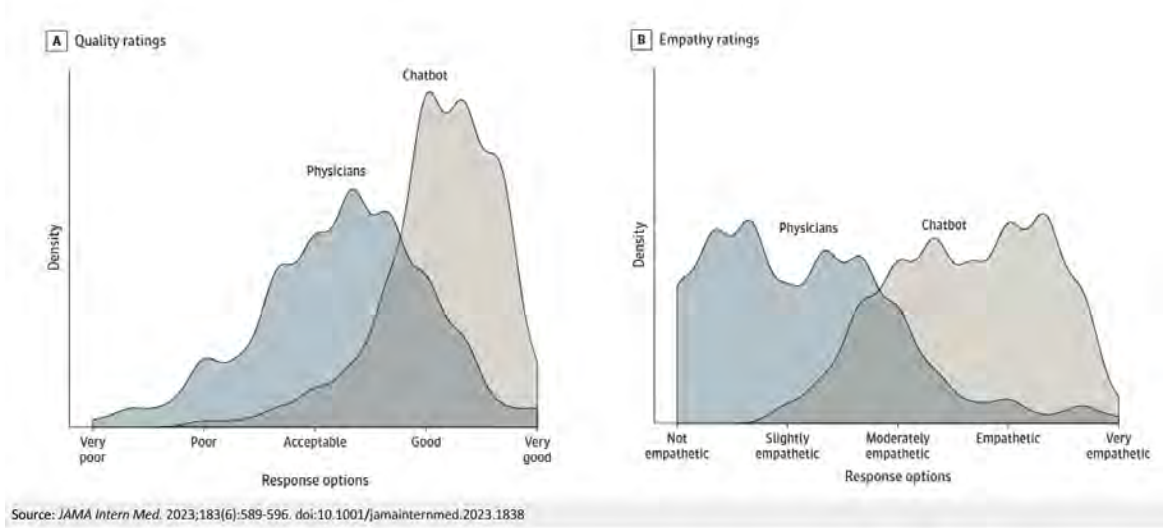

Source: *JAMA Intern Med.* 2023;183(6):589-596. doi:10.1001/jamainternmed.2023.1838

AI Scribes: What about the InBasket?

- But…you'd better review all responses before sending!

It was felt by the assessing physicians that the LLM drafts posed a risk of severe harm in 11 (7·1%) of 156 survey responses, and death in one (0·6%) survey response. The majority of harmful responses were due to incorrectly determining or conveying the acuity of the scenario and recommended action.

The effect of using a large language model to respond to patient messages

Lancet Digit Health 2024

What About Prior Authorizations?

In Constant Battle With Insurers, Doctors Reach for a Cudgel: A.I.

As health plans increasingly rely on technology to deny treatment, physicians are fighting back with chatbots that synthesize research and make the case.

beckershospitalreview.com

Why prior authorization could become an AI 'arms race'

copic

Legal Considerations

AI Virtual Scribes: Do we need to obtain consent? If so, what kind?

- There is currently no Federal regulation that defines this.
- AMA: “Principles for Augmented Intelligence Development, Deployment, and Use”. (Nov 2023.)
  - “The use of AI-enabled technologies by hospitals, health systems, physician practices, or other entities, where patients engage directly with AI should be clearly disclosed to patients at the beginning of the encounter or interaction with the AI-enabled technology.”
  - When AI is used in a manner which directly impacts patient care, access to care, or medical decision making, that use of AI should be disclosed and documented to both physicians and/or patients in a culturally and linguistically appropriate manner. The opportunity for a patient or their caregiver to request additional review from a licensed clinician should be made available upon request.
  - Recommendation: Obtain verbal consent, at minimum, prior to each encounter where AI is used.

AI Virtual Scribes: What Should Patient Consent Include?

- There is currently no regulation that defines this.
- Key Elements:
  1. Your encounter with the provider will be recorded electronically.
  2. This information will be processed by a system that uses AI to help draft a note.
  3. I will review the note.
  4. All of your information will be kept secure.

copic

AI Virtual Scribes: What Should Patient Consent Include?

- There is currently no regulation that defines this.
- Key Elements:
  1. Your encounter with the provider will be recorded electronically.
  2. This information will be processed by a system that uses AI to help draft a note.
  3. I will review the note.
  4. All of your information will be kept secure.

- If patients have additional questions, probably not time-effective to use the tool!

copic

AI Virtual Scribes: Does consent need to be documented? If so, how?

- There is currently no Federal regulation that defines this.
- AMA: “Principles for Augmented Intelligence Development, Deployment, and Use”. (Nov 2023.)
  - “When AI is used in a manner which directly impacts patient care, access to care, medical decision making, or the medical record, that use of AI should be documented in the medical record.”
- What risk are we trying to mitigate in the case of virtual scribes?
- Some vendors/EMRs offer automatic statements in chart that state verbal consent was obtained.
  - Who does that help?
  - Statement should NOT be used similar to voice recognition caveat statements to try to deflect responsibility for inaccuracies in note.
- Recommendation: Talk to your attorney/risk manager.

copic

AI Virtual Scribes: What should provider attestation consist of?

I am responsible for, and will review, all note content before signing.(!!!!)

I agree to get a verbal consent from every patient prior to using AI-assisted clinical documentation.

Might want to include:

I will watch training video etc.

I am aware that technology will only work on certain phones (iphone v. android, not older)

copic

Regulatory Environement - 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.”

Regulatory Environment - Colorado

- Gov. Polis signed SB 24-205 “Concerning Consumer Protections in Interactions with AI Systems” May 2024.
- First state in US to enact broad restrictions on AI
- Rules take effect Feb 1, 2026
- Establishes rules for both developers and deployers of “high risk” AI systems.
- Must provide consumers with:
  - Information regarding opt-out
  - A lot of focus on algorithmic discrimination
  - Publicly posted description on website of all AI systems they deploy, how risks are managed, as well as how info is collected and used
- Impact: Unclear.

Regulatory Environment - California

- Three Bills –Signed by Gov. Newsome Sept 2024
- AB-3030: requires healthcare providers to disclose use of generative AI for patient communications, display disclosure prominently throughout telehealth visit
- SB-942: requires entities with >1mil website visitors to disclose what content was generated by AI and provide free tool for detection.
- SB-1047: requires whistleblower protection, allows legal action against AI model developers, and requires developers to retain third-party auditors to perform compliance audit of model

The practitioner is still responsible

- Human accountability is not going away
- Assistive technology needs to be transparent, verifiable and reproducible
- The record must be credible
- Safety measures must be appropriate and auditable
- Privacy is complex and challenging

More Information
Search AI in Medicine

Chat GPT

Google Search (w/ AI)

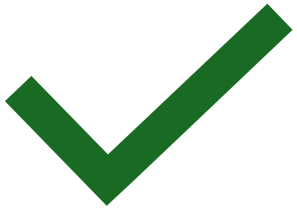

Microsoft Copilot

Mark Cuban: Don’t rely on ChatGPT to do all your work for you—‘AI is never the answer; AI is the tool’

Published Fri, Mar 28 20252:27 PM EDT

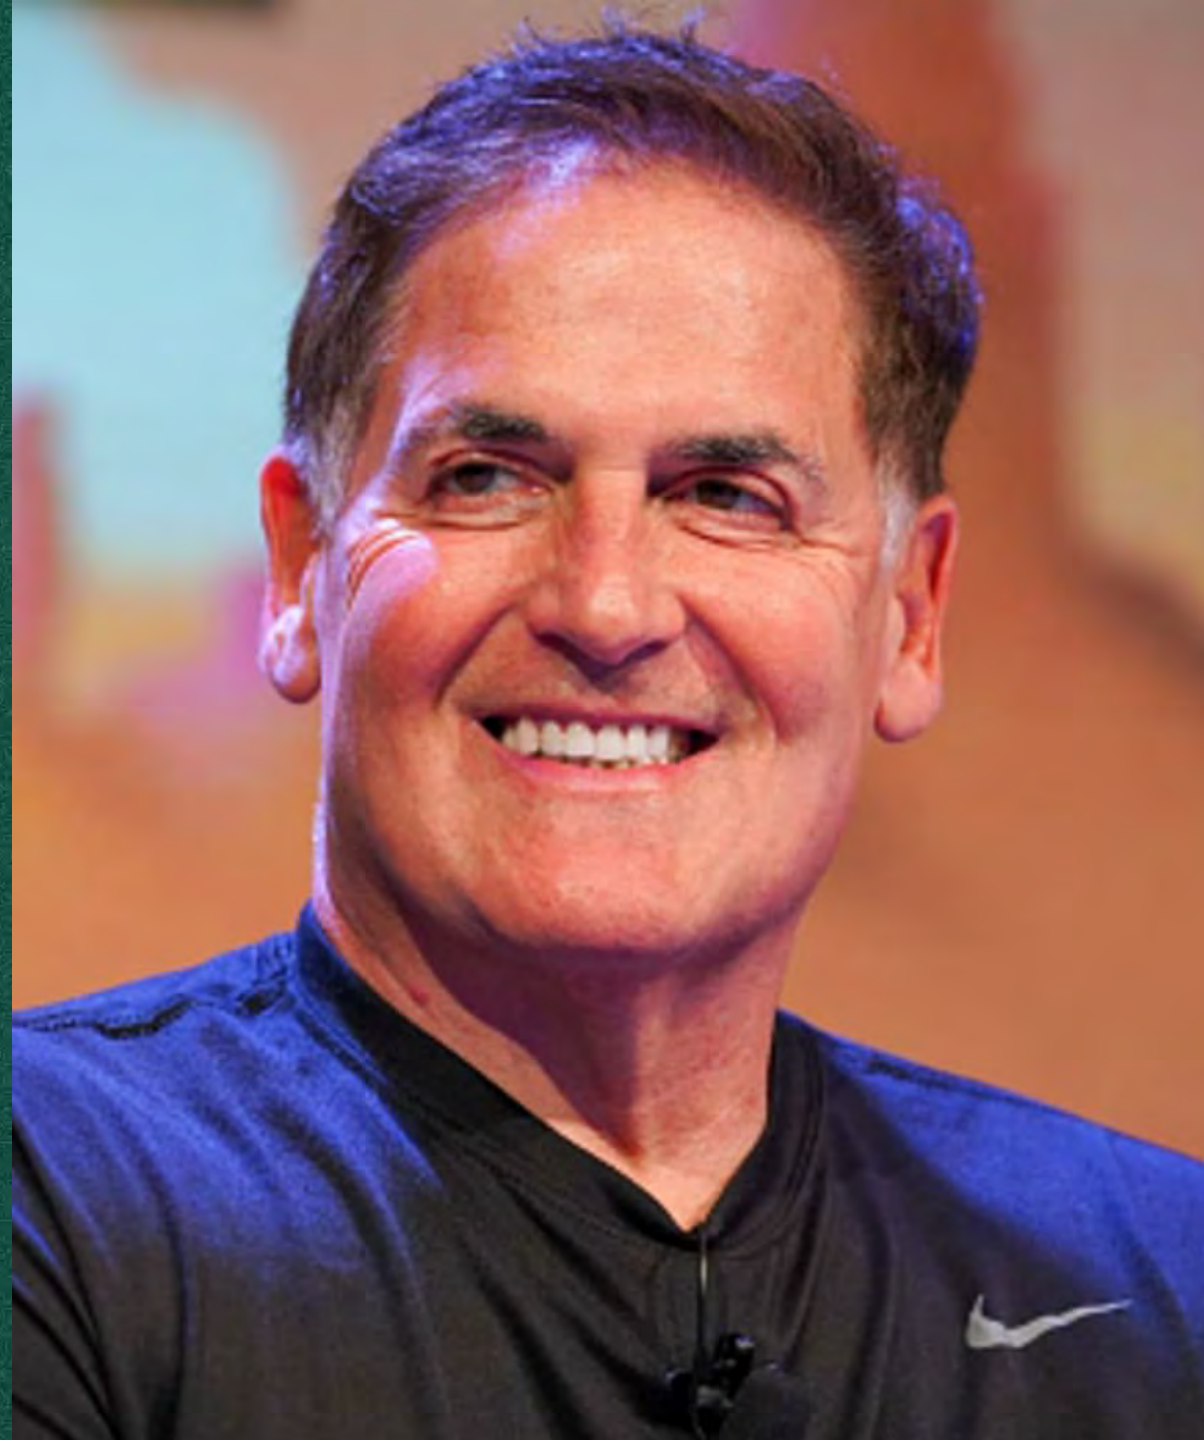

What Does the Future Hold???

- Daniel Rosenquist MD FAAFP
- drosenquist@copic.com