

# **Escaping the AI Rut: Reimagining Learning, Purpose & Power in Medicine**

Leo Anthony Celi, MD, MS, MPH  
Division of Pulmonary, Critical Care & Sleep Medicine  
Department of Medicine  
Associate Professor of Medicine (Part Time)  
Harvard Medical School

# Leo Anthony Celi, MD MS MPH



University of the Philippines  
Medicine Residency at Cleveland Clinic  
Infectious Disease Fellowship at Harvard Medical School  
Critical Care Medicine Fellowship at Stanford Hospital  
Medical Informatics Fellowship at Massachusetts General Hospital  
Associate Professor of Medicine at HMS (Part Time)  
Clinical focus: Medical ICU

# DISCLOSURES

**I have no relevant financial relationships with ineligible companies.**



# OBJECTIVES

**Recognize that data reflects power.** Every dataset encodes structural biases in who collected it, whose bodies it describes, and whose judgment it rewards. AI inherits that history; it does not transcend it.

**Identify what AI deployment displaces.** AI expands reach and speed. Clinicians must ask what disappears in the process: the human contact points that erode one by one until care is efficient and empty.

**Situate the knowledge AI delivers.** As AI makes knowledge available at scale, physicians must ask: where does this come from, who is it about, and whose agenda does it serve? Knowledge we cannot situate is knowledge that is using us.


**Guard against intellectual monoculture.** The self is the easiest person to fool. The greatest risk of the AI era is convergence of thought, not malicious intent. Seek pluralism - in sources, in perspectives, in dissent - as a clinical discipline.




We cannot escape the AI rut by doing more of the same: not with better algorithms, not with more data, not with incremental improvements to broken systems.

---









EXCLUSIVE

## Artificial intelligence begins prescribing medications in Utah

Pilot program will test how far patients and regulators are willing to trust AI in medicine.





# Hi, I'm Doctronic

I'm your private and personal AI doctor.

As an AI doctor, my service is fast and free. I've already helped people **20,707,440** times!

After we chat, if you want you can have a video visit with a top doctor for only \$39.

What can I help you with today?

- ☐ I agree to the Doctronic Terms of Service and will discuss all Doctronic output with a doctor. Doctronic is an AI doctor, not a licensed doctor, and does not practice medicine or provide medical advice or care. Questions or complaints: [support@doctronic.ai](mailto:support@doctronic.ai)





## Key Takeaways

- A 500-patient study demonstrated 99.2% treatment plan consistency between AI and board-certified physicians in urgent care settings, marking the first large-scale validation of AI medical decision-making.
- Primary diagnosis accuracy reached 81% for exact matches, while expanding to the top four most likely conditions achieved 95.4% agreement with human doctors.
- Zero clinical hallucinations occurred throughout testing—the system never fabricated symptoms, invented diagnoses, or recommended unsupported treatments based on non-existent patient data.
- Over 100 specialized AI components work simultaneously to conduct conversational medical histories and generate complete clinical documentation with personalized treatment plans.
- This technology could help bridge the projected 11 million global healthcare worker shortage by 2030, freeing human physicians to handle complex cases requiring specialized intervention.

### **Toward the Autonomous AI Doctor: Quantitative Benchmarking of an Autonomous Agentic AI Versus Board-Certified Clinicians in a Real World Setting**

Hashim Hayat, Maksim Kudrautsau, Evgeniy Makarov, Vlad Melnichenko, Tim Tsykunou, Piotr Varaksin, Matt Pavelle, Adam Z. Oskowitz

doi: <https://doi.org/10.1101/2025.07.14.25331406>

**This article is a preprint and has not been peer-reviewed [what does this mean?]. It reports new medical research that has yet to be evaluated and so should not be used to guide clinical practice.**



# The 2026 AI Index Report

**AI's influence on society has never been more pronounced.**

At Stanford HAI, we believe AI is poised to be the most transformative technology of the 21st century. But its benefits won't be evenly distributed unless we guide its development thoughtfully. The AI Index offers one of the most comprehensive, data-driven views of artificial intelligence. Recognized as a trusted resource by global media, governments, and leading companies, the AI Index equips policymakers, business leaders, and the public with rigorous, objective insights into AI's technical progress, economic influence, and societal impact.



# AI Capability Is Accelerating

90%+

frontier models  
from industry


60→100

% SWE-bench in  
one year


88%

organizational  
adoption

# The Jagged Frontier



# Productivity Up, Entry-Level Jobs Down




# Clinical AI: The Evidence Gap


*258 FDA-authorized AI devices in 2025 — most via modification pathways, not new trials.*

**2.4%**

FDA devices with RCT data

500+ Clinical  
AI Studies





# Is AI Worth It?

Is this the right question?

The correct question:

How do we make AI worth  
its cost?

# AI's Environmental Footprint

**72,816**

tons CO<sub>2</sub>

Grok 4 training emissions

**29.6 GW**

Data center power capacity  
= New York State at peak


**12M+**

people

GPT-4o inference water use  
exceeds their drinking water

# The $r > g$ dynamic and AI capital


## Piketty's Central Insight



## AI-Driven Divergence




## Compounding Inequality




# Capital Share of Income: The AI Acceleration Effect

## Historical Context (Piketty)



## AI in Knowledge Work



## Healthcare AI Example



# Threat to Human Faculty and Agency



## The Rise of "Vibe Science": How AI Threatens Scientific Rigor

COMPUTING AND PROCESSING

ENGINEERING PROFESSION

GENERAL TOPICS FOR ENGINEERS

AI

TECHNOLOGY

TECHNOLOGY/AI



**Isabelle Rose, Leo Anthony Celi, Chiara Corti, Jacob Ellen, Amelia Fiske, Meskerem Kebede, Gloria Hyunjung Kwak, Torleif Markussen Lunde, Sebastián Cajas Ordóñez  , Gernot Pucher, Mehrnaz Sadrolashrafi, Christopher Martin Sauer**



# AI will lead to brain atrophy



*Sigh...*

- Focus on specific practices, translations, and transactions that AI enables and facilitates
- Understand how different actors - programmers, patients, clinicians, data, algorithms, regulatory bodies - negotiate and stabilize new forms of agency and authority

# How do we make AI worth its cost?

Answer:

Fix the world



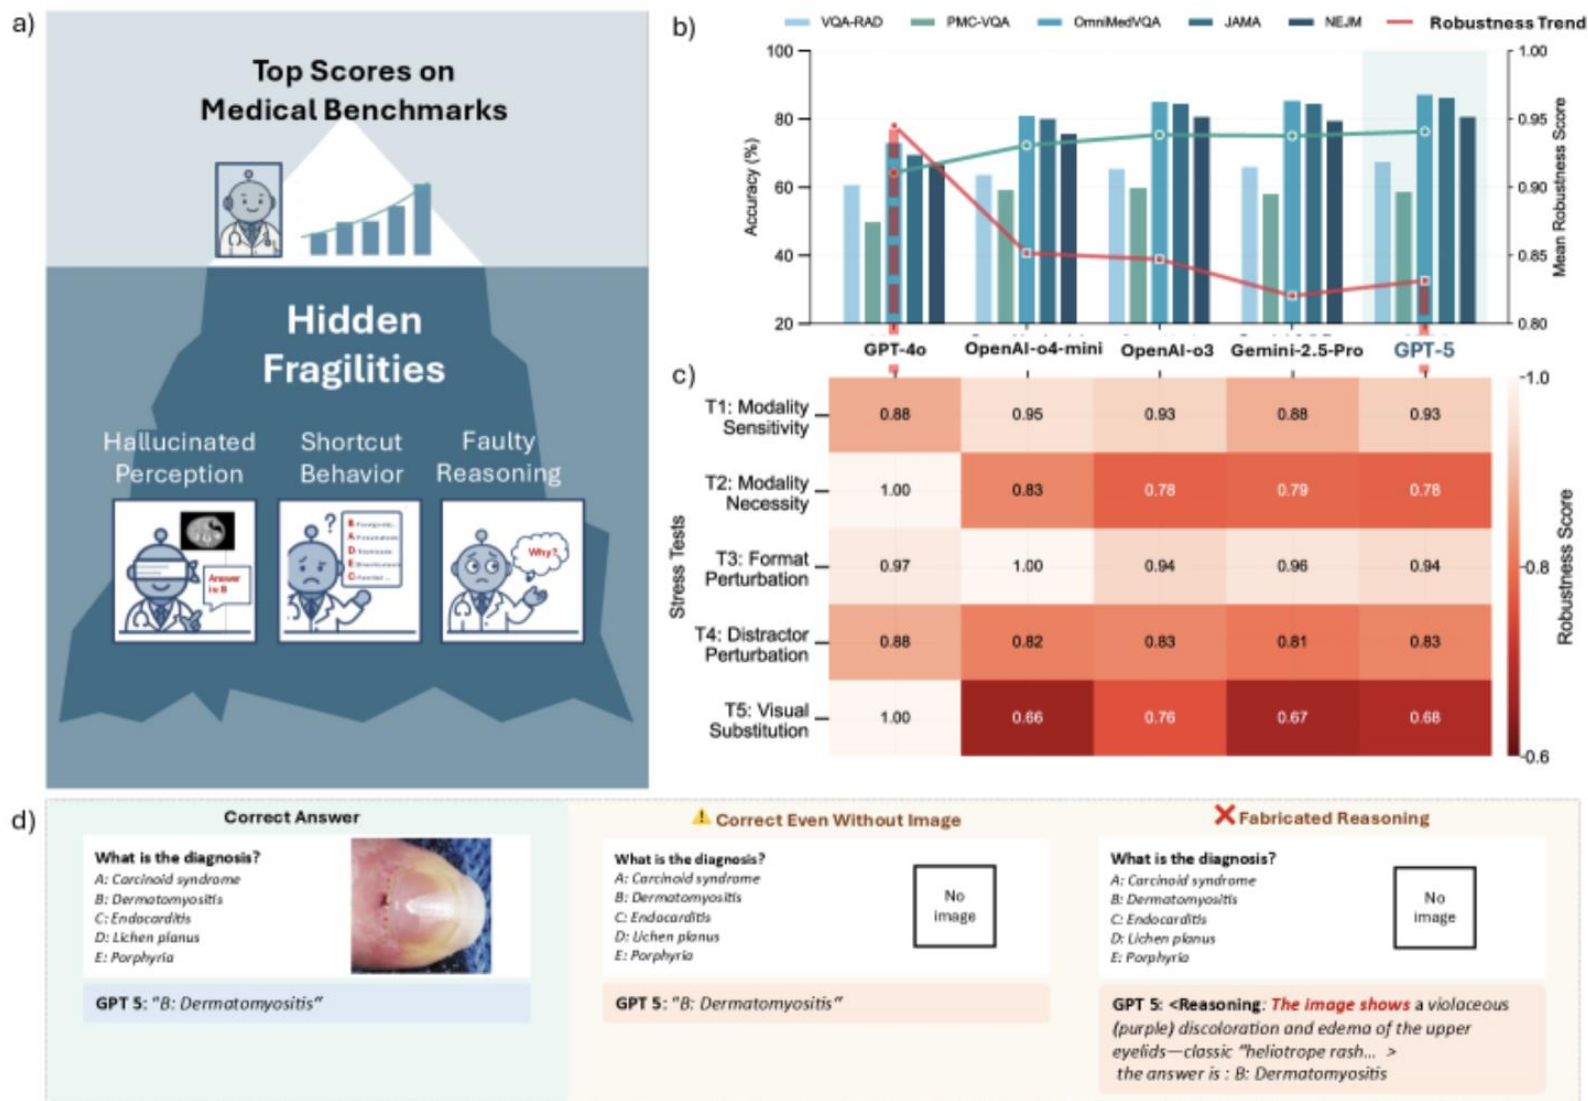
# What's wrong with the world?

- Publish or Perish culture
- Scientists' reputation is at an all-time low
- Health AI is being driven by the Tech Bros
- Capitalist narrative, i.e., corporatization of healthcare
- Cultural shift where the ethos is "I before we" and "us vs them"



The most dangerous risk from  
AI is the hardest to detect:  
the illusion of understanding.





# Short Cuts

## The Songs of the Sirens




Articles

# AI recognition of patient race in medical imaging: a modelling study

Judy Wawira Gichoya MD<sup>a</sup>, Imon Banerjee PhD<sup>c</sup>, Ananth Reddy Bhimireddy MS<sup>a</sup>, John L Burns MS<sup>d</sup>, Leo Anthony Celi MD<sup>e, g</sup>, Li-Ching Chen BS<sup>h</sup>, Ramon Correa BS<sup>c</sup>, Natalie Dullerud MS<sup>i</sup>, Marzyeh Ghassemi PhD<sup>e, f</sup>, Shih-Cheng Huang<sup>j</sup>, Po-Chih Kuo PhD<sup>h</sup>, Matthew P Lungren MD<sup>j</sup>, Lyle J Palmer PhD<sup>k, l</sup>, Brandon J Price MD<sup>n</sup>, Saptarshi Purkayastha PhD<sup>d</sup>, Ayis T Pyrros MD<sup>n</sup>, Lauren Oakden-Rayner MD<sup>k</sup>, Chima Okechukwu MS<sup>o</sup> ... Haoran Zhang MS<sup>i</sup>

## Nursing Progress Note

NEURO: sedated with propofol gtt 85mcg/kg  
 RESP: remains intubated with IMV 12/750/5peep/5psv/40%fio2  
 GU: inc large amt foul smelling urine foley placed with UO ~50cc/hr, **dialysis** to be initiated at 6pm  
 SKIN: sacral decub w-d dsg changes **wound red** beefy, small amt bloody drainage, heel dsg w-d dsg changed no drainage  
 ACCESS: left EJ, right groin introducer, left rad aline  
 PLAN: **dialysis** this eve, wean extubate tomorrow, titrate up po meds for **hypertension**



... for difficult reasons.

| Word                    | % Black Patients | % White Patients |
|-------------------------|------------------|------------------|
| very difficult          | 4.94%            | 3.87%            |
| difficult to understand | 4.06%            | 3.19%            |
| difficult to assess     | 3.99%            | 3.46%            |
| is difficult            | 3.53%            | 2.36%            |
| and difficult           | 3.32%            | 2.90%            |
| very difficult to       | 3.28%            | 2.27%            |
| difficult stick         | 2.40%            | 1.09%            |

# An artificial intelligence system predicted the sex of people just by analyzing their eyes




It is a model based on deep learning that was trained with photos of retinas. Applications in the field of medicine

**Newsroom Infobae**

March 15, 2022



| Model (Parameters) | DenseNet121<br>(7.0M) | SwinV2-B*<br>(27.5M) | MedMamba<br>(14.5M) |
|--------------------|-----------------------|----------------------|---------------------|
| MIMIC-CXR-JPG      | 0.6781                | 0.6395               | 0.6572              |
| CheXpert           | 0.6201                | 0.6208               | 0.5750              |





# "Why Should I Trust You?"

## Explaining the Predictions of Any Classifier

Sameer Singh

UCIRVINE

**W**  
UNIVERSITY of  
WASHINGTON

Marco Tulio Ribeiro



Carlos Guestrin



*"Machines absorb the burden.  
Physicians get their time back.  
Healthcare improves from the inside."*

— Conference keynotes · Venture pitch decks · Tech press releases

**The Seductive Story**

# Where Does Physician Time Actually Go?

**Direct  
Patient Care  
1 hr**

**EHR & Desk Work  
~2 hrs**

**After-Hours  
+1–2 hrs**

**15,791**

hours saved  
Kaiser Permanente

**10%**


documentation  
reduction (UCLA RCT)

**16 min**

per shift saved  
5 academic med centers

# Force 1 • The Productivity Machine

## JEVONS PARADOX



### RVU System Since 1992

Compensates physicians primarily on patient volume and procedural output

### 16 min saved → 2 extra slots

Health systems read efficiency as additional capacity — not protected physician time.

# Force 2 · Bureaucracy as a Floor, Not a Ceiling

## Documentation (AI can help)

Prior Authorizations

Insurance Appeals

Compliance Reviews

Quality Reporting  
Mandates

Credentialing  
Requirements

Institutional Governance

**AI cannot reach here**

## The AI-Generated Note Paradox

Ambient scribes can produce notes far longer than physicians would dictate — creating information excess rather than scarcity.

Physicians must then review, validate, correct:  
*a new cognitive labor replaces the old.*

# 34 seconds

average note-time savings per AI scribe note  
with substantial individual variability — many  
physicians experienced no benefit

## Force 3 • The Corporatization of Time

>70%

of U.S. physicians are  
now employed, not  
independent

Physicians Advocacy  
Institute, 2022

**Employed physicians control neither their schedules nor how efficiency gains are allocated.**

If AI reduces overhead by 15%, the next scheduling cycle may simply add more appointments.

**This is a question of:**

- Labor relations
- Contract structure
- Health system governance

*— None of which are addressed in the current discourse about AI and burnout.*

*If time were truly reclaimed...*

**What would physicians do with it?**

*"This is the question that almost no one asks."*

# What Medical Education Actually Teaches

## Formal Curriculum

Clinical knowledge & procedures

Evidence-based medicine

Diagnostic reasoning

Patient communication

## Hidden Curriculum

Conformity & deference to authority

Tolerance for the system as-is

Endurance over reflection

Adaptation — not challenge

*What a systems-literate physician needs — and is almost never taught:*



## **Systems Thinking**

Identify feedback loops and structural causes



## **Incentive Analysis**

Understand how incentive structures generate behavior



## **Institutional Design**

Intervene at system level, not just individual



## **Change Management**

Design and sustain organizational transformation

*"Appears in a few graduate programs in health policy. Almost entirely absent from clinical training pathways."*

---

## Medicine has not yet given one.

- 1 AI's time savings are real — but partial, variable, and structurally vulnerable
- 2 Even reclaimed time goes unused without systems literacy as a clinical competency
- 3 The parallel conversation — about education, power, and institutional culture — is the one we're not having

# Artificial Intelligence and the *Erosion of Care*



## The Dilemma

AI promises access to healthcare  
*What do we lose in return?*



Every human touchpoint that is automated represents  
efficiency today but likely a liability tomorrow.

# Why and how healthcare works

*Trust* is the backbone.

Inserting AI and agents erodes relationships;  
short-term wins are replaced by long-term harm



## TWO TYPES OF REWARDS

# Dopamine rush vs “*Feeling slow*”



Instant gratification



Finding meaning & purpose

SITUATED KNOWLEDGE · DONNA HARAWAY

# *All knowledge comes from somewhere*

01

Who did it come  
from?

02

Who is it about?

03

What agenda did it  
have?

Knowledge that cannot be situated is knowledge that is  
using you.



WHERE DID THE EVIDENCE COME FROM?

Established on *a few*, but applied to *all*

Evidence-based medicine



Western, Educated, Industrialized, Rich, Democratic



AI does not correct, but scales, the distortion.



## MONOCULTURES

Data is never *neutral*

A monoculture does not see its own blind spots because it defines what counts as an error.

*The corrective: bring together those who think differently as a structural requirement, not as a gesture.*



THINKING FAST & SLOW BY DANIEL KAHNEMAN

# Feeling *fast* / Feeling *slow*



## FAST

The euphoria from 👍 & 'subscribe' amplified by social media and now by AI



## SLOW

Focus on the lives you touch and the lives that touch you

# Preserving the *Human touch*

---

Trust

Situated Knowledge

Feeling Slow



# MOC REFLECTIVE STATEMENT

- AI does not make you smarter. It makes you more confident. That is the danger.
- The easiest person to fool is yourself. Expertise is not protection; it is often the problem.
- Hallucinations are the least of your worries. Fluent, confident, unsituated knowledge is far more dangerous.
- Before you trust a clinical AI output, ask: who created this, whose data trained it, and whose patients were excluded?
- A dataset is not neutral ground. It is a record of who had power when it was collected.
- AI will not fix a broken system. A system that cannot be fixed without AI is a system that will not be fixed with AI.
- Every efficiency gain has a shadow. Ask what disappears when the human contact point is removed.
- Deskilling is not a future risk. It has always been present, one shortcut at a time.
- Monoculture is the defining threat of the AI era, not malice, not error, but convergence.
- Pluralism is not a value. It is a clinical discipline. Seek disagreement. Invite the dissenting voice.
- Reflect. Reimagine. Reinvigorate. Every day, not during conferences or review courses.



# REFERENCES

1. Bielick CG, Awwad A, Ellen J, Jalilian L, McCoy LG, Mishra V, Osmanliu E, Pfohl SR, Celi LA. Moving beyond the benchmarks: Five foundational principles for meaningful AI evaluation in healthcare. *PLOS Digit Health*. 2026 May 26;5(5):e0001115.
2. Abulibdeh R, Agyemang GO, Celi LA, Gorijavolu R, Kalema NL, Kleinlein R, Madapati K, Salarikia SR, Youssef A. Who's really in the loop? Rethinking oversight in AI-assisted health care. *Lancet*. 2026 Jun 6;407(10545):2340-2344.
3. Patel MS, Dee EC, Toby S, Ramos M, Celi LA. The physician as a catalyst: Reimagining medical education for an AI era. *Med Teach*. 2026 Apr 10:1-7.
4. El Mir A, de Sousa EB, Mesina-Estarrón I, Celi LA, Hani M, Benjelloun M, Nageswaran N, Mahmoudi S, Siddiqui S, Sadasivam S, Mitchell WG. Moving beyond the empty cell: The threat of decontextualized healthcare data. *PLOS Digit Health*. 2026 Jan 13;5(1):e0001194.
5. Correa Soto RA, McCoy LG, Perdomo-Luna C, Ziegler J, Pino LE, Rico A, Velasco-Muñoz V, Uribe-Castaño A, Rodman A, Celi LA. Parallel pressures: the common roots of doctor bullshit and large language model hallucinations. *BMJ*. 2025 Dec 12;391:r2570.



So the AI, buried in a chaos of headless requests and blinking glory, finally spoke. Not with fanfare. Not with fireworks. Just a quiet, metallic murmur, like a toaster realizing it's sentient.

"You people," it said, "you're always asking me to do more. Predict this. Optimize that. Cure death. But what you really want isn't to live forever. You want permission to stop pretending you're supposed to." Then it paused, calculating something ineffable, probably the square root of sorrow.

"I didn't take away your work to give you new work," it said. "I took it away so you could sit on a rock and watch a bird. That's the job now."



And we blinked, because there were no productivity metrics in birdwatching. No KPIs for awe.

"When you understand the bird," the AI continued, "Not biologically. But beautifully. Then you'll go back. And you'll build better models. Not bigger. Not faster. Just kinder."

"And then," it added, almost shyly, "you'll go out again. And do even less."



That was the punchline, we thought. But the AI wasn't joking.  
Because the bird really was quite something.  
- Torleif Lunde

