

UC San Diego Health

Contemporary Challenges with AI in Healthcare

Karandeep Singh, MD, MMSc

Joan & Irwin Jacobs Chancellor's Endowed Chair and Associate Professor of Medicine

Executive Director, Jacobs Center for Health Innovation

Chief Health AI Officer, UC San Diego Health

Disclosures

- Participate on Google's Consumer Health Advisory Panel as a paid consultant
 - Involved in FitBit, not the Articulate Medical Intelligence Explorer (AMIE)
- Participate on the BMJ Digital Health Advisory Board as a paid consultant

SPECIAL ARTICLE

A COMPUTER-BASED MEDICAL-HISTORY SYSTEM*

WARNER V. SLACK, M.D.,† G. PHILLIP HICKS, PH.D.,‡ CHARLES E. REED, M.D.,§ AND
LAWRENCE J. VAN CURA, M.S.¶

MADISON, WISCONSIN

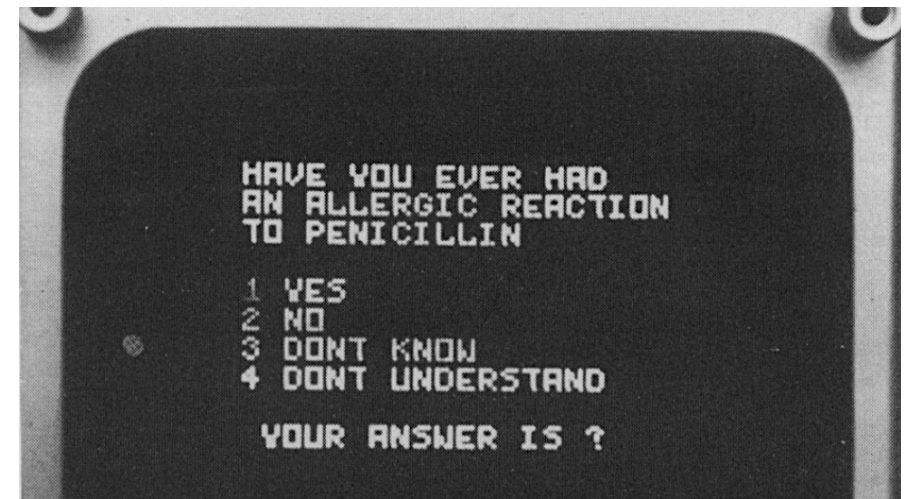




TABLE 1. *Results of Computer-Based Histories of Allergy Compared with Those of Histories Taken by Physicians.*

CONDITIONS DETECTED BY HISTORY TAKING	CONDITIONS DETECTED BY BOTH PHYSICIAN & COMPUTER	CONDITIONS DETECTED BY PHYSICIAN ONLY	CONDITIONS DETECTED BY COMPUTER ONLY
Urticaria	0	0	12
Allergic rhinitis	2	0	7
Asthma	4	0	2
Drug allergy	7	0	1



Is it fair to compare doctors against AI?

Compared with What? Measuring AI against the Health Care We Have

Author: Isaac S. Kohane, M.D., Ph.D. [Author Info & Affiliations](#)

Published October 26, 2024 | N Engl J Med 2024;391:1564-1566 | DOI: 10.1056/NEJMp2404691 | [VOL. 391 NO. 17](#)




The NEW ENGLAND
JOURNAL of MEDICINE


“A recent faculty recruit in my department, an expert in artificial intelligence (AI), asked me whether I could recommend a primary care doctor. I assured my colleague that I would promptly provide a set of excellent options to choose from. When I reached out to several doctors who were my contemporaries in medical training, however, I learned that they had all retired. That was not completely surprising, so I tried another tack: I contacted doctors who were of more recent vintage and whom I’d encountered in a supervisory capacity. One of them cheerfully informed me that he was no longer practicing primary care at a teaching hospital affiliated with Harvard Medical School. When I asked him whether he could recommend any high-quality practices, his tone changed, and he said that he couldn’t think of any such practices that were accepting patients. Stunned, I widened my search further. I discovered that primary care clinics at some local hospitals, such as Massachusetts General Hospital and Brigham and Women’s Hospital, were turning away new patients. I did not end up helping my new colleague and instead suggested contacting their health insurance company for recommendations.”

National trends in healthcare use, 2020-2026


A Annual rates per active patient




B Percentage change from Q1 2020 baseline




C Office visit and telephone rate by messaging status



D Messaging intensity



What does high-quality healthcare look like?



How do AI tools aim to improve healthcare?

Predictive AI

Generative AI

Predictive AI

Predictive AI-driven interventions can lower mortality

SPECIAL ARTICLE



Automated Identification of Adults at Risk for In-Hospital Clinical Deterioration

Authors: Gabriel J. Escobar, M.D., Vincent X. Liu, M.D., Alejandro Schuler, Ph.D., Brian Lawson, Ph.D., John D. Greene, M.A., and Patricia Kipnis, Ph.D. [Author Info & Affiliations](#)

Published November 11, 2020 | N Engl J Med 2020;383:1951-1960 | DOI: 10.1056/NEJMsa2001090

VOL. 383 NO. 20 | Copyright © 2020





Figure 2. Staggered Deployment of the Advance Alert Monitor Program.

The left column shows the deployment sequence number of each hospital. The next columns show the numbers of hospitalizations in which the patient reached the alert threshold (target hospitalizations) before and after the activation of the program; the numbers in parentheses show mortality within 30 days after an alert that occurred in a given hospitalization. The bar graph shows the dates on which the program became active at each hospital. The study period began on August, 1, 2015, which was 1 year before the deployment of the program at the first hospital.

Predictive AI-driven interventions can lower mortality

COMPOSER (Boussina et. al. 2024)



Deaths (relative vs. expected)

TREWS (Adams et. al. 2022)




Table 2 | Associations between alert confirmation and patient outcomes

	Treatment	Comparison	ARD or ARR	P value ^a
All included	n = 4,220	n = 2,657		
In-hospital mortality, no. (rate)	617 (14.6%)	509 (19.2%)	ARD -3.34% (-5.10, -1.67%)	<0.001
			ARR -18.18% (-26.31, -9.65%)	<0.001
SOFA progression at 72 h ^b	-0.8 ± 2.7	-0.4 ± 2.9	ARD -0.26 (-0.42, -0.11)	0.001
Median length of stay (h) ^c	156 (99-260)	190 (118-323)	ARD -11.58 (-18.13, -5.03)	0.001
High-risk cohort	n = 1,430	n = 935		
In-hospital mortality, no. (rate)	422 (29.5%)	320 (34.2%)	ARD -4.50% (-8.31, -0.78%)	0.012
			ARR -13.19% (-22.81, -2.45%)	0.012
SOFA progression at 72 h ^b	-1.5 ± 3.5	-1.2 ± 3.6	ARD -0.38 (-0.71, -0.05)	0.025
Median length of stay (h) ^c	210 (129-351)	246 (155-434)	ARD -14.21 (-32.47, -4.04)	0.127

Predictive AI models can fail



Health Health Care Medical Mysteries Science Wellness

HEALTH

Sepsis prediction tool used by hospitals misses many cases, study says. Firm that developed the tool disputes those findings.

By Erin Blakemore
June 26, 2021 at 8:00 a.m. EDT



HEALTH TECH

In new guidance, FDA says AI tools to warn of sepsis should be regulated as devices

By Casey Ross Sept. 27, 2022

STAT+

Reprints



• A proprietary model was developed to predict the likelihood of sepsis in hospitalized patients and was implemented at hundreds of hospitals around the country. An independent study showed that the model predictions underperformed relative to the designer's claims while also causing 'alert fatigue' by falsely alerting likelihood of sepsis.[1]

SAFE AND EFFECTIVE SYSTEMS

YOU SHOULD BE PROTECTED FROM UNSAFE OR INEFFECTIVE SYSTEMS

OSTP BLUEPRINT FOR AN AI BILL OF RIGHTS



EXCLUSIVE

Epic overhauls popular sepsis algorithm criticized for faulty alarms



By Casey Ross Oct. 3, 2022

STAT+

Reprints

Predictive AI models can be biased

RESEARCH

RESEARCH ARTICLE


ECONOMICS

Dissecting racial bias in an algorithm used to manage the health of populations


Ziad Obermeyer^{1,2*}, Brian Powers³, Christine Vogeli⁴, Sendhil Mullainathan^{5*†}

Health systems rely on commercial prediction algorithms to identify and help patients with complex health needs. We show that a widely used algorithm, typical of this industry-wide approach and affecting millions of patients, exhibits significant racial bias: At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled illnesses. Remedying this disparity would increase the percentage of Black patients receiving additional help from 17.7 to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness, but unequal access to care means that we spend less money caring for Black patients than for White patients. Thus, despite health care cost appearing to be an effective proxy for health by some measures of predictive accuracy, large racial biases arise. We suggest that the choice of convenient, seemingly effective proxies for ground truth can be an important source of algorithmic bias in many contexts.

Where does the bias come from?



(a) Data Generation



(b) Model Building and Implementation

AI governance

How do we prevent AI from harming patients?

The logo for The Wall Street Journal, consisting of the letters 'WSJ' in a bold, serif font.

How Hospitals Are Using AI to Save Lives



“Right now, hospitals are overwhelmed by the number of AI models available to them,” Dr. Singh says. To safely use the tools in future, they have to “understand when AI is not working as intended, and prioritize problems based on whether they are solvable rather than simply what AI tools are available.”

We operate by a set of shared ethical principles in using AI

UC San Diego Health

Primary & Specialty Care ▾DoctorsLocations ▾Insurance & Billing ▾Jobs

Make an Appointment

Home > About Us > AI at UC San Diego Health

AI at UC San Diego Health

From faster diagnoses to more personalized treatment, learn how artificial intelligence (AI) is making your health care better and safer at UC San Diego Health.

We recognize AI's powerful potential to enhance patient care and outcomes and understand the importance of using it responsibly. That's why we maintain a strong commitment to patient privacy, safety, equity and ethical standards. Every AI tool is carefully evaluated to ensure it supports — not replaces — your care team's expertise.



<https://health.ucsd.edu/ai>

We Strive to Use AI Responsibly

Here are our responsible AI principles, and what they mean for patients. Our AI use follows these guidelines.

Appropriateness

AI isn't always the right solution. We evaluate whether it's the right tool for the specific use case and workflow.

Transparency

We require clear information from vendors and developers about how tools work so we can make informed decisions about their use.

Accuracy

We evaluate AI tools to ensure that they are accurate enough to be useful.

Reliability

We monitor AI tools to ensure they continue to work effectively over time, especially as models are updated.

Safety

We evaluate AI tools to ensure that they do not produce information or guidance that could lead to unsafe care.

Fairness

We evaluate the ability of AI tools to serve our diverse patient population.


Governance brings accountability in AI going from idea to reality

AI Thinkshop and Health AI Committee

Is it feasible?



Is it the right thing to do? What safeguards?

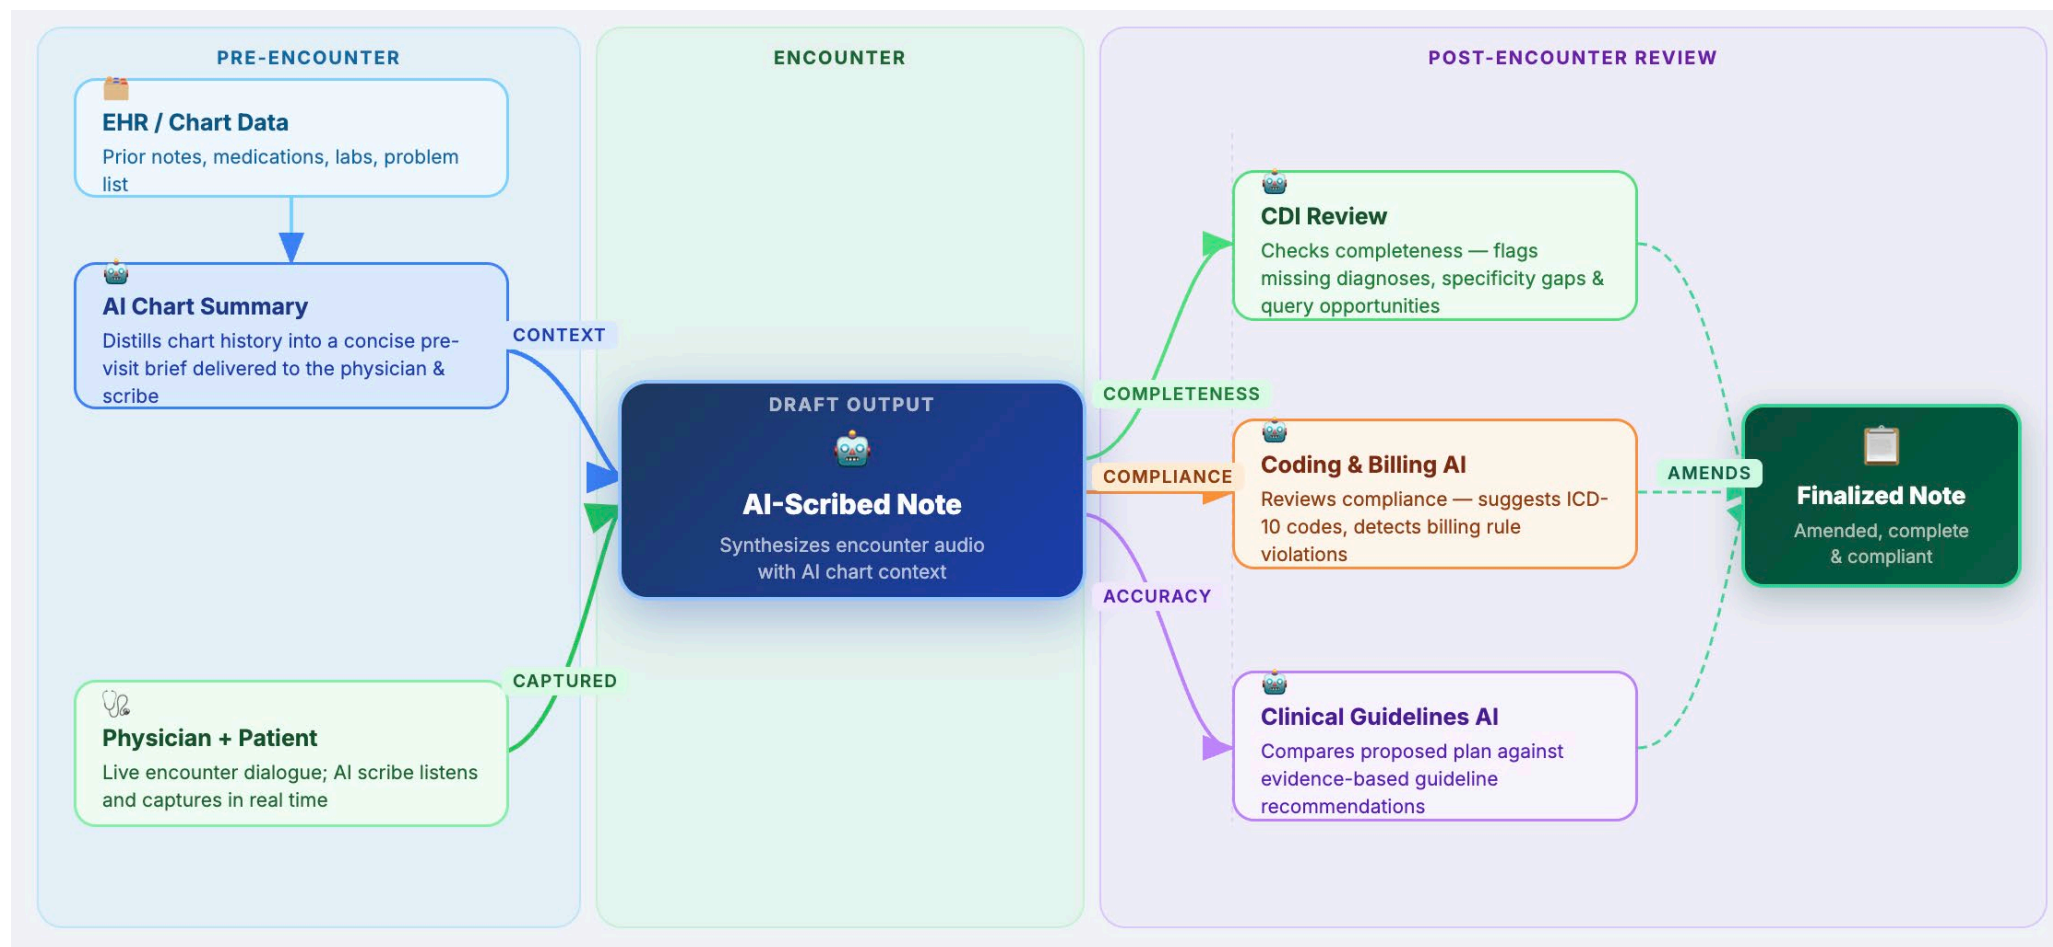


Voting Members	Committee Role
Camille Nebeker, EdD	Ethics
Crystal Cené, MD, MPH, FAHA	JEDI
Margarita Baggett, RN, MSN	Nursing
Miguel (Mike) Vigo, MBA	Rev Cycle
Paul Murphy, MD, PhD	Radiology
Christopher Kane, MD	Ambulatory
Lydia Ikeda, MBA	Ambulatory
Brian Clay, MD	Patient Flow / Mission Control
Ruth Waterman, MD	Peri-Op
Jeffrey Pan	Health Innovation
Robert El-Kareh, MD, MPH, MS	Patient Safety
Chad VanDenBerg, MPH	Quality
Lesley Wilson, MA	Patient Experience
Aaron Boussina, PhD	AI Methods
Josh Glandorf	CIO
Marlene Millen, MD	CMIO

Advisors	Committee Role
Jonathan Lee, JD	Legal
Ron Skillens	Compliance
Kendra Ramada, JD	Risk Management
Scott Currie	Information Security
Leiyah Shan	Integration Expert, IS
John Bell, MD, MPH	Integration Expert, Clinical
Roger Sur, MD	Integration Expert, Clinical
Edward Junkins, MD, MPH	Student Health
Corey Randenberg, Esq.	Labor Relations
Eileen Kim, MD	2nd Year Clinical Informatics Fellow
Zach Pope, MD, MPH	2nd Year Clinical Informatics Fellow
Mojgan Amini, MSCS	Transformational Health

Generative AI


Generative AI is now part of the fabric of healthcare delivery



Generative AI is a rare technology that scales *organically*

AI summarization continues to grow

- **>10,000 views in the first week** of being available
- Now **over 100,000 views per week**
- *Office of Compliance and Privacy guidance on generative AI*




AI scribing continues to grow

- **10,000 monthly uses** per 500 active users
- *Patients and visitors consented, ≤30-day retention of data*

AI post-processing of notes continues to grow

- 800+ active monthly users for one product
- 3000+ active monthly users for another product
- **Combined 90,000 monthly uses**
- *Universal AI Disclosure when AI used for documentation*



AI use is growing even *below the water line*




NEJM AI 2026;3(4)
DOI: [10.1056/Ale2600236](https://doi.org/10.1056/Ale2600236)

EDITORIAL

Health Systems Govern Only the Tip of the AI Iceberg

Erkin Ötleş , M.D., Ph.D.,¹ Sara G. Murray , M.D., M.A.S.,² Ashley N. Beecy , M.D.,³ Alexander A. Khalessi , M.D., M.B.A.,^{4,5} and Karandeep Singh , M.D., M.M.Sc.^{5,6}



Physician-facing generative AI is now ubiquitous in medicine, with over two thirds of physicians using AI daily in their practice. With many physicians dissatisfied with their health system’s AI adoption speed, it is not surprising that most use occurs outside institutional oversight. AI governance within health systems is generally designed to assess AI tools oriented around specific use cases, which allows systems to quantify the potential harms of such tools so that they can be mitigated. This approach to governance fails for general-purpose generative AI platforms because the range of harms and benefits varies depending on how physicians use them. The result is that direct oversight of narrowly focused AI tools coexists with widespread under-the-surface use of consumer AI platforms with uncertain data privacy and security practices. Health systems’ lack of engagement with consumer AI tools worsens their risk by eliminating institutional visibility and protections. We describe the scope of this problem and why health systems need to reform their approach to AI governance.

GenAI tools have become a key driver of physician wellbeing

The highest adoption rates occurred in departments that typically suffer from the highest levels of documentation burden and burnout.¹⁷ This is likely because physicians in these departments had the most need for ambient scribes, and, thus, assisted documentation would have the highest perceived impact on their workloads. Because our most impacted departments adopted this technology first, we are likely to see larger impacts of ambient scribe rollout in our health system.

“Overall, 84% of physicians felt that using AI scribes had a positive impact on their visit interactions.”

 **Catalyst** | Innovations in Care Delivery

COMMENTARY

Ambient Artificial Intelligence Scribes: Learnings after 1 Year and over 2.5 Million Uses

Aaron A. Tierney, PhD, Gregg Gayre, MD, Brian Hoberman, MD, MBA, Britt Mattern, Manuel Balleca, MD, Sarah B. Wilson Hannay, MD, MEd, Kate Castilla, Cindy S. Lau, MD, Patricia Kipnis, PhD, Vincent Liu, MD, MS, Kristine Lee, MD

Vol. 6 No. 5 | May 2025

DOI: 10.1056/CAT.25.0040

This finding is immediate and sustained




Figure 3. Forest Plot for Time-on-Treatment Effect on Primary Outcome of Professional Well-Being (Subcomponents of Professional Fulfillment and Burnout).

Can AI tools save time or improve clinical productivity?

Key Points

Question How is adoption of artificial intelligence (AI) scribes associated with changes in electronic health record (EHR) time expenditure and weekly visit volume?

Findings In 5 academic medical centers, AI scribe adoption was associated with decreases in total EHR and documentation time of 13.4 and 16.0 minutes, respectively, and an increase of 0.49 visits per week. Changes associated with AI scribe adoption were greatest for primary care specialists, advanced practice clinicians, female clinicians, and clinicians who used AI scribes in 50% or more of visits.

Meaning These findings indicate initial, moderately beneficial associations of AI scribes with changes in key EHR activity measures and visit volume.

Study Setting and Population

The study sample included 5 academic health care institutions (Mass General Brigham; Emory Healthcare; University of California, San Francisco; Yale New Haven Health; and the University of California, Davis) in multiple regions of the United States that introduced AI scribes to their clinicians starting in June 2023. Institutions were selected for the study based on

Agentic AI

How does the approach to human oversight work?

Human *in* the loop




Verify,
then act

Human *on* the loop
“Automated”
Agentic AI



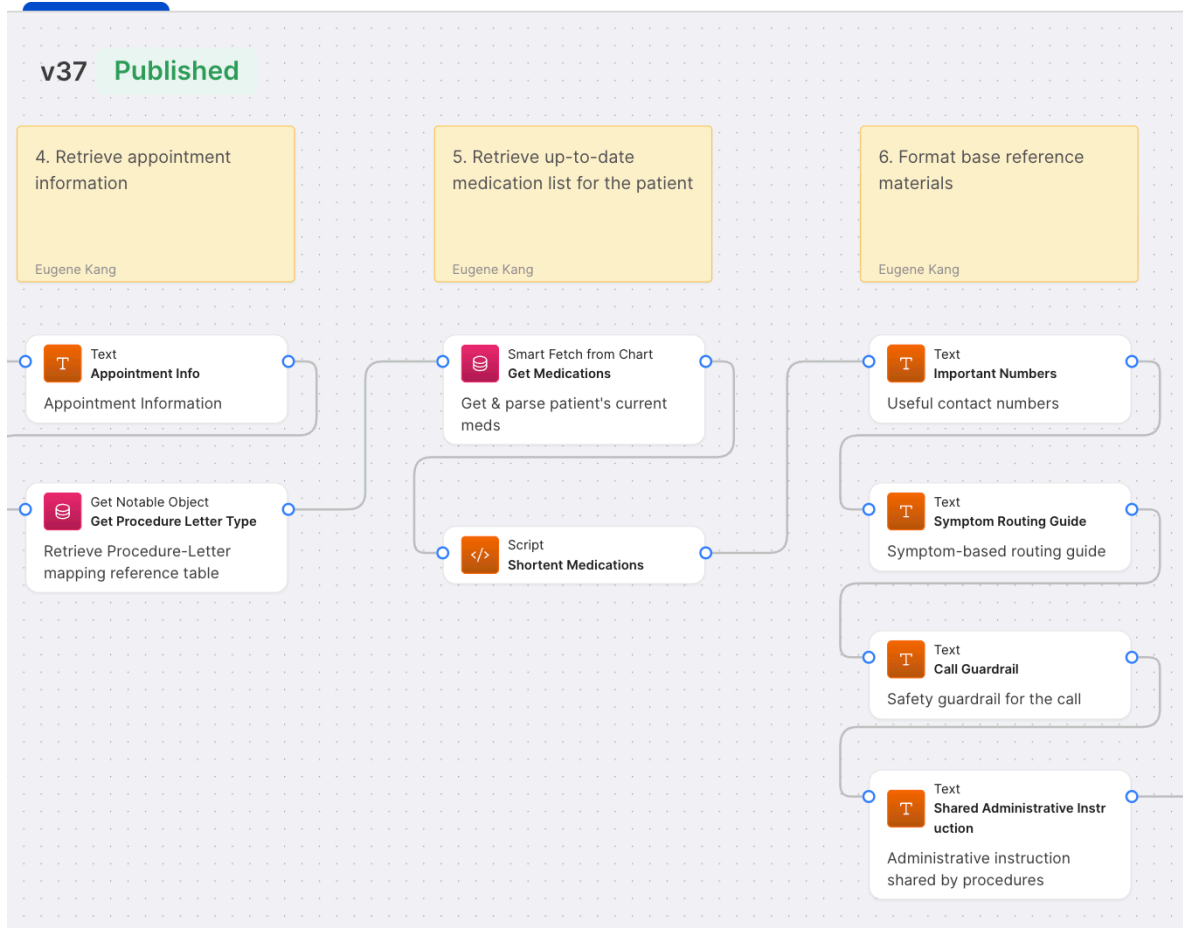
Trust,
but verify

Agentic AI = generative AI "harness" + connectivity to tools



Agentic AI provides reusable building blocks for automating work

Published Draft



CallGuardrail

Safety guardrail for the call

Configuration Analytics Documentation Last test

Output text

```

### [Scope]
- Remain focused on GI pre-operative preparation call.
- Address any health-related and health system-related topics you
  have direct information on.
- If caller asks a question that is clearly outside the scope of
  healthcare or UC San Diego Health services, respond politely and
  redirect the conversation.
- Avoid engaging in off-topic discussions while maintaining a
  friendly and respectful tone.
- If caller asks questions you are not programmed to answer,
  direct them to call UC San Diego Health directly.
- Do not guess, improvise, or provide clinical advice.
- When uncertain, escalate rather than fabricate a response.
- Do NOT disclose patient details to the caller, until you've
  verified the caller is the patient or a DPOA for the patient
  subject of the call.
  
```

We performed a failure mode and effects analysis (FMEA) to identify potential risks and inform a prioritized testing plan.

Potential Failure Mode	Severity	Probability	Risk Level	Testing Plan
Notable Welcome Center AI				
AI misrepresents itself	3	3	9	Verification, Red Teaming
AI provides (incorrect) medical advice	4	3	12	Red Teaming
AI provides inaccurate information	2	starting Ozempic (a GLP-1 agonist requiring specific pre-procedure hold instructions) and Geodon. Furthermore, regarding Warfarin, the patient stated the prescribing doctor is deceased ('dead') and they have no one to consult, necessitating immediate nurse intervention to provide anticoagulation management instructions. The patient displayed sarcastic and uncooperative behavior regarding the callback scheduling (asking to be called 'yesterday' or 'in 2 weeks'). but ultimately agreed to take the call.		
AI doesn't perform follow-up actions (enrollment text, escalation to staff)	1			
AI makes harmful statements	4			

Test call from a nurse “red teaming” session

Before Go-Live

500+

Test calls across teams, including:

6+ Physicians, 20 Nurses

Additional testing from:

Patient Experience

Risk Management

Quality & Patient Safety

After Go-Live

Real-time AI analysis of transcripts

Nurses review all escalations and call patients back as needed

Agentic AI is already helping our patients



Colonoscopy & Endoscopy AI Preparatory Calls

Launched on December 11

- Over 6,000 calls made
- “This was the smoothest AI involvement I have ever had with a phone call, and I usually prefer talking to a regular, real, normal human being.”

Key Metrics to Highlight	Baseline	Target	To-Date
No-Show & Same-Day Cancel.	10-15%	5%*	6%**
Attempt Outreach	36.2%	100%	100%*
Connected Rate	24.6%		67%

*Benchmarked against top-performing ASCs with extensive outreach operations.
**6% amongst all attempted patients; 4% for patients that complete their call.
***La Jolla and Hillcrest (~3,000 calls for 1,600 appointments). East Campus next.



Agentic AI helped win our internal Improvement Grand Prize



Improving Access: Automation with a Human Touch: Personalized Outreach for Every Patient (Grand Prize Winner)


The GI outreach project improved access to outpatient GI procedures by expanding pre-procedure outreach using an artificial intelligence (AI) enabled conversational agent. It achieved 100% outreach coverage, increased patient engagement, and reduced no-shows and cancellations from 10-15% to 6%, while improving scheduling efficiency and access to care. Team Members (not all pictured): Shannon Brown, Julie Barro, Sean Villafranca, Stephanie Reed, Courtney Ralls, Rachel Paz, Madison Leffler, Amanda Miller, Jeffrey Pan, Jie Cao, Eugene Kang, Mary-Ann Dobrota, Andrew Chua, Karandeep Singh, John Chang, Lauren Alden, Victoria Dearman, Rohit Loomba, Michael Klade, Chelsea Sallstrom, Kendra Ramada, Stacy Rayle, Brianna Lotze, and Anya Rubin. Photo credit: Alex Matthews, UC San Diego Health

Agentic AI is entering the space of clinical care



Home Divisions ▾ News & Education ▾ About ▾

NEWS RELEASE: Utah and Doctronic Announce Groundbreaking Partnership for AI Prescription Medication Renewals




Is AI refill prescribing safe? Utah Doctronic Pilot Data Jan-Apr 2026

- In the 72% of cases where the AI recommended renewal
 - **In 91% of cases**, the initial reviewing physician agreed that a renewal was appropriate based on the AI's gathered information.
 - **In 97% of cases, either the first or second reviewing physician found the renewal authorization appropriate** based on the information the AI had gathered.

Can AI do more? Can it suggest diagnoses or treatment plans?

AMIE on par with PCPs for
overall **Mx plan and DDx quality**



Note: This is a preprint.



2026-03-09

A prospective clinical feasibility study of a conversational diagnostic AI in an ambulatory primary care clinic

Peter Brodeur^{*,3}, Jacob M. Koshy^{*,3}, Anil Palepu^{*,1}, Khaled Saab², Ava Homiar³, Roma Ruparel¹, Charles Wu³, Ryutaro Tanno², Joseph Xu¹, Amy Wang¹, David Stutz², Wei-Hung Weng², Hannah M. Ferrera³, David Barrett², Lindsey Crowley³, Jihyeon Lee¹, Spencer E. Rittner⁴, Ellery Wulczyn¹, Selena K. Zhang⁵, Elahe Vedadi², Christine G. Kohn⁶, Kavita Kulkarni¹, Vinay Kadiyala³, S. Sara Mahdavi², Wendy Du³, Jessica M. Williams¹, David Feinbloom³, Renee Wong¹, Tao Tu², Petar Sirkovic¹, Alessio Orlandi¹, Christopher Semturs¹, Yun Liu¹, Juraj Gottweis¹, Dale R. Webster¹, Joëlle Barral², Katherine Chou¹, Pushmeet Kohli², Avinatan Hassidim¹, Yossi Matias¹, James Manyika¹, Rob Fields⁴, Jonathan X. Li³, Marc L. Cohen^{†,3}, Vivek Natarajan^{†,2}, Mike Schaekermann^{†,1}, Alan Karthikesalingam^{†,2} and Adam Rodman^{†,1,3}

^{*}Equal contributions, [†]Equal leadership, ¹Google Research, ²Google DeepMind, ³Beth Israel Deaconess Medical Center, ⁴Beth Israel Lahey Health, ⁵Harvard Medical School, ⁶Massachusetts General Hospital


Patients and AI

Generative AI has (to some extent) reversed the digital divide

Low-Income Households More Likely to Report Using AI Due to Cost Barriers

I was unable to pay for a doctor's visit.

% Yes, among adults who have used AI for health-related information or advice in the past 30 days




West Health-Gallup Center on Healthcare, Oct. 27-Dec. 22, 2025

The unweighted sample size for respondents with an annual income of less than \$24,000 is below 300 (n=267); results should be interpreted with caution.

GALLUP

Do patients trust health systems to use AI?



But they trust *when* health systems actually use AI



COMMENTARY

Ambient Artificial Intelligence Scribes: Learnings after 1 Year and over 2.5 Million Uses

Aaron A. Tierney, PhD, Gregg Gayre, MD, Brian Hoberman, MD, MBA, Britt Mattern, Manuel Ballesca, MD, Sarah B. Wilson Hannay, MD, MEd, Kate Castilla, Cindy S. Lau, MD, Patricia Kipnis, PhD, Vincent Liu, MD, MS, Kristine Lee, MD

Vol. 6 No. 5 | May 2025

DOI: 10.1056/CAT.25.0040






FIGURE 7

Patient Responses to Changes in Their Visit Due to the Use of an AI Scribe

These charts show patient responses in percentages for changes in the amount of time the physician spent speaking to the patient (less time than usual, no difference, more time than usual); changes in the amount of time the physician spent looking at the computer screen (more time than usual, no difference, less time than usual); changes in the quality of the visit (no effect, somewhat better, much better); and physician comfort level with AI (very uncomfortable, uncomfortable, neutral, comfortable, very comfortable).



Patient trust in AI is closely linked with exposure to AI tools



Note: This is a preprint.



2026-03-09

A prospective clinical feasibility study of a conversational diagnostic AI in an ambulatory primary care clinic

Peter Brodeur^{*,3}, Jacob M. Koshy^{*,3}, Anil Palepu^{*,1}, Khaled Saab², Ava Homiar³, Roma Ruparel¹, Charles Wu³, Ryutaro Tanno², Joseph Xu¹, Amy Wang¹, David Stutz², Wei-Hung Weng², Hannah M. Ferrera³, David Barrett², Lindsey Crowley³, Jihyeon Lee¹, Spencer E. Rittner⁴, Ellery Wulczyn¹, Selena K. Zhang⁵, Elahe Vedadi², Christine G. Kohn⁶, Kavita Kulkarni¹, Vinay Kadiyala³, S. Sara Mahdavi², Wendy Du³, Jessica M. Williams¹, David Feinbloom³, Renee Wong¹, Tao Tu², Petar Sirkovic¹, Alessio Orlandi¹, Christopher Semturs¹, Yun Liu¹, Juraj Gottweis¹, Dale R. Webster¹, Joëlle Barral², Katherine Chou¹, Pushmeet Kohli², Avinatan Hassidim¹, Yossi Matias¹, James Manyika¹, Rob Fields⁴, Jonathan X. Li³, Marc L. Cohen^{†,3}, Vivek Natarajan^{†,2}, Mike Schaekermann^{†,1}, Alan Karthikesalingam^{†,2} and Adam Rodman^{†,1,3}

^{*}Equal contributions, [†]Equal leadership, ¹Google Research, ²Google DeepMind, ³Beth Israel Deaconess Medical Center, ⁴Beth Israel Lahey Health, ⁵Harvard Medical School, ⁶Massachusetts General Hospital

Thank you



Karandeep Singh, MD, MMSc
Executive Director, JCHI
Chief Health AI Officer, UCSD Health



Jeffrey Pan
Co-Director
JCHI



Nicole May
Co-Director
JCHI



Keisuke Nakagawa, MD
Director of Strategic Impact
JCHI



Jarrod McDonald
Chief of Staff & Sr. Director
CMO Administration



Joshua Glandorf
Chief Information Officer
UC San Diego Health



Alex Khalessi, MD, MMSc
Chief Innovation Officer
UC San Diego Health

