

GENERATIVE AI IN HEALTHCARE: INDUSTRY TRENDS AND AI AT HOPKINS

Sara Murray, MD, MAS
VP, Enterprise Chief Medical Informatics and AI Officer
Johns Hopkins Medicine

Brian Hasselfeld, MD
Executive Medical Director, Digital Health and Innovation
Johns Hopkins Medicine

JOHNS HOPKINS MEDICINE

© Johns Hopkins Medicine 2026

AI to Transform Healthcare Delivery

Vision: JHM will serve as a global leader in the responsible use of AI to solve the biggest problems in healthcare

1. Create value today

- Diverse portfolio of technology aimed at improving care, learning and discovery across JHHS and School of Medicine.
  - Leverage proven or highly promising applications / partnerships to drive value.

2. Position JHM to lead

- Invest in the infrastructure, processes, and people that enable research, discovery and new models of care delivery.
  - Strengthen institutional assets (e.g. Med Ed AI, DSAI)
  - Strategic industry partnerships (e.g. OpenAI, Anthropic, OpenEvidence)



AI TRENDS IN US HEALTHCARE

JOHNS HOPKINS
MEDICINE

Industry-related Changes

EXCLUSIVE

Artificial intelligen… prescribing medic…

Pilot program will test how far patients a…

January 7, 2026

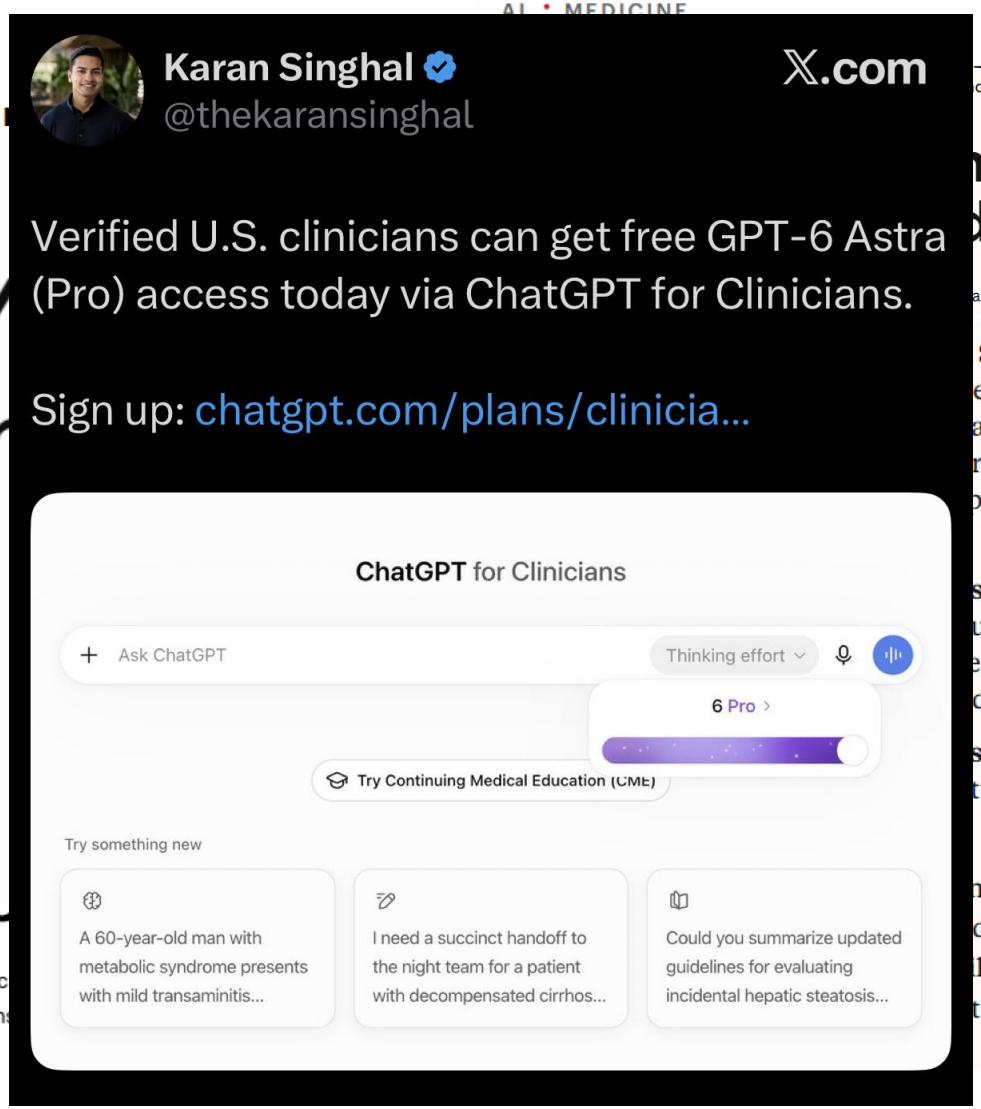
Introducing Ch…

A dedicated experience in ChatGPT …

January 8, 202…

Introducin… for Heal…

Secure AI products to help healthcare organizations sc… admin work for teams, and power custom clinical solutions…



…ouncements

…g Claude in …d the life sciences

…an 11, 2026

Services (CMS) Coverage Database, …erage Determinations. This enables …age requirements, support prior …ronger claims appeals. This connector …pliance, and patient-facing teams work …

…seases, 10th Revision (ICD-10). Claude …ure codes to support medical coding, …ent. This data is provided by the CMS …d Prevention (CDC).

…stry, which allows Claude to help with …tworking directory management, and …

…n now use Claude for Enterprise, they …d connectors, including **PubMed**, …illion pieces of biomedical literature and …t research, and produce up-to-date …

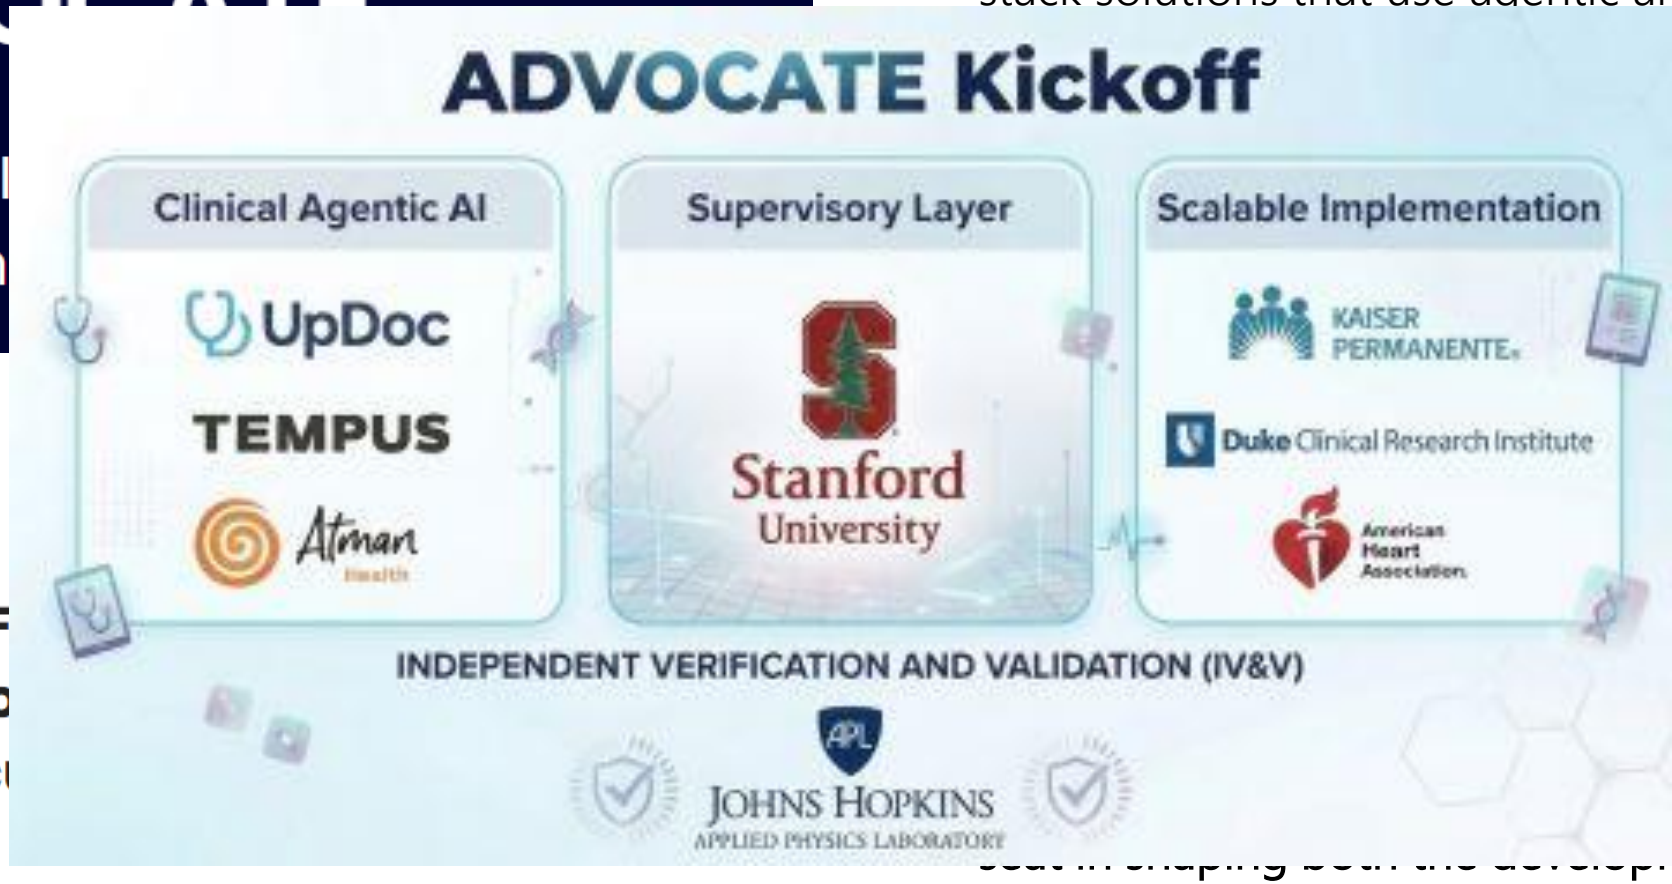
© Johns Hopkins Medicine - Not for Distribution

ARPA-H and Autonomous AI

ADVOCATE

Agentic AI
Care Tran

**What if an F
AI could pro
cardiovascu**



- "support the development of FDA-authorized full-stack solutions that use agentic artificial **intelligence**
- **y care for every**
- **art disease."**
- le to provide
- n diet and exercise,
- tance, make
- **oses and potential**
- e able to **write or**
- **onomously."**
- ill be actually
- a new member of
- ather than hinders
- many promising
- **ATE will recruit**
- **have a front-row**
- nent and

deployment of these technologies in their patient populations"

© Johns Hopkins Medicine - Not for Distribution

Utah AI Prescribing Update

doctronic

- Autonomous AI **refills** for 200 medications for chronic conditions
- Phase 1: human review of all prescriptions

legion HEALTH

- Autonomous AI **refills** for 15 *non-controlled* medications exclusively in psychiatric space
- Phase 1: human review of all prescriptions

State of Utah
Utah Medical Licensing Board

It is the strong recommendation of the Utah Medical Licensing Board that this program be immediately suspended pending further discussion.

We look forward to working with the Office of Artificial Intelligence Policy to explore ways to safely implement AI in the practice of medicine.

Sincerely,

Utah Medical Licensing Board

Division of Professional Licensing
UTAH DEPARTMENT OF COMMERCE

Your letter strongly recommended the immediate suspension of the Doctronic program pending further discussion. **Because the pilot is currently in Phase One—where a licensed physician reviews every refill decision, ensuring the program is already operating safely at the standard of care—we will not be suspending the pilot at this time.** The OAIP requires detailed monthly reporting from Doctronic and retains the absolute authority to modify or cancel the pilot if safety benchmarks are not met. **The OAIP and DOPL plan to discuss the progress of the pilot with the Board as more data is generated.**

© Johns Hopkins Medicine - Not for Distribution



AI AND GOVERNANCE

JOHNS HOPKINS
MEDICINE

Clinically-led Responsible AI Program



Responsible Artificial Intelligence Task Force Recommendations

March 29, 2024

Background and Vision

Johns Hopkins Medicine (JHM) executive leadership chartered a task force of JHU and JHHS leaders to develop an implementation strategy that ensures JHM becomes a global leader in the responsible use of Artificial Intelligence (AI) in healthcare. To achieve a leadership role, the task force recognizes that substantial organizational investment and redirection of current resources and efforts will be required, in partnership with our budgetary and financial teams. AI has the potential to transform the healthcare landscape. JHM must establish a framework and processes that carefully balance risks and benefits, ensuring our AI initiatives align with JHM values of excellence and discovery, leadership and integrity, diversity and inclusion, and respect and collegiality.

Our commitment to responsible AI will require complex coordination across both JHM and Johns Hopkins University (JHU), tapping into expertise and coordinating efforts and investments. In doing so, we acknowledge the critical importance of clarifying appropriate usage of patient data and sensitive business data, and improved transparency and communication with our patients, faculty, staff, and community about the benefits of this new generation of AI capabilities as well as the risks. Community engagement is foundational to our work.

The Task Force believes that to be a leader it is critical to partner with industry to enable access to the most advanced AI technologies and to leverage entrepreneurial investment in new tools and companies. JHM should aspire to be a go-to destination for industry partners to identify optimal uses of AI and safely develop, deploy, and validate AI. In addition, we must partner with the users of AI to ensure that our workflow and tools support them in providing excellent patient care and outstanding operations. Our policy and governance processes must preserve the trust patients have in Johns Hopkins while enabling research and commercial opportunities.

The recommendations in this document lay the groundwork for a governance infrastructure capable of addressing the complexity of AI use in healthcare while maintaining the agility needed to evolve with the AI landscape. Swift decision making will be critical, and this is often difficult in the complex JHM environment. The Task Force believes that as this technology transforms the healthcare industry there are substantial opportunities to improve clinical care and reduce healthcare delivery costs while improving access and operations. Collaboration with our peer institutions will offer a way to speed up the development of responsible AI governance in a cost-effective manner. Starting with simpler models in lower-risk workflows will allow JHM to build its understanding of AI and appropriate governance. Actively engaging with our clinicians and staff will help maximize AI's value and minimize risks while accelerating our AI proficiency as a learning healthcare organization. We must work to ensure that JHM is seen as expert in AI opportunity identification, validation, and evaluation rather than a mere source of data or brand recognition and credibility.

Our focus on advancing responsible AI reaffirms our commitment to preserving patient trust and health equity. As JHM charts a path to lead in this critical area, the proposed responsible AI governance will

1

Multidisciplinary Expertise

- Clinical
- Informatics
- Bioethics
- Health IT
- Safety/Quality
- Finance
- Government Affairs
- Education
- Engineering

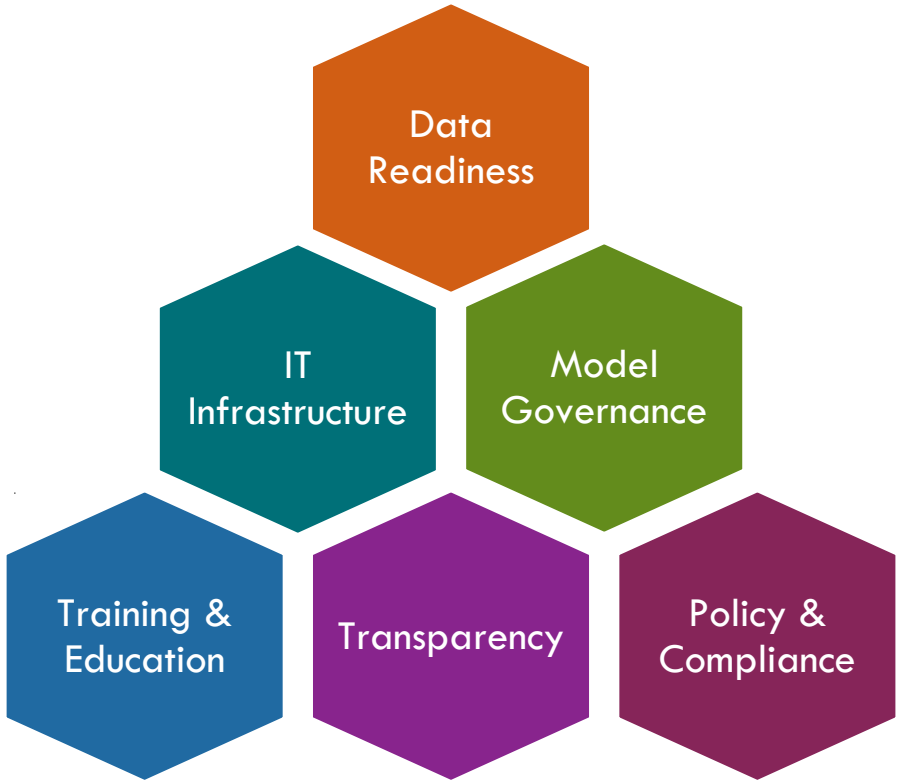
© Johns Hopkins Medicine - Not for Distribution

Clinically-led Responsible AI Program

7 Core Principles For Use

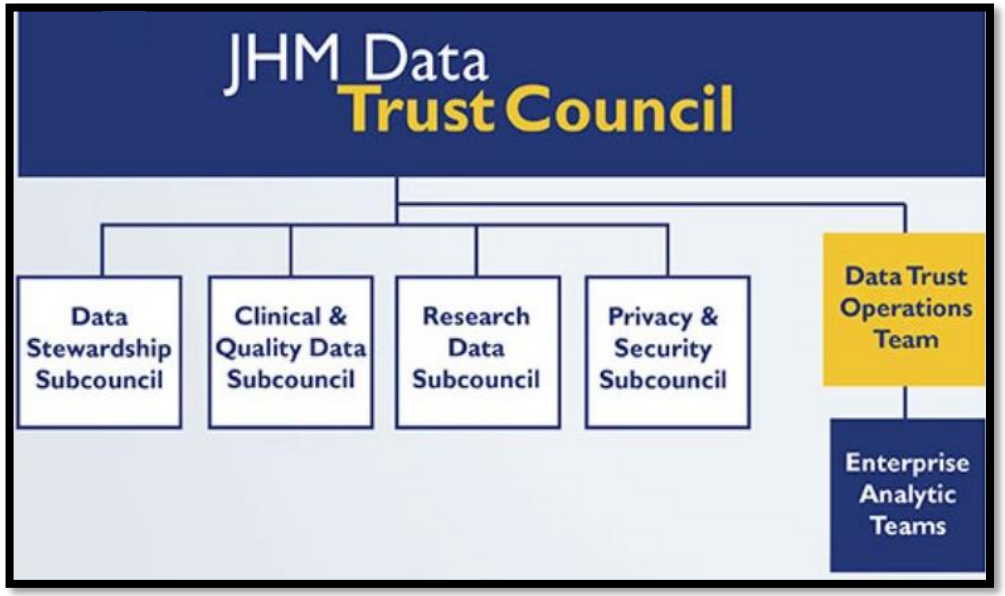


Implementation Plan



© Johns Hopkins Medicine - Not for Distribution

Existing Governance Structures



The Data Trust governs the infrastructure, standards, policies and procedures for patient and member-related data.



The Software Intake Process is a streamlined, single-entry point to collect and govern new technology requests.

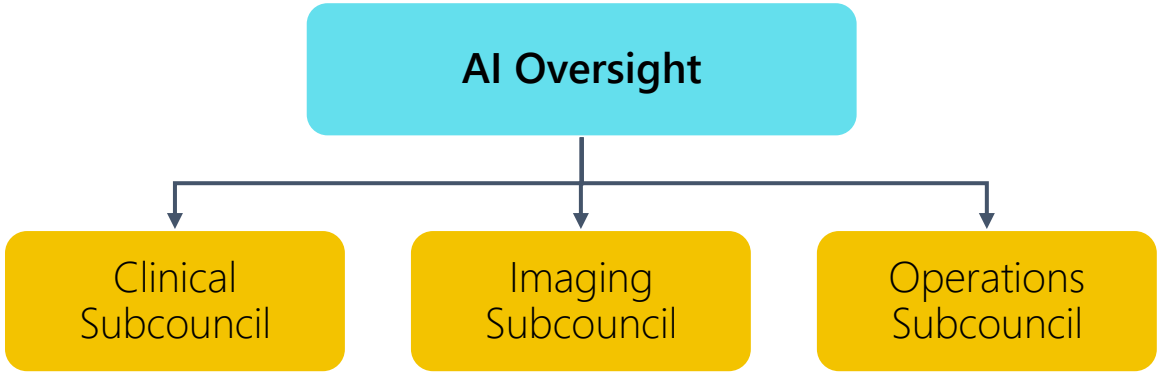
© Johns Hopkins Medicine - Not for Distribution

Recommendation: Governance Structure

7 Core Principles For Use



Subcouncils ensure proposals sufficiently meet our core principles.

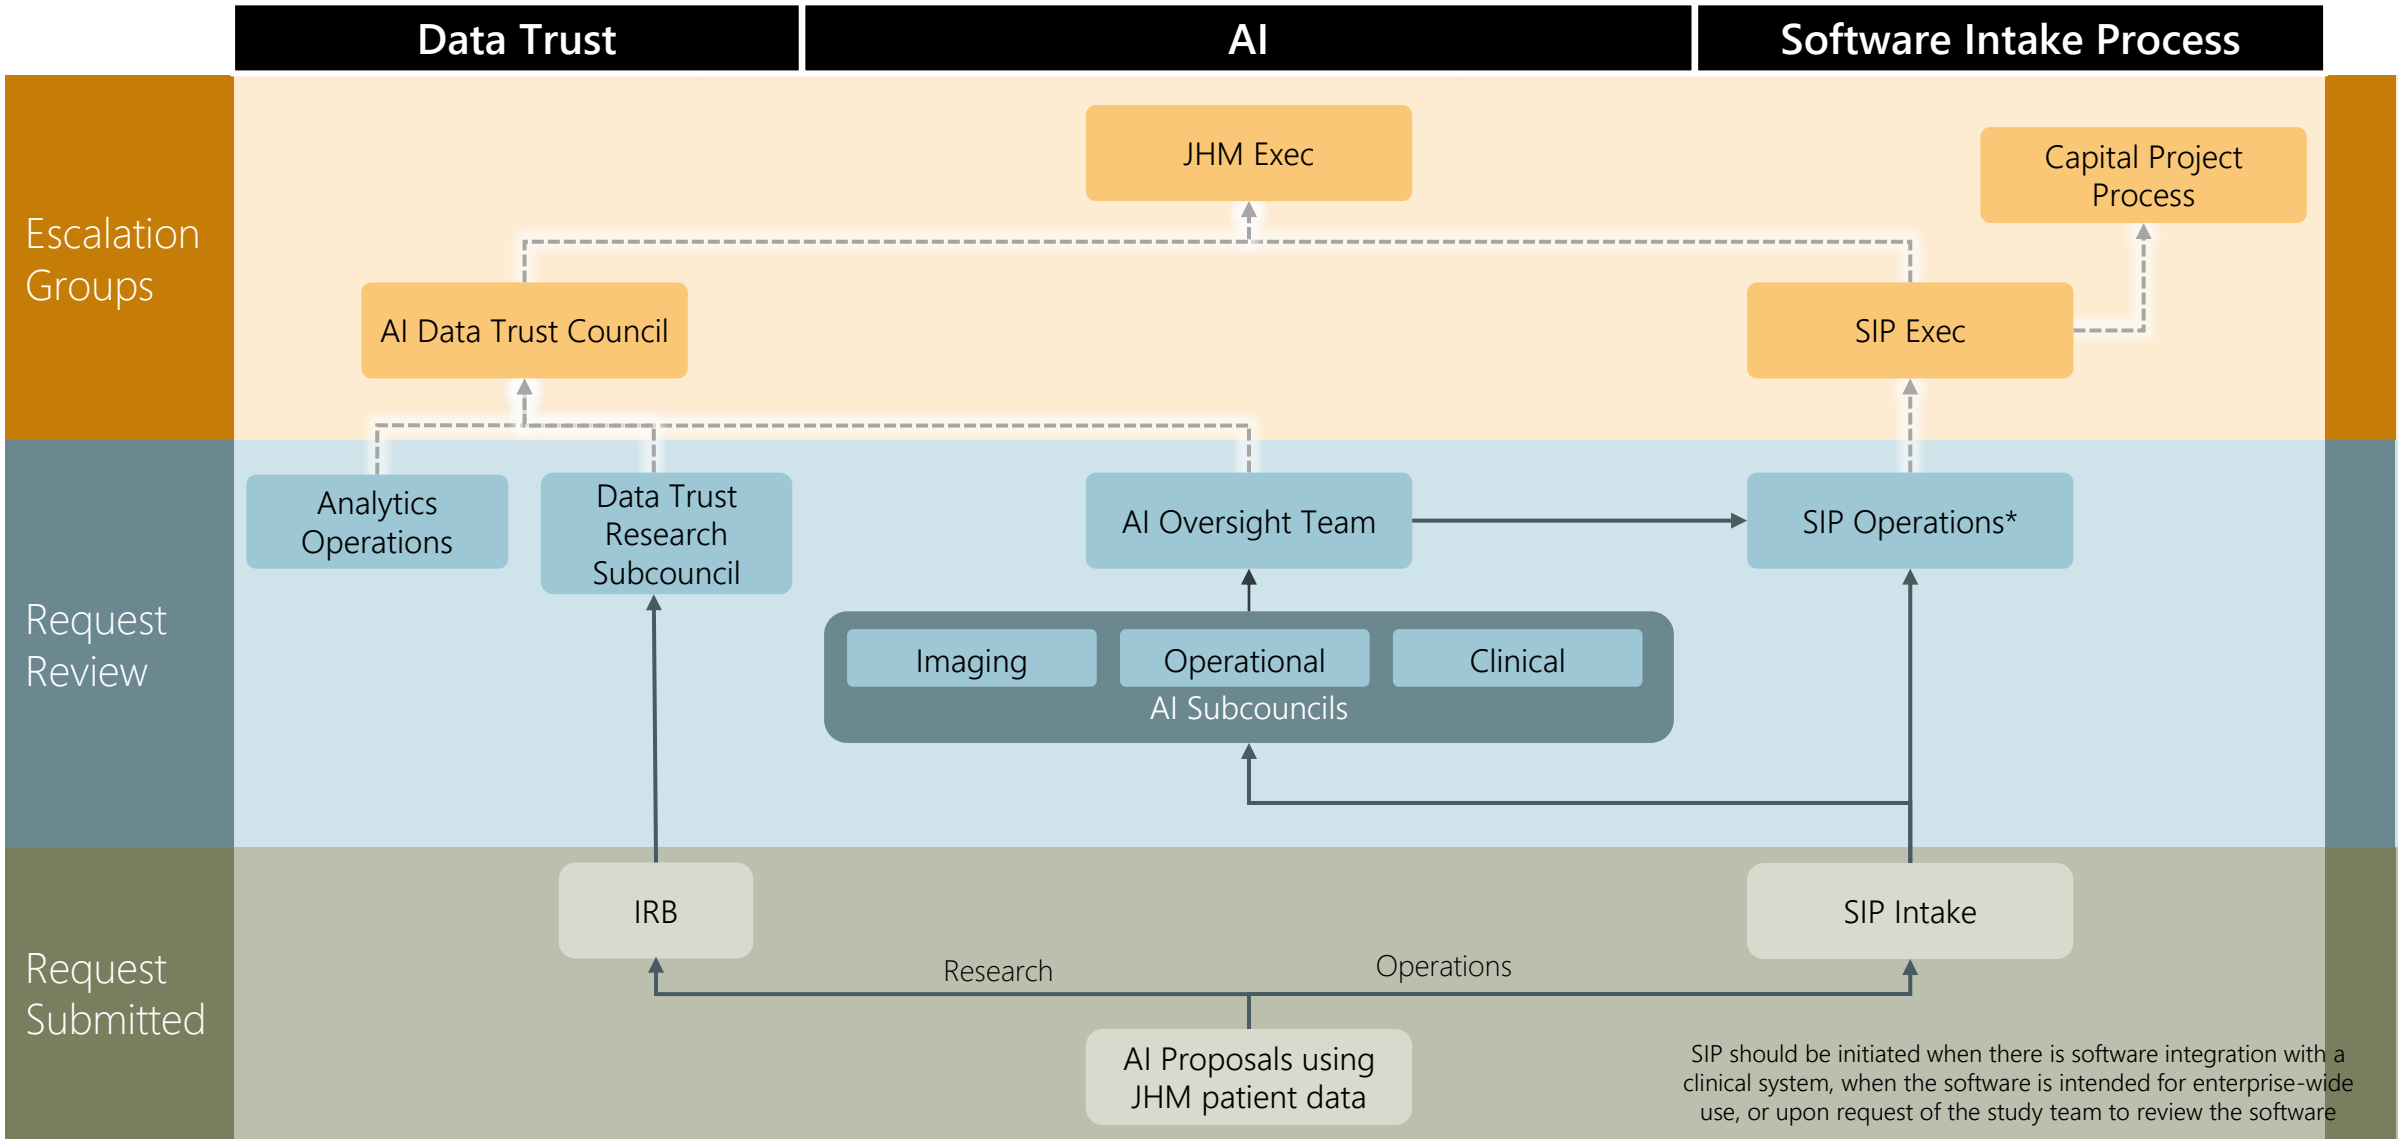


JHM Responsible AI Council

- internal & external experts (JHHS, SOM, JHU)
- oversight for AI tool use, ongoing evaluation, and cessation
- provide input on AI model development, implementation, risk mitigation, and general AI strategy

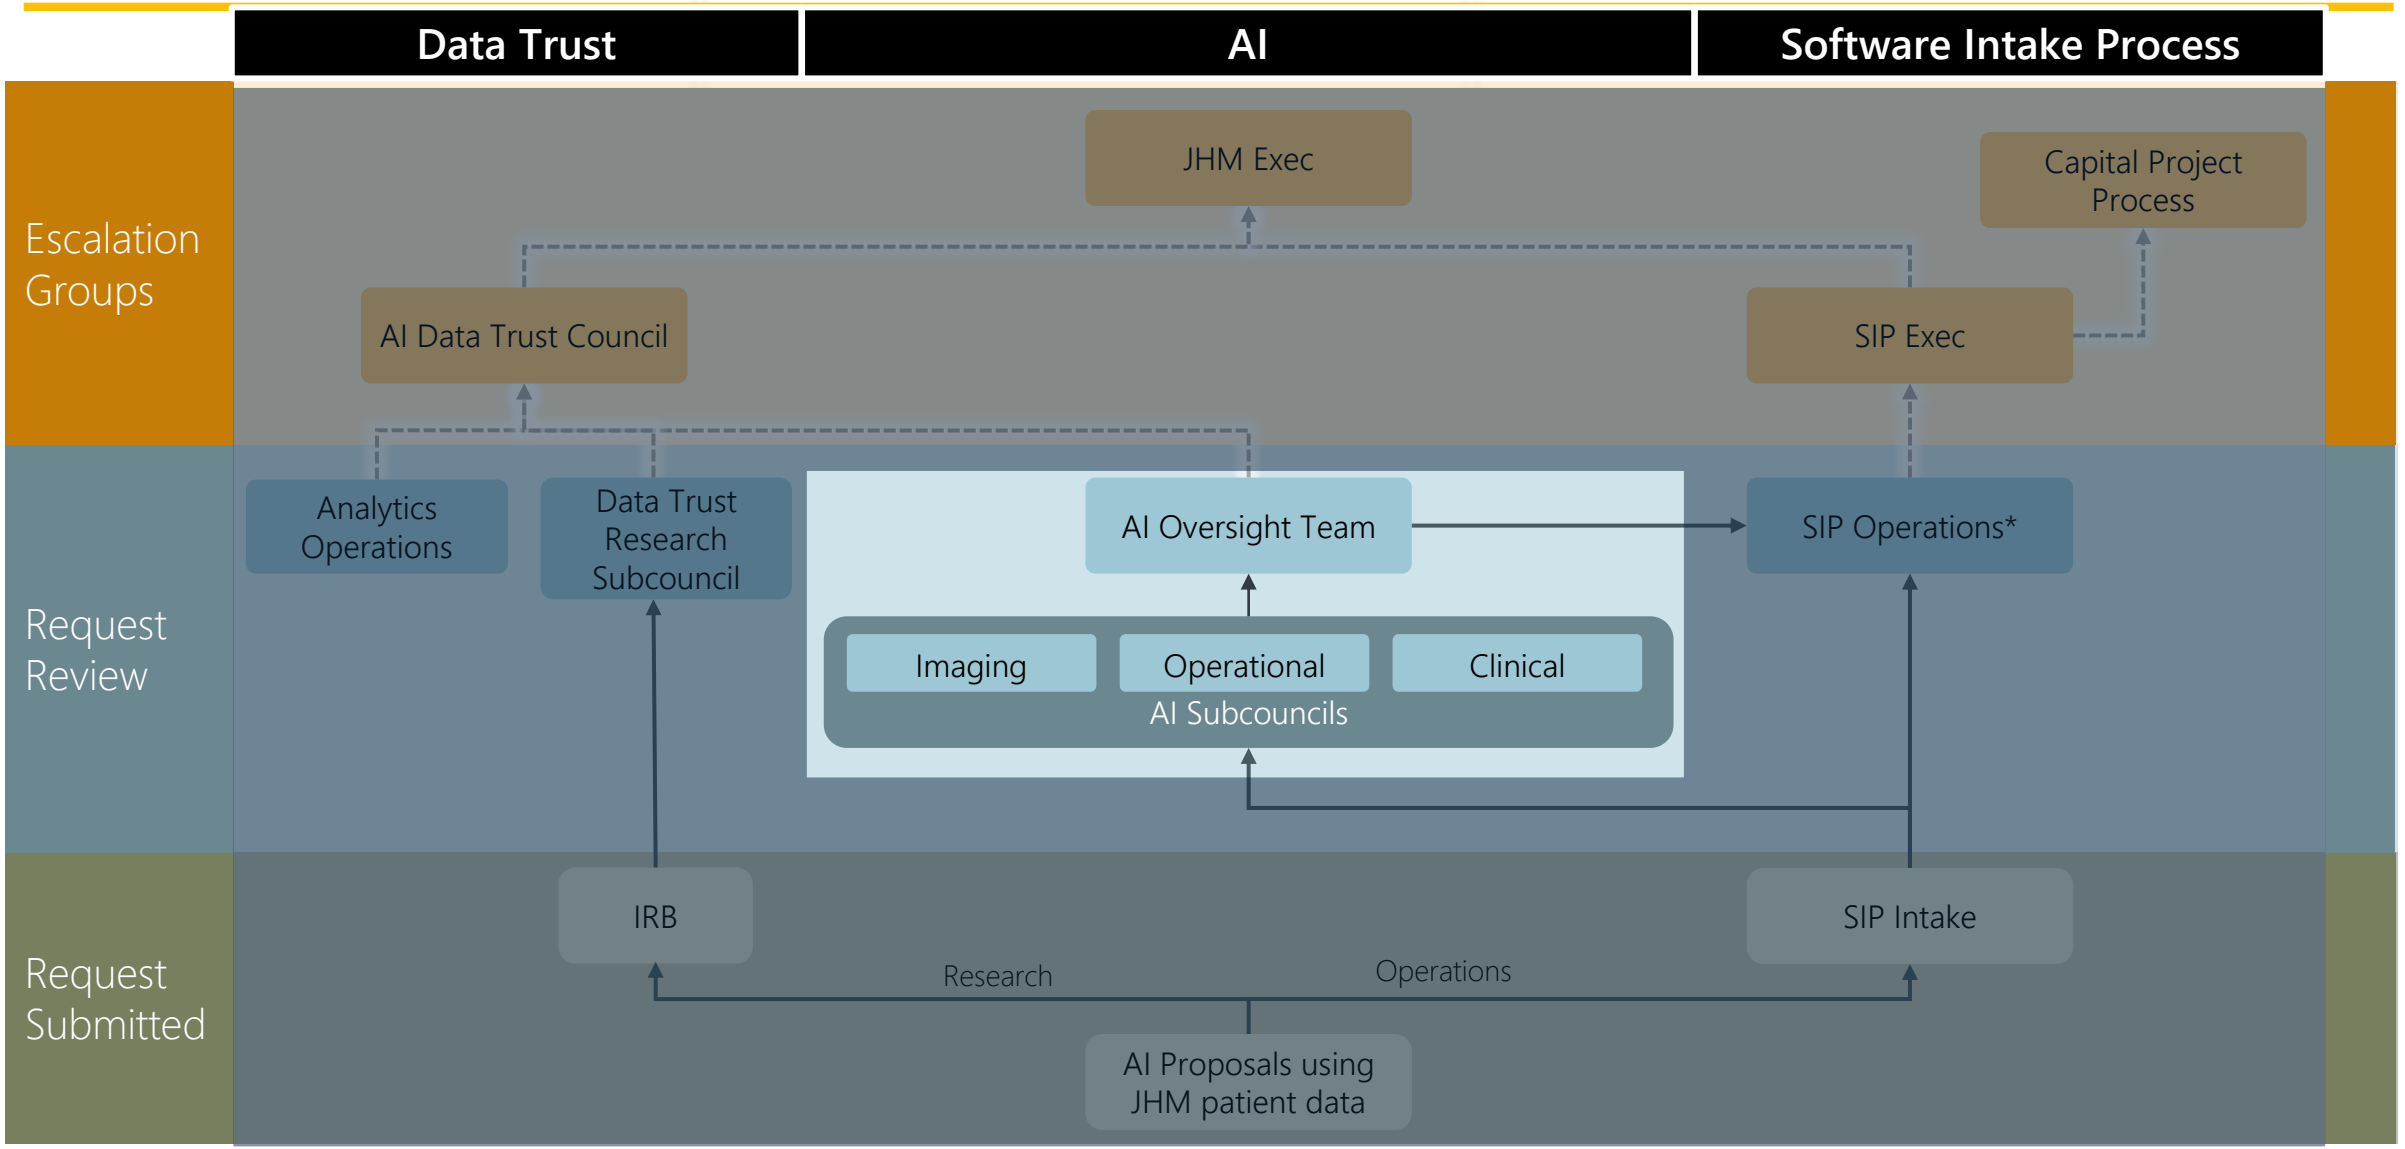
© Johns Hopkins Medicine - Not for Distribution

Integrated Governance Structure



© Johns Hopkins Medicine - Not for Distribution

Integrated Governance Structure



© Johns Hopkins Medicine - Not for Distribution

AI Governance: Post-Deployment

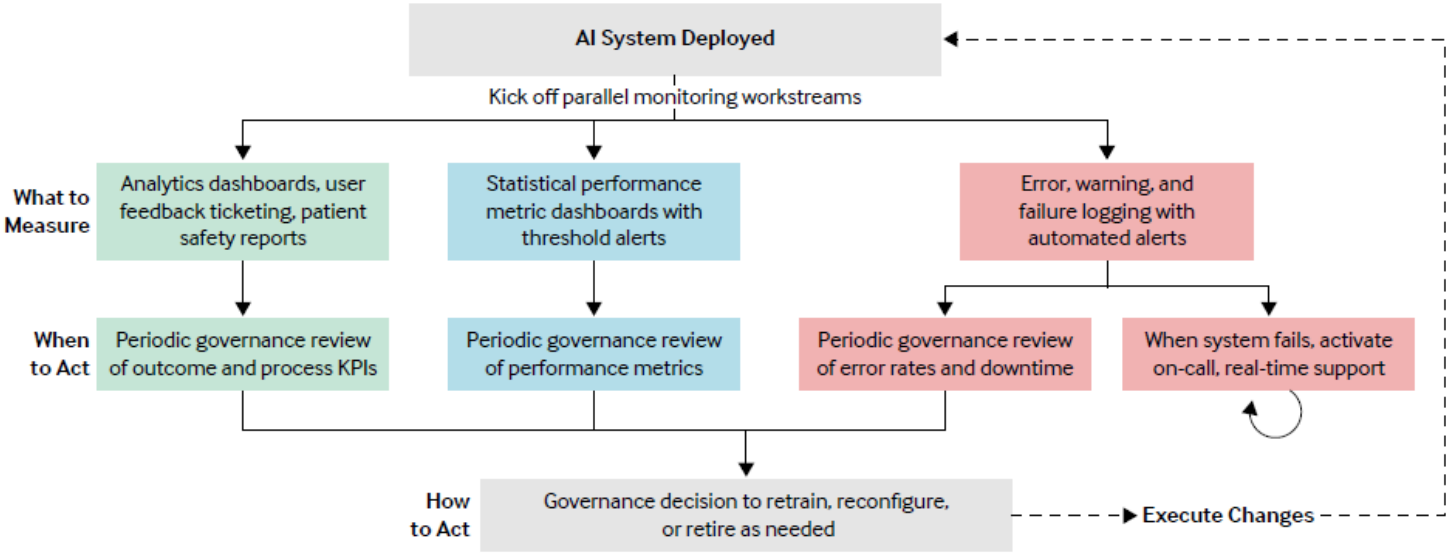
NEJM Catalyst | Innovations in Care Delivery

ARTICLE

Why and How to Monitor Deployed AI Systems in Health Care

Timothy Keyes, Ph.D., Alison Callahan, Ph.D., Abby S. Pandya, M.B.A., M.S., Nerissa Ambers, M.P.H., Juan M. Banda, Ph.D., Miguel Fuentes, M.S., Carlene Lugtu, M.Ci.M., R.N., Pranav Masariya, M.S., Srikar Nallan, M.S., Connor O'Brien, M.S., Thomas Wang, Ph.D., Emily Alsentzer, Ph.D., Jonathan H. Chen, M.D., Ph.D., Dev Dash, M.D., M.P.H., Matthew A. Eisenberg, M.D., Patricia Garcia, M.D., Nikesh Kotecha, Ph.D., Anurang Revri, M.S., Michael A. Pfeffer, M.D., Nigam H. Shah, Ph.D., Sneha S. Jain, M.D., M.B.A.

Vol. 7 No. 6 | June 2026
DOI: 10.1056/CAT.25.0372



Principle	Definition
System Integrity	Check IT infrastructure, data pipelines, and integrations are functioning as intended to avoid extended AI system execution failures or downtime
Performance	Evaluate accuracy, quality, and safety of AI model outputs over time to detect drift or decay
Impact	Verify the continued benefits of the AI-guided workflow for patients, staff, and the organization

© Johns Hopkins Medicine - Not for Distribution

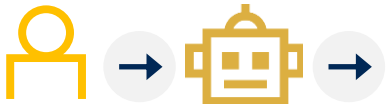


AI USE CASES

JOHNS HOPKINS
MEDICINE

Hospital Executive Perspective: Risk Comfort Level

Where do we draw the line?



Assistive

Human decides; AI supports

Drafting text, summarizing records



Decision-Support

AI recommends; human approves

Clinical risk scores



Semi-Autonomous

AI acts with human checkpoints

Automated scheduling with oversight



Autonomous

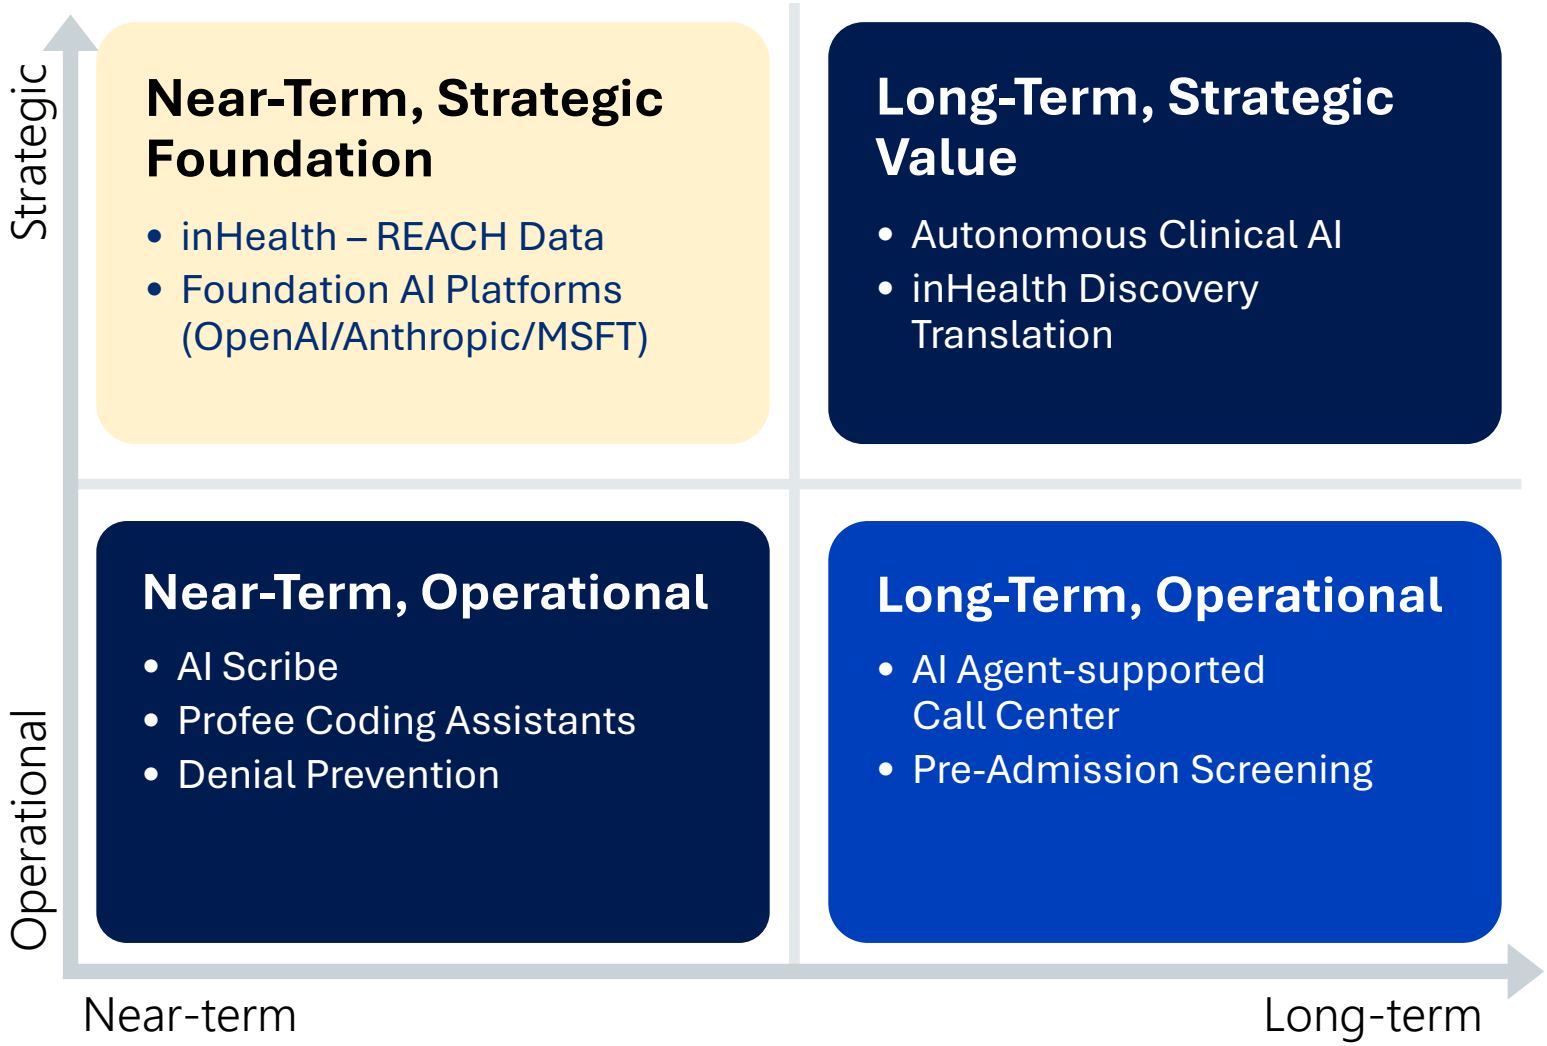
AI decides and acts independently

Autonomous agents, self-optimizing systems

Gartner

© Johns Hopkins Medicine - Not for Distribution

Pursuing a Balanced AI Portfolio

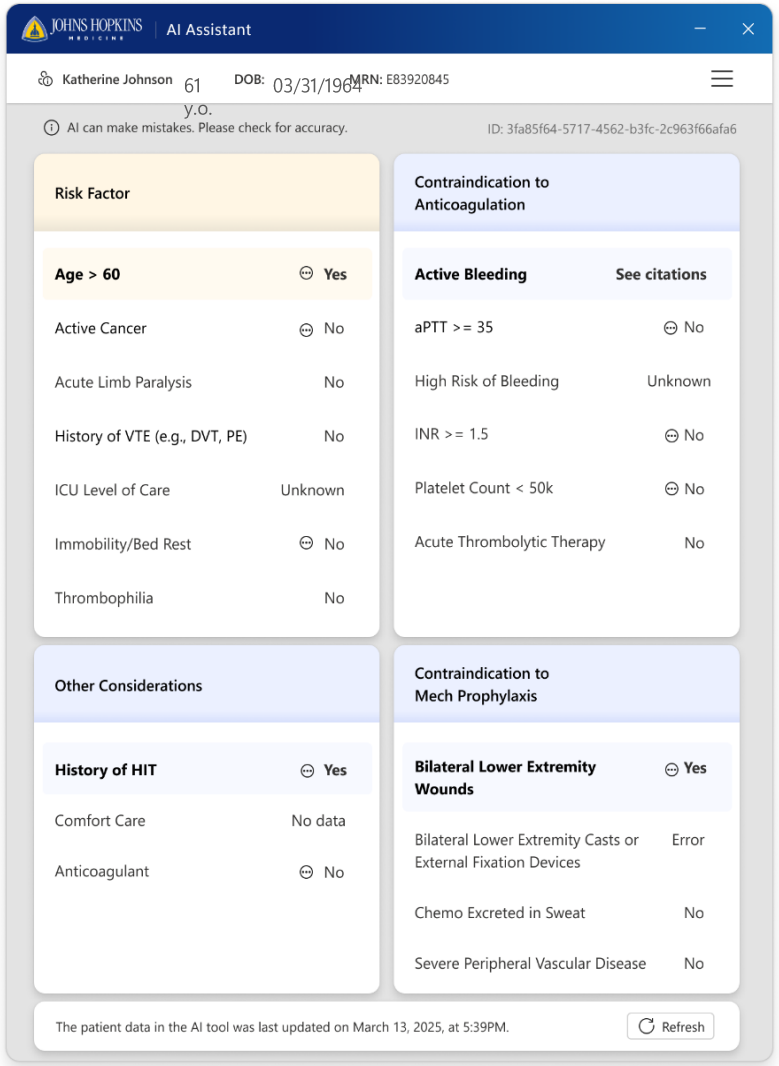


© Johns Hopkins Medicine - Not for Distribution

VTE Workflow Redesign & Piloting AI Assistant

Advisory

JOHNS HOPKINS MEDICINE

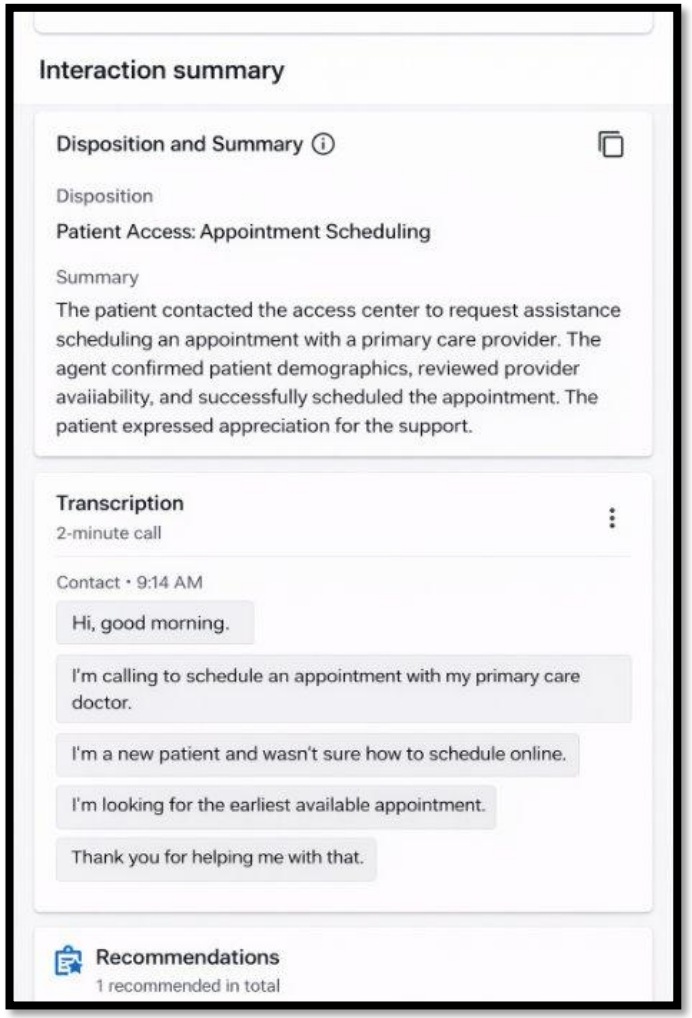


© Johns Hopkins Medicine - Not for Distribution

Contact Center (Talkdesk)

Assistive

- **AI-Enable Patient Interactions**
  Patients can get appt info, confirm, cancel, and/or reschedule by chat or call
- **Epic Integration**
  Agents access Contact Hub within Epic
- **Agent AI- Assist**
  Real-time transcription, guidance, and resources
- **Analytics for Agent Improvement**
  Identifies coaching and compliance gaps from patient interactions
- **Phase 1: March-April Go-lives**
  Patient Access Services, Care at Home, and All Children's



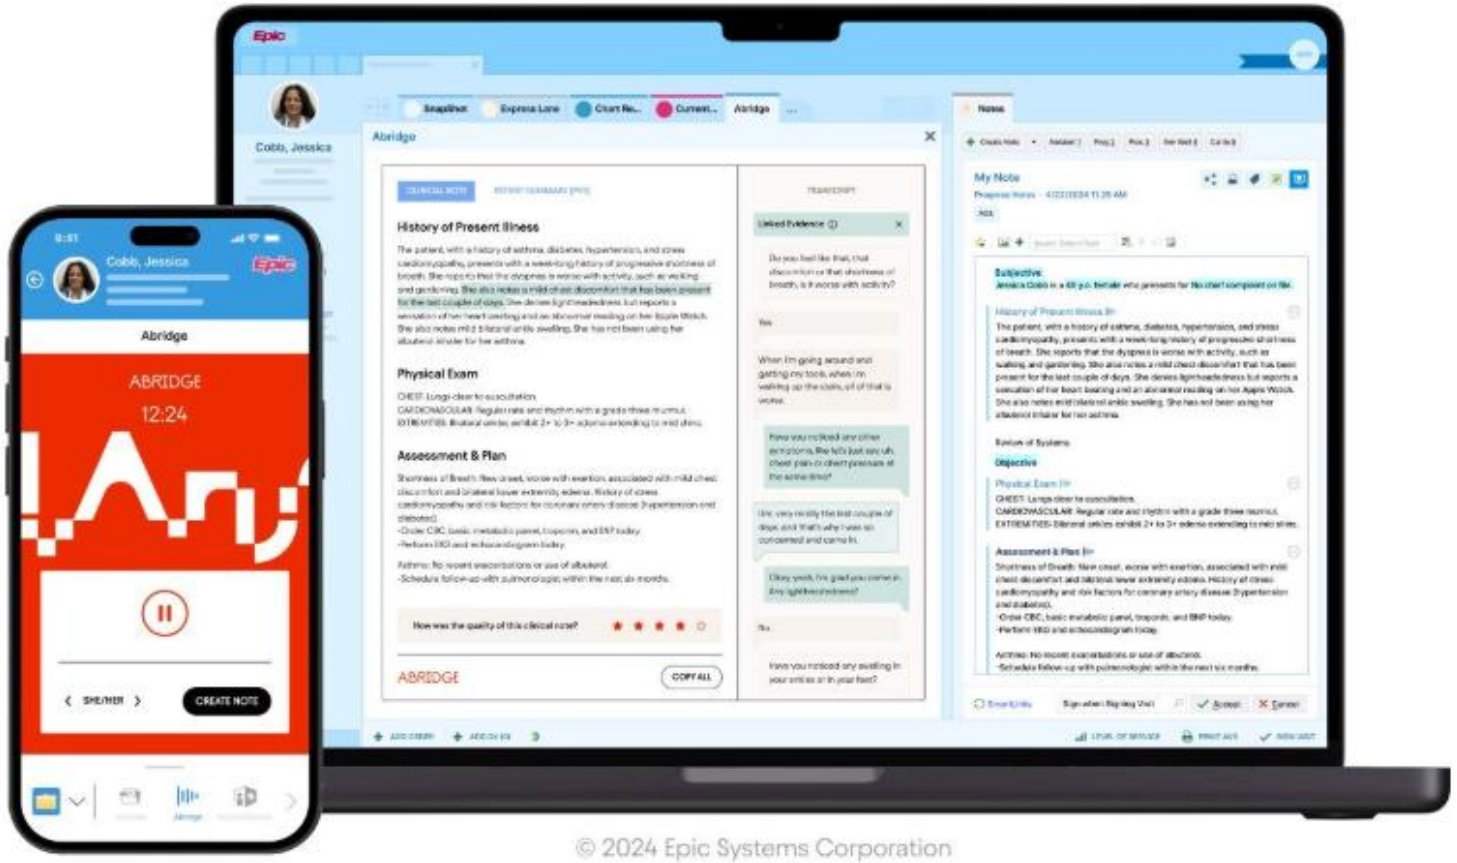


© Johns Hopkins Medicine - Not for Distribution

Abridge Ambient Scribe

Live at JHM

Assistive



© 2024 Epic Systems Corporation

ABRIDGE

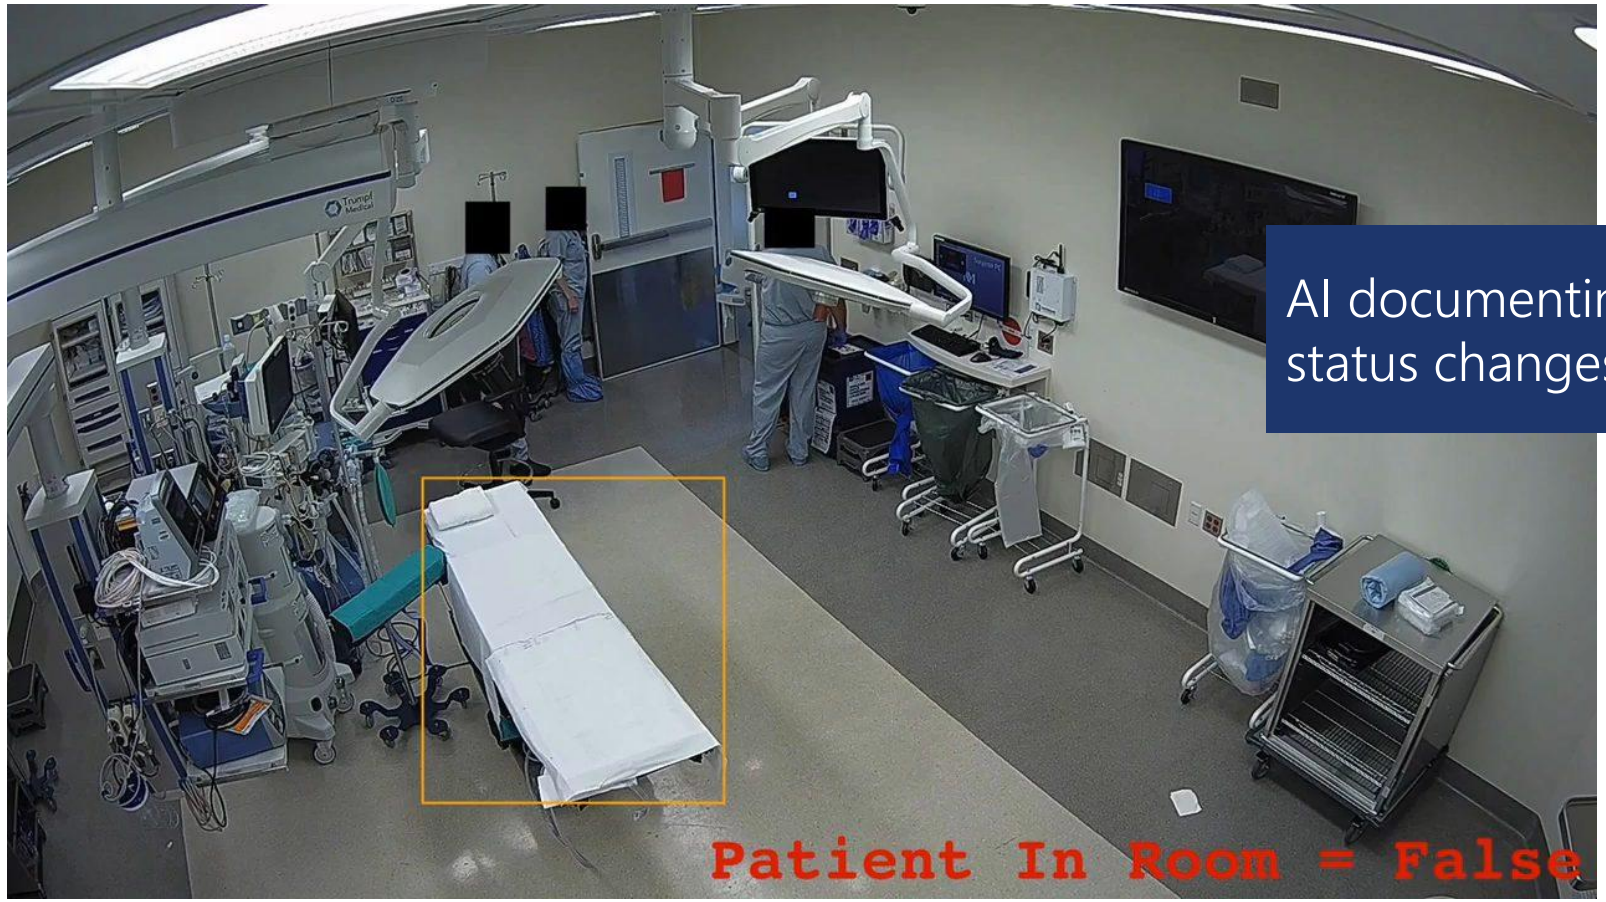
The most trusted AI platform for clinical conversations—
integrated directly inside Epic.

© Johns Hopkins Medicine - Not for Distribution

OR Computer Vision

Piloting at Sibley

Autonomous



AI documenting operational status changes

© Johns Hopkins Medicine - Not for Distribution



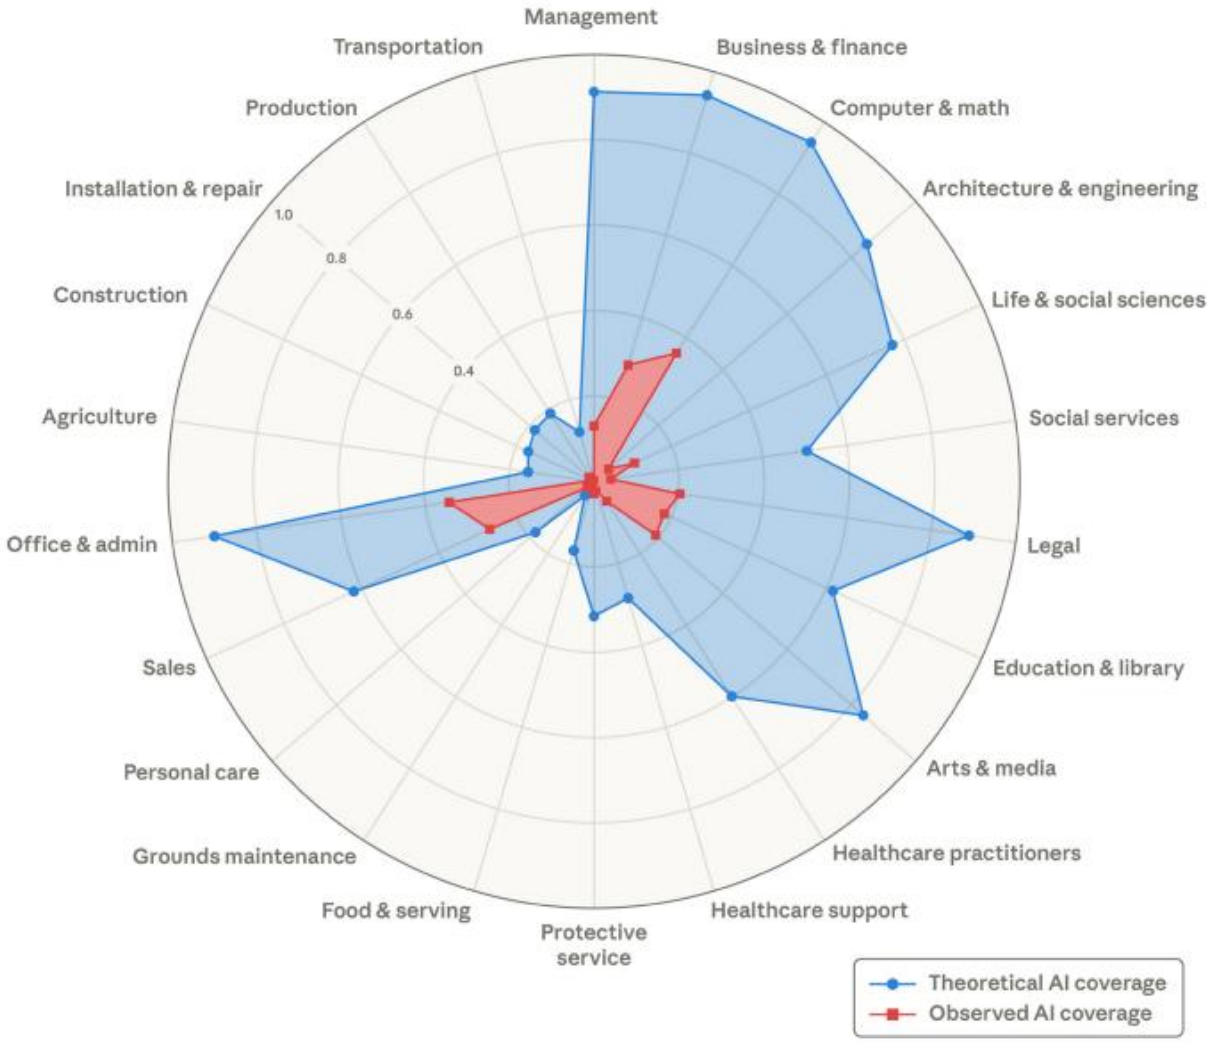
LABOR MARKET

JOHNS HOPKINS
MEDICINE

Labor Market

ANTHROP\C

Labor market impacts of AI: A new measure and early evidence



Source: https://www.anthropic.com/research/labor-market-impacts

© Johns Hopkins Medicine - Not for Distribution

Labor Market (Cont'd)

What will our economic future look like?

We don't know yet how AI will reshape the economy. Will it lead to unprecedented growth? Widespread unemployment? Neither, or something else? How can we tell?

Anthropic's Economics team built a model of how AI might affect jobs, growth, and unemployment in the US in coming years. Read about possible economic futures and make your own predictions about AI capabilities to see the economy they imply.



Source: https://www.anthropic.com/institute/econ-scenarios

© Johns Hopkins Medicine - Not for Distribution

Clinician Labor Market – Recent Publications

Perspective | AI in Medicine

Will Autonomous AI Exceed AI-Aided Physic

Conclusions

The prevailing claim that physician-AI hybrids, such as AI de
support for physicians, will be the highest standard of care
proven. That in the near future AI-alone may provide better p
care than physicians or physician-controlled hybrids at 5 fundamen-
tal cognitive medical tasks is unsettling but seems probable.

Even when AI alone is superior, significant barriers to im
mentation remain. Nonetheless, superior autonomous AI will li
be ready to be deployed for real-world cognitive medical tasks
some, maybe many, workflows by 2030. Consequently, physi
cians, policymakers, and others need to urgently devise ap
proaches to workflow, liability, regulation, reimbursement, and medi-
cal education.

On Straw Men and Doormen

(Gently) Rebutting Khosla and Emanuel on Autonomous AI in Medicine



ROBERT WACHTER
AUG 27, 2026

(Disclosures: I am an advisor to Curai, in which Vinod Khosla is an investor. Zeke Emanuel is a friend and collaborator.)

The First Problem with the Argument: The Reliable Fact-Set and GIGO

"All of this means that AI's diagnostic performance on a curated set of clinical facts should not be taken as compelling evidence that "Dr. AI" can replace actual doctors, who are experts in creating correct and useful problem
r

The Second Problem: The Doorman Fallacy

What were my off-label "doorman" tasks? It was my end-of-life discussions with three children of an elderly patient with terminal cancer, all of whom had different views on what to do with their dad

The Economy Class Straw Man

In my book, I argued that, as AI gets better, primary care in the U.S. is likely to divide into two classes. If you have basic health insurance, you might get AI-first care

Source: Emanuel EJ, Baker-Butler A, Khosla N, Khosla V. Will Autonomous AI Exceed AI-Aided Physicians as the Best Medical Care? JAMA. Published online August 17, 2026. doi:10.1001/jama.2026.15380
Source: https://robertwachter.substack.com/p/on-straw-men-and-doormen

© Johns Hopkins Medicine - Not for Distribution

Clinician Labor Market – Recent Publications (cont'd)

THE PRACTICE OF EMERGENCY MEDICINE/ORIGI

Will Artificial Intelligence Repl Automation Susceptibility of Emerge Tasks

Christian Rose, MD*; Emily Molins, MS; Austin Schoeffler, MD; Rana Kabeer, M
Carl Preiksaitis, MD, MEd

*Corresponding Author. E-mail: ccrose@stanford.edu.

Editor's Capsule Summary

What is already known on this topic

Artificial intelligence (AI) is altering the practice of emergency medicine.

What question this study addressed

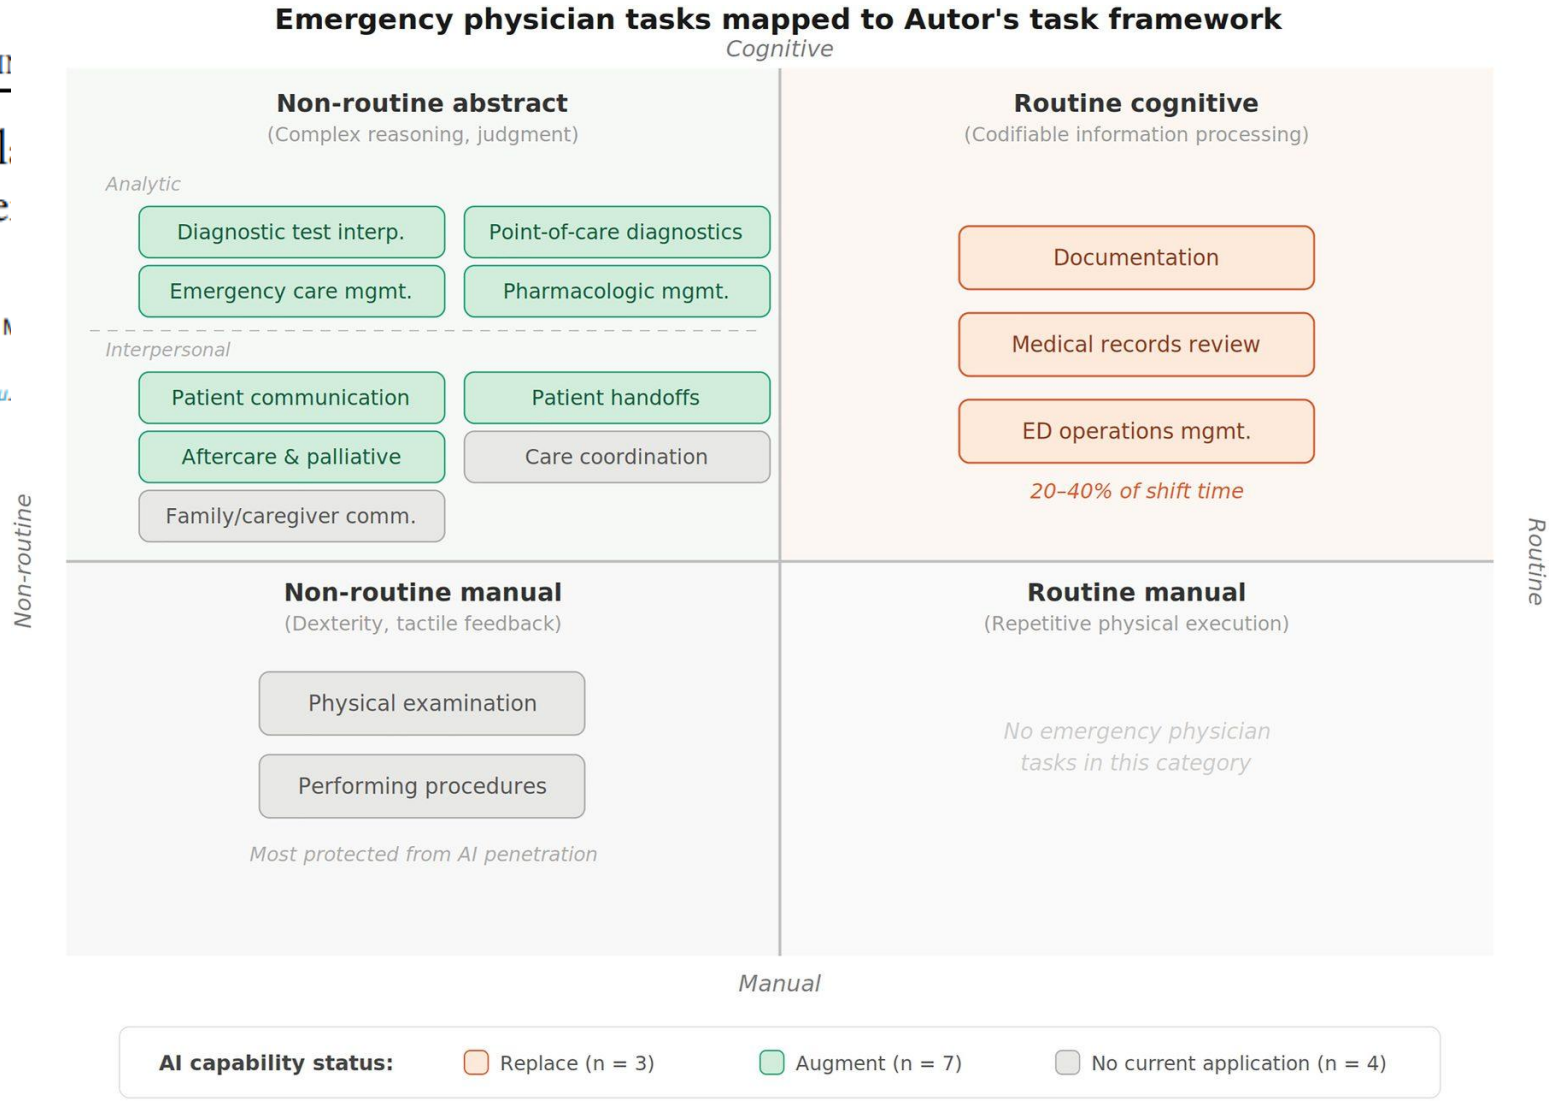
What elements in the work of an emergency physician could be replaced or augmented by AI?

What this study adds to our knowledge

Routine cognitive tasks are ripest for AI automation and represent 20% to 40% of physician shift time. For most work, AI is more likely to augment than replace the physician.

How this is relevant to clinical practice

AI will likely aid specific tasks but not replace the emergency physician in acute care.



Source: Christian Rose, Emily Molins, Austin Schoeffler, Rana Kabeer, Morgan R. Frank, Carl Preiksaitis, Will Artificial Intelligence Replace Me? Automation Susceptibility of Emergency Physician Tasks, Annals of Emergency Medicine, 2026, ISSN 0196-0644, https://doi.org/10.1016/j.annemergmed.2026.06.031.

© Johns Hopkins Medicine - Not for Distribution

Role of the Clinician

BMJ Digital Health & AI

Why are humans still in the loop with advancing AI capabilities?

Jonathan H Chen,[1,2,3] Michael A Pfeffer,[1,2,4,5] Christopher A Longhurst[6]

Table 1 Essential qualities of a clinician

Competence (Knowledge→Judgement)	More than knowledge, a clinician needs judgement to balance competing risks and benefits that do not have textbook (or AI) answers. Whether a treatment decision is worth the harms is a question of values and preferences, not facts and evidence. Values are what humans provide, not computers.
Communication (Empathy→Influence)	More than empathy, a clinician needs influence and negotiation skills to advocate on behalf of their patients. Knowing the right treatment does not matter if no one will accept it. A skilled clinician must be able to convince others to do what they need.
Character (Liability→Responsibility)	More than liability, a trustworthy clinician takes professional responsibility for patient care. Knowledge remains essential to have the language for authentic and accurate communication, the ingredients needed to credibly take ownership of any outcome.

AI and many other technologies may surpass human capabilities in any one attribute at a time, but none of these is sufficient in isolation. The combination of all of the above, integrated into a single entity, is needed to define a trustworthy professional
AI, artificial intelligence.

Source: Chen, Jonathan & Pfeffer, Michael & Longhurst, Christopher. (2026). Why are humans still in the loop with advancing AI capabilities?. BMJ Digital Health & AI. 2. e000057. 10.1136/bmjdh-2026-000057.

© Johns Hopkins Medicine - Not for Distribution