Artificial Intelligence in Primary Care: From Hype to Practical Impact

Steven Lin, MD
Clinical Professor
Chief, General Primary Care
Director, Stanford Healthcare AI Applied Research Team
Division of Primary Care and Population Health
Stanford University School of Medicine
Stanford, Calif.

5 June 2025
11:45 a.m. – 12:45 p.m. ET

Presenter

Steven Lin, MD

Clinical Professor

Chief, General Primary Care

Director, Stanford Healthcare AI Applied Research Team

Division of Primary Care and Population Health

Stanford University School of Medicine

Stanford, Calif.

Steven Lin, MD

Steven Lin, MD is a practicing family physician, Chief of General Primary Care at Stanford University School of Medicine, and the President of the Society of Teachers of Family Medicine. He is the Founding Director of the Stanford Healthcare AI Applied Research Team, a primary care-focused AI implementation center partnered with 50 organizations across industry, academia, non-profit and government to bring leading-edge AI technologies from “code to beside” in support of the Quintuple Aim. He was a James C. Puffer/American Board of Family Medicine Fellow with the National Academy of Medicine and is the author of 500 scholarly works and conference presentations.

Disclosures

- Dr. Steven Lin is the Principal Investigator on sponsored research administered by Stanford University with Google and Omada Health. He is an advisor for Codex Health and Gaia Health. These companies were in no way involved in the conception, development, editing, review, or approval of this presentation.
- The views expressed in this presentation are those of the author and do not necessarily reflect the official policy or position of the Department of Defense, nor the U.S. Government.
- This continuing education activity is managed and accredited by the Defense Health Agency, J-7, Continuing Education Program Office (DHA, J-7, CEPO). DHA, J-7, CEPO and all accrediting organizations do not support or endorse any product or service mentioned in this activity.
- DHA, J-7, CEPO staff, as well as activity planners and reviewers have no relevant financial or non-financial interest to disclose.
- Commercial support was not received for this activity.

Learning Objectives

At the end of this presentation, participants will be able to:

1. Explain the basics of artificial intelligence (AI), machine learning, and generative AI.
2. Identify and describe ways AI is being used in primary care.
3. Summarize key challenges in integrating AI into healthcare.
4. Discuss the ethical landscape of AI in clinical settings.
5. Develop a plan to integrate AI in clinical practice.

Let’s Start

With a Poll

What is Your Relationship with AI?

A) AI and I are going steady, and things are getting serious.

B) I know about AI, but we are still in the “getting to know each other” phase.

C) AI? Pretty sure I’ve seen this movie (Terminator) and know how it ends...

Understanding AI/ML and Generative AI

Traditional AI

Artificial Intelligence (AI) - any technology that mimics human cognitive functions

Machine Learning (ML) - AI that improves with data, recognizing patterns

Deep Learning - a subset of ML using neural networks

Generative AI - AI that generates new text, images, audio/visuals or code based on training data

Artificial Intelligence Taxonomy

WEEKLY AI USE CASES FOR ANALYTICS
AI FOR BI ROCKS

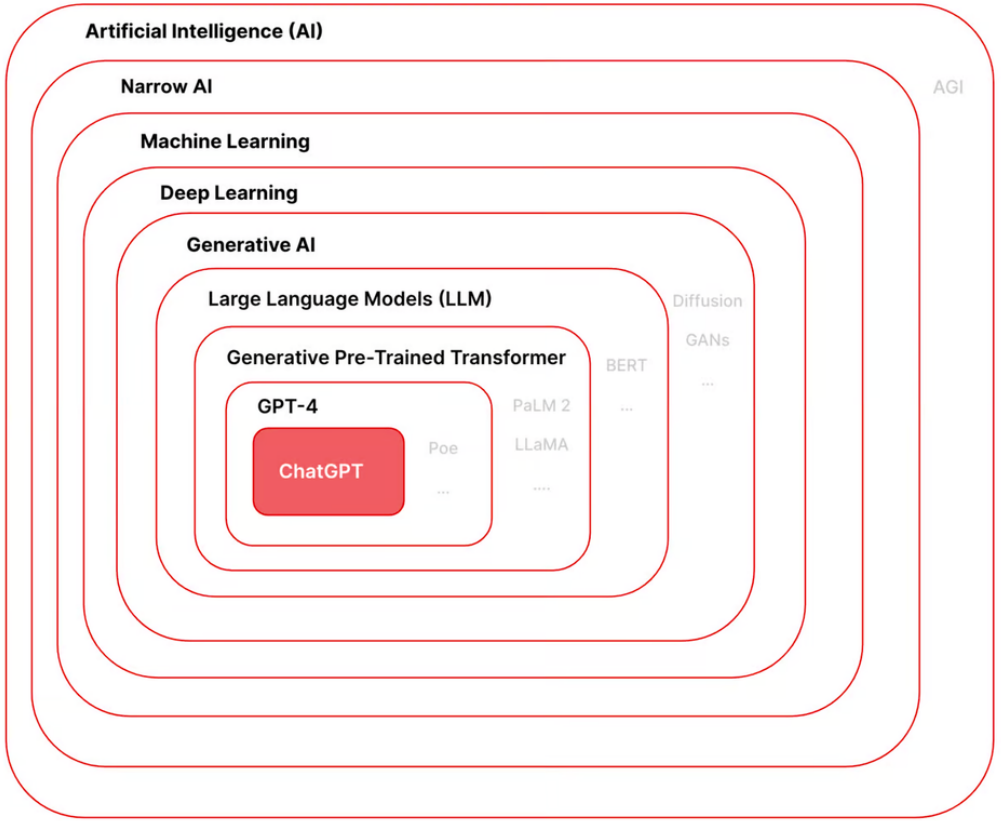

Core Types of Traditional AI Outputs

Prediction - AI estimates future outcomes based on patterns in data

Classification - AI sorts or categorizes data into predefined groups

Examples

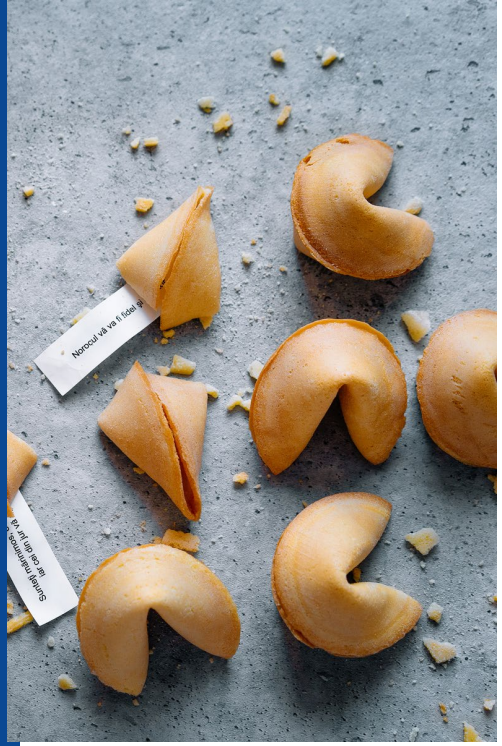

(stock.adobe.com, n.d.)

Prediction Task

Predicting a patient's risk of readmission based on EHR data

Classification Task

Identifying whether a skin lesion is benign or malignant

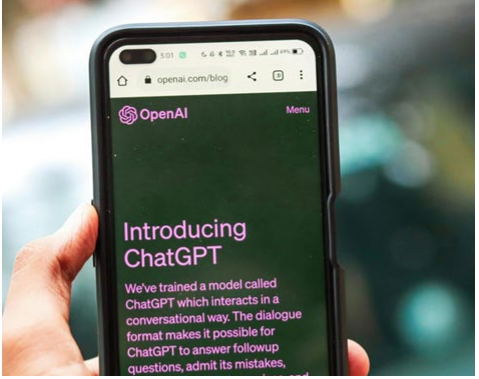

(stock.adobe.com, n.d.)

Text

AI-generated clinical notes or chart summaries

Images

AI-generated or enhanced medical images

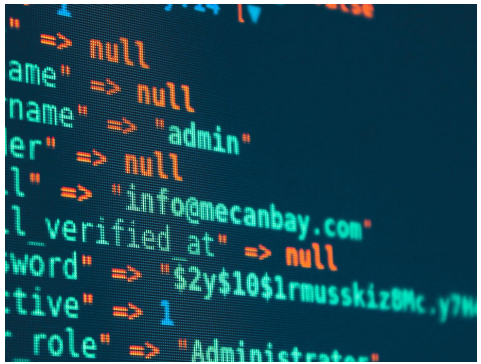

Audio/Visuals

AI-generated education videos and voice synthesis

Code

AI-assisted coding support for electronic health record (EHR) customizations

Core Types of GenAI Outputs

Text Generation - AI creates human-like text

Image Generation - AI produces or enhances medical images

Audio & Video Generation - AI synthesizes spoken language or video content

Code Generation - AI writes or completes programming scripts

Examples

Model Training

Step 1: Data collection and processing

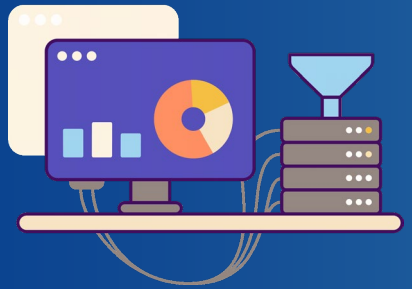

Step 2: Feature extraction pattern recognition

Step 3: Fine-tuning and feedback

https://wonderchat.io/blog/how-long-did-it-take-to-train-chatgpt

https://www.sparxitsolutions.com/blog/everything-you-need-to-know-about-generative-ai-development/

Key Differences Training Generative AI vs Traditional AI

Difference Category	Traditional AI	Generative AI
 Objective	Classify, predict, or detect patterns in existing data	Generate new content based on learned patterns
 Data type/Scale	Typically uses structured, labeled data	Massive-scale, largely unstructured data without explicit labels
 Learning Approach	Typically supervised learning although it *can* be unsupervised	Almost exclusively unsupervised learning approaches
 Compute	Can be trained on smaller datasets with moderate compute power requirements	Massive datasets and significant compute resources due to large-scale data architectures
 Fine-tuning/Updates	Typically requires full retraining to improve performance	Fine-tuned with reinforcement learning or domain-specific training without a full retrain

https://www.sparxitsolutions.com/blog/everything-you-need-to-know-about-generative-ai-development/

Models, models, and more models

Grok 3

xAI

GPT-4o

OpenAI

Llama 3

Meta

Mistral 7B

Mistral

DeepSeek R1

High-Flyer

Claude 3.7

Anthropic

Gemini 2.0

Google

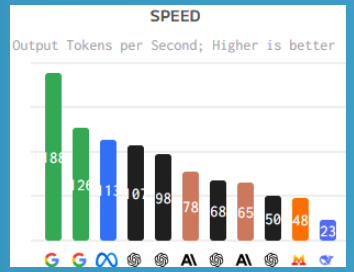

GenAI Leaderboards

Chatbot Area (lmarena.ai)

Artificialanalysis.ai

Tradeoffs

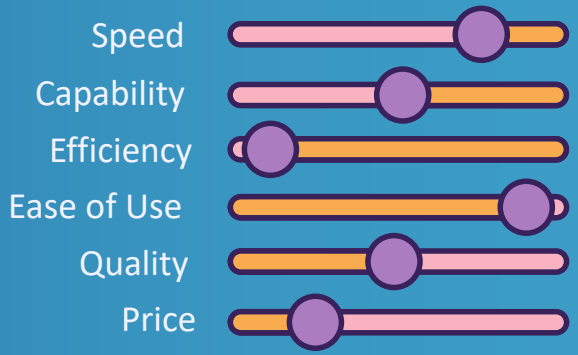

A Case Study

5 GenAI Primary Care Use Cases

Outline

1 AI Chart Review and Summarization

2 AI Assisted Clinical Decision Making

3 AI Translated Personalized Instructions

4 Ambient AI Scribing

5 AI Chatbot for Triage and Medical Advice

Case Study

Use Case 1

The set up

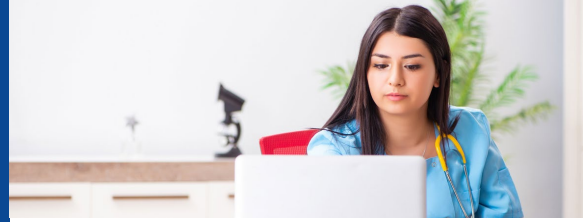

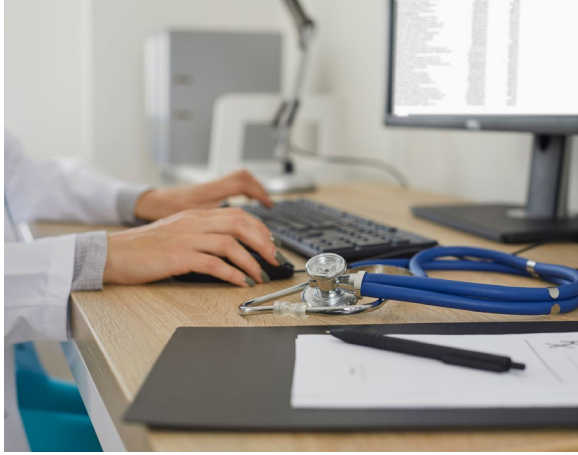

(stock.adobe.com, n.d.)

- You are a **primary care physician** running behind in clinic
- Your next patient is here for a **follow-up visit after a hospitalization**
- You recognize his name (Juan), and vaguely recall that he has diabetes and hypertension, but **you haven't seen him in 3 years**

Case Study

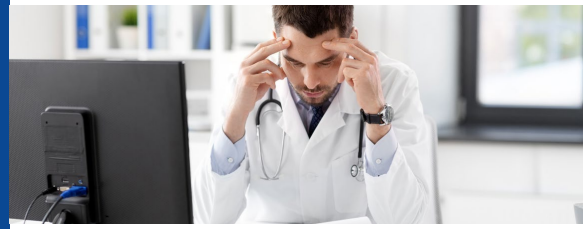

(stock.adobe.com, n.d.)

Use Case 1

The problem

Data Overload and Fragmentation

- **Overabundance and scatter of information** within EHRs is a barrier to patient care and key source of physician burnout
- For every hour physicians spend in front of patients, another hour is spent in front of the EHR; **chart review accounts for the second highest proportion of EHR time** (second to documentation)

(Holmgren et al., 2024)

Case Study

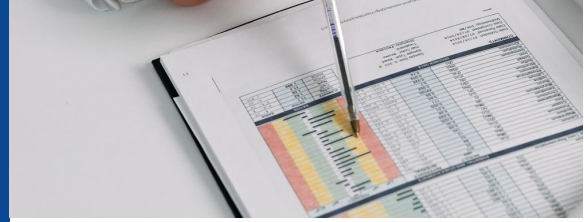

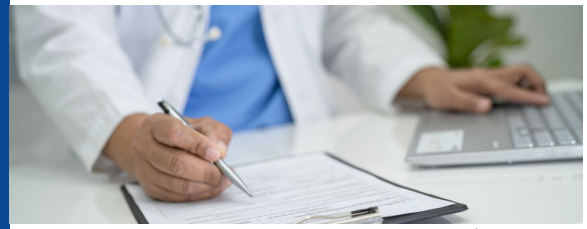

(stock.adobe.com, n.d.)

Use Case 1

The opportunity

AI Chart Review and Summarization

- AI-based systems that **extract, analyze, and summarize** patient clinical information from EHRs can reduce physicians' cognitive burden from information chaos
- Study: **AI helped physicians finish chart review 18% faster** while maintaining high accuracy

(Chi et al., 2021)

Case Study

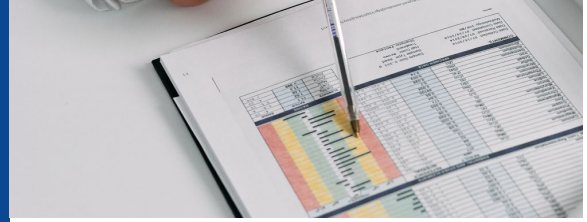

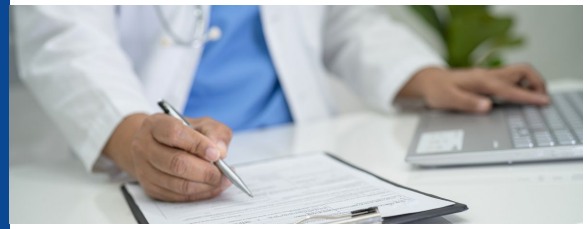

(stock.adobe.com, n.d.)

Use Case 1

The result

- The AI tells you that Juan was hospitalized one week ago with **a new diagnosis of congestive heart failure**; it summarizes the key results of all pertinent labs, imaging, and procedures done
- He was discharged home with carvedilol, lisinopril, atorvastatin and metformin

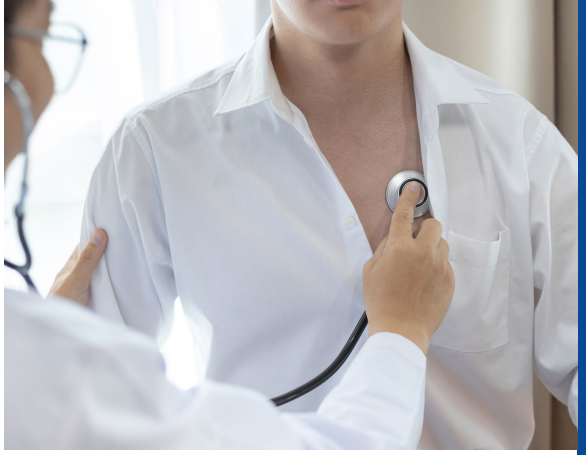

(stock.adobe.com, n.d.)

Case Study

Use Case 2

The set up

- You enter the exam room to see Juan, who appears well; his vital signs, history and physical exam show no signs of fluid overload
- You noted from the AI summary that his **hemoglobin A1C = 8.5%**
- You wonder **which medication is the best to add to his regimen?**

(stock.adobe.com, n.d.)

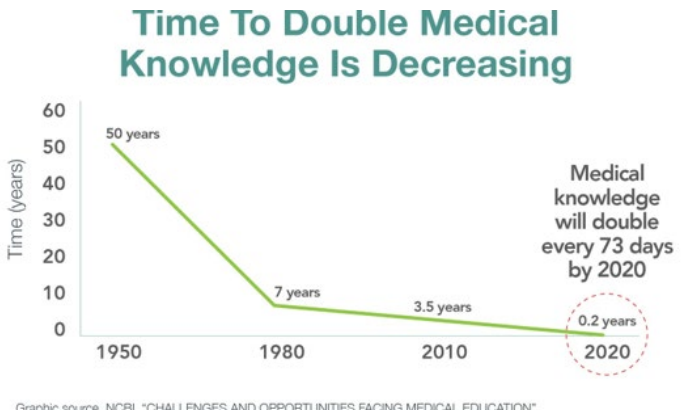

Case Study

Use Case 2

The problem

Rapidly Changing Evidence and Guidelines

- Medical knowledge **doubles every 2.5 months** (in 1950, it was 50 years)
- It takes an average of **17 years for evidence to change practice**
- **Example**: <25% of heart failure patients are getting guideline-directed medical therapy

(Densen, 2011)
(Balakumaran et al., 2019)

(vectorstock.com, n.d.)

Case Study

Use Case 2

The opportunity

AI Assisted Clinical Decision Making

- AI-based systems trained on **up-to-date scientific literature** can provide **on-demand consultation** to answer physician's questions
- AI can also generate **real-world evidence from medical records** to answer questions that can't be answered by studies

(Callahan, 2021)

(vectorstock.com, n.d.)

Case Study

- The AI tells you that, based on recent trials and both the 2022 American Heart Association (AHA)/American College of Cardiology (ACC)/Heart Failure Society of America (HFSA) and 2023 American Diabetes Association (ADA) guidelines, a **Sodium-Glucose Cotransporter 2 (SGLT2) inhibitor is strongly recommended** for heart failure with or without diabetes

- You start Juan on dapagliflozin, and ask him to come back and see you in 3 months

Use Case 2

The result Ai

Case Study

Use Case 3

The set up

(stock.adobe.com, n.d.)

- You're ready to send Juan home, but **he has numerous questions** about his medical conditions, his medications, and what foods he should eat or avoid, etc.

- Your **educational materials don't match** someone with his exact set of conditions/medications, and they **don't come in Spanish** (his first language)

Case Study

(stock.adobe.com, n.d.)

Use Case 3

The problem

Poor Patient Instructions and Education

- Patients only remember **49% of the decisions** made during talks with their doctors
- Written instructions improve recall, but are time-consuming to generate and are often not aligned with patients' **language, literacy level, and unique set of medical conditions/medications**

(Laws et al., 2018)

Case Study

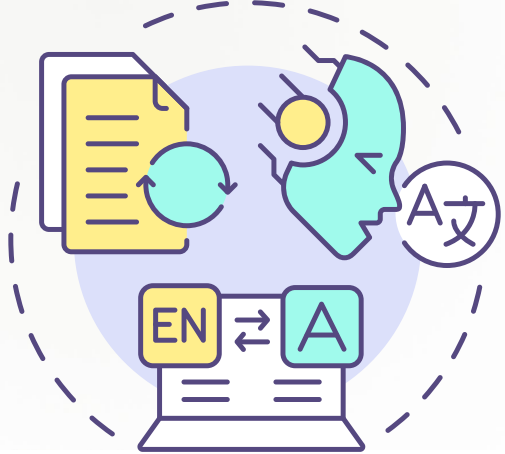

(stock.adobe.com, n.d.)

Use Case 3

The opportunity

AI Translated Personalized Instructions

- AI-based systems that **retrieve** historical instructions, **reason** related medical knowledge, and **refine** language to fit patients' needs can **create personalized instructions** for doctors to edit

- Study: **AI-generated instructions rated more accurate and helpful than humans**

(Liu et al., 2022)

Case Study

(stock.adobe.com, n.d.)

Use Case 3

The result

- The AI generates **personalized after-visit instructions** for Juan: clear, Spanish-language, patient-friendly education for his unique set of conditions, medications, a culturally sensitive diet plan, and a heart failure action plan

- Juan goes home feeling more **confident and empowered**

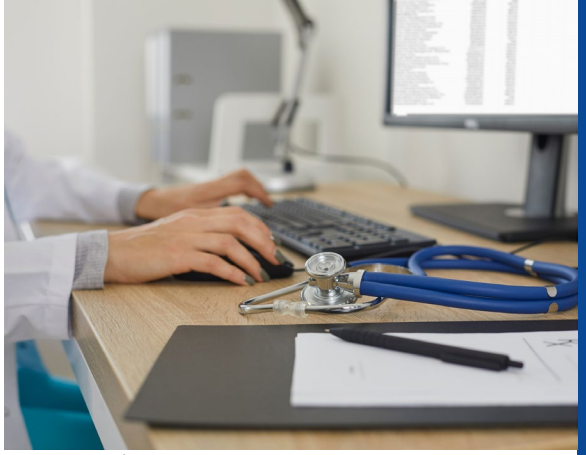

(stock.adobe.com, n.d.)

Case Study

Use Case 4

The set up

- You are finally finished with your clinic for the day
- **You realize you have not written any of your notes**
- It’s your daughter’s birthday and you had promised her that you would be home for dinner

(stock.adobe.com, n.d.)

Case Study

Use Case 4

The problem

Burden of Clinical Documentation

- **Burden of documentation** within EHRs is a key source of physician burnout and lost productivity
- For every hour physicians spend in front of patients, another hour is spent in front of the EHR; **clinical documentation accounts for the highest proportion of EHR time at 36%**

(Holmgren et al., 2024)

(stock.adobe.com, n.d.)

Case Study

Use Case 4

The opportunity

Ambient AI Scribes

- AI systems that **capture** audio data from clinical conversations, **convert** audio to text, **extract** and **summarize** relevant data can **automatically generate notes** for doctors to review and edit

- Study: Physicians using these tools report **improved efficiency** and significant **reductions in burnout**

(Shah et al., 2025)

(stock.adobe.com, n.d.)

Case Study

Use Case 4

The result

- Using a combination of chart summarization and real-time scribing, the **AI drafts all your progress notes**
- You review the drafts, edit them, and sign your notes in minutes
- **You get home in time for dinner**, much to your daughter's delight

Case Study

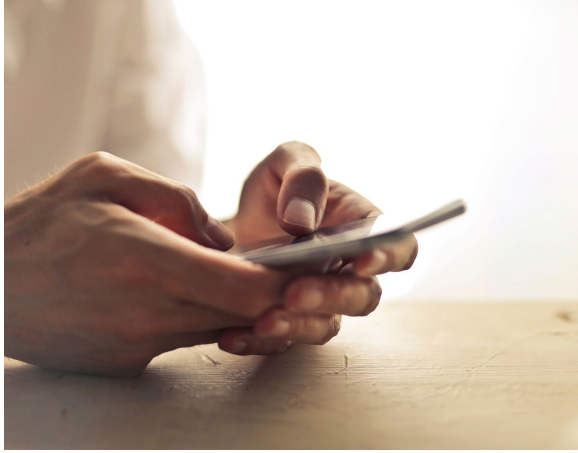

(stock.adobe.com, n.d.)

Use Case 5

The set up

- One month later, **Juan feels unwell** after a weekend of partying and eating BBQ
- He's not sure if he needs to see you, so he messages you using the portal on a Thursday night: **"Hey Doc, it's hard to catch my breath and my feet are swollen. Should I be worried?"**

Case Study

(stock.adobe.com, n.d.)

Use Case 5

The problem

Rising Asynchronous & Continuous Care

- The COVID pandemic hastened the adoption of virtual care, resulting in a **1.6-fold increase in electronic patient messages**
- Current approaches like limiting number of messages patients can send, billing for responses, or delegating responses to less trained staff **limit access to care**

Case Study

Use Case 5

The opportunity

AI Chatbot for Triage and Medical Advice

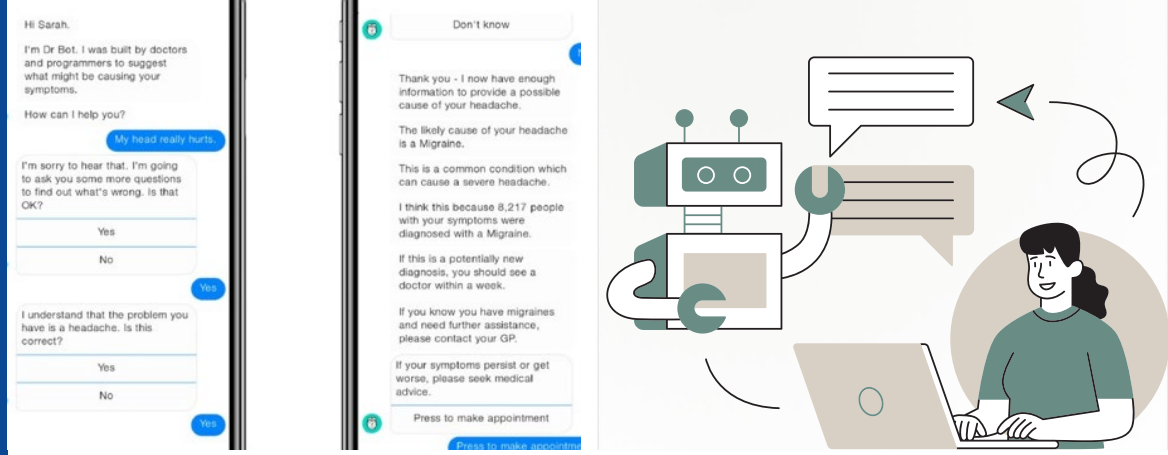

(shutterstock.com, n.d.)

- Large language models can respond to patient messages, **triage** or **provide medical advice** with human-quality text
- Study: ChatGPT responses to medical questions rated **higher quality and more empathetic** than humans
- Study: **Adoption** of AI-generated draft replies reached **20% within 5 weeks,** and providers report significant **reductions in burnout**

(Perlis & Collins, 2025)
(Garcia et al., 2024)

Case Study

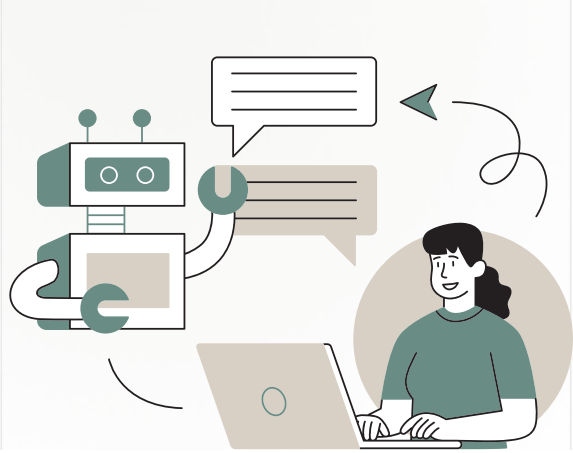

(shutterstock.com, n.d.)

Use Case 5

The result

- The AI chatbot finds out from Juan that he has **gained 5 lbs in the past week and has trouble laying flat to sleep**; it connects Juan to the on-call nurse, who **schedules him a next-day visit**
- You see Juan the next day, treat him with diuretics, and **save him from another hospitalization**

In Summary

5 GenAI Primary Care Use Cases

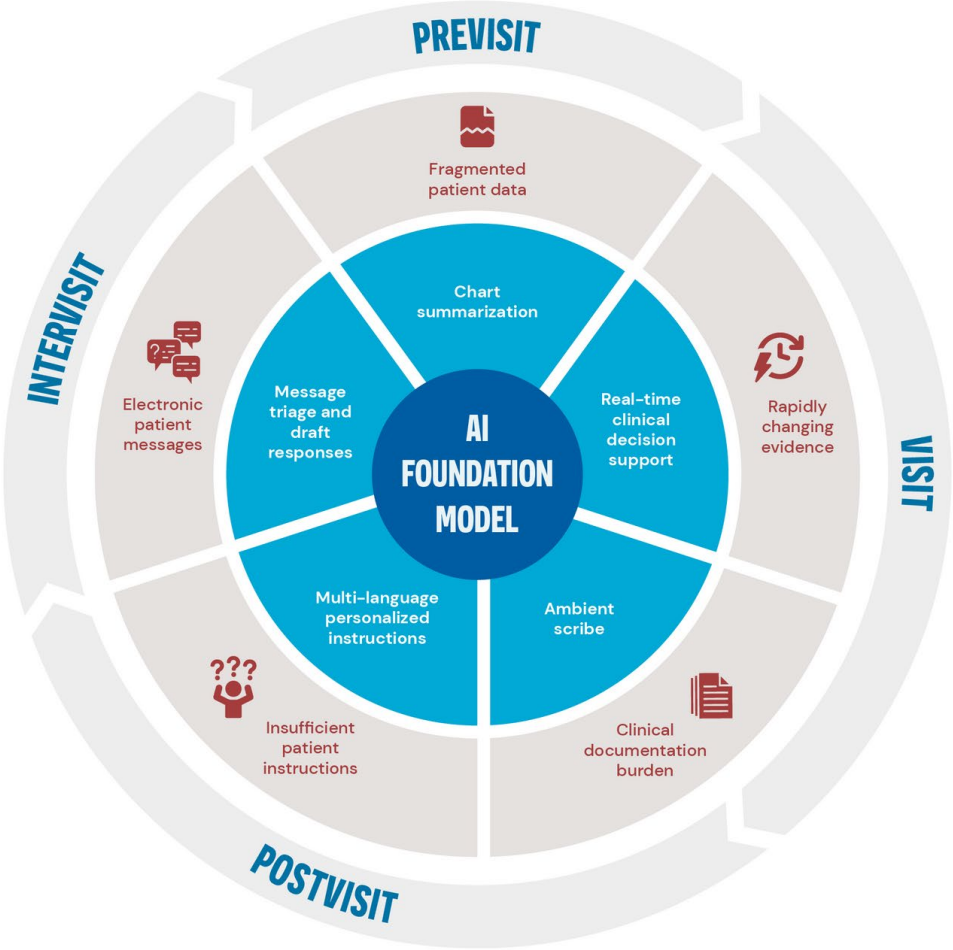

Challenges of the Moment

Five Critical Issues

Evaluating Generative AI Performance

AI Performance Drop from Training to Real-World Use

The Hidden Cost of Maintenance, Data Shifts & Re-Training

Speed of AI Development Moving Faster than Medical Education

AI Development and Its Impact on the Environment

What is Needed

Five Critical Solutions

Clear, Clinically Relevant Benchmarks for Generative AI

Increased Emphasis on Real-World Evaluation

Sustainable AI Governance, Maintenance and Adaptation Plans

Scalable, Personalized, and Evergreen AI Education to Keep Pace

Frameworks for Evaluating Environmental Impact of AI

Ethical Considerations for

People-Centered AI Implementation

Phase	Key Questions	Taking Action
Development & Design	Are we solving the right problem—and who gets a say?	Many tools are developed without diverse teams. *Ask: Were frontline clinicians and diverse patient groups part of the process?*
Training, Testing & FIne Tuning	What data was this trained on—and what was left out?	AI can encode hidden biases from historical data. *Ask: Does this tool work equally well for different patient populations?*
Implementation and Integration	Who has access—and is AI creating new gaps in care?	AI can widen care gaps if it's not accessible to all—language barriers, digital literacy, etc. *Ask: Does this tool improve access for all, or does it shift resources away from those who need them most?"*

Getting Started with AI in Your Practice

Practical Tips

Instead of asking “what AI should I use?”, start by identifying and analyzing a problem and ask “Is AI needed to solve this problem?”

Take a free prompt engineering course to learn how small tweaks in wording can improve AI outputs (e.g. Coursera, Learnprompting.org).

Look for clinical validation studies, real-world case examples, comparison studies, and peer feedback before adopting a tool.

Research and pick a framework for evaluating tools pre-implementation (e.g. Stanford’s FURM, Duke’s AI Playbook, and many others).

Verify that vendors have ISO, HITRUST, or SOC 2 certifications and verify HIPAA compliance.

Begin with low-risk tasks like summarizing research or writing emails using tools like Open Evidence or SciSpace.

Final Thought

AI in medicine isn’t coming—it’s already here. The question isn’t whether to use it, but how to ensure we do so responsibly, keeping humans and outcomes at the center.

Key Takeaways

- Traditional AI analyzes data to predict or classify; generative AI uses learned patterns to create new content.
- AI is already being used along every step of the patient-provider care journey.
- Challenges exist in the real-world implementation of AI, as does serious ethical considerations.
- You can take meaningful action to get started with AI in practice.

References (1 of 3)

Artificial Intelligence Taxonomy. (n.d.). *GPTech Blog*. Retrieved November 12, 2024, from https://www.gptechblog.com/5-diagrams-to-help-you-understand-generative-ai/

Balakumaran, K., et al. (2019). Evaluation of a Guideline Directed Medical Therapy Titration Program in Patients With Heart Failure With Reduced Ejection Fraction. *International Journal of Cardiology*, 292, 77–83. https://doi.org/10.1016/j.ijcard.2019.05.066

Callahan, A., Gombar, S., Cahan, E. M., Jung, K., Steinberg, E., Shah, N. H., & Patel, B. N. (2021). Using aggregate patient data at the bedside via an on-demand consultation service. *NEJM Catalyst Innovations in Care Delivery*, 2(4). https://doi.org/10.1056/CAT.21.0224

Chi, E. A., Chi, G., et al. (2021). Development and validation of an artificial intelligence system to optimize clinician review of patient records. *JAMA Network Open*, 4(7), e2117391. https://doi.org/10.1001/jamanetworkopen.2021.17391

References (2 of 3)

Densen, P. (2011). Challenges and Opportunities Facing Medical Education. *Transactions of the American Clinical and Climatological Association*, 122, 48–58. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3116346/

Everything You Need to Know About Generative AI Development. (n.d.). *Sparx IT Solutions Blog*. Retrieved November 12, 2024, from https://www.sparxitsolutions.com/blog/everything-you-need-to-know-about-generative-ai-development/

Garcia, P., Ma, S. P., Shah, S., et al. (2024). Artificial intelligence–generated draft replies to patient inbox messages. *JAMA Network Open*, 7(3), e243201. https://doi.org/10.1001/jamanetworkopen.2024.3201

Holmgren, A. J., & Sinsky, C. A. (2024). National comparison of ambulatory physician electronic health record use across specialties. *J Gen Intern Med*, 39(4), 738–740. https://doi.org/10.1007/s11606-024-08930-4

How Long Did It Take To Train ChatGPT? (n.d.). *Wonderchat Blog*. Retrieved November 12, 2024, from https://wonderchat.io/blog/how-long-did-it-take-to-train-chatgpt

References (3 of 3)

Kniberg, H. (2023, October 17). *Generative AI in a Nutshell - How to Survive and Thrive in the Age of AI* [Video]. YouTube. https://www.youtube.com/watch?v=2IK3DFHRFfw

Laws, M. B., Barton, A. C., et al. (2018). Factors Associated With Patient Recall of Key Information in Ambulatory Specialty Care Visits: Findings From the PROMISES Project. *PLOS ONE*, 13(2), e0191940. https://doi.org/10.1371/journal.pone.0191940

Liu, F., Yang, B., et al. (2022). Generating accurate and faithful patient instructions. *Advances in Neural Information Processing Systems*, 35. https://proceedings.neurips.cc/paper_files/paper/2022/file/77c08a6e68ae25433f1d117283c0e312-Paper-Conference.pdf

Perlis, R., & Collins, N. (2025). Study finds people prefer AI over clinician responses to questions in the electronic medical record. *JAMA*, 333(9), 738–740. https://doi.org/10.1001/jama.2024.24065

Shah, S. J., Devon-Sand, A., Smith, M., et al. (2025). Ambient artificial intelligence scribes: Physician burnout and perspectives on usability and documentation burden. *Journal of the American Medical Informatics Association*. https://doi.org/10.1093/jamia/ocad084

Questions?

How to Obtain CE/CME Credits

2025 JUNE CCSS: Evidence-Based Approaches for Advancing Excellence in Primary Care

Credits are awarded by session. To claim CE/CME credit or certificate of attendance for the session(s) you attend, you must register by 4:00 p.m. ET on June 6, and then you must complete the course evaluation and posttest for each session by 11:59 p.m. ET on Thursday, June 19, 2025.

1. Visit the main event page at https://www.dhaj7-cepo.com/content/2025-jun-ccss to register for the live event or to log in to your account if already registered.
2. On the main event page, select the "Get Started" tab (located in the menu below the event title on the desktop and at the bottom of the page on mobile devices). Note: This tab will not appear unless you are registered and logged in to your account.
3. Under the "Get Started" tab, scroll down to a session you attended and select "Claim credit."
4. Proceed to take the evaluation and posttest to obtain your certificate after the session has ended.

All completed courses and certificates are available in your account. Refer to your Pending Activities for sessions you have yet to complete. You must complete the required course items by Thursday, June 19 to receive credit.

Questions? Email DHA, J-7, CEPO at dha.ncr.j7.mbx.cepo-cms-support@health.mil.