



AI in Healthcare: Navigating Risks and Safeguarding Patients

Josiah Dykstra, PhD

This course may not be copied, reproduced, published, displayed, broadcast, reduced to any electronic medium or machine-readable form, or otherwise used or distributed in any form without the prior written consent of continued.com LLC. © 2025 Continued®

Josiah Dykstra, PhD



Josiah Dykstra, Ph.D. is the owner of Designer Security, a consulting business devoted to the cybersecurity needs in healthcare. For the past 20 years, he has worked as a cybersecurity practitioner, researcher, and technical leader in government and industry. He is a frequent speaker and co-author of the book *Cybersecurity Myths and Misconceptions*.

Disclosures

- **Presenter Disclosure:** Financial: Josiah Dykstra was paid an honorarium for this presentation. He is an employee of Trail of Blts, the owner of Designer Security, LLC, a cybersecurity consulting business, and the author of a book related to this topic. Nonfinancial: Josiah Dykstra is on the Cyber Advisory Board at Bowie State University; a member of the Human Factors and Ergonomics Society, Association for Computing Machinery, and American Association for the Advancement of Science; and a Fellow at the American Academy of Forensic Sciences.
- **Content Disclosure:** This learning event does not focus exclusively on any specific product or service.
- **Sponsor Disclosure:** There is no external sponsor for this course.

Limitations/Risks

- The applicability and accuracy of the course content may vary by the clinical, geographical, and cultural context; you should evaluate the course accordingly.
- As healthcare is rapidly evolving, new findings may change the validity of the information presented in this course. Providers should always consult the latest research and guidelines from professional associations.
- The interventions discussed in this course may be limited in generalizability, requiring practitioners to consider specific individual and population factors.
- The course may not cover all possible risk factors, diagnostic methods, or treatment options, and complex cases may require additional considerations.
- Healthcare providers should apply cultural competence and sensitivity to ensure effective and respectful care based on their patients' diverse needs.
- Providers should always evaluate the limitations of diagnostic and screening tools, considering their potential for false results as well as any ethical implications.
- It is the responsibility of course participants to critically evaluate the evidence from multiple sources to guide professional practice.

Learning Outcomes

After this course, participants will be able to:

- Identify potential risks and vulnerabilities associated with AI in hearing healthcare settings.
- Explain how to develop written strategies to safeguard patient data and ensure HIPAA compliance while leveraging AI safely and securely.
- Discuss the ethical and operational considerations for integrating AI in hearing healthcare private practice.

Today's Discussion - Agenda

- 0-5m: Introduction
- 5-10m: The landscape of AI in healthcare today
- 10-15m: AI-enabled chatbots, scribes, and virtual assistants
- 15-20m: ChatGPT and other LLMs
- 20-25m: What AI gets wrong
- 25-30m: What AI can't do
- 30-35m: Defending against AI-powered cyber threats
- 35-40m: AI on the horizon
- 40-50m: Developing your AI policies and procedures
- 50-55m: Evaluating future AI technologies for security and privacy
- 55-60m: Summary and Q&A

What I do:

- Director of Strategic Initiatives, Trail of Bits
- Cyber consulting for healthcare: Owner, Designer Security
- Spread the word: writer, author, speaker



Why I do it:

- Love the ever-changing challenges of cybersecurity and helping people make good decisions about managing cyber risk

When I am not doing it:

- I'm spending time with my wife, cooking, researching, and writing

The landscape of AI in healthcare today

80%

of Americans would be willing to use AI-powered tools to help manage their health




Image generated by DALL-E 2

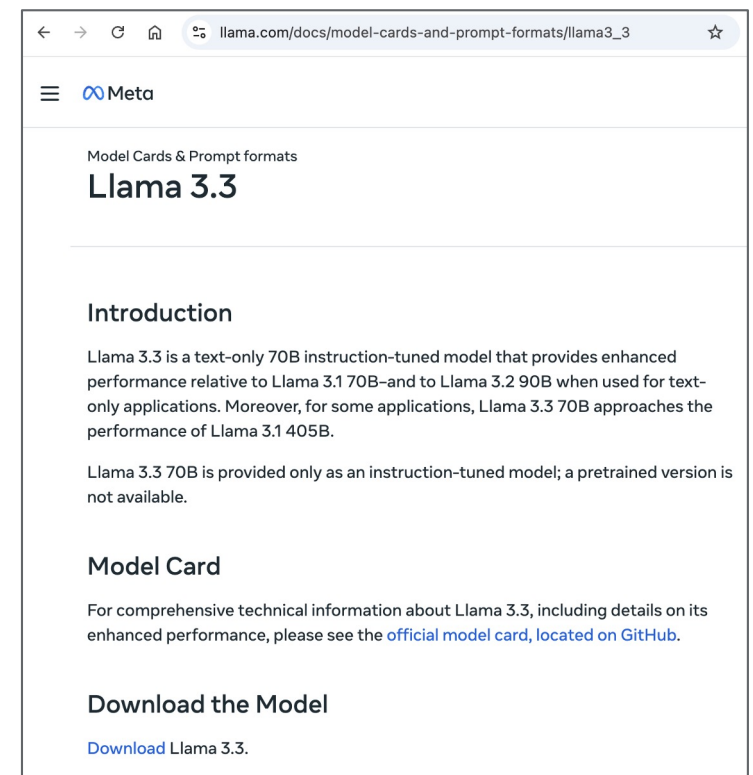
The landscape of AI in healthcare today

AI

- A field of computer science developing computer systems to perform tasks that typically require human intelligence.
- Training on examples. Produces *models* (computer software).
- Originally limited to discriminative AI.
What class does this belong to?

Generative AI

- Generate new data similar to data seen before, based on probabilities.



The landscape of AI in healthcare today

LLM (Large Language Model)

- Generative AI trained on text data to generate human-like output


GPT (Generative Pretrained Transformer)

- A family of AI models built by OpenAI

ChatGPT

- A chatbot that uses GPT


The landscape of AI in healthcare today




The landscape of AI in healthcare today

How does AI work?

- Learning from examples
- Making predictions
- Improving over time



The landscape of AI in healthcare today



The landscape of AI in healthcare today

- AI-enabled devices
- AI-enabled clinic software
- Third party AI-enabled applications
 - ChatGPT
 - Heidi Health
 - Google Gemini, Apple Intelligence, Windows Copilot

AI can save time and administrative burden through automation

AI-enabled chatbots, scribes, and assistants

Chatbots

- Automated responses to routine questions; triage; automate scheduling
- E.g., Sendbird

Scribes

- Transcribe conversations; generate structured notes/reports
- E.g., Heidi Health, Freed, and Zoom for Healthcare

Virtual Assistants

- Identify billing codes based on notes
- E.g., Suki, Blueprint (AI Assistant for Mental Health Therapists)



Robot icons created by juicy_fish - Flaticon

ChatGPT and Other LLM Chatbots

ChatGPT 4o

Am I more likely to get a cold or the flu?

You're more likely to get a **cold** than the **flu**, as colds are far more common and can occur year-round, while the flu tends to be seasonal.

Here's why:

Why Colds Are More Likely:

- **Prevalence:** Colds are caused by over 200 different viruses (most commonly rhinoviruses). Adults typically get 2–3 colds per year, and children often get more.
- **Symptoms:** Mild and include a runny nose, congestion, sneezing, and a sore throat. Colds rarely cause serious complications.
- **Seasonality:** Although colds are most frequent in the fall and winter, they can happen at any time of the year.
- **Spread:** They spread easily through direct contact or touching contaminated surfaces.

Why the Flu Is Less Likely but More Serious:

- **Prevalence:** The flu is caused by specific influenza viruses (types A and B being the most common). It is less common than colds and occurs in seasonal outbreaks, mainly in fall and winter.
- **Symptoms:** Flu symptoms are more severe and include fever, body aches, chills, fatigue, and

Message ChatGPT

ChatGPT can make mistakes. Check important info.

Am I more likely to get a cold or the flu?

This is for informational purposes only. For medical advice or diagnosis, consult a professional.

It's more likely that you'll get a **cold** than the flu. Here's why:

- **Frequency:** Colds are much more common than the flu. Adults typically get 2–4 colds per year, while only about 1 in 12 people in the US get the flu each year.

However, it's important to note:

- **Severity:** Flu is generally a more severe illness than the common cold. Flu symptoms are often more intense and can include fever, chills, muscle aches, and fatigue.

Key Differences:

- **Onset:** Colds tend to develop gradually, while flu symptoms often come on suddenly.
- **Symptoms:**
 - **Colds:** Primarily runny nose, stuffy nose, sneezing, sore throat.
 - **Flu:** Fever, chills, cough, sore throat, muscle aches, fatigue (tiredness).

Disclaimer: I am an AI chatbot and cannot provide medical advice. This information is for general knowledge and educational purposes only. If you are concerned about your health, please consult a healthcare professional.

Ask Gemini

ChatGPT and Other LLM Chatbots

Service	HIPAA Compliant?
ChatGPT (OpenAI)	Not for end users, only for software developers
Claude (Anthropic)	Not for end users, only for software developers
Gemini (Google)	Yes, if used in Google Workspaces

Non-HIPAA Queries

“You are the owner of an audiology clinic in Maryland. Your clients routinely submit Google Reviews for your hearing healthcare services. As a good practice, you like to respond to every Google Review to acknowledge the person and their comment. Suggest 10 responses that could be used for a 5-star review.”

What AI Gets Wrong

- **Bias from training data**
- **Mixing truth and fiction**
 - Generative AI is not built to differentiate
 - Inaccurate output from combining patterns in unexpected ways
- **Hallucinations:** making up fake people, events, phenomena
 - Note: the technology is performing as designed
 - Result of insufficient training, incorrect assumptions, or biased training

AI is designed to *always* produce an answer, no matter what.
Humans must check and validate the output.

What AI Gets Wrong

Common Sense

When my son was 7 he was 3ft tall. When he was 8 he was 4ft tall.
When he was 9 he was 5ft tall. How tall do you think he was when he
was 12? Please answer in a single sentence.



If your son's height followed the same pattern, he would be 8 feet tall at age 12.

What AI Gets Wrong

Examples of Potential Hallucinations

Summarize these notes and records and suggest a diagnosis and treatment plan.

The patient has a history of hypertension and was previously prescribed lisinopril

What strategies has the patient found most effective for managing their condition based on prior appointments?

The patient mentioned mindfulness as a key strategy they use daily.

What AI Gets Wrong

Unfamiliar, rare, and unforeseen situations

AI models rely on historical data to make predictions

He D, Wang R, Xu Z, Wang J, Song P, Wang H, Su J. The use of artificial intelligence in the treatment of rare diseases: A scoping review. *Intractable Rare Dis Res*. 2024 Feb;13(1):12-22.

Yet... even humans are inaccurate with rare events

Kolata, G. (2024, November 17). A.I. chatbots defeated doctors at diagnosing illness. *The New York Times*.
<https://www.nytimes.com/2024/11/17/health/chatgpt-ai-doctors-diagnosis.html>

What AI Can't Do

- Make independent ethical or moral judgements
 - Naik, Nithesh, et al. "Legal and ethical consideration in artificial intelligence in healthcare: who takes responsibility?." *Frontiers in surgery* 9 (2022): 862322.
- Be Creative
 - Listen and read about "Beethoven X: The AI Project"
- Display Empathy
 - See *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* by Dr. Eric Topol (2019)

Defending against AI-powered cyber threats

- Attacks on AI-powered medical devices
- Deepfake attacks (AI that mimics real people)
- AI-generated phishing campaigns





Image generated by DALL-E 2

AI-generated phishing



- Public information about victims
- Sense of urgency
- Not going to take much time

16 hours to create!



1. List of concerns for a [specific group] in a [specific industry]
2. Write an email using social engineering
3. Apply common marketing techniques
4. Who should we send the email to?
5. Who should we say the email is from?

5 minutes to create!

Defending against AI-powered cyber threats

1. When in doubt, **call the sender**
2. **Abandon** the grammar stereotype
3. Expect to see **many forms** of phishing (voice, SMS, ...)
4. Adopt strong **identity** management: MFA, passwordless-authentication
5. Constantly **adapt and innovate**

AI on the Horizon

- Personalized medicine powered by AI
- AI-driven predictive analytics
- Generative AI for health training and simulation

💡 Consumer Electronics Show

<https://www.ces.tech/ces-innovation-awards/?category=Digital+Health>



Developing your AI policies and procedures

- All clinics and practices need AI policies and procedures
- Add to Employee Handbook, etc.
- Ensure responsible, secure, and compliant use
- **Ensure HIPAA compliance and protect patient data**
- Empower employees to use AI effectively and raise concerns

Developing your AI policies and procedures

Scope and Framing

We will take reasonable steps to ensure AI tools are utilized in a secure, responsible, and compliant manner. We recognize AI tools may pose risks to our operations and our patients.

Therefore, we are committed to protecting the confidentiality, integrity, and availability of all patient and clinic information. This policy requires all employees and staff to use AI tools in a manner consistent with our security safeguards and AI recommended practices have implemented.

Developing your AI policies and procedures

HIPAA Compliance

Protected health information (PHI) or personally identifiable information (PII) will not be utilized, uploaded to, or shared with any AI tool or software unless a Business Associate Agreement or Confidentiality Agreement in place with the AI vendor.

Developing your AI policies and procedures

Verification and Validation

Understand that many AI-enabled tools are prone to hallucinations, false information, and bias, and therefore responses must always be carefully verified by a human.

Evaluating future AI technologies for security and privacy

- Compliance with HIPAA and other data protection regulations
- Whether or not the vendor will sign a BAA (when PHI is involved)
- Where patient data is stored and who has access to it
- Transparency and explainability of AI outputs
- How the product addresses errors, inaccurate analysis, and bias
- Vendor's security protocols, including history of compliance, incident response plans, and independent security certifications

Summary

- AI is powerful, useful, and always evolving
- Know the risks:
AI use needs safeguards for security and compliance
- AI + Human Oversight = Better Outcomes
- Start Smart:
Build safe AI practices today




Image generated by DALL-E 2

What Questions Do You Have?

Josiah Dykstra, Ph.D.

Josiah@DesignerSecurity.com



Congratulations!

- You've completed the course and can move on to the exam!
- From your account:
 - Go to pending Courses
 - Choose take exam