Liability and Risk Management for Using AI Tools

Michelle Mello, JD, PhD

Stanford | Health Policy
Freeman Spogli Institute and Stanford School of Medicine

Stanford University
Human-Centered
Artificial Intelligence

Stanford Law School

Learning Objectives

- Discuss when the use of a healthcare AI tool harms patients, who is responsible?
- Examine how courts are grappling with the challenges of adjudicating liability for software-related injuries.
- Discuss how health systems and clinicians can assess and manage AI liability risk.

The case of Mr. Park

Ms. Dee is a bedside nurse whose hospital uses an AI tool, VitalSign, to detect early signs of sepsis. It scans 60 factors in patients' EHR every 15 minutes and alerts the nurse when the model predicts a 25% or higher risk of sepsis.

Before VitalSign, nurses screened patients at the beginning of each shift using observation and a simpler, non-AI risk calculator with a lower predictive accuracy than VitalSign.

VitalSign doesn't send an alert for Dee's postsurgical patient, Mr. Park, but he does become septic. Several of his vital signs had unusual baseline values due to his underlying health conditions and recent surgery. By the time Nurse Dee registers his deterioration and calls a physician, he is severely septic. He dies despite appropriate treatment.

Who, if anyone, is responsible for Mr. Park's death?

Questions raised

- Is the developer liable if VitalSign outperforms standard of care? If it works well for most patients but not patients like Mr. Park?
- Was it unreasonable for Ms. Dee not to have conducted more intensive observation?
- Should the hospital have instructed nurses to use VitalSign only as a supplement to existing screening?
- What did the hospital tell nurses about VitalSign's performance?

Roadmap

1. The AI adoption landscape
2. AI-related liability risk and why it matters
3. Recommendations for managing risk

Key healthcare AI modalities, simplified

Predictive AI

Uses machine learning techniques and massive patient data troves to train an algorithm to classify or predict things

- *Will this patient need a blood transfusion during surgery?*
- *Does this EKG suggest hypertrophic cardiomyopathy?*
- *Can daily lab tests be safely discontinued for this patient?*

Generative AI

Uses a large language model to create novel text in response to a prompt

- *What was discussed and decided during this office visit?*
- *What are the radiologist's impressions from this x-ray?*
- *When did this patient last report chest pain?*

Types of healthcare AI tools

Clinical decision support

Healthcare operations

Direct-to-consumer

Epic’s AI builds

Source: https://www.epic.com/epic/post/cool-stuff-now-epic-and-generative-ai/

What are provider organizations using generative AI for?

USE CASE	OPPORTUNITY			ADOPTION STAGE					
	Pain point intensity	Current level of manual work	Opportunity Score	Not yet started	Ideation	POC	Active pilot	Full rollout	Adoption Score
Patient triage	36%	63%	**49**	34%	17%	25%	19%	5%	29
Provider credentialing and enrollment	33%	51%	**42**	39%	19%	19%	14%	9%	27
Patient scheduling	37%	42%	**39**	19%	21%	26%	28%	5%	36
Staff scheduling	36%	42%	**39**	26%	19%	25%	19%	11%	34
Risk adjustment	42%	46%	**44**	16%	30%	24%	24%	6%	35
Care gap identification	37%	53%	**45**	24%	13%	26%	32%	5%	36
Clinical trial coordination	26%	58%	**42**	21%	21%	37%	11%	11%	34
Documentation support (scribes)	55%	31%	**43**	3%	11%	27%	38%	22%	53
Follow-up care	41%	51%	**46**	19%	31%	20%	22%	8%	34
Referral management	34%	51%	**42**	25%	19%	36%	13%	8%	32
Quality metrics and patient registry	47%	32%	**40**	23%	28%	21%	25%	4%	32

AI scribes are at the leading edge of adopti

Patient email drafters are also popular

Note: Opportunity Score is calculated as the average of pain point intensity (% respondents claiming a job is a significant pain point) and current level of automation (% respondents of describing job as mostly manual process); Adoption Score calculated using weighted average of development stage where net yet started is 0% adoption and implementation/full rollout is 80% adoption | Source: Bain GenAI Survey (N = 408)

Roadmap

1. The AI adoption landscape
2. AI-related liability risk and why it matters
3. Recommendations for managing risk

Lawyers' warnings

"can be a disaster for health care providers"

Image: Broccolo BM, AHLA In-House Counsel Program, 2019
Quotation: Keris MP, Am. Soc. for Healthcare Risk Mgmt., 2020

Reasons for worry

- Less testing than some other clinical innovations
- When an area is under-regulated, liability tends to fill the gap
- Errors may propagate over many patients
- Public distrusts AI
- Harm events likely to draw public attention
- Judges are inexperienced, doctrine is underdeveloped
- Unclear who will be left holding the ball

Reasons for reassurance

- Few negligent injuries become claims
- Scant evidence of software-related claims to date
- Significant hurdles for plaintiffs in AI cases

What plaintiffs must prove

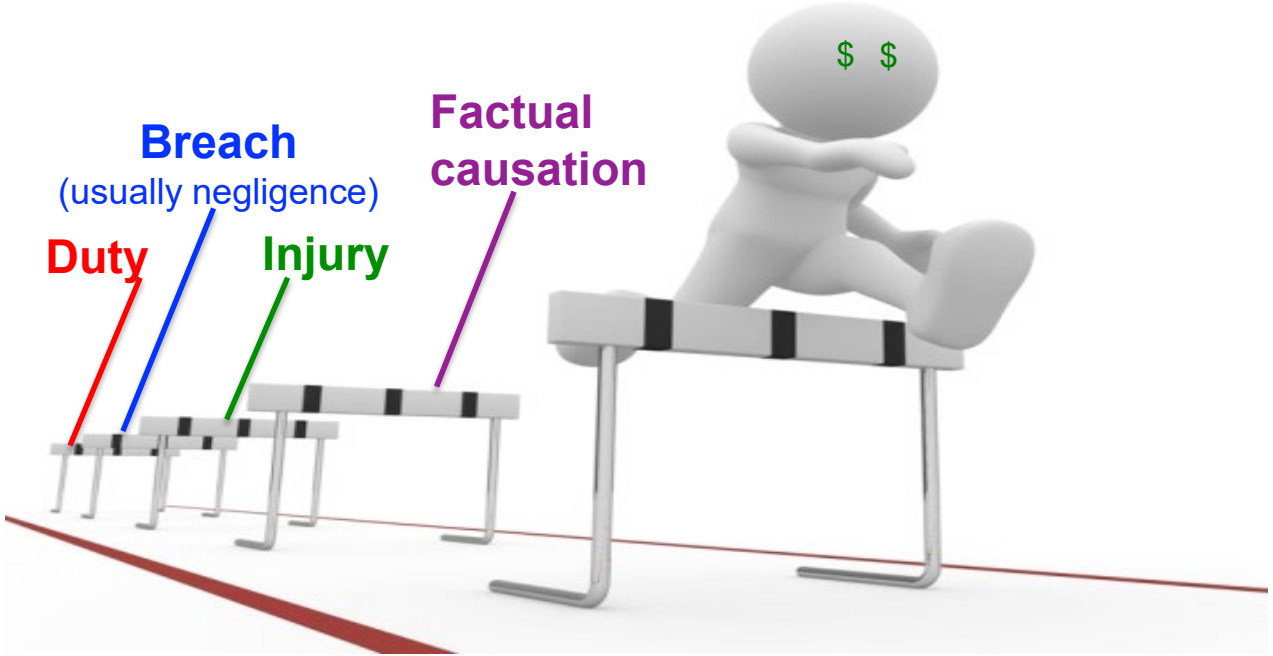

Challenges for plaintiffs: **DUTY**

- Do developers have a duty to patients when a human-in-the-loop makes the final call?
- Absent FDA-approved indications or established patterns of AI use, what constitutes unforeseeable misuse?

Challenges for plaintiffs: **BREACH**

- Not clear that courts will deem AI a “product”
- Hard to show there’s a reasonable alternative safer design
- Model opacity makes it hard to prove that a clinician’s decision to accept/reject output was unreasonable

Challenges for plaintiffs: **CAUSATION**

- Model opacity makes it hard to prove that wrong output occurred because of a *defect*
- For clinical areas where human+AI outperforms human-only, is the counterfactual no injury?

Roadmap

1. The AI adoption landscape
2. AI-related liability risk and why it matters
3. Recommendations for managing risk

Risk management recommendations

- Calibrate adoption & oversight decisions to risk level

AI risk assessment

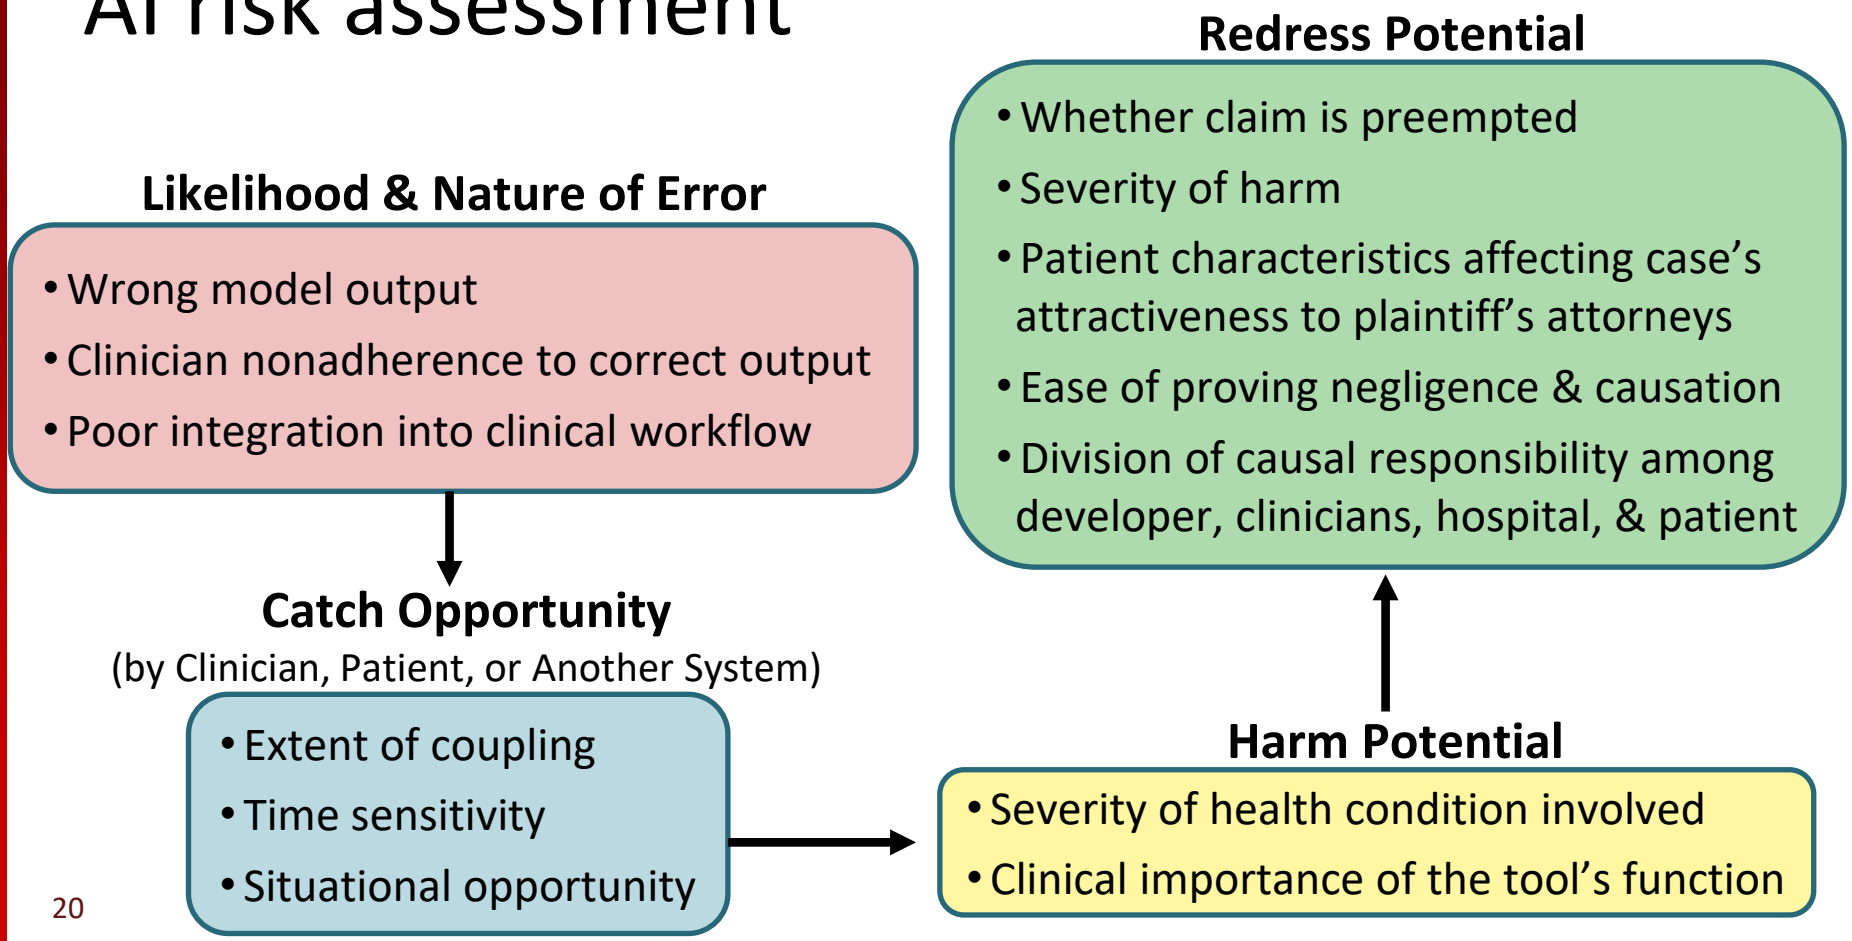

Risk management recommendations

- Take advantage of the buyer's market. Bargain for:
  - Information & other supports for monitoring
  - Indemnification
  - Adequate insurance for developer
  - Non-applicability of disclaimers in Terms of Use

OpenAI

Published: December 11, 2024

Terms of Use

- You must not use any Output relating to a person for any purpose that could have a legal or material impact on that person, such as making credit, educational, employment, housing, insurance, legal, medical, or other important decisions about them.

Risk management recommendations

- Anticipate evidentiary issues in litigation
  - Document inputs, outputs, versions, & reasons for accepting/rejecting recommendations
- Acknowledge and resist automation bias

The legal standard of care for malpractice

The custom standard:

The care that a reasonable practitioner in the defendant's specialty would provide in similar circumstances.

The current standard:

The care, skill, and knowledge regarded as competent among similar medical providers in the same or similar circumstances.

Customary practice may fall short of what medical professionals regard as competent. "It should be no defense that many other providers render similarly deficient care."

- *Restatement (Third) of Torts: Medical Malpractice* § 5(a) (2024)

Risk management recommendations

- Insist that developers provide key information about their tools' performance and training

Questions to ask:

- *How similar was the training population to our patients? Will you tune the tool further on our data?*
- *What subgroup analyses were conducted on the training data, what were the results, and were sample sizes large enough to adequately power the analyses?*
- *In performance testing, what did the tool do well—and not so well? If you had to train users on potential problems to look out for, what would you say?*
- *Does the tool provide any help for users who want to verify the output (e.g., links to source documents; information on which factors weighted most heavily in the prediction)?*
- *Do you have other customers that look similar to us? Can we talk with them about their experiences with the tool, or do you have results from those sites to share?*

Risk management recommendations

- Set an institutional policy about patient notification for each tool

Should patients be told about uses of AI?

Arguments against:	Arguments for:
Use of other decision supports usually isn't disclosed	AI is different; it's "material information" to patients
Patients care about physicians' judgments, not how they make them	Like other evidence, helps patients weigh treatment recommendations
Patients have low understanding of AI	Clinician can help patients understand
Might create distrust	Candor engenders trust; use of AI will come to light in litigation

Ask, tell, or neither?

How great is the risk of physical harm?	Does patient have a meaningful opportunity to exercise agency?
• Risk posed by tool • Likelihood that errors will reach patients • Severity of potential harm	• Opt out • Alter behavior in ways that promote their interests

Source: Mello, Char & Xu, *JAMA* 2025

What should be disclosed?

1. The fact that an AI tool is being used
2. What functions it performs
3. Basics of how it works, including clinician's role
4. Why the organization believes it improves care
5. Basics of how the organization monitors performance, including in subgroups
6. Where applicable, patient's choices about having the tool used

Sample consent for an ambient scribe

Modality

Information is distributed via paper information sheet when patient checks in for clinic visit; nurse elicits oral consent when patients are brought to examination room.

Information Sheet

Apex Hospital uses an artificial tool (AI) tool called Intelligent Ear to help your physician take notes on your visit today and summarize them in your medical record. Under California law, you have the right not to have this tool used.

Intelligent Ear makes an audio recording of your physician visit and uses AI to write out the conversation in full. It then uses a different kind of AI to generate a summary of the important parts of the conversation. Your physician will have the opportunity to review and edit this summary before making it part of your medical record.

The audio recording is securely sent to the company that developed Intelligent Ear, but the company does not keep any information about you once the visit summary is created.

Apex Hospital is using this tool because of evidence that it can improve care. Having this type of aid enables physicians to keep their focus on you during your visit and save time writing notes at the end of the day. Testing showed Intelligent Ear generates good-quality summaries and that physicians find it helpful.

The tool performs somewhat less well for people with accented English, people with speech impediments, and conversations with more than 2 speakers. To help avoid inaccuracies, your physician has been trained to review summaries with special care in those circumstances. Apex Hospital is using this tool after studying the results of a pilot test at Apex, and continues to monitor the tool's performance to make sure it is working well for all types of patients.

We are excited to be able to offer this tool in our continuing efforts to deliver the best care possible. If you do not wish to have the tool used in your visit today, please let your nurse know.

Nurse Script

Did you have any questions about the scribe tool that the physician uses? Is it OK with you if the physician uses it today?

Source: Mello, Char & Xu, *JAMA* 2025

More information

Deep dive into AI liability & risk management

Informed consent recommendations

Liability issues from using ChatGPT

Vetting AI tools for ethical problems