

Issue Brief: Liability for Clinical AI

October 2026

As AI assumes an expanding range of functions in clinical care, questions about responsibility for patient harm are becoming more urgent. Existing liability frameworks were not designed for technologies in which developers, deploying organizations, and physicians exercise different degrees of authority and control over clinical functions and decisions. Responsibilities may overlap, but each party remains accountable for the risks it is positioned to understand, prevent, or mitigate. Aligning responsibility with these factors is essential to realizing AI's potential to improve patient care.

Physicians identify liability as one of their [most serious concerns](#) about adopting AI. In the AMA's 2026 Physician Survey on Augmented Intelligence, [more than one-third of respondents](#) selected clear liability frameworks as the regulatory action most likely to increase their trust in AI and their likelihood of adopting it. Respondents also cited postmarket surveillance, including adverse event reporting, and oversight by governing bodies as ways to build confidence.

[AMA policy](#) holds that liability for AI-related harm should rest with the parties best positioned to know the risks and to prevent or mitigate them. This brief outlines how the AMA Center for Digital Health and AI (the Center) is applying that principle to physicians, developers, and those who deploy AI, recognizing that these roles may overlap. It focuses on AI systems that influence patient care, including those operating within clinical workflows or supporting decisions elsewhere in the care pathway. This brief does not address AI systems used directly by patients outside clinician-directed care. Increasingly autonomous AI raises additional questions addressed under "What's Next."

Key Issues

In the current liability landscape, the following issues are of primary concern to the AMA:

- **Physicians may be held responsible for errors they cannot reasonably detect.** Physicians may be expected to rely on AI systems they were not informed were in

use or had no authority to select or decline. Meaningful physician oversight depends on conditions physicians often do not control. These include what developers disclose, how interfaces present outputs, whether physicians have the authority to intervene, and whether clinical workflows allow sufficient time and information for review. Physicians often cannot inspect a model, see how it was validated, or know how it performs on their own patients.

- **Deployers' AI-specific responsibilities remain unsettled.** Healthcare organizations that deploy AI control many of the conditions that shape its use. These include tool selection, vendor contracts, interface configuration, staffing, training, time for review, and post-deployment monitoring. Healthcare organizations already owe duties of care for equipment acquisition, staffing, physician credentialing, oversight, and policies, though the scope of those duties varies by state. What remains largely unsettled is how those duties apply to AI. Open questions relate to qualification or credentialing of AI systems, local testing, version updates, user training, safety and outcomes monitoring, incident response, and decisions about when to take a tool out of service.
- **It is unclear when product liability applies to AI.** Developers control a tool's design, development, labeling, and warnings. Product liability law can hold manufacturers responsible for defective designs and inadequate warnings. Whether those rules reach

software, and which AI systems they cover, is unsettled and varies by jurisdiction. A tool's regulatory status does not settle its legal status.

- **Contract terms and disclaimers can shift responsibility without reflecting control over risk.** Contracts between developers and deploying organizations may allocate financial responsibility for AI-related harm through indemnification provisions, liability limits, and other terms. Patient- and physician-facing terms may also restrict dispute resolution options or limit legal claims for injuries. These arrangements can place responsibility on parties with little ability to identify, prevent, or mitigate the harm.
- **Some patients may have no clear path to recovery.** Not every patient injury results from a failure for which a party is legally responsible, even when AI is involved. But when harm results from defective technology, negligent implementation, or other actionable failures, responsibility can be difficult to establish because it is distributed across parties, obscured within systems, or limited by contract. Patients harmed by these failures should have meaningful avenues for redress.

AMA Policy and How the Center Is Applying It

AMA Policy [H-480.939](#) states that “liability and incentives should be aligned so that the individual(s) or entity(ies) best positioned to know the AI system risks and best positioned to avert or mitigate harm do so through design, development, validation, and implementation.” The Center applies this policy by aligning accountability with each party's authority, access to information, and ability to prevent or mitigate harm. In practice, that means:

- **Physicians** are responsible for the clinical judgment they can meaningfully exercise.
- **Developers** are responsible for what they design and build.
- **Deployers** are responsible for where and how AI is put to work.

Where control is shared, responsibility is shared, but it does not become interchangeable, and no party's obligations disappear because another party also bears some.

This principle is central to the Center's broader approach to [clinical AI governance](#). Clinical authority must be earned through evidence and maintained through demonstrated performance, with authority and accountability kept aligned.

The strength of the evidence should inform the authority an AI system is permitted to exercise, the oversight it requires, and the responsibilities of those who develop, deploy, and use it. Where evidence is insufficient or does not apply to the intended clinical setting, deployment should be reconsidered. Physicians should not be asked to compensate for those limitations through additional review or documentation.

The considerations below reflect the Center's current thinking on how these principles apply across the development, deployment, and clinical use of AI. We reference formal AMA policy where applicable. They represent the Center's emerging analysis, offered to advance the conversation and to guide future work, including evidence-gathering, stakeholder engagement, and policy development.

1. Enable meaningful physician oversight

Physicians need the right information, at the right time, in a form they can act on. Meaningful physician review depends on access to sufficient information to evaluate an AI system's output in the relevant clinical context, alongside the time, training, and authority to act on it. These conditions cannot be assumed simply because a physician is presented with an AI-generated result. Clear expectations about physician responsibility when AI systems fail are also essential to supporting appropriate reliance on these technologies and building trust in their clinical use.

- **Account for what a physician could reasonably detect.** A physician's responsibility for AI-related harm should reflect their ability to review an AI system's output, identify potential errors, and intervene to prevent or mitigate harm. That ability depends on the information available to them, the clinical workflow, and the physician's actual authority to act. Consistent with AMA Policy [H-480.931\(5\)\(b\)](#), physicians should not be held liable for the performance of an AI-enabled technology when they do not know or have reason to know of concerns about its quality and safety. Two examples:

— *Triage error: the physician cannot see the AI decision.* If a triage model deprioritizes a case, the physician sees only an order of patients. The physician may see that patient hours later with no way of knowing the system caused the delay.

- *Upstream data error: the physician cannot verify the inputs.* An AI system synthesizing records from multiple providers and/or patient inputs may generate an output based on missing laboratory results, outdated diagnoses, or incomplete historical information. The output may appear clinically plausible, even when it is based on incomplete or inaccurate information that the physician cannot reasonably identify or verify.
 - **Establish expectations for reasonable reliance.** Physicians should be able to rely appropriately on AI systems that have demonstrated adequate performance for their intended clinical functions and have been responsibly implemented, without assuming responsibility for risks they cannot reasonably identify or mitigate. Whether reliance is reasonable should reflect the supporting evidence, clinical risks, conditions of deployment, and the physician's ability to meaningfully evaluate and act on the system's output. Physicians who reasonably rely on AI systems should have the benefit of strong liability protections. Physicians retain their duty to patients and their authority to exercise independent clinical judgment.
 - *Conditions supporting reasonable reliance:* Reasonable reliance is supported when AI systems perform a defined clinical function, have demonstrated performance for that function in comparable settings, provide outputs physicians can meaningfully evaluate and act upon, and undergo changes through controlled, documented processes. Regulatory authorization can also support a finding of reasonable reliance when an AI system is used within its authorized intended use and under conditions consistent with its validation.
 - *Conditions limiting reasonable reliance:* Reasonable reliance may be limited when an AI system's performance has not been adequately characterized for the clinical function at issue, or when errors cannot reasonably be identified through the available information and clinical workflow. A fabricated detail in a clinical summary, for example, can look entirely consistent with the record the physician reviewed. Vigilance alone cannot substitute for evidence and design.
 - **Physicians should be free to decline unproven AI.** Physicians should retain the authority to decline AI systems whose clinical validation, usefulness, or place in the standard of care remains uncertain, without being penalized for exercising their clinical judgment.¹

Meaningful physician oversight cannot be established through labeling, regulatory authorization, or institutional policies alone. It must be supported by evidence, information, clinical workflows, and the authority to intervene.
- Define developers' responsibilities**
- Developers play a central role in determining how AI systems perform, what risks they introduce, and whether those risks can be identified and mitigated throughout their use.
- **Developers are responsible for what they design and build.** Developers make the design and development decisions that shape how AI systems function, the risks they introduce, and whether users can detect and respond to errors. These include decisions about training and validation data, system architecture, intended use, interfaces, and the information made available to users. Real-world performance also depends on how a system is deployed and used. Even so, developers remain responsible for the risks arising from what they control.
 - **Limit developers' ability to shift responsibility to physicians.** Physician oversight should not serve as a blanket justification for transferring responsibility or liability for AI-related harm. Warnings and labeling alone cannot establish the conditions necessary for meaningful physician oversight. Physicians also need usable information about intended use, performance, limitations, and the conditions under which outputs can be reliably evaluated and acted on. As AMA Policy [H-480.931](#) at (2)(a)(ii) holds, "transparency requirements should be designed to meet the needs of the end users. Documentation and disclosure should enhance patient and physician knowledge without increasing administrative burden."
 - **Duties continue after deployment.** Developer responsibilities extend throughout the lifecycle of an AI system, proportionate to their control over its design, performance, and associated risks. They

¹ See AMA Policy H-480.939 at (6): "Physicians should not be penalized if they do not use AI systems while regulatory oversight, standards, clinical validation, clinical usefulness, and standards of care are in flux."

continue to the extent the developer retains control over the technology, has access to performance and safety information, or can identify and mitigate emerging risks.

2. Set clear expectations for deployers

Deployers may include health systems, hospitals, physician practices, and other organizations responsible for introducing AI into clinical care and have responsibilities for the selection, implementation, oversight, and ongoing use of clinical AI. These responsibilities are distinct from those of developers and physicians, and they complement them. If multiple parties contribute to an AI-related harm, liability should reflect each party's actual contribution and capacity to prevent it. Shared responsibility must never become a way to avoid liability.

- **Institutional duty of care.** Deployers should have fit-for-purpose processes for:
 - evaluating and selecting AI systems;
 - validating their performance in the intended clinical setting;
 - monitoring safety and performance;
 - responding to emerging risks.

Training should prepare physicians to use technology as intended. Disclosures should be calibrated to the setting, so that physicians are not left to filter excessive or extraneous information. These processes should be tailored to clinical function, the consequences of error, and the organization's ability to oversee and intervene.

- **Modifications and updates.** Deployers are responsible for the implementation decisions and modifications within their control. Liability exposure should reflect the nature of the change: routine configuration differs from a material modification that alters a system's function or performance. Deployers should understand how updates are managed, including which changes they control, and establish appropriate processes for assessing their implications for clinical use and patient safety.
- **Information flow and corrective action.** Deployers should ensure that safety and performance information moves among developers, clinicians, and appropriate oversight bodies. They should also have processes for responding to adverse events and for modifying, restricting, or discontinuing use when warranted.

3. Preserve patients' path to recovery

Patients harmed through the use of clinical AI should have meaningful avenues for seeking accountability and appropriate remedies. As responsibility for clinical functions is distributed across developers, deployers, and physicians, existing liability frameworks may not adequately address every circumstance in which harm occurs. Protecting physicians who rely appropriately on AI must go hand in hand with clear expectations for developer and deployer accountability. Otherwise, patients may be left without recourse when harm results from failures for which those parties are responsible.

- **Developer responsibility for defects.** Patients harmed by defective AI systems should have meaningful avenues to pursue remedies from responsible developers. Liability frameworks should account for the developer's role in creating and controlling the relevant risks, including whether recovery without proof of negligence may be appropriate in some circumstances. Contracts should not allow developers to transfer responsibility for defects within their control to physicians or deployers, and other parties should remain accountable for their own independent failures.
- **Patient-facing disclaimers.** Disclaimers, consent forms, and other patient-facing terms should not be used to eliminate or improperly restrict patients' legal remedies for injuries resulting from defective AI systems or other actionable failures in their care.

What's Next

As clinical AI capabilities advance and their use expands, liability frameworks must evolve so that accountability stays aligned with clinical authority and control. The AMA Center for Digital Health and AI will work with physicians, patients, developers, healthcare organizations, insurers, legal experts, and policymakers on six priority areas:

- **Increasingly autonomous AI.** As AI systems assume more consequential clinical functions with less direct physician oversight, how should liability for patient harm be allocated? Key questions include authorization to perform clinical functions and the respective obligations of developers and deploying organizations.
- **Automation bias and meaningful oversight.** How should liability frameworks account for foreseeable automation bias on AI systems, particularly those that

perform reliably most of the time? What responsibilities do developers and deployers have to design systems and clinical workflows that support meaningful physician oversight?

- **Patient information and consent.** When should patients be informed that AI is being used in their care, and when is informed consent appropriate? How should expectations evolve as AI assumes more consequential clinical functions, and how can patients understand who remains accountable for their care?
- **Information sharing and safety reporting.** What information about AI performance, limitations, errors, and adverse events should developers and deployers share with physicians, patients, regulators, and one another? What mechanisms are needed to identify emerging risks and support timely corrective action?
- **Practical implementation and economic sustainability.** How can liability and oversight expectations be right-sized for practices and health systems with different resources and capabilities? What training, infrastructure, clinical time, and sustainable payment or financial support are necessary to make these responsibilities feasible without creating new barriers to access or innovation?

- **Data privacy, cybersecurity, and liability.** How should liability for patient harm arising from AI-related privacy or cybersecurity failures be allocated among developers, deployers, and physicians? How should existing legal obligations and each party's authority, access to information, and ability to prevent or mitigate harm inform that allocation?

These priorities will guide the Center's work to convene stakeholders, address evidence gaps, shape policy, and develop practical approaches to accountability as clinical AI evolves. We welcome input from physicians, patients, developers, healthcare organizations, insurers, legal experts, and policymakers on the practical challenges, evidence needs, and potential solutions that should inform this work.

Conclusion

As AI assumes increasingly consequential functions in clinical care, accountability must evolve alongside clinical authority. Physicians should be responsible for the clinical judgment they can meaningfully exercise; developers and deployers for the risks within their control; and patients must have meaningful avenues for redress when harm results from actionable failures. Getting this right is essential to earning the trust of both patients and physicians, enabling responsible innovation, and translating AI's capabilities into better patient care.



Center for
Digital Health & AI