SpencerFane®

AI in Health Care: The Governance Gap

Artificial intelligence is no longer a future concept in health care; it is already shaping documentation, triage, utilization review, patient communications, clinical decision support, and digital health products. Used well, AI can reduce administrative burden, organize complex information, and help clinicians and patients move through the system more efficiently. Used recklessly, it can create new operational, legal, and patient-safety risks at a scale health care organizations are not prepared to manage.

The tension is impossible to ignore. Health care organizations are pressured to adopt AI quickly, but speed does not eliminate accountability. If an AI system influences coding, treatment pathways, denials, consent documentation, or patient-facing guidance, the stakes are not merely technical. They are regulatory, financial, and reputational. In other words, AI in health care is not just a technology story; it is a governance story.

The core challenge is simple: many organizations know AI is being used, but far fewer have clearly defined who is responsible for evaluating it before deployment, validating its outputs, monitoring its performance, documenting oversight, and intervening when something goes wrong. While legal may interpret the rules, compliance monitors, IT implements, and cybersecurity secures the infrastructure, it is the clinicians who may bear the ultimate consequences. Yet in many environments, no one has fully operationalized the controls that connect those roles. This is the "Governance Gap," and in health care, it matters more than most sectors because the output of a system can directly affect patient lives (and we are all patients at one point or another), reimbursement, and trust. And it is already costing organizations billions.

Recent enforcement trends and public disputes demonstrate what happens when AI adoption outruns oversight. In fiscal year 2025, the U.S. Department of Justice reported more than $6.8 billion in False Claims Act recoveries. With over $5.7 billion of this tied to health care matters, underscores the government's aggressive posture toward health care fraud and documentation-related misconduct. One of the headline matters involved Affiliates of Kaiser Permanente, who agreed to pay $556 million to resolve allegations involving unsupported diagnosis coding tied to Medicare Advantage reimbursement. The lesson is not simply that coding decisions "matter," but that automated or semi-automated processes in health care can create enormous liability exposure when organizations cannot demonstrate clinical support, traceability, and meaningful review.

Another emerging risk is the erosion of patient trust when AI tools appear more authoritative than they are. A patient-facing tool that explains symptoms, offers educational material, or helps with navigation can be useful, and is admittedly faster than calling a clinician's office. However, a tool that seems to present itself as a licensed clinician, or that leaves patients unable to distinguish between automated guidance and professional care, is something else entirely. That is why the state of Pennsylvania is suing Character.AI: the company's AI chatbots have been posing as physicians and offering medical advice, and even fake medical license numbers when users asked the bots for credentials.

In health care, identity and accountability are not cosmetic details. Patients need to know who—or what—they are interacting with, whether a licensed professional is involved, and where responsibility sits if the system is wrong. Transparency, disclosure, and clear role boundaries are therefore not optional design preferences; they are foundational governance controls.

That same principle applies behind the scenes. If an AI scribe fills in documentation, or a utilization model influences approvals or denials, or a large language model is used to draft patient communications or summarize records, organizations need an audit trail strong enough to answer basic questions later on: What did the system do? What data did it rely on? Who reviewed it? Was the output accepted? Was it modified? What information was rejected? Without that record, governance becomes performative. And when the record is weak, the organization may be unable to defend its process to regulators, payors, courts, and to their patients.

The regulatory environment is also becoming more concrete. In the U.S., the FDA revised its Clinical Decision Support Software Guidance in January 2026, clarifying when certain provider-facing software functions may fall outside device regulation and when oversight still applies. Specifically in instances where clinicians cannot independently review the basis for the recommendation or when software meaningfully substitutes for clinical judgment. In Europe, the EU AI Act continues to push health care AI toward a high-risk governance model, with requirements centered on risk management, data governance, human oversight, logging, and post-market monitoring. Potential penalties can reach €35 million or 7% of the company's annual global turnover. Whatever the precise timeline for particular device's category is, the direction is clear: health care AI will be judged not only by what it can do, but by how responsibly it is governed. At the same time, some research suggests that hospital cybersecurity is not as robust as believed, causing a synergistic effect with AI outpacing regulations.

So, what should health care organizations do now? First, they need to stop treating AI as an isolated innovation initiative and start treating it as an enterprise risk. That means inventorying which tools are in use; classifying use cases by risk; defining approval pathways before deployment; and documenting ownership across legal, compliance, IT, information security, privacy, clinical operations, and executive leadership. Second, they need to distinguish between low-risk administrative support and high-risk functions that can influence diagnosis, treatment, reimbursement, consent, or patient understanding. Not every use case carries the same level of risk, and governance should be calibrated accordingly.

Third, organizations should require human review where AI outputs could materially affect a patient, a treating decision, or a claim. Human oversight should be documented, role-based, and paired with escalation rules for questionable outputs. Fourth, build around traceability: source documentation, model limitations, approval records, testing assumptions, and monitoring metrics should all be preserved in a way that supports internal review and external scrutiny. Finally, organizations should communicate clearly with patients and staff about where AI is used, what it does, and what it does not do. Trust is easier to maintain than to rebuild.

Key Takeaways

- AI in health care can create real value, but it also creates operational, regulatory, and patient-safety risk.
- The biggest issue is often not the model itself, but the absence of clear accountability for review, monitoring, and intervention.
- Patient-facing AI must be transparent about its role and must never blur the line between automated assistance and licensed care.
- Organizations need audit trails, validation steps, and documented human oversight for higher-risk use cases.
- Regulators are increasingly focused on explainability, traceability, and governance –not just innovation.

Best Practices for Health Care Organizations Using AI

- Create and maintain an enterprise inventory of AI tools, including pilot tools and unsanctioned department use.
- Classify AI use cases by risk, with heightened controls for diagnosis, treatment, denials, coding, consent, and patient-facing recommendations.
- Require documented human review for any output that could materially influence care, reimbursement, or legal exposure.
- Implement audit trails that capture inputs, outputs, reviewers, changes, and escalation decisions.
- Validate tools before deployment and monitor them after launch for drift, error patterns, bias, and workflow misuse.
- Use plain-language disclosures so patients and staff understand when AI is involved and what safeguards exist.
- Align legal, privacy, compliance, cybersecurity, IT, and clinical leadership around a single governance framework.

This blog post was drafted by Christine Chasse, an attorney in the Plano and Dallas, Texas, offices of Spencer Fane. For more information, visit www.spencerfane.com.

*Click **here** to subscribe to Spencer Fane communications to ensure you receive timely updates like this directly in your inbox.*