

HEALTH LAW WEEKLY

October 23, 2026

Board Oversight of AI in Health Care: New Technology, Familiar Duties

Tracy E. Miller, Brach Eichler LLC

Charles E. Binkley, Hackensack Meridian Health

Artificial Intelligence (AI) has initiated a transformative period in health care delivery as in other sectors of society. What distinguishes health care are the unique ethical and legal obligations owed to patients and how these will apply to AI-enabled clinical care and operations. Although health systems have used early forms of AI—such as rules-based risk scores—for decades, more powerful tools including machine learning and agentic AI, pose new demands in assessing both potential benefits and risks. Important questions about accuracy, impact, and human oversight are surfacing with particular salience for health care delivery.

The far-reaching changes prompted by AI in health care delivery call upon boards to exercise their fiduciary duty to oversee AI adoption and use. As boards seek to carry out this responsibility, they must grapple with the distinct challenges AI presents: the rapid pace of change, the expansive potential benefits, the scope of risk, and the fact that AI cuts across all aspects of mission and operations from clinical practice to administration and workforce, with uncertainty about its efficacy and impact. This article examines those challenges and provides suggestions for how boards can oversee the integration of AI into health care treatment and operations, recognizing that the expectations for board oversight of AI are still evolving and will vary by health care setting.

Even among similar organizations, the approach boards take will be shaped by the existing board structure and practices and by the interaction with executive management. Notably, the term “AI governance” is widely used in relation to the systemic oversight of AI at the management level—the structures, processes, and approaches for day-to-day management of AI. This is distinct from and not the board’s responsibility but requires a robust dialogue with executives managing AI and a well-delineated crosswalk of AI governance to board oversight. Similar to other areas of governance, board oversight of AI should be carried out with a clear demarcation of duties between directors and management, with the board responsible for oversight but not day-to-day operations, which are the purview of management.

Copyright 2026, American Health Law Association, Washington, DC. Reprint permission granted.

Legal Obligations in Overseeing AI

In overseeing AI, boards will be held to long-established legal standards and expectations for board duties. The technology itself, however novel, does not change the framework for board obligations, although it requires an examination of how board education, committee functions, and reporting should be refashioned to encompass AI. AI requires the board to gain new competence and determine how it can serve as an independent source of oversight and judgment for the organization in the face of a technology with an accelerating pace of change.^[1]

Boards of directors must meet two core fiduciary duties, the duty of care and the duty of loyalty. In many states, case law or statutes establish a third duty of mission or obedience, requiring that boards further and preserve the mission of the organization and adhere to governing documents and applicable laws. The duty of care requires board members to act in good faith with a reasonable degree of diligence and attention in carrying out board duties. As established in the seminal case of *In re Caremark International, Inc. Derivative Litigation*, as part of these duties, board members must act in good faith to ensure that an effective system exists to gather information and report to the board about risks to the organization.^[2] The duty of care is bounded by the business judgment rule, which recognizes that board members cannot be held accountable for a decision they make, even if it proves wrong or harmful to the organization, as long as they act in good faith, without conflict, and with the requisite degree of care.^[3] Inherently, the rule recognizes that boards acting responsibly will at times make decisions with imperfect information. As boards embrace AI as part of their oversight responsibility, with the uncertainty of both expansive opportunity and risk, the business judgment rule creates broad latitude for board decisions. Given AI's reach across fundamental values touching on quality, workforce impact, and strategy, AI use should be closely aligned with the duty of obedience to mission.

Building Blocks of Board Oversight

Board Education

To carry out the fiduciary duty to oversee AI, boards must have a general understanding of AI's strategic role and capabilities and how these capabilities will shift over time. Board education is essential, with more extensive education for board members who serve on the Technology Committee or a specially designated AI Committee. All board members should share a basic AI fluency. More specific education, especially related to higher-risk models, should be provided as needed. Members will inevitably bring to the board knowledge and experience with AI from other arenas. This existing knowledge does not replace the need for formal health care-related AI education. Rather, such education should augment the knowledge that board members already possess.

The board should also receive regular updates at full board meetings and relevant committee meetings.^[4] Annual board retreats could offer a valuable opportunity to focus intensively on the organization's current use of AI, uses on the horizon, and the evolving regulatory framework.

Board Composition and Structure

No one blueprint exists for board composition or committee structure to oversee AI. Ideally, boards will be able to recruit new members that add to the board's capacity and confidence in tackling AI oversight. The board should assess how needed expertise is best provided to the board overall and to board committees through some combination of membership, internal staff, and external consultants. Reporting cadence and content to the full board and to board committees should also be reviewed to address AI effectively.

As the board evaluates its oversight of AI, it will also be important for the board to consider whether to integrate AI concerns into existing committees or to create a separate AI Committee. For boards that have a Technology Committee, that Committee may be well-positioned to oversee AI with added members or competencies to evaluate the core areas of operations AI will impact—quality, workforce, compliance, mission/ethics, and those relevant to each organization's operations. Members from other standing committees, Finance, Quality, Audit and Compliance, and Human Resources, could serve on the Technology Committee or on a separate Committee established to oversee AI. Board committees can include individuals with AI expertise, even non-board members, consistent with state corporate law.^[5] Given the significant potential overlap in committee responsibility, the Governance Committee should review and revise the charters for each board committee so that AI oversight is clear to both board members and executive staff that must report to the board and board committees.

Roadmap of Existing AI Uses

A roadmap of current AI uses within the organization across all areas of operations is a valuable tool and starting point for board oversight. This roadmap can inform the board's stewardship over the range of AI tools: those that are built internally, purchased from vendors, integrated into existing software, or are part of research and education at academic medical centers. How is AI being used, with what risk controls, and at what cost? How are the models validated, assessed for impact, and monitored post-deployment? Which uses are under consideration for adoption and in what timeframe? Boards should understand the goals for each use, metrics to evaluate performance, projected cost, and measures to mitigate risk. This information could be presented in a chart that is updated each time the board considers AI. Capturing and tracking the information in this way will help boards focus on the highest cost and highest risk AI uses, while prioritizing patient safety and quality.

Current AI Landscape in Health Care

A strategic assessment of the current landscape for AI in the health care industry is essential information for boards as they weigh their organization's strategy and priorities. This assessment is most relevant by health care sector. For example, health systems, academic medical centers, and health plans are leaders in AI adoption. Their uses of AI, available options, and competitors' use of AI that may drive their decisions will differ from the way large physician practices, hospice, home care, and other providers use AI to address the distinct quality, patient engagement, financial, and operational challenges they face. Boards should be advised of the predominant current use cases, as well as those likely to emerge in the foreseeable future, in their sector of the health care industry to inform both long and short-term planning.^[6] Knowledge about how competitors are using AI, in addition to the overall trends in the health care industry, will help boards better assess the risk not only of proceeding with AI adoption but also of failing to do so. Boards should be updated periodically about emerging knowledge and the evidence-base for AI use cases.

Board Oversight: Key Areas for Consideration

Mission and Strategic Alignment

While management will generally have greater AI expertise than board members, the board can still provide critical guidance to preserve and further mission, core values, and strategic alignment in AI use. In the first instance, the board can set the tone for the organization's AI strategy, welcoming AI innovation while also signaling the need for rigorous risk assessment and mitigation. In overseeing AI adoption, boards should especially consider the potential impact on patient care and experience, clinician engagement, and the physician-patient relationship. A successful AI strategy will be measured in large part by patient and clinician trust, without which the most promising AI tool will never meet its full potential.

AI must be integrated into human systems for use by and for individuals. Failure to consider the human impact of AI, especially on the normative clinician-patient relationship, could derail the best devised strategy. Moreover, those relationships will remain integral to compassionate, patient-centered care as AI clinical use expands. Whether the board draws upon the Encyclical by Pope Leo XIV in Catholic health care organizations, guidance from other faiths, or humanitarian values, it can seek to preserve those human qualities as a touchstone of health care delivery.^[7]

AI's impact on employees and human resources practices may also be far-reaching, ranging from recruitment to job satisfaction and retention. The board has an obligation to safeguard workforce culture and loyalty as organizational assets and should identify the specific actions, whether workforce restructuring or use of AI to replace human judgment in hiring and employee evaluation, that would require board reporting and

Copyright 2026, American Health Law Association, Washington, DC. Reprint permission granted.

approval.^[8] Overall, the board should consider the ethical, legal, and reputational risks such changes may entail.

More broadly, boards can set priorities for AI aligned with goals for the organization in mission, clinical care, business opportunity, and operations. Given the cost and complexities of adopting health AI, board-set priorities would enable the organization's leadership to target their time and efforts accordingly. One process for input on AI strategy would be to present categories of AI use to the board, e.g., clinical or finance, with analysis of the core benefits and risks of each, and ask the board to rank which categories the organization should prioritize.

Leadership and Accountability

As part of its oversight responsibility, boards should seek a plan from management for incorporating competence in AI within the organization to carry out AI initiatives effectively. Equally important is the need for management to identify leadership and clear lines of accountability for AI at each critical stage of adoption—evaluation, approval of a new use, implementation, and monitoring. AI will cut across many competencies in health care delivery, all of which should be integrated into the structure and process to initiate and manage AI use cases.

One important vehicle for board oversight of AI selection and adoption is a multi-disciplinary AI Governance Committee at the management level with leadership representation from all critical areas of expertise: operations, medical and nursing care, IT, legal, compliance, finance, security, privacy, human resources, and ethics and/or mission.^[9] The board could charge the AI Governance Committee, working with the CEO or another executive, with ensuring that needed expertise is engaged in AI evaluation and approval, contracting, implementation, and risk assessment and mitigation. While risk assessment and mitigation should be an essential function of the AI Governance Committee, the board should set an expectation that the Committee will take a broad view of AI governance, giving due consideration to the many stakeholders who will be impacted by AI.

Assessing Benefit—The Question of Return on Investment

A fall 2025 poll conducted by the National Association of Corporate Directors found that board members perceived uncertainty about AI return on investment (ROI) as the single leading barrier to AI adoption.^[10] Given the significant financial cost of AI, the projected and actual ROI for each use proposed and implemented is critical information for the board. This information should be updated regularly so that the board can gain insight on AI's capacity and cost effectiveness. As with other investments, boards should conduct a rigorous assessment of major capital commitments for AI,^[11] understanding that the direct unit of measurement may not be financial. Although risk tolerance is

subjective, an objective process for evaluating both the absolute risk and ROI of an AI tool allows the board to be transparent in weighing the trade-offs involved.

For instance, licensing Ambient AI Scribes for clinicians is costly, and the overall time savings for an individual clinician is modest and seldom enough to increase patient volume.^[12] Nonetheless, many clinicians have come to rely on Ambient AI Scribes and report increased wellness when using the tool.^[13] It might be challenging for a health system to hire and retain clinicians without offering Ambient AI Scribes. Where is the ROI? In this case, it could be measured by clinician recruitment, retention, and well-being, far less tangible than adding one more patient to a clinician's schedule.

The board should consider the risks of any AI tool as part of a model's ROI valuation to inform a holistic assessment of its overall impact. The risks of *not* using AI may also be critical as the standard of care shifts.^[14] For instance, a commonly used radiology triage tool can identify an intracranial hemorrhage on a CT scan and flag the study for immediate evaluation. While there is no direct avenue available for reimbursing the use of this AI tool, a health system or hospital that has not deployed it may lag in the standard of care, posing potential liability.

Overseeing the Risks Posed by AI

Overseeing risk is a core board function and should be considered in parallel with ROI when evaluating an AI tool. With AI, boards will need to understand the nature and scope of the risks the technology presents, including the risks to patient safety and quality, cybersecurity, data privacy, compliance, workforce impact, ethical issues, and reputational. This general knowledge will enable boards to ask questions, seek information, and assess AI options, although each major AI initiative presented to the board for action should be evaluated individually for risk. AI regulatory risk is marked by the absence of a unified federal framework and evolving state laws that focus on a host of issues including consent, privacy, high-risk clinical use, AI use in coding and billing, and transparency.^[15] Moreover, existing laws and regulations, including the Health Insurance Portability and Accountability Act, Food and Drug Administration regulations, state licensure and professional standards, and laws barring employment discrimination, apply to AI. Boards should receive ongoing updates as the regulatory framework for AI changes.

Many health systems have adopted a standardized process for evaluating the risks of AI tools from multiple perspectives including patient safety and quality, legal, data acquisition and integrity, cybersecurity, clinician engagement, scalability and maintenance, and model type.^[16] A standardized risk assessment process adds rigor, ensures consistency, and helps document that the health system has exercised due diligence in vetting AI tools. Implicit in the assessment is the board's relative risk tolerance generally, and AI-related risk tolerance specifically. For instance, a board focused on competitive advantage may accept a higher level of risk in approving AI

tools, skewing the AI portfolio toward generative and agentic tools that have less human oversight. Another organization may be relatively less risk tolerant and only procure tools that have direct human oversight.

As a starting point, boards should understand the existing structures, policies, and practices to mitigate risk within the organization and the revisions and expansion required or implemented to encompass AI. Part of this assessment should include an understanding of how AI expertise has been embedded in each functional area, ranging from medical and nursing leadership, IT, legal and compliance functions, cybersecurity, and data privacy to mission/ethics. AI will require not only new competencies of management but also strong collaboration across disciplines and departments to safeguard the organization's interests.

As boards oversee AI risk, they should consider the following areas:

- **Policies and Procedures.** The board should establish an expectation that policies are being revised or created to address AI across the spectrum of clinical care, operations, and administration before any use is implemented.
- **Patient Safety and Quality.** While AI poses risks across many areas of health care operations, the risks to patient safety should receive special attention from the board. Boards should be informed about how use cases are vetted for any risk to the quality of care for patients.
- **Data Governance.** AI demands access to massive quantities of data for training purposes, can generate new intellectual property from data access, and facilitates the flow of data from one vendor to others. The board should seek a plan from management about policies and practices to protect patient, employee and company data in new transactions/arrangements and data sharing in prior arrangements at increased risk due to AI. Arrangements to monetize patient data or allow its use for training AI, even with de-identified data, should also come to the board given the risks posed to patient privacy, trust, and reputation.
- **AI Model Evaluation and Selection.** AI model selection requires cross-disciplinary cooperation and rigorous analysis of AI vendors and tools to preserve capital investment and facilitate successful integration into existing practices and systems.^[17] The board should seek confirmation that a process to evaluate and accept vendors with the capacity to meet the demands of AI acquisition has been implemented.
- **Liability.** AI has generated fundamental unresolved questions about liability. Although it is clear that physicians can face liability and disciplinary action for fully delegating clinical decisions to AI,^[18] the law is unsettled about whether AI developers and health systems will be held legally accountable when AI contributes to patient harm. Boards should understand this as another risk in AI adoption.
- **Audit and Monitoring.** To carry out the duties imposed by *Caremark*, the board should ensure that there is a process in place to audit and monitor AI use, with

serious events akin to patient safety sentinel events or major security breaches reported to the board.[\[19\]](#)

Conclusion

Beyond the more limited role set for board oversight by case law and statute, boards often serve a critical function in overseeing strategy, technology, capital investment, and other key areas. AI challenges boards to determine how they can best fulfill this role in relation to a technology that is complex, rapidly developing, and presents both uncertain promise and risk across fundamental aspects of an organization's health care mission and operations. The board, like the broader organization, should assess the strengths and weaknesses of its current practices and chart a path forward to meet the challenges that AI poses.

About the Authors

Tracy E. Miller, JD, is Counsel at Brach Eichler where she advises health systems, physicians, and other health care clients about governance, transactions, compliance, and contracts. Prior to joining Brach Eichler, she was EVP & General Counsel of a health system where she oversaw all legal matters, including advice regarding governance, transactions, system policies, AI, data privacy, cybersecurity, and regulatory affairs. She has extensive experience advising boards and executives, assisting them to realize their business, strategic, and operational goals. She previously served as Partner and Deputy Chair of the Health Care Practice at a law firm in New York City and as Executive Director of a New York State task force on ethics in medicine. She is the former Chair and current member of the Executive Committee of the Health Law Section of the New York State Bar Association and a member of the Advisory Council of The Hastings Center for Bioethics. She has published extensively in leading legal, health policy, and bioethics publications.

Charles E. Binkley, MD, FACS, is Director of AI Ethics and Quality at Hackensack Meridian Health and Professor of Surgery at Hackensack Meridian School of Medicine. He chairs the School of Medicine's AI Task Force and the Faculty Assembly, as well as the health system's Clinical AI Champions Group. Dr. Binkley is a nationally recognized expert in medical artificial intelligence and translational bioethics. He serves as an advisor for local and national organizations, including the Greater New York Hospital Association, the Institute for Healthcare Improvement, the American College of Surgeons, the Society of University Surgeons, the Association of American Medical Colleges, and The Hastings Center for Bioethics. Dr. Binkley's scholarly work focuses on the ethical integration of AI into clinical care, with particular emphasis on clinical decision-making, value-driven AI implementation, human-AI teaming, and accountability for harm. He is the author of *Encoding Bioethics: AI in Clinical Decision Making*.

[1] See Dylan Sandin et al., *Director Essentials: Implementing AI Governance*, Nat'l Ass'n of Corporate Directors (Sept. 2025) (NACD), <https://www.nacdonline.org/all-governance/governance-resources/governance-research/director-faqs-and-essentials/implementing-ai-governance/>.

[2] *In re Caremark Intern. Inc. Deriv. Litig.*, 698 A.2d 959 (Del. Ct. Ch. 1996); See also *In re Boeing Co. Deriv. Litig.* C.A. No. 2024-1210-MTZ, slip op. at 1-2, 21-22 (Del. Ch. Aug. 13, 2026); *Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019) (emphasizing the importance of the duty of care in relation to “mission critical” risks). For general discussion of board duties, see Tracy Miller and Valerie Gutmann, *Changing Expectations for Board Oversight of Health Care Quality: The Emerging Paradigm*, 2 J. of Health & Life Sciences 43-49 (July 2009).

[3] See, e.g., *In re Caremark* at 967; *In re Boeing* at 20-22, 43-44.

[4] For a discussion and recommendations for reporting to the board, see Sandin *supra* note 1, at 7-8, 14-15; Rob Gerberry and Michael Peregrine, *Health Care Corporate Governance, Board Oversight of AI, Part Two-What Does a Framework Look Like?* AHLA's Speaking of Health Law Podcast (Dec. 12, 2025), <https://www.americanhealthlaw.org/content-library/ahla-daily/2025/december/12/ahla-speaking-of-health-law-podcast-health-care>; Holly Gregory, *AI and the Role of the Board of Directors*, Harvard Law School Forum on Corporate Governance 6-7 (Oct. 2023), <https://corpgov.law.harvard.edu/2023/10/07/ai-and-the-role-of-the-board-of-directors/>.

[5] Depending on state law, committees that include non-board members are advisory only, and cannot exercise delegated board authority. See for example, N.Y. Not-for-Profit Corp. Law § 712(e) (McKinney 2026).

[6] For further discussion of board consideration of AI and business strategy, see Gregory *supra* note 4, at 1-2.

[7] Pope Leo XIV, *Magnifica Humanitas: On Safeguarding the Human Person in the Time of Artificial Intelligence*, The Holy See (May 25, 2026), <https://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html>.

[8] Rob Gerberry and Michael Peregrine, *Health Care Corporate Governance: Emerging Importance of the Human Capital Committee*, AHLA's Speaking of Health Law Podcast (Sept. 1, 2026), <https://www.buzzsprout.com/221709/episodes/19730564-health-care-corporate-governance-emerging-importance-of-the-human-capital-committeeHumanCapital>.

[9] For discussion of different AI governance approaches in health care systems see, Valerie Parker et al., *AI Governance in Health Systems: Aligning Innovation, Accountability, and Trust*, Duke-Margolis Inst. For Health Policy (Oct. 2024), <https://healthpolicy.duke.edu/sites/default/files/2024-10/AI%20Governance%20in%20Health%20Systems.pdf>; see also American Medical Association, *Steps Forward; Governance for Augmented AI* (Apr. 2025), <https://edhub.ama-assn.org/steps-forward/module/2833560>.

[10] Sandin *supra* note 1, at 9 (finding that 32% of directors who responded to an NACD survey identified “uncertain ROI/improvement” as a barrier to AI adoption).

[11] Sandin *supra* note 1, at 12 (noting that “Capital allocation is an important area of strategic oversight for boards” and recommending that boards scrutinize large AI expenditures).

[12] Lisa Rotenstein et al., *Changes in Clinician Time Expenditure and Visit Quantity With Adoption of Artificial Intelligence-Powered Scribes: A Multisite Study*. 335(16) J. Am Med. Ass’n. 1408 (2026).

[13] *Id.*

[14] Maliha George et al., *Artificial Intelligence and Liability in Medicine: Balancing Safety and Innovation*, 99(3) The Milbank Quarterly 629 (2021), <https://doi.org/10.1111/1468-0009.12504>.

[15] For discussion of federal and state AI regulations, see David Peloquin et al., *Paging Dr. Algorithm: Navigating the Regulatory Landscape for Agentic AI in the Health Care Industry*, AHLA’s Health L. Weekly (Feb. 2026), <https://www.americanhealthlaw.org/content-library/health-law-weekly/2026/february/20/paging-dr-algorithm-navigating-the-regulatory-land>; Akin Gump Strauss Hauer & Feld, *Artificial Intelligence in Clinical Decision-Making: Regulatory Roadmap and Reimbursement Strategies* (Dec. 11, 2025), <https://www.akingump.com/en/insights/alerts/artificial-intelligence-in-clinical-decision-making-regulatory-roadmap-and-reimbursement-strategies>; Ryan Brown, *No AI Law, No Problem: Pennsylvania Applies Its Medical Practice Act to an AI Chatbot*, Duane Morris Health Law Blog, (May 8, 2026), <https://blogs.duanemorris.com/healthlaw/2026/05/08/no-ai-law-no-problem-pennsylvania-applies-its-medical-practice-act-to-an-ai-chatbot/>.

[16] Charles Binkley et al., *An early pipeline framework for assessing vendor AI solutions to support return on investment*, 8(1) NPJ Digit Med. 368 (2025).

[17] *Id.* at 1-3.

[18] Fed'n of State Med. Bds., *Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice* at 5 (Apr. 2024), <https://www.fsmb.org/siteassets/advocacy/policies/incorporation-of-ai-into-practice.pdf>.

[19] *In re Caremark supra* note 2, at 968.