Chambers AND PARTNERS

CHAMBERS GLOBAL PRACTICE GUIDES

Digital Healthcare 2026

Definitive global law guides offering comparative analysis from top-ranked lawyers

USA – California: Trends and Developments

Eric Fish and Paul Deeringer
Hooper Lundy & Bookman

USA – CALIFORNIA

Trends and Developments

Contributed by:
Eric Fish and Paul Deeringer
Hooper Lundy & Bookman

Hooper Lundy & Bookman is the largest law firm in the US dedicated exclusively to healthcare. With over 60 professionals across six offices – Los Angeles, San Francisco, San Diego, Boston, Denver, and Washington, DC – HLB serves clients in all 50 states and two US territories from major health systems and hospitals to medical groups, managed care organisations, digital health companies, and medical device manufacturers. HLB's focus on healthcare enables a depth of expertise which is unmatched in general practice. The firm has shaped landmark appellate decisions, recovered billions of dollars for provider clients, and influenced federal and state healthcare legislation and regulation. Its capabilities span Medicare and Medicaid reimbursement, fraud and abuse defence, transactional work, compliance, litigation, and government relations, offering clients a seamless, multidisciplinary approach. Consistently ranked among Chambers USA's top health law firms nationally and in California, Massachusetts, and Washington, DC, HLB is trusted by the industry's most respected providers as a true long-term business partner.

Authors

Eric Fish is a leading digital health attorney whose practice sits at the intersection of regulation, policy, and innovation. Based in Washington, DC, he advises healthcare clients at HLB on compliance, privacy, enforcement risk, and the regulatory implications of AI and emerging technologies. Eric brings rare depth to his practice: as former chief legal officer of the Federation of State Medical Boards, he led the drafting of the Interstate Medical Licensing Compact and developed landmark guidance on AI in clinical settings. Clients value his ability to translate complex regulatory challenges into clear, actionable strategy, making him a trusted partner as the healthcare landscape continues to evolve.

Paul Deeringer is a partner in HLB's San Francisco office, advising health systems, medical groups, and health technology companies on complex transactions and regulatory matters, including digital health. He returned to the firm after almost 14 years in senior leadership at John Muir Health, a USD2.6 billion California integrated health system, most recently as senior vice president and chief strategy officer. That particular operating background gave Paul first-hand insight into how healthcare organisations evaluate and adopt digital health tools, and he combines legal and regulatory depth with executive judgment to deliver practical, efficiently scoped advice on technology-driven transactions and partnerships.

Hooper, Lundy & Bookman, P.C.

401 9th Street NW
Suite 550
Washington D.C. 20004

Tel: (+1) 202 905 2163 (E Fish)
(+1) 415 875 8514 (P Deeringer)
Email: efish@hooperlundy.com
pdeeringer@hooperlundy.com
Web: www.hooperlundy.com

Contributed by: Eric Fish and Paul Deeringer, **Hooper Lundy & Bookman**

AI Is Scrubbing In: Is Your Governance Ready?

What healthcare leaders should do now to govern AI before risk outpaces innovation

- AI is no longer a future-state issue for healthcare organisations; it is already embedded in clinical, operational, revenue cycle, documentation, and patient-facing workflows, often ahead of formal governance controls.
- Healthcare organisations should move quickly to inventory AI use across the enterprise, identify "shadow AI" activity, and establish a governance framework that addresses data use, safety reporting, bias assessment, validation, and workforce training.
- AI oversight should have board-level visibility and practical authority. Effective AI governance requires multidisciplinary councils with authority to approve, oversee, escalate concerns regarding, and retire AI tools, rather than councils limited to providing recommendations.
- Existing and future vendor agreements should be reviewed for AI-specific risk allocation, including data ownership and permitted use of patient data, model training restrictions, validation obligations, audit rights, liability caps, indemnity, and remedies for algorithmic errors.
- Clinical AI tools require heightened controls because clinicians remain responsible for patient care decisions. Policies should require human oversight, documentation of AI-assisted decision-making, training, and procedures for hallucinations, adverse events, and unauthorised tool use.
- AI governance is now a transactional, investment, and regulatory readiness issue. Organisations that build accountable programmes will now be better positioned for vendor negotiations, diligence, accreditation activity, enforcement scrutiny, and future AI-related partnerships or transactions.

Introduction

Artificial intelligence is rapidly reshaping every facet of the healthcare system, unlocking efficiencies in administration, and fundamentally altering care delivery. AI-enabled software is being used to streamline back-office operations, support billing functions, facilitate reimbursement requests and appeals, and respond to patient enquiries. Ambient documentation tools are being employed to document clinical encounters and populate patient charts in an effort to find ways to reduce administrative burdens on clinicians. Predictive algorithms are powering tools utilised by clinicians to guide decision-making. Consumer-facing applications are also helping patients understand, and in some cases independently treat, their conditions.

However, the development of proper oversight and governance frameworks has failed to keep pace with AI adoption. This regulatory lag significantly elevates the risk profile for organisations utilising these technologies. While industry data indicates that over 90% of healthcare organisations deploy AI in some capacity, fewer than half have implemented comprehensive strategies and formal policies to govern its use or have a dedicated environment for testing prior to deployment.

This maturity gap in governance and risk mitigation is affecting a disproportionate number of lower-resourced and rural hospitals. While large academic medical centres can have staff-dedicated AI governance committees and utilise teams of data scientists to audit algorithms for bias, many healthcare systems and independent hospitals often lack the capital and personnel for robust oversight. When adverse events occur, these structural disparities will be scrutinised as contributing factors to patient harm or noncompliance with state or federal regulations.

Governance, therefore, should be thought of as strategic imperative for the legal and ethical deployment of artificial intelligence in healthcare settings. Proper governance policies must be operationalised through a comprehensive framework that explicitly tailors oversight mechanisms to the unique parameters of the setting in which AI is used. Furthermore, organisations must proactively mitigate external liability by negotiating stringent vendor contracts that strictly define data usage rights, algorithmic performance standards, and indemnification obligations.

This shift is increasingly financial as well as operational. Health systems are no longer only licensees of AI tools – they are emerging as investors, pilot sites, and co-development partners for AI vendors. As a result, effective AI governance is more than a compliance exercise; it is a critical component of deal readi-

ness and will be increasingly important in satisfying the diligence requirements of investors, acquirers, and strategic partners.

Essential elements of an AI-ready governance strategy

Irrespective of size or the extent to which AI is integrated into care and operations, health systems should view the creation of robust AI governance structures as a strategic imperative, rather than a passing compliance measure. Properly designed and actively refined, these governance structures can transform AI from a potential operational liability into a secure, scalable clinical asset.

The essential elements of any healthcare organisation's strategy should focus on five major areas:

- system-wide AI governance frameworks;
- secure and transparent data use;
- institutionalised safety and oversight protocols;
- reducing the impact of bias on tool performance; and
- properly ensuring that personnel are suitably educated to use and understand AI.

1. Establishment of effective system-wide governance framework

A robust AI strategy mandates a formalised, risk-based governance framework overseeing the complete lifecycle of AI applications, from vendor selection and deployment through retirement. This framework must permeate all organisational functions, encompassing direct patient care, daily operations, and administrative support. Achieving this integration requires the dissolution of departmental silos and the formation of a multidisciplinary governance council. This body must include technical and cybersecurity experts alongside frontline providers, administrative personnel, executive leadership, and patient advocates. Furthermore, governing boards must incorporate AI utilisation and outcome metrics into their routine oversight, affording algorithmic performance the exact fiduciary scrutiny applied to financial performance.

In practice, the governance council must be vested with genuine decision-making authority and empowered to act efficiently and decisively. It should report into an existing board committee structure (typically finance or audit) with defined budget authority for governance activities, rather than operating as a standalone advisory body with no claim on institutional resources.

2. Securing data and ensuring transparency

Robust data privacy policies are more than something necessary for regulatory compliance. With concerns about how data are collected and used prevalent among patients, proper policies are foundational to maintaining patient trust. Healthcare organisations must operate with complete transparency, systematically educating staff and patients regarding the collection, utilisation, and potential repurposing of clinical data for AI development.

Operationally, this requires the following.

- Defining explicit permitted uses of exported data, including rights concerning monetisation and model outputs.
- Enforcing the principle of least privilege by limiting data exports to the absolute minimum required for a vendor's specific function.
- Prohibiting the re-identification of de-identified patient data through binding contractual clauses.
- Fortifying third-party agreements with comprehensive audit rights and strict compliance obligations.

3. Institutionalising safety and oversight protocols

Maintaining human-in-the-loop oversight over use of AI in patient encounters requires cultivating an institutional culture that encourages staff to voluntarily report adverse patient safety events, hallucinations, and performance issues. Capturing these occurrences in internal reporting systems enables the early detection of issues with AI tools. Well-designed reporting practices proactively insulate healthcare systems against legal vulnerabilities and support sustained innovation.

4. Mandating bias assessment and localised validation

The reliability of an AI tool depends directly upon the integrity and representativeness of its training data. Models trained on datasets unrepresentative of a hospital's specific demographic generate severe risks of

misdiagnosis, inequitable care delivery, and operational failure. Given that individual patient characteristics frequently dictate treatment efficacy, algorithmic equity constitutes a fundamental component of patient safety. Proper AI governance requires rigorous bias assessments prior to deployment. It also necessitates a critical evaluation of the underlying training data, forcing institutions to determine whether the tool underwent bias detection during development and whether the model is adequately tuned to the localised demographics of the specific community served.

These validation and bias-testing obligations carry real cost, and organisations should account for that cost in total cost of ownership at the procurement stage, not treat it as an unbudgeted afterthought once a tool is already in use.

5. Implementing targeted workforce education

Governance frameworks should also include comprehensive education of staff to facilitate the safe clinical use of AI. Organisations must establish clear, documented authorisation protocols defining role-based access permissions for specific AI applications. Prior to authorising providers or staff to use AI to support clinical care, they should demonstrate functional competency regarding the tool's capabilities and limitations, alongside a thorough understanding of institutional AI policies. Enterprise-wide educational initiatives that establish a foundational knowledge base will ensure the consistent and safe application of AI technologies across the workforce.

Using governance to address risk in vendor contracts

With AI becoming embedded as a standard capability in both administrative and clinical tools, it is paramount to rethink the approach to dealing with vendors of AI tools. Traditional software-as-a-service (SaaS) agreements are insufficient to capture the unique risks of generative AI and machine learning. At a minimum, both existing and future contracts should align with organisational governance structures to ensure proper monitoring of the risks associated with the tool. Patient-facing, clinical tools create more risk of harm and, as such, contracts should be reviewed more frequently and in more detail than administrative tools.

This analysis becomes more complex where the health system's relationship with the vendor extends beyond a purely transactional one. Health systems increasingly take equity stakes in AI companies, serve as pilot or proof-of-concept sites, or co-develop tools alongside a vendor – arrangements that can place the system in the position of investor, customer, and data source simultaneously. Where any of these roles intersect with a vendor that also receives referrals or business from system-affiliated physicians, the arrangement warrants a Stark Law and Anti-Kickback Statute remuneration analysis, including independent fair market value support for the equity terms, data-sharing value, or pilot-site compensation involved.

As an initial step towards proper governance, healthcare organisations should conduct a systematic review of existing contracts with third-party vendors. Many existing agreements will lack provisions addressing emerging standards regarding data access, ownership, or permissible use in connection with AI technologies. Many of these agreements may also preclude vendor accountability through legal disclaimers, suffer from limited representations and warranties and severe limitations on vendor liability, and have other limitations on available remedies. Early identification and remediation of existing contractual gaps will mitigate legal and operational exposure and better position an organisation for future negotiation. For existing contracts coming up for renewal or extension, such timing may present a natural opportunity to address these issues. For contracts still early in their term and/or with auto-renewal provisions, if the vendor is unwilling to affirmatively address these issues through a contract amendment, it may be worth evaluating whether existing contractual language contains material change or other provisions that could be read broadly enough with respect to newfound AI capability integration to force a renegotiation.

As part of the procurement process, it is key to require disclosure of detailed information about the testing and validation of AI tools, the evaluation of bias, and probe whether the willingness of the vendor to refine the tool once deployed in the organisation's ecosystem. Governance policies should also establish non-negotiable boundaries tailored to AI. This includes:

- customised liability caps that specifically address algorithmic errors;
- strict data residency and usage rights, which prevent vendors from using protected health information to train commercial models without consent;
- an uncapped intellectual property indemnity that serves as a protection against AI tools generating outputs that infringe on the intellectual property of others; and
- proactive monitoring with regular validation and testing, comparing AI outputs to known sets of performance data, assessing use-case relevant outcomes, and user confidence in the tools.

Tailoring governance for AI used in clinical settings

The adoption of AI in clinical practice necessitates that any governance policy provides sufficient human accountability. Legally and ethically, clinicians remain entirely responsible for the use of, and any outcomes associated with, AI-generated treatment or documentation.

To mitigate risk to the clinician and employing organisation, policies should require proper documentation of how AI was factored into the delivery of care. Should a physician's prescribed course of treatment align with the AI's recommendation, the clinical reasoning validating this concordance must be explicitly documented. Conversely, if a practitioner deviates from an AI-delineated pathway, they must document the clinical rationale for that choice. Such deviations are typically justified when adherence to the algorithm's recommendation would fail to uphold the established standard of care. While the proliferation of AI offers substantial clinical utility, the failure to apply independent clinical judgment to its outputs exposes practitioners to significant professional liability, can compromise patient safety, and may run afoul of a developing set of state laws requiring human oversight of all AI uses in healthcare.

Where AI tools measurably change physician productivity or throughput – eg, by reducing documentation time in ways that increase capacity for patient volume – organisations may need to revisit compensation plans built around productivity or wRVU benchmarks. Organisations should confirm that any resulting compensation remains commercially reasonable and consistent with fair market value, and should avoid incentive structures that could be read as rewarding AI-influenced referral or utilisation patterns. Where an AI clinical decision-support tool functionally substitutes for physician judgment, organisations also should consider whether its use raises corporate-practice-of-medicine concerns, particularly for management service organisations (MSOs) or other non-physician-owned entities deploying the tool on behalf of affiliated providers. This is a particularly acute concern given the recent resurgence of corporate practice enforcement, with regulators increasingly scrutinising arrangements that allow non-physician entities to influence clinical decision-making.

Physician alignment structures should extend this same rigour to medical staff governance. Privileging and credentialing processes should treat competency with system-approved AI clinical tools the way they treat any other significant new technology or procedure, requiring documented training and, where appropriate, tool-specific sign-off through the medical staff office rather than leaving adoption to informal practice. Professional liability carriers are also beginning to ask about AI governance and physician training as part of underwriting, and a documented alignment and credentialing framework can support more favourable coverage terms. This trend supports greater coordination of such work with the organisation's risk-management and insurance functions, not just its clinical and legal teams.

Rapid adoption of ambient medical scribes also requires proper governance protocols to sidestep the legal issues that are emerging. These tools capture clinician-patient conversations and generate draft clinical notes for clinician review. Preliminary study findings suggest that these tools may reduce documentation burden and improve patient-provider communication, indicating potential to mitigate physician burnout. However, if a physician fails to catch and correct a false or hallucinated clinical detail generated by the scribe, traditional liability still attaches to the human provider whose signature is on the chart. These tools also have complex issues surrounding patient consent and data management, and organisations deploying the tools have been subject to lawsuits around the failure to obtain consent.

An equally pressing vulnerability to address in any governance policy is the unauthorised use of generative AI tools that have not been approved. Referred to as "shadow use" of AI, the use of consumer-grade, non-HIPAA-compliant tools – often on personal devices, by clinicians and administrative staff to summarise patient histories, draft appeal letters to insurers, or translate discharge instructions – creates immense legal risks. These risks include greater vulnerability for data privacy breaches, compliance blind spots, and an open door for fraud, waste, and abuse events.

Mitigating these risks associated with AI deployed in a clinical context requires governance policies to address these issues but also to provide for continuous clinician education on both the capabilities and the strict limitations of AI in a clinical setting. Also essential is a reminder that use of AI does not absolve the clinician from any associated ethical or professional responsibilities.

AI governance as a transactional and investment imperative

The scale of capital now flowing into healthcare AI makes this a transactional issue as much as a compliance one. US digital health companies raised USD14.2 billion in 2025, a 35% increase over 2024, with AI-enabled companies capturing 54% of all funding (see 2025 Year-End Digital Health Funding Overview: A Tale of Two Markets, Rock Health (12 Jan 2026)). This pace continues in 2026, with USD7.4 billion invested across 244 deals in the first half of the year, including 20 megadeals of USD100 million or more from just 19 companies, representing 45% of all capital invested (see H1 2026 Funding and Market Overview: Durable Roots, Shifting Routes, Rock Health (13 July, 2026)). Health systems are participating in this capital formation directly, not only as customers. A number of large systems now operate dedicated venture funds focused on digital health and AI, including Mass General Brigham's AI and Digital Innovation Fund and venture arms affiliated with Providence Health and MemorialCare (see 2026 Digital Health Investors Directory, digitalhealth.com). Health systems including Corewell Health, Memorial Hermann, Duke, and Intermountain Health have also appeared as co-investors alongside traditional venture capital in recent AI-enabled care delivery financing rounds (see Marissa Plescia, 4 Notable Health Tech Funding Announcements in June, MedCity News (29 June 2026) (as aggregated by Digital Health VC Hits $7.4B in H1 2026, MarketScale (2026)).

This dual role changes the governance calculus. Building on the contractual analysis above, the way an investment itself is structured – as a passive minority stake, an active joint venture with board or information rights, or a straight commercial pilot – carries different Stark Law and Anti-Kickback Statute considerations, and should be driven by the underlying business relationship rather than administrative convenience. Governance should also ensure that the investment decision and the clinical procurement decision run through separate channels and separate documentation, so that neither influences the other.

These same capital trends make AI governance a diligence and valuation issue in M&A and affiliation transactions. Digital health M&A activity rose 61% to 195 deals in 2025, and the pace has continued into 2026 (see 2025 Year-End Digital Health Funding Overview: A Tale of Two Markets, Rock Health (12 Jan 2026)). As acquisitions of AI-enabled platforms and of provider organisations with established AI programmes become more common, a target's governance maturity – its vendor contract terms, its bias-testing practices, its documentation of clinical oversight – is becoming a real driver of deal risk and valuation, and not a technical matter to resolve after closing.

Governance now addresses tomorrow's issues

The ability to manage the integration of AI through these immediate actionable steps will be critical for all health systems. It is likely that, over the next 18 months, there will be greater regulatory focus on such things as AI governance requirements, vendor procurement policies, and more stringent enforcement of data privacy requirements. With AI heavily automating high-stakes areas like claims processing and denials management, proactive strategies establishing proper oversight over procurement and use of AI tools will

be essential to mitigate exposure to false claims and fraud, waste, and abuse actions.

As with any operational efficiency claim, projected labour savings or productivity gains from AI tools should be independently validated before they are built into budgets or financial forecasts, with finance leadership involved alongside legal and clinical stakeholders in that validation.

Furthermore, organisations that invest in AI governance today may be better positioned if such frameworks become embedded in future accreditation and regulatory requirements. In June 2026, The Joint Commission launched its Responsible Use of AI in Healthcare Certification programme. This voluntary designation recognises healthcare organisations that have implemented appropriate governance structures, risk management safeguards, monitoring processes, and workforce education to support the responsible use of AI. Although voluntary, the programme reflects a broader trend toward formalised accountability for AI deployment in healthcare. As regulators, payors, and accrediting bodies continue to develop AI oversight mechanisms, certifications of this type could become an increasingly important benchmark and may ultimately inform future regulatory or accreditation standards.

Conclusion

Artificial intelligence is no longer a future consideration for healthcare organisations. It is already embedded in clinical workflows, administrative operations, revenue cycle functions, and patient engagement strategies. The question facing health systems is not whether AI will be used, but whether it will be deployed under a governance framework capable of managing the legal, ethical, operational, and patient safety risks that accompany it.

As federal and state policymakers, accrediting bodies, and enforcement agencies increase their focus on AI oversight, healthcare organisations should expect heightened scrutiny of procurement practices, data governance, validation protocols, and human oversight requirements.

Properly designed AI governance mitigates legal risk, empowers care delivery, and supports business operations. A clear framework for evaluating tools before deployment and monitoring performance after implementation is critical to this effort, as are policies that establish clear expectations of both users and vendors. Appropriately tailored governance frameworks can provide a stable foundation for building systemic trust in healthcare AI while minimising associated risks.

Approached this way, governance is not only a risk-mitigation measure, but also a value-creation and deal-readiness function, positioning organisations to pursue AI-related investments, partnerships, and transactions from a place of strength.

CHAMBERS GLOBAL PRACTICE GUIDES

Chambers Global Practice Guides bring you up-to-date, expert legal commentary on the main practice areas from around the globe. Focusing on the practical legal issues affecting businesses, the guides enable readers to compare legislation and procedure and read trend forecasts from legal experts from across key jurisdictions.

To find out more information about how we select contributors, email Luke.Wilson@Chambers.com