NAVIGATING RISKS IN THE PROCUREMENT OF THIRD-PARTY AI TOOLS

by Robert Kantrowitz, Shubha Lakshmanan, Valerie Rock, and Sean Sullivan

Robert Kantrowitz
(robert.kantrowitz@kirkland.com) is a Partner at Kirkland & Ellis LLP in New York, NY.

Valerie Rock
(vrock@pyapc.com) is a Principal at PYA P.C. in Atlanta, GA.

Shubha Lakshmanan
(slakshmanan@waudcapital.com) is a Senior Director of Compliance & Privacy – Portfolio Operations at Waud Capital Partners in Chicago, IL.

Sean Sullivan
(sean.sullivan@alston.com) is a Partner at Alston & Bird in Atlanta, GA..

Artificial intelligence (AI) is reshaping healthcare, offering solutions that span clinical, operational, and administrative domains. As adoption accelerates, compliance professionals face the dual challenge of harnessing AI's benefits while managing multifaceted risks, especially when these tools are sourced from third-party vendors.

This article explores the evolving regulatory landscape, potential risks in development and deployment of third-party AI tools, and key risk mitigation strategies in the following areas:

1. Due diligence of AI vendors/developers
2. Contract negotiation with vendors/developers
3. Continuous auditing and monitoring of third-party AI tools

Notable use cases of AI in healthcare

AI's integration into healthcare is no longer theoretical. Providers leverage AI-driven tools to enhance diagnostics, personalize treatment, streamline workflows, and improve patient engagement. For example, AI-powered diagnostic systems have demonstrated the potential to outperform human radiologists. AI is also optimizing administrative functions by reducing documentation time and appointment no-shows, automating eligibility and prior authorization determinations, and targeting preventive care.

Beyond diagnostics, AI supports real-time patient monitoring, automates billing and claims management, and even assists in drug discovery and clinical trial recruitment. Generative AI and large language models are increasingly used for clinical documentation, patient communication, and knowledge management, offering both efficiency and new compliance challenges.

Rapidly evolving regulatory landscape

This rapid innovation is occurring against a backdrop of evolving regulation. The regulatory environment is dynamic, with federal agencies like the U.S. Food and Drug Administration (FDA), U.S. Department of Health and Human Services, and the Federal Trade Commission issuing guidance on transparency, risk management, and patient safety. States are also enacting their own AI laws, such as California's SB1120, which mandates human oversight and fairness in AI-driven healthcare decisions. In the U.S., the FDA has approved nearly 1,000 AI-driven medical devices, with more anticipated.[1] Internationally, frameworks like the European Union AI Act set additional standards for transparency and accountability.

Healthcare organizations must also comply with longstanding general privacy laws applicable to healthcare, such as HIPAA, General Data Protection Regulation, and the California Consumer Privacy Act, which impose strict requirements on the use and sharing of patient data, especially when handled by third-party vendors.

At the same time, the Trump administration's executive order "Removing Barriers to American Leadership in AI"[2] revoked the Biden administration's "Safe, Secure, and Trustworthy AI."[3] Trump's executive order emphasizes deregulation and innovation but leaves organizations largely on their own to manage transparency, accountability, ethics, and compliance with existing regulatory risks.

To further complicate this regulatory maze, litigation and enforcement are increasingly setting precedents under the Unfair and Deceptive Practices Act and similar laws at both the federal[4] and state levels,[5] focusing on the representations that companies make to the public regarding their AI products.

Potential risks of third-party AI tools

While AI offers transformative potential, it also introduces significant risks, particularly in the use of third-party solutions:

- **Patient safety and efficacy**: AI-powered diagnostic or treatment tools are only as reliable as the data and algorithms behind them. If a third-party AI system is trained on incomplete, outdated, or nonrepresentative data, it may generate inaccurate results. For example, an AI tool that underperforms for certain demographic groups could miss early signs of disease or recommend inappropriate interventions. Placing undue trust in AI recommendations without sufficient human oversight can potentially endanger patient safety and increase malpractice liability for the provider.
- **Fraud, waste, and abuse**: AI tools are increasingly used by both payers and providers to code medical procedures, recommend treatments, obtain prior authorizations, and evaluate appropriateness of billed charges. While these tools can save time on documentation, mistakes and inaccuracies can cause improper coding decisions and medically unnecessary procedures or prescriptions, thus leading to fraud and abuse risk under payer rules and federal and state false claims acts.

AI-powered diagnostic or treatment tools are only as reliable as the data and algorithms behind them.

- **Data security and privacy**: AI models thrive on data. The more data the model ingests, the more accurate and useful it can become. However, healthcare data is particularly sensitive, and privacy risks are magnified with large language models and generative AI. Personally identifiable information, protected health information, and other sensitive data that are fed to AI models should remain private and secure. Generally, AI outputs should not produce personal information (unless required or authorized). AI algorithms could also potentially collect and analyze patient data without explicit consent, especially when derived from electronic health records (EHRs) or collected passively through wearable technology.

- **Bias and health disparities**: Depending on the training data used, AI models can exhibit certain biases towards different populations. Algorithms trained on lighter-skinned patients have been reported to fail to accurately diagnose skin cancer in darker-skinned patients and demonstrate similar biases against historically underserved, poorer populations, despite often having the greatest healthcare needs.
- **Hallucinations**: AI hallucinations can recommend incorrect dosages or invent side effects and lead busy physicians to order unnecessary or even harmful treatments. Texas's enforcement in 2024 against a Dallas-based healthcare technology company demonstrates the real-world consequences of hallucinations.[6]
- **Legal exposure**: Recent class action lawsuits have alleged that insurers' use of AI to automate claim denials violates good faith obligations and consumer protection laws.[7] Lack of explainability, transparency, or human oversight can further increase liability for both third-party vendors and companies that deploy their products.
- **Explainability, transparency, and trust**: Third-party AI models can sometimes function as "black boxes" without a clear understanding of how they produce outputs. This can be problematic in clinical contexts where explainability is essential for treatment decisions and reimbursement. If a provider cannot understand how an AI tool arrived at a diagnosis, it may be difficult to defend malpractice claims and regulatory inquiries. As AI becomes more embedded in healthcare workflows, organizations must also manage unique operational, legal, and ethical risks around explainability. Explainability refers to an AI system's ability to communicate how it reaches a specific outcome. In clinical settings, reliance on opaque or black box models can complicate validation, erode trust, and introduce safety and efficacy concerns, even when the outputs appear statistically accurate.
- **Integration and workflow disruption**: AI tools may not integrate well with existing EHRs, workflows, and legacy technology, especially for smaller providers with limited IT resources and staffing. As a result, AI adoption can lead to disruptions, delays, and complications in day-to-day operations, and potentially even lead to clinical risk.
- **Billing and reimbursement**: Virtual health assistants, in both human and bot form, are becoming more prevalent. Bringing healthcare to the patient improves access to care and reduces expensive hospitalizations. However, legacy reimbursement models often do not recognize or cover these innovative and beneficial services, even as they become increasingly common.

For example, an AI virtual assistant that uses motion tracking to monitor a patient's at-home rehab exercises and provide feedback may not be reimbursed by payers, requiring providers to charge the patient directly for this service. In fact, payers may offer these services directly to their members, or patients may utilize third-party apps on their phones that they pay for directly, regardless of the treatment plan.

As providers and their patients increasingly use these tools to reach their treatment goals, payers will need to determine how to value these services and whether the same services should be valued differently when performed by a human or an AI model. Similarly, because the AI bot is not a licensed and credentialed provider, then a physician or other practitioner should generally be responsible for the service and monitor the accuracy and efficacy of the tool. While the AI performs the service, provider organizations should monitor reimbursement trends and determine how to compensate the physician for that oversight and supervision. Likewise, the organization should ensure its professional liability insurance covers issues that could arise from errors in the AI.

Providers should consider payer coverage, state licensure requirements, physician compensation arrangements, malpractice coverage and cost, as well as the evolution of reimbursement models when evaluating AI- or bot-based virtual health services. Communicating clearly to payers and patients the services being provided and the responsible party for the cost mitigates the risk of false claims liability and potential refunds to payers.

The complexity of these risks demands robust compliance oversight, and considering the AI governance strategies described below, to mitigate risks that are inherent to advanced AI systems used in healthcare.

Vendor due diligence

To address these challenges, organizations should adopt a structured AI vendor due diligence process under their risk management framework. This five-phase approach includes: (1) determining needs, alignment, and risks; (2) conducting vendor selection and due diligence; (3) establishing clear contracting terms; (4) executing onboarding and implementation; and (5) continuously monitoring and assessing deployed tools.

Before engaging and performing due diligence on vendors and their respective AI tools, organizations should identify clinical or operational goals and ensure alignment with broader strategic priorities. Understanding the applicable goals and priorities will guide the due diligence process, including the type and level of risks to consider and the scope and depth of the review. For example, procuring AI for clinical support will generally have greater and different risks (e.g., patient safety) than for administrative functions (e.g., operational costs and efficiencies). The due diligence process should, in turn, reflect these variations in risks and needs and will aid parties in understanding the scope of due diligence at the outset. A detailed request for proposal should be issued, outlining technical, ethical, and regulatory requirements.

At a high level, the due diligence may cover both the tool and the AI vendor's posture with respect to data sensitivity (from a privacy, security, trade secret, and intellectual property (IP) perspective), regulatory exposure (e.g., whether the vendor has received proper approvals or meets the applicable legal requirements), and downstream impact (e.g., on patients or business operations). Through these lenses, procuring organizations would evaluate the model development process, data sources, bias and related mitigation strategies, and explainability and validation methods, among other aspects. For example, the procurer would inquire about the rights the vendor received to use the relevant data, what and how often tests were performed, and if and how the vendor is able to validate or explain the tool's outputs.

Contract negotiations with AI vendors

Negotiating contracts with AI vendors is not unlike contracting with other Software as a Service (SaaS) providers; however, certain contractual provisions become even more important when those contracts involve software that can replicate human thought processes, such as clinical support. While issues such as payment terms, dispute resolution, basic representations and warranties, and termination rights may look similar to other SaaS agreements, healthcare providers and their counsel should carefully consider the following issues when contracting with AI vendors:

Scope of services, description of the AI solution, and performance metrics

Many AI solutions are customized to the buyer's specific business and needs. For this reason, the contract should clearly define the scope of work, including the specific AI solutions to be provided, the expected deliverables, and the timeline for implementation. If the developer will be training the AI model during implementation, consider incorporating milestones to ensure that it is performing as expected and to monitor accuracy, reliability, and efficiency. To further mitigate risk, consider including contract remedies such as credits, unpaid extensions, or termination.

Testing and incident management

Expect AI vendors to provide regular performance reports and continuously accept feedback and complaints from customers; vendors should ensure that feedback will be used to refine and improve the model. The vendor should independently test the AI model for accuracy, bias, fairness, and representativeness of the AI outputs, ensuring that both the AI and the output do not infringe or misappropriate third-party IP rights or violate applicable laws. Many AI vendors will accept contract provisions that require the vendor to engage a reputable third-party testing firm to conduct annual reviews. If any "AI issues" are discovered, then the vendor should be obligated to notify the customer, assist the customer in preventing or rectifying such issues (especially to the extent that they degrade the services or could potentially harm the healthcare provider or its patients), and take appropriate steps to notify any other potentially affected persons.

Privacy and security

One of the key differences between advanced AI models and other software solutions is the vast amounts of data needed to develop and operate the AI. The AI must ingest massive amounts of information, and the more data available, the more accurate and effective the AI will be. However, in a healthcare setting, that data is typically shielded by various federal and state privacy protections, including HIPAA, consumer privacy laws, and wiretapping laws. However, the data pools needed to feed AI models magnify the security risks.

Various mechanisms exist to mitigate these privacy and security risks, including strong representations, warranties, and covenants about security measures and IT infrastructures (including for their subcontractors), as well as indemnification provisions for data breaches. Likewise, from a privacy perspective, a combination of business associate agreements (BAAs) and data protection agreements (DPAs) express patient authorization or consent, and use of de-identified data can allow the AI model to have the necessary data to be effective. However, the contract should prohibit the AI vendor from attempting to re-identify data and allow the data to be used only for the express purposes described in the patient's consent, BAAs, and DPAs.

IP and confidentiality

Generative AI is different from other software solutions in that the more it is used, and the more data it is given, the more it knows, learns, and improves. But what if a provider does not own the rights to the data inputs fed into the AI model? Or what if that data has certain additional privacy protections? In many cases, the provider may not be allowed to turn that data over to the AI model. Sensitive data includes prompts containing proprietary business information, asking that it analyze the data to perform a task or evaluate a competitive strategy, radiology images, and patient voice recordings used in ambient listening technology. While the AI vendor may want access to this data to train its model, providers should be wary of who owns the data, what privacy laws may restrict its use, and what rights the AI vendor has to use the input data (including to train its models).

Similarly, output data can be just as valuable, and while AI vendors may request access rights to the AI outputs, this data should be owned by the end user: the healthcare provider customer. Given the vendors' lack of control over the inputs and questionable ownership rights of the AI model training data, some AI vendors avoid granting ownership rights of the outputs to their customers. In these cases, the AI vendor may include a provision that the vendor does not own the output data without expressly granting ownership of the data to the customer, but will still insist on the customer granting the vendor a license to that data for monitoring, quality assurance, and continued AI model development and training.

Healthcare providers should push for full ownership rights over the outputs and limit vendors' access and use of the outputs to specific purposes, such as providing technical support or services to the customer. Consider granting only time-limited access rights, such that the AI vendor may retain both inputs and outputs only for the duration of the user's session or for only seven days. While the AI model may not be able to "learn" from the customer's data, it will retain the most recent inputs and outputs for a short period, after which the provider will be responsible for maintaining any data in a separate system.

Finally, healthcare providers should obtain representations and warranties, as well as express indemnification, from the AI vendor that the AI solution and the data used to train the solution do not and did not infringe on any third-party rights. While the AI vendor may be unwilling to expressly grant ownership of the AI outputs to its customer, the vendor should at least indemnify the customer from any claims that the AI model used data without appropriate data rights.

Training and support

The vendor should be obligated to assist in adequate training and support to effectively use generative AI technology. Contracts should specify the training programs and resources that vendors will provide, including ongoing support and troubleshooting. Also, consider incorporating service level agreements that define expectations, including model uptime, response times, the quality of support services, and remedies for noncompliance.

Changes in regulatory landscape

Finally, given the fast-paced evolution of generative AI and regulatory oversight of AI, contracts should address processes for monitoring the legal landscape, communicating changes to all parties, and amending the contract to maintain its primary purpose while complying with updates to state or federal laws.

AI tool auditing and education

After conducting due diligence on the tool, providers should verify that it performs as expected and is subject to appropriate oversight; they should also continue to monitor and assess their compliance and ethical use. Oversight should be layered, from the user level to departments like IT and compliance.

Clinical decision support tools are a helpful example to understand the importance of auditing and education. Providers may adopt them independently for EHR integration or use versions offered by EHR vendors. These tools are generally used for decision support and not as the decision-maker.

The tool should integrate with the EHR using a standard electronic format (e.g., HL7) and allow providers contemporary access to structured and unstructured clinical data from medical records (via natural language processing). It should rely on current medical literature and emphasize extraction over excessive summarization. Citations should link to original sources, and treatment rationale should be included to ensure traceability and explainability.

Physicians should be trained on methods of effectively asking the tool questions, otherwise known as queries or prompts. "Prompt-engineering" refers to crafting questions that guide the AI tool to identify accurate and concise information from large datasets. Ongoing training should cover prompt engineering and how to detect inaccuracy, bias, or drift from the appropriate meaning. Instruction should also reinforce emotional quotient, bedside manner, sensitivity to social determinants of health, and cultural context. Regularly survey physicians on usability, accuracy, bias, reliability, and integration with workflows.

Medical directors and clinical leadership should oversee tool use from clinical and ethical perspectives, such as through peer review programs. Reviews should assess the level of reliance on the tool, including adherence to or deviation from standards of care and applicable guidelines (e.g., FDA). Audit trails may be needed to evaluate the tool's recommendations against the physician's decision-making.

IT and HIPAA oversight teams should monitor security, encryption, data storage and access, open vs. closed data sources, penetration testing, and user risks. Security updates may change data handling or encryption methods, which should be reviewed regularly—at least annually—and during updates. Any changes should be reported to the legal, compliance, or contracting teams to assess related risks.

Conclusion

As AI becomes increasingly integrated into healthcare, organizations must balance innovation with responsibility. The use of third-party AI tools introduces complex regulatory, operational, and ethical challenges that demand proactive governance. By implementing structured due diligence, robust contracting, continuous monitoring, and comprehensive training, healthcare providers can harness AI's potential while safeguarding patient safety, data privacy, and legal compliance. CT

Endnotes

1. Vijaytha Muralidharan et al., "A scoping review of reporting gaps in FDA-approved AI medical devices," *NPJ Digital Medicine* 7, no. 1 (2024): 273, https://www.nature.com/articles/s41746-024-01270-x.
2. The White House, "Removing Barriers to American Leadership in Artificial Intelligence," Presidential Action, January 23, 2025, https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/.
3. Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, Executive Order 14110, 88 Fed. Reg. 75,191 (Nov. 1, 2023), https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
4. Ken Paxton, Attorney General of Texas, "Attorney General Ken Paxton Reaches Settlement in First-of-its-Kind Healthcare Generative AI Investigation," news release, September 18, 2024, https://www.texasattorneygeneral.gov/news/releases/attorney-general-ken-paxton-reaches-settlement-first-its-kind-healthcare-generative-ai-investigation.
5. Federal Trade Commission, "FTC Announces Crackdown on Deceptive AI Claims and Schemes," news release, September 25, 2024, https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes.
6. Ken Paxton, Attorney General of Texas, "Attorney General Ken Paxton Reaches Settlement in First-of-its-Kind Healthcare Generative AI Investigation."
7. Lauren Clason, "AI, Algorithm-Based Health Insurer Denials Pose New Legal Threat," Bloomberg Law, April 8, 2025, https://news.bloomberglaw.com/daily-labor-report/ai-algorithm-based-health-insurer-denials-pose-new-legal-threat.

Takeaways

- Artificial intelligence (AI) in healthcare offers transformative benefits across diagnostics, administration, and patient engagement, but also introduces significant risks, especially when sourced from third-party vendors.
- Regulatory oversight is rapidly evolving, with new guidance and laws, making compliance a moving target.
- Key risks include patient safety, data privacy, algorithmic bias, and legal liability, particularly under fraud and abuse laws.
- Effective risk mitigation requires a structured framework for vendor due diligence, contract negotiation, and ongoing auditing and education.
- Clear communication with payers and patients about AI-driven services and cost responsibilities is essential to reduce reimbursement issues and legal exposure.