AMERICAN HEALTH LAW ASSOCIATION

February 20, 2026

Paging Dr. Algorithm: Navigating the Regulatory Landscape for Agentic AI in the Health Care Industry

February 20, 2026

David Peloquin, Ropes & Gray LLP | **Hillary Kalay**, University of California |
Allison Trimble, BJC Health System | **Jake Fulton**, Ropes & Gray LLP

Looking for a new doctor? An AI agent can find the physician and schedule your first appointment. Need a prescription re-filled? An AI chatbot can process the refill for you. Seeking prior authorization for a procedure? The AI agent will help your health plan determine whether you will receive care. Having trouble finding a therapist? The chatbot will see you now.

Over the past several years, artificial intelligence (AI)—and more recently, AI acting independently in an "agentic" manner—has rapidly proliferated in the health care industry with use-cases first reaching core administrative functions and now, increasingly, clinical functions. This growth has understandably drawn interest and investment by stakeholders seeking to harness its efficiency-generating potential and the attention of lawmakers at all levels of government aiming to limit the risks AI poses to patients' health, safety, and privacy.

This article provides a definition of agentic AI and surveys a variety of health care industry use-cases in both clinical and administrative applications. It then turns to a discussion of the legal and regulatory responses to these developments in the United States, detailing tensions as the Trump administration seeks federal preemption of nearly all state-level AI laws and states continue to enact laws regulating AI to address myriad concerns including health, safety, consumer protection, and civil rights. Following a discussion of the status of AI regulation, the article briefly considers enforcement action and industry guidance, then concludes with commentary on how health care industry stakeholders can build AI governance structures and update workflows to maintain compliance and flexibility as the regulatory landscape continues to evolve.

Agentic AI in Health Care

Agentic AI refers to a subset of AI systems designed to handle complex workflows and autonomously take actions or sequence tasks in real time with limited or no human prompting, distinguishing them from more passive, prompt-driven assistants like chatbots. Agentic AI tools are characterized by capabilities beyond traditional language modeling or other AI models that are designed for the completion of a single task in a routine manner.[1] For instance, agentic models are designed with features such as the ability to operate with autonomy and adaptiveness to changing settings, or sequencing and completing a number of tasks to achieve unique goals provided varying context without human input or intervention.[2] Due to this comparatively high level of independence, agentic AI tools have begun to be developed for and adopted in a number of industries, including finance, accounting, software development, and health care.[3]

In the health care industry, these tools are deployed in administrative settings, with the goal of reducing administrative burden in terms of both time and cost on health care providers, payers, and patients. For instance, appointment scheduling platform Zocdoc has launched a voice AI agent that interfaces directly with patients and resolves roughly two-thirds of scheduling calls without human intervention, lightening the load of human customer service representatives and thereby reducing wait times for patients.[4] As another example, health tech company Cedar has employed a voice agent trained on its billing data that patients can call to address questions about their medical bills.[5]

Agentic AI systems are distinguished from generative AI by their capacity to operate with greater autonomy. Agentic systems pursue goals through multi-step planning and execution, make decisions, and take independent actions across connected systems without continuous human guidance.[6] In the health care context, the distinction can be illustrated through clinical documentation and prior authorization workflows. Ambient AI scribes—which listen to patient-clinician conversations and

generate clinical notes for physician review prior to entry into the electronic medical record—are examples of generative AI, because they produce content based on input and require human validation before action is taken.[7]

By contrast, agentic AI systems are being piloted to assist with prior authorizations by autonomously comparing payer requirements to patient information, prompting physicians to gather missing documentation, and completing and tracking submissions across payer and provider systems.[8] They are also being used to assess medical images to identify conditions and make recommendations for clinical care.[9] These and other clinical use-cases for agentic AI have been identified as opportunities for providers to free up time for more complex patient care activities; however, many critics of AI in the clinical context are skeptical of AI's place in the exam room as such instances may test the line between augmentation of clinical capacity and substitution of professional judgment. Recognizing both the potential benefits of agentic AI—including enhanced diagnostic accuracy, greater operational efficiencies amid rising costs and reduced funding, and relief from clinician burnout—and the importance of preserving professional judgment, industry stakeholders are actively exploring, developing, and implementing these tools in the health care context. Consistent with the AI boom in other industries, lawmakers have begun to regulate in this area.

Federal Response: Preemption to Avoid a "Patchwork" of AI Regulations

To date, Congress has not enacted legislation to regulate AI broadly. While there have been a few proposed bills in the 119th Congress related to the use of AI in health care, they represent narrow regulatory efforts and have not progressed toward being enacted.[10] Rather, the majority of activity related to AI regulation by the federal government has come from the Executive Branch. The Trump administration has signaled a clear preference that AI should be regulated by the federal government rather than the states, but it has been less clear about the details of such regulation. In July 2025, the White House released a report titled *Winning the Race: America's AI Action Plan*,[11] which included three pillars: Accelerate AI Innovation, Build American AI Infrastructure, and Lead in International AI Diplomacy and Security. The first listed group of policy aims is to "remove red tape and onerous regulation." This aim includes rejecting the Biden administration's Executive Order on AI,[12] which the report characterizes as "hinder[ing] AI development or deployment" and which President Trump rescinded on the first day of his second term, and withholding AI-related federal funding from states with "burdensome" AI regulations.[13]

Since the July 2025 report, the Trump administration has taken further actions to promote its preferred deregulatory and anti-regulatory approach. The President has twice urged Congress to pass legislation placing a ten-year moratorium on states' ability to enact and enforce AI statutes and regulations. The House version of the One Big Beautiful Bill Act[14] included such a moratorium, but the provision was rejected by the Senate, which voted to remove the provision by a 99-1 margin.[15] The Trump administration sought to have the moratorium provision included in an early version of the 2026 National Defense Authorization Act, but it was ultimately excluded from the bill.[16] The Trump administration is likely to continue to push for Congress to enact legislation expressly preempting state AI laws and regulations, though no bills including the ten-year moratorium or similar provisions are currently pending.

After multiple failed attempts to enact preemptive legislation, the Trump administration has pursued alternative avenues to limit the reach of state AI regulation. On December 11, 2025, the administration issued an Executive Order on Ensuring a National Policy Framework for Artificial Intelligence (Order). [17] The Order advocated for a single national framework enabling the United States to win the AI race instead of a "patchwork of 50 different regulatory regimes" complicating compliance and otherwise imposing cumbersome requirements.[18] To disincentivize further state AI legislation and enforcement of existing state laws, the Order directed the U.S. Attorney General to establish an AI Litigation Task Force to challenge state AI laws that "unconstitutionally regulate interstate commerce."[19] Further, the Order directs the Secretary of Commerce to evaluate all existing state AI laws and withhold remaining funding for the federal Broadband Equity Access and Deployment (BEAD) Program from states deemed to have onerous AI laws; moreover, the Order encourages other agencies to withhold discretionary grant funding from such states or condition the grant of such discretionary funding on states' agreements not to enact new AI laws or not to enforce existing AI laws. [20]

Aside from its strategic litigation and funding provisions, the Order also calls for three courses of action in line with the administration's preemption aims. First, the Order requires the Chairman of the Federal Trade Commission to issue a policy statement on the application of the FTC Act's prohibition on unfair or deceptive acts to AI models, including under which circumstances state laws requiring alteration to AI models' truthful outputs are preempted by the FTC Act.[21] Second, the Order also requires the Federal Communications Commission Chairman to initiate a proceeding to determine whether to adopt a federal reporting and disclosure standard for AI models that preempts conflicting state laws. [22] Third, the Order directs the Special Advisor for AI and Crypto and the Assistant to the President for Science and Technology to "prepare a legislative recommendation establishing a uniform Federal policy framework for AI that preempts State AI laws,"[23] suggesting that the administration understands that federal legislation provides the clearest path forward for a uniform federal AI policy.

While the Trump administration has pushed repeatedly for preemption and regulatory uniformity, the exact contours of federal legislation establishing a national AI policy are unclear in terms of both the breadth and depth of the resulting regulatory regime. State regulation to date—which is discussed in the section below—has exemplified a variety of approaches to AI regulation in terms of both the types of entities directly regulated by the laws (e.g., developers of AI tools versus entities implementing or otherwise using AI) and the level of abstraction addressed by the laws (e.g., high-level risk-based approaches versus laws directly regulating certain conduct, industries, or specific AI use-cases).

AI-related Executive Orders and publications from the Executive Branch have provided few indicia of these considerations and other regulatory aims, and the administration's stated goal of reducing restrictions on AI development and integration across economic sectors is likely to result in limited federal regulation. As such, the implications of a federal AI policy on health care organizations are uncertain at this point, though stakeholders should continue to monitor for developments in this area.

As a notable and recent example of its approach to AI, in January 2026, the Food and Drug Administration (FDA) published revised final guidance on clinical decision support (CDS) software that reflects a deregulatory approach to AI.[24] The guidance sets forth a new enforcement discretion

policy for CDS tools that provide a single clinically appropriate recommendation but otherwise meet all statutory criteria for CDS software to be exempt from FDA regulation as a medical device. Under the prior 2022 guidance, such CDS tools would not be subject to enforcement discretion.[25]

Additionally, the Trump administration has taken a number of steps to integrate AI into the day-to-day workflows of the federal government across all federal agencies and departments. The Department of Health and Human Services and the Centers for Medicare & Medicaid Services (CMS) have both released documents outlining their plans to integrate AI into agency workflows as part of a wider and ongoing bipartisan effort[26] for the federal government to integrate AI to generate efficiency and cost savings. Though most agency AI initiatives are internally facing, the technology is already being rolled out in a manner that directly implicates beneficiaries of federal health care programs. In January 2026, CMS launched the Wasteful and Inappropriate Service Reduction (WISeR) innovation model—a pilot leveraging AI along with human review to expedite prior authorization for Medicare beneficiaries in six states—to reduce clinically unsupported care in services susceptible to fraud, waste, and abuse.[27] Given that human involvement in the review process remains, the WISeR model does not appear to be entirely agentic; however, the WISeR rollout has drawn scrutiny as it expressly places AI in the prior authorization decision-making chain, an area many states have already sought to regulate or limit, and could lead to additional denials of care as it compensates participating companies on the basis of averted expenses. It is too early to determine the full extent of AI utilization in the Medicare context, though CMS may expand the WISeR program or launch other pilot AI programs in the coming months and years. As such, stakeholders should continue to monitor the outcome and potential replication of WISeR and other AI-related initiatives.

State Regulation and Response to Federal Preemption Push

While the federal government has taken few steps to regulate AI affirmatively, states have moved rapidly to address AI's risks and benefits through a variety of regulatory approaches with differing implications on health care stakeholders in terms of both the entities regulated and the type of regulation. For instance, some states, like Colorado, have broadly imposed a reasonable care framework on developers and deployers of AI across industries, requiring risk assessment for a range of potential use-cases rather than prohibiting outright or regulating specific use-cases. This approach is broadly reminiscent of the omnibus, risk-based approach to regulating AI pursued by the European Union (EU) through the EU AI Act.[28] Other states, like Nevada, have foregone the general risk-based approach in favor of banning or requiring safeguards on specific use-cases that legislators believe have already generated, or pose a significant risk of generating, harm. To date, state laws have not made explicit legislative distinctions between agentic and non-agentic AI, though laws targeting specific AI use-cases tend to include in their scope agentic AI even if they do not specifically label it as such.

Colorado was among the first states to enact state legislation regulating the use of AI broadly. Colorado's AI Act applies to both developers and deployers (users) of "high-risk" AI systems, defined as those making or being a substantial factor in making a "consequential decision," including decisions that have a material legal or significant effect on the provision or denial of health care services.[29] While the law does not use the term "agentic AI," it seems likely that many agentic AI systems would be considered "high-risk" AI systems under the law given their ability to act autonomously with minimal

human prompting. The Act requires such deployers to maintain risk management policies and programs, follow certain reporting frameworks, and disclose use of the AI system to consumers interacting with the AI system.

Several states have adopted the risk framework characteristic of the Colorado AI Act in their own legislation; however, certain state laws pose no direct requirements on health care organizations or only apply to a subset of organizations. For instance, the Responsible AI Safety and Education (RAISE) Act,[30] signed into law in New York on December 19, 2025, includes requirements such as developing safety protocols, performing periodic independent audits of AI systems, and engaging in certain reporting requirements for safety incidents. However, the law applies only to "large developers" of AI systems, defined as those whose lifetime computing spend in training certain AI models exceeds an aggregate $100 million, with at least $5 million spent on training a particular AI model.[31] The spending threshold applies to AI models developed internally or procured through a third party, though computing costs for training performed by third party developers does not count toward the purchasing entity's computing cost threshold. Due to the high spending threshold and the Act's exemption for accredited colleges and universities engaged in academic research,[32] many health care organizations do not immediately fall under the purview of the Act. Nevertheless, because the Act's definition of a "large developer" is based upon funds spent training AI models cumulatively, rather than annually, hospitals and other health care organizations may come within the scope of the bill as they continue to train AI tools and reach the aggregate spending threshold.

Other states regulate health care organizations directly, rather than through the use of a risk framework, and focus on discrete use-cases for AI in the health care context centered around three themes: disclosures to patients, prohibitions on misrepresenting AI as a licensed professional, and limits on AI supplanting clinician or payer decision-making—each of which becomes more complex as agentic AI systems capable of autonomous, multi-step actions enter clinical workflows. For instance, California and Utah require disclosure when AI is used in clinical communications unless a human reviews the output;[33] Texas mandates disclosure when AI is used for diagnosis or treatment.[34] California has additionally banned AI from misrepresenting itself as a licensed medical professional; [35] Nevada has a similar prohibition for misrepresentation of licensure as a mental or behavioral health provider and requires licensed provider oversight when chatbots are used.[36] Finally, numerous states, including California, Texas, and Illinois, either directly ban the use of AI tools from independently reaching decisions regarding medical necessity or clinical treatment, or require provider review of AI outputs used in the clinical decision-making process.[37] Other activity-based regulations will doubtless come to pass as new use-cases for AI, and particularly agentic AI, continue to take hold in the health care industry, and as states formulate, experiment with, and borrow responses to such uses. Given that it can operate relatively autonomously with minimal human input, agentic AI is likely to trigger the highest level of regulation under these state laws.

As agentic AI functionality continues to proliferate across industries, states may pursue a blended approach incorporating aspects of both current trends, combining broad, risk-based approaches with additional discrete laws or regulations enacted on an ad-hoc basis for use-cases meriting explicit prohibitions or safeguards. For instance, while risk-based laws like New York's may impose strict safety and reporting protocols on developers, there is limited regulatory oversight of, for example, a small therapy practice that procures a developer's AI tool and employs it as an unsupervised AI chatbot

claiming to be a licensed therapist—an activity Nevada and other states have prohibited due to reported instances of harm to patients. It stands to reason given the heavy regulation of the health care industry that many of these ad-hoc laws will address the use of AI in the health care sector. While such laws may not use the term "agentic AI" by name, they are likely to regulate more heavily agentic use-cases given the autonomous nature of such systems and attendant potential impact on patient care.

Although, as discussed above, the Trump administration supports full preemption of state AI regulation and has threatened litigation and funding cuts to states that enact new AI laws and continue to enforce existing AI laws, both red and blue states alike have indicated that they will continue regulating AI in the wake of the recent Executive Order. For example, New York Governor Kathy Hochul signed the state's recent AI bill into law following the Order, and Florida Governor Ron DeSantis criticized the President's move, continuing to push for an AI "bill of rights" in his state.[38] Statehouses nationally are replete with bills addressing AI, including proposed legislation addressing health care, and legislators continue to introduce bills and move toward enacting additional AI legislation despite the Trump administration's preemption push. Some states have proposed legislation modeled after or similar in key respects to laws in other states, including Missouri, where a bill regulating AI in mental health care is substantially similar to Nevada's law addressing the same.[39] Other state legislatures continue to propose novel approaches to regulating AI in health care, including, for example, a bill in Massachusetts that would require health insurance carriers both to disclose the use of AI in the claims review process and report information on AI training data to the state department of insurance.[40]

Sub-regulatory Guidance and Enforcement

We could locate no pronouncements from state medical boards regarding the use of AI in a clinical context, and there have been few or no enforcement actions against physicians for inappropriate use of AI in the provision of care. That said, some state laws delegate authority to professional licensing boards regarding both enforcement and the adoption of rules for safe and ethical AI use. For instance, California grants state professional licensing boards authority to pursue an injunction or restraining order for violations of the ban on misrepresenting an AI agent as a licensed medical provider.[41] Texas statutes require providers to review medical records created with AI in a manner consistent with medical records standards developed by the Texas Medical Board;[42] however, to date, the relevant Texas Medical Board regulations have not been updated to reflect AI-specific requirements.[43] More broadly, where state AI laws delegate enforcement or standard-setting authority to medical boards—as California and Texas have done—boards may need to engage in formal rulemaking under their statutory authority before they can effectively enforce or implement such requirements.[44] We may expect more activity in the coming years, particularly if state medical boards face complaints and are asked to undertake disciplinary actions against licensed professionals stemming from negative patient outcomes related to the use of agentic AI tools. Enforcement actions may also extend beyond individual practitioners to hospitals and corporate entities under corporate practice of medicine doctrines, which prohibit corporations from practicing medicine or unduly influencing physician decision-making, and under theories of vicarious liability or negligent credentialing when institutions deploy AI systems that result in patient harm.[45]

Enforcement by state attorneys general has also been limited to date, though it may pick up in 2026 and beyond as new laws come into effect and violations occur. To date, state AG activity has included a 2024 settlement between the Texas AG and Pieces Technologies, a company whose product drafts AI-generated clinical summaries and documentation within electronic health records for multidisciplinary care teams, stemming from allegedly misleading advertising statements regarding the products' accuracy and hallucination rates.[46] The California AG in 2025 released a legal advisory report outlining potential liability resulting from the use of AI in health care as it relates to the state's consumer protection, civil rights, competition, and patient privacy laws, but we could locate no other AG guidance specific to the health care industry.

Several national medical organizations and other industry groups have issued guidance or otherwise offered insights on the use of AI in clinical settings in a manner that largely mirrors the legislative safeguards that states have been adopting. For example, the Federation of State Medical Boards issued a report[47] offering several best practices for state medical boards' governance of AI use, including requiring transparency and disclosure when AI is used, education for providers on responsible use of AI tools, ensuring equitable access to the benefits of AI for all patients, maintaining a focus on patient privacy and data security, and oversight of how AI use interacts with the legal contours of the "practice of medicine" in their jurisdiction. The American Medical Association similarly reported on principles for AI development and use, advocating for a risk-based approach for the integration of AI into both clinical and non-clinical aspects of health care,[48] and the Joint Commission (JC) and Coalition for Health AI (CHAI) has published guidance[49] based on industry standards for the development, deployment, and use of AI in health care, centering on patient privacy, safety, and data security, as well as training, risk assessment, and quality monitoring.

Take–Aways for Health Care Stakeholders on Navigating AI Regulations

To date, use-cases for AI, including agentic AI, in the health care context have developed rapidly and regulators have worked nearly as rapidly to impose new laws and policy pronouncements addressing the use of AI, particularly in health care. While a great deal of attention is focused on federal preemption efforts, it is highly unlikely that full federal preemption of state laws regulating AI will come to pass, and such preemption would likely not be instantaneous if it were to occur. As the battle over federal preemption continues, state legislatures remain highly active in experimenting with new regulations of AI, which, as discussed above, generate significant and various compliance requirements for health care organizations as a function of the types of entities and specific activities regulated. In 2026 and beyond, state legislatures will doubtless pursue additional regulation, likely with particular focus on higher-risk activities in the health care context—for example, where AI, including agentic AI, is involved in the clinical decision-making process, or where the use of AI tools in workflows creates a potential high risk of harm to patient safety or privacy. Health care organizations should continue to monitor for such new laws and pronouncements, as well as actions taken by state licensure boards and other enforcement entities as existing laws are implemented.

Beyond monitoring for new laws and enforcement actions, health care organizations should take action to ensure compliance with existing laws and, where operating in multiple jurisdictions, take care to design programs that are flexible to accommodate the variety of regulatory approaches taken by the

states thus far. The insights and guidance published by industry groups to date provide a good starting point as health care organizations undertake these compliance initiatives.

Health care organizations seeking to deploy agentic AI solutions should consider establishing comprehensive governance frameworks to guide evaluation, implementation, and ongoing oversight. These governance efforts should begin with the identification of when agentic AI solutions are being evaluated and should include a classification system for potential uses based on risk tiers that take into consideration the intended use cases and the specific harms or inefficiencies the tool aims to address.

Clinician involvement and understanding is essential throughout the evaluation process. Organizations should engage clinician users early in the agentic AI journey to discuss the workflows that will be affected by the proposed solution, recognizing that tools appropriate in one setting, such as the emergency department, may be ill-suited for another, such as primary care. Clinicians should also assess the clinical accuracy of the proposed solution, determining what percentage of accuracy the solution must demonstrate for providers to be comfortable with its use and establishing mechanisms for ongoing monitoring.

Professional judgment and human oversight are also important considerations. Organizations should evaluate the ability of clinicians to exercise professional judgment when using the solution, which will inform the overall potential for clinical risk. The level of human involvement may vary depending on the potential risk posed by the solution and applicable legal requirements. This includes ensuring that clinicians have the ability to review the basis for decisions made by the agentic AI, to override the agentic AI when clinically appropriate, and to make their own determination regarding use of the agentic AI tool based upon the assessed level of risk.

Equal access for all patients should also be evaluated. Organizations should consider including a provider or clinician advocate in the evaluation process to assess potential bias or discrimination effects of the agentic AI solution after testing on training data. Additionally, organizations should document how clinical workflows and decision-making points will be preserved to prevent deskilling that may accompany use of the solution and to ensure backup procedures exist in the event of solution failure.

Technical integration also warrants careful analysis. Organizations should consider identifying whether the agentic solution will interact with any other systems or AI agents, as a single failure point could create a compounding effect if the solution is heavily integrated with multiple other systems or agents. Organizations must also develop the capability to determine when disclosure of AI use to patients and other stakeholders is required based on applicable laws.

Operational readiness requires ongoing education and support. Organizations should consider establishing AI education requirements for all personnel, potentially as part of annual training requirements. Support for operational deployment of agentic AI tools in the clinical setting should include standard operating procedures, scripting for physicians, instructions on obtaining consent where legally required, frequently asked questions, and directions on where to send patients who have additional technical questions regarding matters such as data storage and retrieval. Physicians should clearly understand what information will be included in the official medical record, what disclosures patients will see, and how to respond to patient questions.

Finally, organizations should consider addressing ancillary considerations, including any third-party AI company's access and use rights to data, as well as indemnity and insurance provisions applicable to the agentic AI tool. Organizations that follow the above steps will be better positioned to take advantage of the promises of AI use while minimizing the risk of harm.

About the Authors

David Peloquin, JD, is a partner in Ropes & Gray's health care group in Boston, who advises clients on a wide range of legal and regulatory issues in the area of clinical research, data privacy, provision of health care services, and related activities. David guides clients through service arrangements, transactions and investigations involving complex regulatory questions arising under the Common Rule and FDA regulations, data privacy regulations (including HIPAA, U.S. state privacy laws, and GDPR), state and federal fraud and abuse laws, health care licensing requirements, and laws regulating the use of artificial intelligence. David is also a member of the firm's digital health practice and frequently advises clients on the use of digital technologies in research and clinical settings. His clients include academic medical centers, universities, pharmaceutical and medical device companies, biotech firms, clinical laboratories, information technology companies, and investors. He frequently publishes and speaks on issues related to human subjects research, data privacy and digital health.

Hillary Noll Kalay is Managing Counsel at the University of California's Office of General Counsel in Oakland, California, advising UC campuses and medical centers on matters relating to clinical research, data rights, privacy, procurement, and artificial intelligence. Hillary is recognized within the UC system and nationally as an expert in privacy in international and domestic clinical research, and complex data collaborations and AI. Hillary supports UC's leadership, as well as institutional review boards, health care and research compliance, health procurement, research policy, and privacy offices throughout the UC system. Prior to joining UC's Office of General Counsel, Hillary developed and advised on policies relating to clinical trials and human subject research protection for the UC system. Hillary previously practiced intellectual property litigation at national law firms in the San Francisco Bay Area. She frequently speaks on issues relating to health privacy, artificial intelligence, and clinical research. Hillary also holds a Master of Public Policy and a CIPP/US certification.

Allison Trimble is a clinical research, data privacy and technology attorney with several years of experience working in highly regulated industries both in-house and for an Am Law 100 law firm. In her current role as Associate General Counsel for BJC Health System, she serves as a senior strategic legal advisor to C-suite and senior management guiding digital transformation efforts, cutting-edge data privacy, AI and technology issues.

Jake Fulton, JD, is an associate in Ropes & Gray's health care group in Chicago, who advises clients on legal, regulatory, and transactional, matters, including issues related to health care operations, fraud and abuse laws, and laws regulating the use of artificial intelligence. He serves a broad range of health care clients, including hospitals and health systems, pharmaceutical manufacturers, physician practices and other provider organizations.

[1] D. B. Acharya, K. Kuppan, and B. Divya, *Agentic AI: Autonomous Intelligence for Complex Goals —A Comprehensive Survey*, in Institute of Electrical and Electronic Engineers Access, vol. 13, pp. 18912, 18917.

[2] *Id.*

[3] *Id.* at 18919.

[4] Anastassia Gliadkovskaya, *Zocdoc launches voice AI agent Zo to streamline bookings by phone,* Fierce Healthcare (May 1, 2025), https://www.fiercehealthcare.com/health-tech/zocdoc-launches-voice-ai-tool-zo-streamline-phone-visit-bookings.

[5] Heather Landi, *Cedar rolls out AI voice agent to tackle patients' medical billing questions*, Fierce Healthcare (Apr. 29, 2025), https://www.fiercehealthcare.com/ai-and-machine-learning/cedar-rolls-out-ai-voice-agent-tackle-patients-billing-questions.

[6] *Agentic AI vs. Generative AI: 5 Key Differences*, Exabeam, https://www.exabeam.com/explainers/agentic-ai/agentic-ai-vs-generative-ai-5-key-differences/.

[7] Shreya J. Shah et al., *Physician Perspectives on Ambient AI Scribes*, 8 JAMA Network Open e251904 (2025), https://pmc.ncbi.nlm.nih.gov/articles/PMC11933996/.

[8] Heather Landi, *Highmark Health taps Abridge to build 'real-time' prior authorization using* AI, Fierce Healthcare (Aug. 12, 2025), https://www.fiercehealthcare.com/health-tech/highmark-health-taps-abridge-build-real-time-prior-authorization-using-ai.

[9] Pinto-Coelho L., *How Artificial Intelligence Is Shaping Medical Imaging Technology: A Survey of Innovations and Applications*. Bioengineering (Basel). 2023 Dec 18;10(12):1435. doi: 10.3390/bioengineering10121435. PMID: 38136026; PMCID: PMC10740686.

[10] *See* Healthy Technology Act of 2025*,* H.R. 206, 119th Cong. (2025); Maintaining Innovation and Safe Technologies Act, H.R. 193, 119th Cong. (2025).

[11] The White House, *Winning the Race: America's AI Action Plan* (2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.

[12] The Trump administration rescinded President Biden's 2023 Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (EO 14110) along with scores of other Biden-era policies through EO 14148. *See* Exec. Order No. 14,148, 90 Fed. Reg. 8237 (Jan. 20, 2025).

[13] *See supra* note 11, at 3.

[14] One Big Beautiful Bill Act, H.R. 1, § 43201(c), 119th Cong. (2025) (version engrossed in House on May 22, 2025 which did not ultimately become law).

[15] U.S. Senate Committee on Commerce, Science, and Transportation, *Senate Strikes AI Moratorium from Budget Reconciliation Bill in Overwhelming 99-1 Vote* (July 1, 2025), https://www.commerce.senate.gov/2025/7/senate-strikes-ai-moratorium-from-budget-reconciliation-bill-in-overwhelming-99-1-vote/8415a728-fd1d-4269-98ac-101d1d0c71e0.

[16] Julia Shapero, *Here are key tech provisions the NDAA includes and excludes*, The Hill (Dec. 9, 2025), https://thehill.com/policy/technology/5639209-ndaa-ai-preemption-chip-exports/.

[17] Exec. Order No. 14,365, 90 Fed. Reg. 58,499 (Dec. 11, 2025).

[18] *Id.* at § 1.

[19] *Id.* at § 3.

[20] *Id.* at §§ 4-5.

[21] *Id.* at § 7.

[22] *Id.* at § 6.

[23] *Id.* at § 8.

[24] U.S. Food and Drug Administration, *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff* (Jan. 29, 2026), https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software.

[25] U.S. Food and Drug Administration, *Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff* (Sept. 28, 2022).

[26] Though the Biden administration's major AI Executive Order was rescinded in full, it largely supported the integration of AI into agency workflows. *See* Exec. Order 14,110, 88 Fed. Reg. 75191 (Oct. 30, 2023) (rescinded by Exec. Order. 14,148).

[27] Centers for Medicare & Medicaid Services, *WISeR (Wasteful and Inappropriate Service Reduction) Model*, (accessed Jan. 20, 2026), https://www.cms.gov/priorities/innovation/innovation-models/wiser.

[28] Mayur Jariwala, *A Comparative Analysis of the EU AI Act and the Colorado AI Act: Regulatory Approaches to Artificial Intelligence Governance*, 186.30 Int'l J. Computer Applications 23 (Sep. 2024). https://www.ijcaonline.org/archives/volume186/number38/jariwala-2024-ijca-923954.pdf.

[29] Consumer Protections for Artificial Intelligence, Colo. S.B. 24-205 (2024).

[30] Responsible AI Safety and Education Act, N.Y. A.B. 6453A (2025).

[31] *Id.* at § 1420.9.

[32] *Id.*

[33] *See* Cal. Health and Safety Code § 1339.75; Utah Code Ann. § 13-2-12(3)-(5).

[34] *See* Tex. Bus. & Com. Code § 552.051.

[35] *See* Cal. Bus. and Prof. Code § 4999.9(c).

[36] *See* Nev. Stat. Rev. Ch. 433, 629.

[37] *See, e.g.,* Cal. Health and Safety Code § 1367.01(k).

[38] Andrew Atterbury, *'We have a right to do this': DeSantis wants Florida to move ahead with AI policies*, Politico (Dec. 15, 2025), https://www.politico.com/news/2025/12/15/we-have-a-right-to-do-this-desantis-wants-florida-to-move-ahead-with-ai-policies-00690680.

[39] *See* Mo. S.B. 1444 (2025).

[40] *See* An Act relative to AI health communications and informed patient consent, Mass. H.1210 (2025).

[41] *See* Cal. Bus. and Prof. Code § 4999.9(a).

[42] *See* Texas Senate Bill 1188 (SB 1188)

[43] *See* 22 Tex. Admin Code § 163.

[44] Fed'n of State Med. Bds., *Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice* (Apr. 2024), https://www.fsmb.org/siteassets/advocacy/policies/incorporation-of-ai-into-practice.pdf (adopted by FSMB House of Delegates, Apr. 2024) (recommending best practices for state medical boards in governing physician use of AI in clinical care and noting that boards regulate physicians, not tools, suggesting that boards will need to develop standards specific to AI use before pursuing disciplinary actions).

[45] *See* W. Nicholson Price II, Sara Gerke & I. Glenn Cohen, *Liability for Use of Artificial Intelligence in Medicine*, *in* Research Handbook on Health, AI and the Law 150, 154-57 (Barry Solaiman & I. Glenn Cohen eds., Edward Elgar Publishing 2024) (noting that "hospitals could be liable for negligently choosing, implementing, and maintaining" AI systems, and explaining that "[u]nder the doctrine of respondeat superior, an employer is subject to liability for torts committed by employees while acting within the scope of their employment," such that "should a patient have a bona fide malpractice claim relating to a hospital employee's tortious use of AI to direct the patient's care, and the activities were within the employee's scope of employment, then liability may flow to the hospital system").

[46] Dave Murio, *Texas Attorney General, Pieces Technologies Settle Allegations of Inaccurate Generative AI Advertising,* Fierce Healthcare (Sep. 19, 2024), https://www.fiercehealthcare.com/ai-and-machine-learning/texas-ag-pieces-technologies-settle-allegations-inaccurate-generative-ai.

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

[48] Am. Med. Ass'n, *Augmented Intelligence Development, Deployment, and Use in Health Care* (Nov. 2024), https://www.ama-assn.org/system/files/ama-ai-principles.pdf.

[49] Joint Commission and Coalition for Health AI (CHAI), *The Responsible Use of AI in Healthcare* (2025), https://digitalassets.jointcommission.org/api/public/content/dcfcf4f1a0cc45cdb526b3cb034c68c2?v=3edb8a95&_gl=1*4qliez*_gcl_au*OTAzMzg5MjczLjE3NjU4MTgzNjU.*_ga*MTI0MDI5NTc3Ni4xNzY1O

ARTICLE TAGS

Health Information

1099 14th Street NW, Suite 925, Washington, DC 20005 | P. 202-833-1100
For payments, please mail to P.O. Box 79340, Baltimore, MD 21279-0340