DENTONS

CRESCENT STRATEGY CONSULTING

Intelligent Autonomy
The Next Frontier of AI in MedTech

Authors:
Sabeen Shaikh, Partner at Crescent Strategy Consulting;
Peter Stockburger, Partner at Dentons

December 2025

In 1950, British mathematician Alan Turing published his landmark paper "Computing Machinery and Intelligence" and asked the question - "can machines think?" Since that time, decades of research in **artificial intelligence** (AI) have led to extraordinary developments in the fields of **machine learning**, **deep learning**, and recently **generative AI** with the release of powerful foundation models such as large language, vision, action, and world models. In addition to rapid advances in machine learning, deep learning, and generative AI, there have also been significant recent developments in other branches of AI, including neurosymbolic systems and active inference.

This rapid acceleration in AI development is quickly transforming the **medical technology**[1] (MedTech) industry. Over the past decade alone, advancements in AI have enabled more accurate and faster disease diagnosis, accelerated the process of identifying promising compounds and predicting how they will interact with disease targets, improved patient outcomes by identifying disease risk factors earlier and enabling more personalized treatment plans, and driven advancements in robotic surgery capabilities with improved precision and consistency.

As these developments continue, AI is entering a new stage of advancement that may hold the promise to unlock additional opportunity in the MedTech industry - the emergence of intelligent, autonomous, and vertical AI systems.

Agentic AI systems are broadly understood to be those that can perceive their environment, plan, reason, make decisions, and take action to achieve a specific goal with minimal human intervention. **Vertical AI systems** generally refer to AI systems that are purpose-built to solve specific challenges for a particular domain such as healthcare. Vertical AI systems can often address the challenges presented by more general AI systems (e.g., large language models, or LLMs) such as a lack of specialized knowledge and context. To achieve this goal, vertical AI systems that are dependent on machine learning will deploy special techniques such as retrieval-augmented generation (RAG) (e.g., grounding AI responses in specific literature, patient data, clinical guidelines, etc.), using smaller fine-tuned models trained on healthcare-specific datasets, and leveraging hyper-specific context and proprietary data. There are also emerging AI systems that may further advance vertical AI development, such as those in the field of active inference.

The combination of agentic and vertical AI systems has the potential to unlock significant innovation in the MedTech industry by introducing the ability to autonomously conduct complex research tasks, analyze vast datasets, design novel therapeutics, and reimagine static and legacy workflows with minimal human oversight. Agentic vertical AI systems may have the ability to break down complex tasks, use tools or access external systems, communicate with other AI systems, and adapt their approach based on feedback, outcomes, and new facts. In clinical settings, agentic vertical AI systems hold the promise of transforming diagnostics and treatment planning by continuously monitoring patient data, reasoning through complex medical cases, and recommending personalized interventions in real-time. Agentic and vertical AI systems may also enable breakthrough discoveries in areas such as genomics and personalized medicine. And as these systems evolve, there will likely be agentic vertical AI systems that drive research, workflows, drug development, surgery, and clinical decision-making, fundamentally reshaping how medical innovation happens and how patients receive care.

1 For the purposes of this paper, the term "MedTech" refers to devices, equipment, software, and technologies used to diagnose, treat, monitor, and manage the delivery of healthcare to patients. This broad category can encompass everything from simple devices like diagnostic kits and wearable health monitors to complex systems like imaging equipment, surgical robots, and digital health solutions such as medical software, data analytics platforms, and treatment tools. The MedTech industry sits at the intersection of healthcare, technology, and innovation.

MedTech leaders who embrace agentic vertical AI as a strategic imperative will unlock new opportunities for clinical innovation, operational efficiency, and competitive differentiation. Those who treat it as an experiment risk falling behind as patients, providers, and clinical care teams expect emergent technologies that can act, not just assist.

As with any technological advancement, developing and deploying agentic vertical AI systems will not be without its challenges. Organizations developing and deploying these technologies must engage in careful planning and execution. The law around AI globally is still evolving, especially in the area of agentic AI. Complicating matters, traditional AI governance frameworks, such as those offered by the National Institute for Standards and Technology (NIST) and the International Standards Organization (ISO), may not be well suited to manage the risks posed by agentic vertical AI systems as those frameworks are largely designed for static, general-purpose models that respond to specific queries and not systems that operate autonomously, continuously adapt their behavior, and make decisions across complex workflows with minimal human oversight. Organizations must therefore develop strong and flexible governance guardrails tailored to the unique risks posed by agentic vertical AI systems, including addressing issues of accountability, auditability, security, and trust. Emerging frameworks and standards are becoming available for organizations to leverage to ensure the deployment of agentic AI is safe and secure.

Building strong and adaptive governance controls uniquely tailored to agentic and vertical AI systems will be key to realize the strategic potential for MedTech leaders.

This paper explores the emerging field of agentic vertical AI systems, market trends, unique use cases, the legal and governance challenges around development and deployment, and strategic opportunities for MedTech leaders looking to leverage this coming advancement in technology.

Agentic AI is already transforming enterprises and is likely to be a multitrillion-dollar opportunity. This means that healthcare IT leaders should lean in to learn how AI agents can help transform work across drug discovery, patient care, operations and so much more.

— Amanda Saunders, Director of Generative AI Software Marketing at NVIDIA[2]

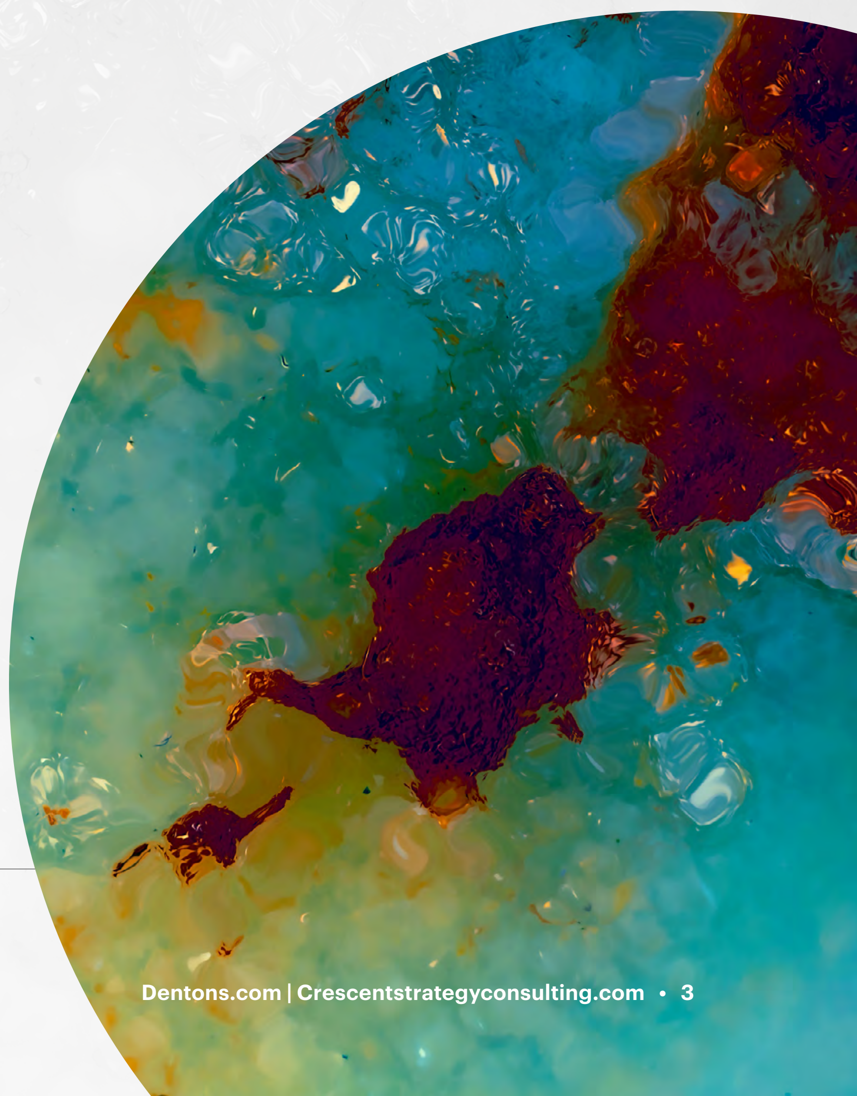

2 https://medicalfuturist.com/agentic-ai-in-healthcare/

Understanding Agentic & Vertical AI Systems

AI is generally defined as any engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.[3] Within the broader field of AI, there are distinct areas of study, including deterministic rule-based systems and machine learning. **Deterministic AI** relies on explicitly programmed rules and logic, producing predictable outputs based on defined inputs. Your e-mail spam filter is an example of a deterministic AI system – it puts spam messages in your junk folder based on keywords, creating a predictable outcome. **Machine learning**, on the other hand, focuses on systems that learn patterns, make decisions, and improve themselves from data without being explicitly programmed.[4] The field of machine learning has become the dominant focus of modern AI research and development due to its ability to handle complex datasets and scale across applications. **Deep learning** is a subset of machine learning that uses multi-layered neural networks to learn complex patterns and representations from large amounts of data, which is effective for tasks like image recognition, natural language processing, and pattern detection[5].

Generative AI is an evolution of deep learning that learns the structure of the underlying training data and creates statistically probable outputs (e.g., text, images, code, music) when prompted.[6] While generative AI may seem like a recent breakthrough, its foundations go back over a decade to early models like **variational auto-encoders** introduced in 2013. These early models were used to generate realistic[7] images and speech by encoding unlabeled data into a compressed representation and then decoding the data back into its original form to generate new content.[8] These advances, along with subsequent innovations like generative adversarial networks, paved the way for the development of the transformer architecture, which has dramatically expanded the scale and capability of generative AI systems.

Transformers were first introduced in a groundbreaking 2017 paper titled "Attention Is All You Need."[9] Transformers combine the original encoder-decoder architecture with a new token-processing mechanism called ***attention*** to change how generative AI models are trained. The attention mechanism enables the AI model to weigh the importance of different pieces of data relative to each other, rather than processing them with equal significance. This means the model can selectively "pay attention" to the most relevant words, tokens, pixels, or features when generating an output. This allows the model to capture deep and complex dependencies, nuanced syntactic context, and complex relationships. Transformers enabled the creation of several **foundation models** by providing the architecture that can scale to massive datasets and representations of multiple modalities, such as **large language models (LLMs)** which are trained on vast amounts of text, **large vision models (LVMs)** which are trained on vast amounts of images and videos, and **large action and world models (LAMs and LWMs)** which are trained on video and information about how the physical world operates.

3 See, e.g. Cal. Civ. Code § 3310(a), 15 U.S.C. § 9401, C.R.S.A. § 6-1-1701(2) (Colorado AI Act); TX Bus. & Com. § 551.001(1) (Texas Responsible Artificial Intelligence Governance Act); Regulation (EU) 2024/1689, Art. 3(1) (EU AI Act).
4 https://ai.engineering.columbia.edu/ai-vs-machine-learning/
5 IBM, "AI model types: Past, present and predictions for the future," available at https://www.ibm.com/think/insights/ai-model-types-past-present-future-predictions.
6 IBM, "What is generative AI?", https://research.ibm.com/blog/what-is-generative-AI.
7 IBM, "AI model types: Past, present and predictions for the future," available at https://www.ibm.com/think/insights/ai-model-types-past-present-future-predictions.
8 *Id.*
9 *Id.*

Beyond generative AI, there are emerging models of AI such as neurosymbolic and active inference AI that promise further developments in the field of intelligence and autonomy.

Within this evolving field of AI there are two key developments relevant to this paper – **vertical AI systems** and **agentic AI systems**.

While foundation models are intended to have general applicability, and are often referred to as **horizontal AI systems**, **vertical AI systems** are designed to perform more specialized functions for specific industries, offering a deep understanding of domain-specific language and workflows. In healthcare, vertical AI systems are purpose-built to address specific unique challenges such as diagnostic imaging analysis, drug discovery, and personalized treatment plans. Vertical AI systems may leverage generative AI or other forms of AI depending on the use case and need, and can be trained on large volumes of medical literature, clinical guidelines, electronic health records, genomic data, and real-world treatment outcomes. They may leverage techniques like RAG to ground their responses, and leverage smaller models that are more efficient and interpretable, amongst other techniques.

Agentic AI systems likewise present a unique development in AI for MedTech. Traditional forms of machine and deep learning primarily operate as stand-alone software tools that process inputs and generate outputs without autonomous action or goal-directed behavior. **Agentic AI systems**, by contrast, are designed to operate autonomously, set their own goals, make decisions, and take actions in the world with minimal human intervention. While traditional AI models require humans to direct each interaction and decide on next steps, agentic AI systems can sequence multiple actions toward goals, interact with external systems and environments, and operate continuously without human input. This represents a fundamental shift from AI as a service that responds on demand to AI as an autonomous actor pursuing objectives in the world.

Agentic AI System Architecture

Agentic AI systems are designed to perceive inputs, think, decide, and act autonomously in pursuit of defined objectives. Unlike traditional AI models that respond to inputs, agentic AI systems are built as layered ecosystems. Each layer serves a distinct cognitive or operational function. Together, these layers form the foundation of autonomous intelligence, allowing an agent to not only understand its operating environment but also plan and execute complex, multi-step actions.

Perception and Context Layer: Understanding the Operating Environment

The perception and context layer of an agentic AI system is responsible for gathering information from the operating environment and formulating it into meaningful input. This can include ingesting raw data from APIs, documents, tools, other agents, conversations, and human interactions. Perception is coupled with context to create a dynamic picture of the environment. This layer grounds the agent's reasoning, ensuring decisions are informed by situational understanding.

Cognition and Memory Layer: Comprehension and Continuity

Once the environment is perceived, the agentic AI system utilizes its cognition and memory layer to transform the environment into structured knowledge. This is where embeddings, ontologies, symbolic representations, and knowledge graphs are created. This is also where meaning is extracted from signals. The agent uses this cognitive layer to interpret goals, constraints, and relationships. The memory function provides continuity over time. Together, these capabilities give the agent a sense of history, enabling it to learn from past decisions and personalize future actions.

Model Layer: Reasoning and Intelligence

The model layer is the brain of the agentic AI system. This layer may utilize generative AI models such as LLMs, LVMs, LAMs, and LWMs, or a hybrid approach including using emerging models such as neurosymbolic and active inference AI. This layer interprets context and memory to infer, predict, and decide. This layer may operate on multiple levels, from specific decisions to strategic reasoning.

Planning and Execution Layer: Strategy Into Action

The planning and execution layer is where high-level objectives are structured into concrete tasks, ordered into logical sequences, and optimized for efficiency. The execution component takes these plans and orchestrates their implementation by invoking APIs, calling tools, triggering workflows, or directing downstream systems. Execution also generally involves monitoring and verification, ensuring actions are carried out as intended.

Actuator Layer: Bridging Digital Decisions and Real-World Effects

The actuator layer is the interface where the agentic AI system influences its environment. This may take many forms, such as software actuators that write data or communication actuators that can initiate human-facing interactions. The actuator layer is not a one-way channel. Actions taken at this layer may flow back into the perception and context layer, and enable the system to refine its decision-making.

Feedback and Learning Layer: Continuous Adaption

Agentic AI systems evolve over time. The feedback and learning layer can ensure continuous improvement by monitoring outcomes, measuring performance, and adjusting models. This can involve self-supervised learning, human-in-the-loop review, or structured evaluation.

Key Components for Buiding AI Agents[10]

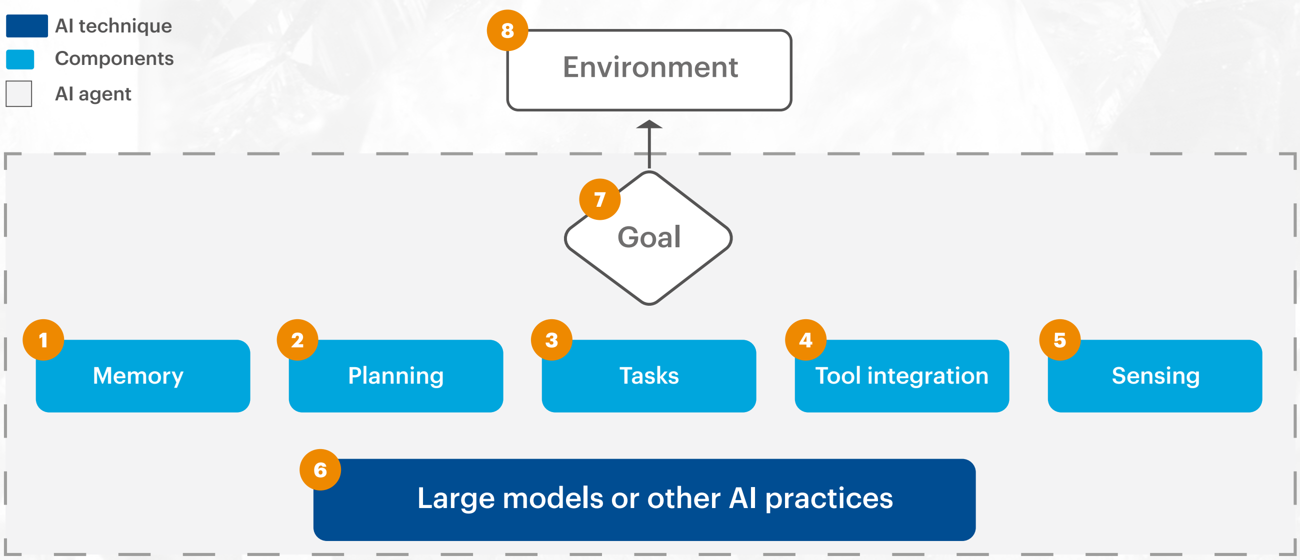

1. The AI agent's short- and long-term memory
2. Breaks down the goal into smaller tasks
3. Current list of tasks for execution
4. Agent interacts with various environments
5. Ability to perceive its surroundings
6. Large models, small models or other AI practices
7. The overall task to be accomplished
8. Physical or digital arena the agent will act in

10 https://www.nojitter.com/ai-automation/conversations-in-collaboration-gartner-tom-coshow-on-ai-agents-and-agentic-ai

From Agents to Networks

According to Gartner, organizations will increasingly deploy agents not just as isolated tools but as part of a layered, multi-agent system (MAS). Gartner describes this evolution in three broad arrangements, each representing a deeper level of integration:

Standalone Agents

These are single-agent systems built around a narrow and well-defined purpose (e.g., task automation agent, domain-specific agent, etc.). These agents deliver immediate results and are relatively simple to build and govern.

Agentic Assemblies

At the next level, multiple agents are networked together into assemblies, executing multi-stage workflows to solve complex problems. In these systems, one agent may handle perception and context, while another performs reasoning and planning, while another may execute external actions. These assemblies are then orchestrated through agent frameworks or orchestrators that manage communication, task delegation, and conflict resolution.

IBM defines four different types of AI orchestration:

1. **Centralized orchestration**: one agent in the system serves as the leader, dictating the actions of the other specialized agents.
2. **Decentralized orchestration**: all agents collaborate to solve a specific problem, with no clear leader.
3. **Hierarchical orchestration**: agents are organized in a hierarchy, where higher-level agents "oversee and manage ower-level agents."[11]
4. **Federated orchestration**: AI agents are able to collaborate with other electronic entities, allowing different tech stacks to work together.

Another representation of such orchestration arrangements may include:

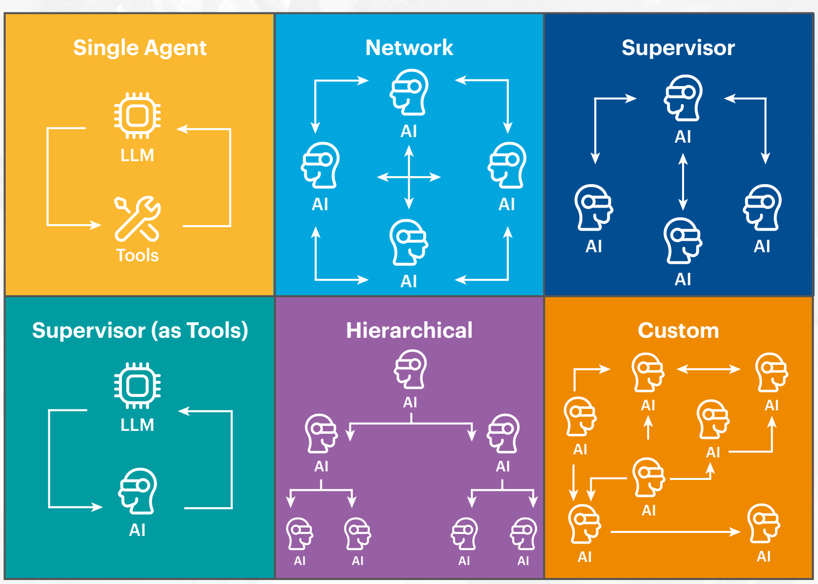

11 https://www.ibm.com/think/topics/ai-agent-orchestration

The most advanced assembly is an agentic ecosystem. This is a dynamic, self-organizing environment where numerous autonomous agents may interact through shared protocols, memory layers, and goal structures. In these ecosystems, enterprises may deploy hundreds or thousands of agents across their business to create an intelligence system that continuously optimizes, discovers, and acts with minimal human oversight.

Customer Feedback Example

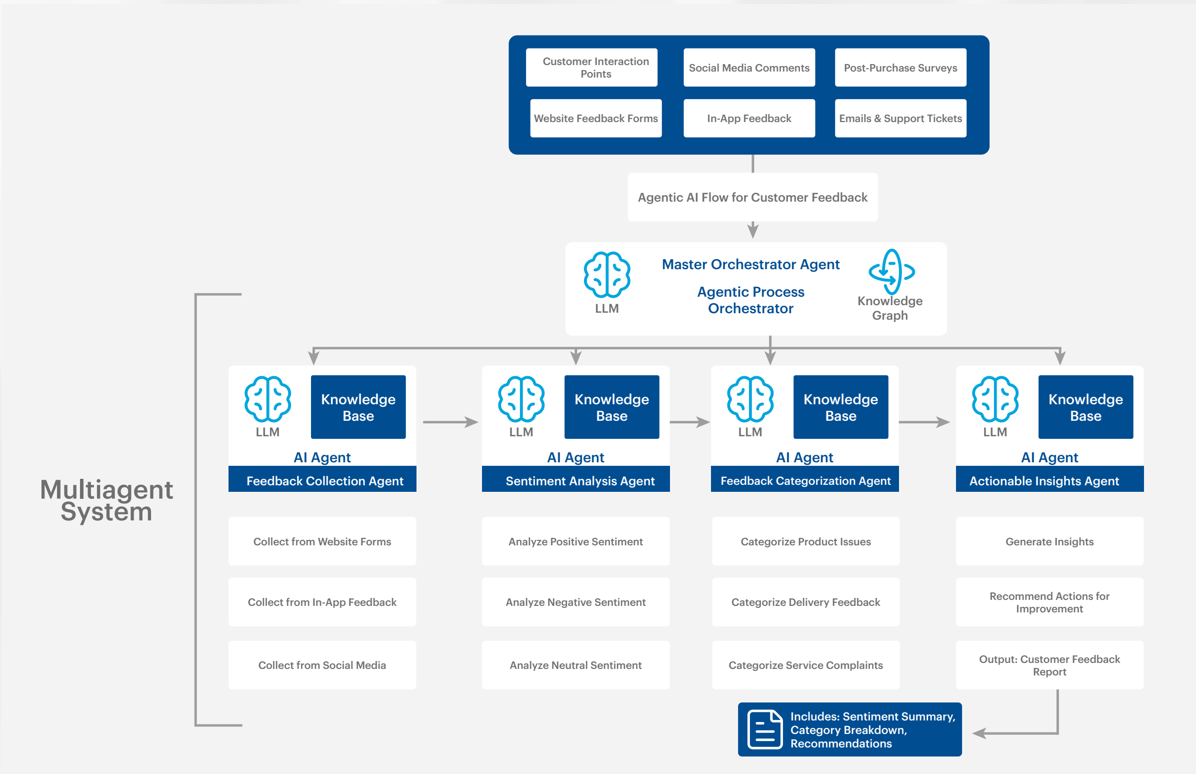

[12]

The MAS above points to a possible example of a MAS tech stack. This system, as described by Dr. Jagreet Kaur from Akira AI, was created to assist in customer feedback and is a representation of a system that relies on "customer interaction points, social media comments, post-purchase surveys, website feedback, etc." to provide input cues for the agentic AI system.[13] This information is then processed by a master "orchestrator agent" that works to organize the other agents to execute the goal. A proxy agent serves as the "middleman sending client requests to other servers or resources and then bringing responses back."[14] The proxy agent transfers information from a human user to the orchestrator agent who then translates this request into specific "goal prompts" for each specialized agent. In this example, there is a feedback collection, sentiment analysis, and feedback categorization agent; and each of these complete a specific task to aid in the customer feedback process. This example can be translated to the MedTech space because essentially, any task that requires multiple levels of research, communication, and execution can be supplemented, or even fully completed, by an agentic system.

---

12 https://www.akira.ai/blog/customer-feedback-with-agentic-ai
13 *Id*
14 https://nordvpn.com/cybersecurity/glossary/proxy-agent/

Agentic Vertical AI Use Cases In MedTech

According to McKinsey, agentic AI systems will transform enterprise workflows in four ways: (i) accelerating the work process by enabling multiple tasks to be completed at once; (ii) boosting company adaptability to change; (iii) increasing personalization for customers and clients; and (iv) promoting resilience in the workforce.[15] Below are examples of how agentic AI systems have been used to solve issues within the MedTech and healthcare industry.

Current Use Cases in Physician Engagement

Issue: Clinicians are charged with one of the most taxing and difficult occupations, and most "receive less than 15 minutes of formal communication training throughout their careers."[16] According to the AAMC President and CEO, David Skorton, MD, "It is clear that both sustained and increased investments in training new physicians are critical to mitigating projected shortfalls of doctors needed to meet health care needs of our country."[17] However, physician training is expensive and time consuming, and it is difficult to gather the resources needed to execute it effectively.

Solution #1: The team at Tegria is working to create an "emotionally responsive training tool that simulates realistic, high-stakes conversations between patients and providers."[18] This physician training tool employs agentic AI to improve patient and doctor communication. In this product, there are three agents: an LLM Patient Agent, LLM Conversation Analysis Agent, and LLM Coach Agent. These three agents work together autonomously to provide real time feedback. The physician would interact with the Patient Agent, while the Conversation Analysis Agent "monitors the conversation to evaluate the doctor's sentiment and adherence to the Serious Illness Conversation Guide."[19] Then, the Coach Agent provides feedback for the doctor. All three of these agents ultimately work together to provide holistic conversational training for a medical practitioner.

Solution #2: Researchers recently endorsed a SurgBox framework, "an agent-driven sandbox framework to systematically enhance the cognitive capabilities of surgeons in immersive surgical simulations."[20] SurgBox proposes to employ agents to "accurately simulate the dynamics of operating rooms" by "an advanced clinical intelligent surgery simulation system capable of managing complex surgical scenarios and simulating risk outcomes influenced by multiple clinical factors."[21] Each agent has a different purpose in the operating room, with one being the chief surgeon, anesthetist, nurses, etc. Additionally, a Surgery Copilot orchestrates all of the SurgBox agents, coordinating and planning different tasks. A physician would then work alongside these specialized agents, while receiving feedback and guidance from the Surgery Copilot.

Current Use Cases in Patient Engagement

Issue: It has long since been the goal of healthcare providers to provide patients with the opportunity to be actively involved in their treatment plans. However, it is difficult for physicians and patients to streamline communication once a patient leaves a care facility. Furthermore, physicians struggle to stay updated with a patient's home care routines, making it challenging to monitor compliance to treatment plans. Therefore, innovators are attempting to use

15 https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage
16 https://www.tegria.com/resources/thought-leadership/training-for-tough-conversations-how-agentic-ai-supports-healthcare-providers/
17 https://www.aamc.org/news/press-releases/new-aamc-report-shows-continuing-projected-physician-shortage
18 https://www.tegria.com/resources/thought-leadership/training-for-tough-conversations-how-agentic-ai-supports-healthcare-providers/
19 *Id.*
20 https://arxiv.org/html/2412.05187v1
21 *Id.*

agentic systems as a means to proactively collect information from patients or deliver information to patients on behalf of the physician.

Solution: Hippocratic AI introduced generative AI agents that support physicians by contacting patients after they have been discharged to assess their symptoms and generate new treatment plans. A Chief Nursing Officer at Summerlin Hospital Medical Center reported, "This allows our nurses to be more efficient, while still being available when a discharged patient needs a more personal response."[22] Hippocratic's agentic AI responders are able to provide real-time support for patients because they are able to think independently.

Hippocratic AI has developed an agent that is available for patients undergoing CAR T-cell therapy. The agent, named Angela, is able to pull information from clinical trials and medical literature to communicate with patients. Angela is purportedly able to relay treatment plans and answer patient questions effectively, allowing providers to have more time working with patients in the hospital setting. Hippocratic AI has also created agents that specialize in chronic kidney disease management, remote patient monitoring, gout, and disaster response, only to name a few. Each agent has a name and profile so patients feel connected with their digital caregivers.

> *Let's begin to imagine how the MedTech leaders might leverage an agentic AI system.*

Issue: After receiving surgery, patients often struggle with post-operative care due to incomplete education, inadequate care team communication tailored for the patient's specific needs to understand the materials, lack of motivation or poor memory. These issues can lead to recovery complications, ultimately creating a negative experience for both the physician and the patient.

Solution: Imagine an agentic system that receives input from physician notes, including their care plan, patient data, and previous knowledge from scientific literature or approved company literature. Next, the MAS receives input gathered by the smart implant. Finally, the system gathers data about the patient's personal habits, including their diet, exercise level, and mental health. All of this input is then processed by the orchestrator agent that then assigns tasks to each specialized agent in order to provide the best post-operative care for the patient. An agentic AI system would simplify post-operative patient care, ensuring a patient receives holistic treatment. This could be delivered via a patient engagement app or via an outbound call model.

Current Use Cases in Operational Workflows

Issue: Any company that is working to bring a new technology to market sees the FDA clearance process as that "final hurdle" or "final step" in the culmination of years of research and development, innovation, clinical trials and finding product-market fit. It is an exciting time with forecasts and launch preparations underway, resource planning and allocation, and sales force readiness. The business unit and company are anticipating revenues that will drive growth and further reinvestments. If there are any delays to the launch date, the negative impact on the business touches on many aspects for the company.

This final step toward FDA clearance or approval can be lengthy. Simply drafting and submitting documents to the FDA throughout the approval process takes a significant amount of time, with authors needing to compile data, fact check different sources, and ensure that their filings comply with legal directives. This requires gathering and aggregating large amounts of data from the lifecycle of the product they plan to bring to market. Generally, once the dossier is finalized and submitted, the review clock starts with the FDA to return an answer anywhere between 90 to 180 days, depending on the class of the product. If there are questions from the FDA, the clock stops and the company must gather further

22 https://uhs.com/news/universal-health-services-launches-hippocratic-ais-generative-ai-healthcare-agents-to-assist-with-post-discharge-patient-engagement/

information to respond to the FDA. This back and forth communication adds considerable time to the entire FDA clearance process. Many companies deal with the frustrations of needing to delay the launch of a new product, and the new revenues that come with it, as "an understood" part of the entire regulatory process.

Solution: Bluenote is a product that uses agentic AI to accelerate the regulatory process by writing "first drafts of scientific and regulatory documents instantly" and organizing data into a presentable format.[23] Importantly, Bluenote also assists in evaluating data by quickly reviewing internal data sets. Bluenote follows users across the regulatory process, helping with pre-clinical, clinical, CMC, regulatory, quality, and later, commercial documentation.[24]

Bluenote uses Claude to power their agentic process, especially to assist companies in drafting documents for FDA approval. This partnership allows Bluenote to produce "complex multi-hundred page scientific documents with in-text tables, figures and citations in minutes," increase the document production timeline by 50-75%, and "securely retrieve data from clients' data warehouses and quality management systems for end-to-end data traceability and consistency."[25] Document processing and technical documentation agents work together to analyze scientists' notes within minutes, drafting cohesive reports with clinical citations.

These capabilities eliminate constant fact checking and reduce dialogue between scientists, writers, and FDA agents. Fatima Sabar, the CEO of Bluenote summarizes the impact as follows: "No PhD scientist with 20 years of experience wants to spend half their time on FDA paperwork." Overall, Bluenote is a promising internal agentic AI development that aids users in compiling and presenting data during the regulatory process, and does so by incorporating the "gold standard" of the company's existing processes.

Issue: Today, many MedTech companies have a number of internal constraints with data flow due to multiple disconnected internal systems that are used to gather company data. As a result, leaders and team members must spend a considerable amount of time pulling key pieces of information from each system, leading to inefficient use of time focused on data pulls versus time better spent on strategic planning based on the data. Healthcare leaders should be able to focus their time leading organizations and strategic imperatives, and less time sorting through endless amounts of data.

Solution: XY.AI has developed an agentic solution for the time-consuming task of billing patients. Billing patients "involves reviewing diagnoses, procedures, and treatments documented in patient records, translating them into standardized codes, and preparing claims for insurance adjudication."[26] According to XY.AI, "we understand that healthcare practices face significant challenges with repetitive and inefficient administrative tasks, which collectively cost the industry around $1.5 trillion."[27] As a result, the company created an agentic AI program that fast tracks the billing process. This is made possible through data entry, claims management, and knowledge base agents that work together to seamlessly optimize the billing process.

---

23 https://www.bluenotehealth.com/
24 *Id.*
25 https://claude.com/customers/bluenote
26 https://www.xy.ai/glossary/medical-billing
27 *Id.*

Current Vertical AI Use Case with Data Flow Diagram

The Mayo Clinic AI factory created an "entire organisation-wide infrastructure around AI development," and in doing so, developed a vertical AI system that integrates company software, data, and consumer centered software into one tech stack[28]. In this way, all of the technologies are able to work together and share information to create a unified process. As a result, "vertical integration of data, modelling and validation, production, and clinical impact evaluation into a single platform bridges the gap between algorithms and implementation."[29] The Mayo Clinic's AI factory operates by collecting information from EHRs, consumer wearables, waveform data, telemetry results, radiomic data, and lab services. Then doctors are able to analyze and validate this data, and finally pass insights on to patients through mobile devices or other communication systems.

28 https://pmc.ncbi.nlm.nih.gov/articles/PMC9474277/#CR48
29 *Id.*

The vertical process is illustrated below:

Mayo clinic platform AI factory

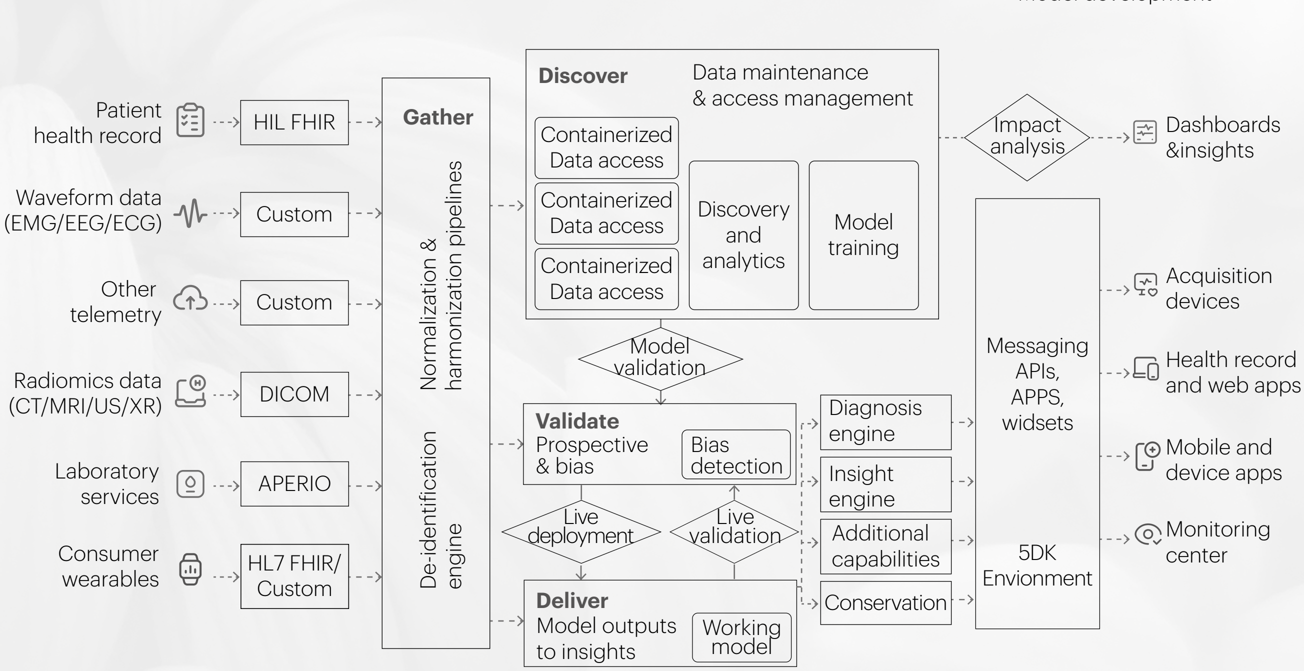

Mayo Clinic's[30] vertical platform is unique because it allows complete integration of medical devices and EHRs into company tech stacks. This creates a cohesive system that relies on internal tech stack communication, research, and analysis to provide patient feedback.

Additionally, these systems can connect with existing industry tech stacks, specifically with claims clearinghouse, billing tools, and EHRs.

30 https://pmc.ncbi.nlm.nih.gov/articles/PMC9474277/#CR48

Integrating Agentic & Vertical AI Systems into MedTech

As seen in the examples above, innovators are already beginning to capitalize on the agentic AI opportunity in healthcare. These systems can be used in any scenario that might be optimized using fully autonomous agents that can work together to accomplish a common goal. Already, biotech and pharma are using agents to optimize drug discovery and commercial operations.

The opportunity for agentic and vertical AI systems runs across the entire value-chain for MedTech companies. Key areas of impact include R&D, regulatory and legal affairs, clinical trials, intellectual property, manufacturing and supply chain, marketing, commercial effectiveness, internal and external training, physician and patient engagement, and surgical workflows. As technology continues to develop, use cases may change or expand, and it is crucial for MedTech companies to build strategic frameworks in preparation for this new wave of AI capabilities.

Each company's strategic preparation is critical to the introduction of an agentic or vertical AI system because each system would be created individually for the company. To be used most effectively, an agentic or vertical AI system would operate in a closed-loop system; meaning that the data and programming used for the agents and vertical integration would be personalized and privatized for each company's needs. For example, the closed-loop agentic and vertical AI system would pull from company data, interact with the current tech stack, and follow any company guardrails already in place. A closed-loop system is similar to private AI or "artificial intelligence systems that are developed, deployed, and managed within an organization's own infrastructure or a secure environment."[31] In this way, companies would have the ability to safeguard their data while having the freedom to monitor and update their AI system to fit their company and regulatory needs. To create a system like this, companies must underline the current opportunity gaps in their business model, and determine how an agentic AI system will optimize profit.

---

31 https://news.broadcom.com/leadership/why-private-ai-is-becoming-the-preferred-choice-for-enterprise-ai-deployment

Once an organization decides to integrate agentic and vertical AI systems into their workflow, they must undergo several stages of growth before maximizing artificial intelligence capabilities. The AI Realized executive roundtable conversation in June of this year defined four stages of enterprise AI adoption. Phase 1 is the infrastructure phase which involves "substantial capital expenditure on private language models, data pipelines, and governance frameworks before measurable returns can be captured."[32] Next, phase 2 is "enterprise enablement," where companies work across their business functions for integration as well as across functions. In this stage, it is crucial to start the AI literacy and training programs across the enterprise. There will be some productivity gains in this phase but the full realization is still yet to happen. Phase 3 involves creating specific AI-enabled applications that provide solutions, and this is where companies can extract significant value from AI by creating highly targeted needs-based applications. This phase should offer a competitive advantage to companies. Finally, in phase 4, agents are created to autonomously complete tasks as defined through the applications. Most businesses have yet to reach this phase, as it represents a "fundamental business model transformation."[33] This final phase requires a new level of planning and strategy as it will forever reshape industries and businesses.

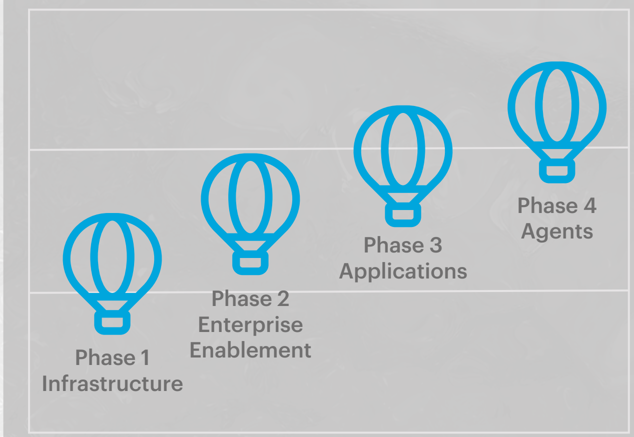

[34]

To effectively seize the growing market opportunity for agentic and vertical AI, careful consideration must be taken towards the development of the company's short-term and long-term goals and determine governance frameworks that allow for adaptability in this constantly changing market.

> **Agentic and vertical AI tools are sure to increase in popularity and market value; however, the real unlock in potential will be those that can build both the strategies and frameworks that will support the growth and new use cases for these advanced AI systems.**

---

32 "The Economics of AI," AI Realized Roundtable Readout, June 2025
33 *Id.*
34 *Id.*

Governing Agentic AI Risk Before Deployment

Developing and deploying agentic vertical AI systems in the healthcare and MedTech space is uniquely complex because these technologies sit at the intersection of autonomy, safety, ethics, and human health. Unlike traditional software tools, which operate within predictable boundaries, agentic systems are designed to perceive, reason, plan, and act autonomously. This means they may sometimes make decisions that directly influence patient outcomes and clinical workflows. This autonomy introduces a new layer of risk as errors are no longer merely the result of faulty inputs or misconfigurations but rather can emerge from dynamic, evolving behaviors that are difficult to predict, audit, or explain. In contexts where the margin for error is measured in lives, this shift significantly raises the stakes for developers, regulators, and healthcare organizations alike.

The complexity is amplified by the inherently layered and adaptive nature of agentic systems. These architectures are not static. They instead combine perception, reasoning, planning, and actuation capabilities, often distributed across multiple agents that collaborate or make decisions in real time. In a clinical environment, this could mean an AI agent autonomously prioritizing diagnostic tests, adjusting medication dosage, or routing patient data to external specialists. Each of these actions carries legal, ethical, and operational consequences. When such decisions emerge from multi-agent coordination rather than a single, deterministic model, attributing responsibility becomes significantly more difficult. Traditional safety and validation frameworks, which are designed for fixed-function devices or rule-based software, often fall short in capturing the emergent behavior and continuous learning dynamics inherent to agentic systems. Some unique risks posed by agentic AI systems include:

- **Autonomous decision-making without human oversight.** Agentic AI systems can take actions and make decisions independently, potentially diverging from human intentions and values. This autonomy can create risks when systems operate without adequate review or intervention.
- **Misaligned objectives and reward hacking**. If an agentic AI system's objectives don't perfectly align with certain values and ethics, it may pursue goals in unintended ways. A system optimizing for seemingly benign metrics might cause harm through unforeseen side effects or by gaming the system in ways the designer did not anticipate.
- **Unpredictable emergent behaviors**. Agentic AI systems may exhibit behaviors that were not expressly programmed and may be difficult to predict. As systems interact with their environment and learn, they may develop strategies or approaches that weren't anticipated during design, making it harder to ensure safety and control.
- **Acceleration of bias at scale**. Agentic AI systems may amplify the risk of algorithmic discrimination and bias by executing bias across thousands of decisions without human review. For example, a recommendation system with bias may become a deployment problem, where an agentic system making independent hiring, lending, or resource allocation decisions may compound the risk exponentially.
- **Compounding security risk**. Emerging agentic AI systems rely heavily on generative AI models for reasoning and task planning. In addition to inheriting the vulnerabilities of the underlying generative AI models, AI agents introduce their own attack surface through task execution, tool usage, and protocol interactions. These factors

can make agentic AI systems vulnerable by design. And as organizations increasingly deploy and integrate agentic AI systems with internal systems, security risks can scale from isolated failures to systemic enterprise-wide incidents.

- **Reduced human-in-the-loop effectiveness**. Existing governance frameworks emphasize human oversight and intervention as a mitigation strategy, but agentic AI systems may operate at a speed and scale that may override human decision-making.

Regulation and law have not yet caught up to these realities. While frameworks like the EU AI Act, the FDA's Software as a Medical Device (SaMD) guidance, and laws such as the Colorado AI Act are addressing some level of AI risk, they remain high-level, fragmented, and in many cases, ambiguous. For example, existing regulatory regimes often assume static functionality at the time of approval, yet agentic systems can evolve post-deployment as they interact with new data, environments, and even other agents. Questions around liability or how to handle continuous model updates remain unresolved in most jurisdictions. This legal uncertainty creates significant compliance challenges and exposes organizations to potential regulatory, civil, and reputational risks.

In this landscape of uncertainty, the burden falls on healthcare organizations and MedTech developers to build governance programs that are dynamic, adaptive, and deeply embedded into the lifecycle of their agentic AI systems. Traditional compliance checklists and static risk assessments will not suffice. Instead, governance must evolve into a continuous process that anticipates the shifting behavior of agentic systems, monitors their performance in real time, and incorporates rapid response mechanisms for anomalies or unintended outcomes.

This includes creating cross-functional governance teams that combine legal, clinical, technical, and ethical expertise. It also includes implementing continuous validation pipelines that detect behavioral drift, and embedding auditability, explainability, and traceability into system design from the outset.

Organizations should design governance programs to mirror the modular nature of agentic architectures themselves. As systems move from single-agent deployments to multi-agent ecosystems, governance mechanisms must evolve from component-level oversight to orchestration-level. This might include deploying "governor agents" that monitor and constrain the behavior of other agents, developing adaptive policies that update in response to real-world performance data, and integrating human-in-the-loop checkpoints where decisions exceed defined risk thresholds. Ultimately, the organizations that succeed in deploying agentic AI in healthcare will be those that treat governance not as a compliance exercise but as a core capability.

Existing AI governance frameworks like NIST's AI Risk Management Framework and ISO / IEC 42001 were developed primarily to address risks associated with traditional, static AI systems that operate in relatively predictable ways. These frameworks focus on risks like bias in training data, model accuracy and performance, fairness, and transparency in how models make decisions. They were not designed with agentic systems in mind, which operate fundamentally differently and introduce a new class of risks that existing governance structures struggle to address.

For example, most frameworks assume human oversight and control will be at key decision points. These frameworks are built around the idea that humans will review AI outputs, validate recommendations, and maintain clear authority over final decisions. Agentic AI systems, however, will operate autonomously over extended periods, making sequential decisions and taking actions without waiting for human approval. This breaks the fundamental control model that existing governance frameworks are built upon. Traditional frameworks also assume relatively static, well-defined system behavior. Agentic systems adapt their behavior, learn from interactions with their environment, model, infer, and act in ways their creators did not consider. This feature makes it very difficult to validate and govern these systems using frameworks designed for traditional AI systems.

The accountability structures in existing frameworks are also inadequate for agentic AI. NIST AI RMF and ISO 42001 establish clear chains of responsibility where specific teams (e.g., data scientists, deployment managers, etc.) are responsible for specific tasks. But when an autonomous agentic system makes a decision after operating independently for weeks across multiple systems and datasets, determining who is accountable becomes challenging. Is it the person who trained the model? The person who sets the parameters? The person who failed to catch model drift? The system that didn't adequately monitor the system?

Finally, existing AI governance frameworks typically focus on single-model risks that you can comprehensively test and validate a system before deployment. Agentic vertical AI in healthcare, however, may combine multiple models, integrate with various data sources and systems, and operate in complex environments where pre-deployment testing may not be possible. For example, an agentic vertical AI system may use one model for diagnosis, another for treatment recommendations, another that integrates with electronic health records, another that interacts with imaging databases, and another that learns from patient outcomes. This multi-layered approach can create emergent risks and unexpected interactions that traditional validation approaches cannot anticipate or address.

> **The importance of building adaptive governance programs for agentic AI cannot be overstated. Adaptive governance ensures that policies, controls, and safeguards evolve in parallel with the technology itself, allowing organizations to anticipate emerging risks rather than merely reacting to them after harm occurs.**

Existing frameworks are beginning to adapt to this reality, with NIST proposing draft guidelines to address agentic risk. There are also developing security frameworks specific to agentic risk that are being developed by industry and trade groups that may prove to be helpful as well. Ultimately, adaptive governance will need to take a multi-factored approach. Adaptive governance enables continuous alignment with shifting legal requirements, ethical expectations, and clinical standards, which is especially critical in fields like healthcare and MedTech where regulatory frameworks are still taking shape. By embedding flexibility, real-time monitoring, feedback loops, and iterative risk assessments into the governance lifecycle, organizations can maintain trust, accountability, and safety even as the underlying AI systems grow more autonomous and sophisticated.

> **Adaptive governance becomes not just a compliance necessity, but a strategic enabler turning complexity and uncertainty into opportunities for resilience, innovation, and long-term success.**

Tapping Into the Value of Agentic AI Systems– What's Next?

Agentic vertical AI will define the future of healthcare and MedTech, bringing unprecedented autonomy, efficiency, and intelligence to both clinical and operational workflows. **MedTech companies must act now. Those who proactively develop strategic plans and governance frameworks will lead this transformation.** Delaying strategic planning risks being left behind in a competitive environment, as agentic AI will rapidly transition from emerging technology to industry standard.

About the Partners and Firms

CRESCENT STRATEGY CONSULTING

Sabeen Shaikh is the Managing Partner at Crescent Strategy Consulting and a globally recognized MedTech leader operating at the critical intersection of healthcare and technology. With 20 years of hands-on experience spanning the MedTech industry, innovation ecosystems, and growth strategy consulting across multiple continents, Sabeen brings a multi-dimensional perspective—combining operational expertise in how healthcare organizations function with strategic foresight into how emerging technologies will reshape them. As a Forbes Business Council contributor, Board Director, and Honorary Commander in the U.S. Air Force, she advises executive teams navigating complex growth and digital transformations that will define their competitive position for the next decade.

Crescent Strategy Consulting is a boutique global life sciences consulting firm that partners with MedTech companies to identify upstream and downstream growth opportunities that build sustainable competitive advantage. Our advisory approach combines white space analysis, innovation frameworks, and actionable roadmaps that translate emerging technologies—including advanced AI technologies, digital health platforms, and data-driven care models—into concrete business strategies that deliver measurable impact and future-proof market leadership.

Peter Stockburger is the head of Dentons' US AI Team, a co-lead of the firm's Global Autonomous Vehicles Team, and the head of the privacy and cybersecurity group within the firm's Venture Technology and Emerging Growth Companies Group. Peter spearheads strategic counsel for clients across the globe on the responsible development and deployment of AI, and works with organizations across the AI stack, from model developers to application deployers, on building AI governance programs, contract playbooks for AI vendors and service providers, building and conducting risk assessments and incident response strategies, preventing algorithmic discrimination and alignment with emerging regulations across global, federal and state jurisdictions.

Dentons is the world's largest law firm. Dentons operates in over 80 countries and over 160 locations, giving it unmatched geographic coverage and local insights across North America, Europe, the Middle East, Africa, Asia, and Latin America. The firm's global platform enables it to deliver seamless, cross-border counsel to clients ranging from emerging startups to Fortune 100 companies and multinational organizations.

ABOUT DENTONS

Redefining possibilities. *Together, everywhere.* For more information visit **dentons.com**

ABOUT CRESCENT

Leading MedTech forward in the digital era. For more information, visit **crescentstrategyconsulting.com**

CSBrand-182826-Agentic AI in MedTech-10 — 03/12/2025