Avoid the Pitfalls of Healthcare AI

What leaders must do differently

June, 2026

Public

Purpose statement

This document provides an overview of how artificial intelligence (AI) can deliver sustained value. It is intended solely to help you assess the business benefits of integrating AI into clinical and operational workflow.

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

This document may include some forward-looking content for illustrative purposes only. Some products and features discussed are indicative of the products and features of a prospective future launch in the United States only or elsewhere. Not all products and features discussed are currently offered for sale in the United States or elsewhere. Products and features of the actual offering may differ from those discussed in this document and may vary from country to country. Any timelines contained in this document are indicative only. Timelines and product features may depend on regulatory approvals or certification for individual products or features in the applicable country or region.

ORACLE

Table of contents

Artificial intelligence (AI) holds significant promise to improve healthcare efficiency, reduce costs, and enhance patient outcomes. Yet many healthcare AI initiatives fall short—not because the technology is insufficient, but because deployment is misaligned with real-world operations. For healthcare leaders, the challenge is ensuring AI delivers sustained value. Achieving that goal requires integrating AI into clinical and operational workflows, built on high-quality data—not confining it to isolated pilots.

Understanding how AI works

Traditional software operates according to explicit rules where every possible decision and action must be defined in advance. Whereas AI systems learn patterns from data to generate predictions or content. This approach enables AI to perform complex tasks such as understanding language, imaging analysis, and pattern detection.

Increasingly, AI models are multimodal, integrating structured data, clinical text, and images, critical in healthcare, where meaningful insights often depend on combining sources of data.

There are different types of AI. **Predictive AI** identifies patterns and forecasts likely outcomes, as in identifies abnormalities in radiology images. **Generative AI** creates new content such as text, code, or images. In healthcare, generative AI can help clinicians summarize patient records, draft discharge summaries and referral letters, and support research through rapid literature review. A newer form of AI, called **Agentic AI**, goes a step further. It can plan, make decisions, and take self-directed actions within boundaries set by human rules.

Together, all these forms of AI can reduce people's manual, repetitive work, while providing data and information to support human decision-making.

Compute

Computing infrastructure is a foundational component of enterprise AI. Modern AI models require significant computational power to train and operate, performing vast numbers of mathematical calculations across large datasets. These workloads are typically handled by graphics processing units (GPUs), which are deployed within data centers that provide the power, cooling, and networking required for AI at scale.

Many organizations access this infrastructure through cloud platforms, enabling them to leverage advanced computing capabilities without building their own facilities. This combination of specialized hardware and cloud infrastructure has accelerated AI development and made it economically feasible to deploy AI across enterprise systems, including high-performance use cases such as genomic analysis, drug discovery simulations, and large-scale population health modeling.

Bolted on versus built in

While AI models are powerful, creating a standalone solution is fundamentally insufficient and risky. Isolated point solutions operate outside core workflows and include standalone AI scribes or analytics tools that identify high-risk patients but can't take any action. These "bolted-on" solutions show limited impact because insights remain disconnected from workflows that are managed by EHR, enterprise resource planning (ERP), human capital management (HCM), and other systems which have the necessary regulatory approvals and compliance built in. These bolted on solutions do not and could compromise the accuracy, reliability, and usefulness of AI, which depend heavily on both the data used during training and the data available at the time of inference.

Standalone AI solutions lack awareness of real-time patient data, local clinical workflows, organizational policies, or continuously evolving medical knowledge unless explicitly connected to those systems. As a result, outputs may be incorrect. More importantly, standalone models lack longitudinal context and are subject to confidently generating responses that appear plausible but are not factually grounded in patient relevance or institutional data.

In healthcare, this is not a minor technical issue but a critical safety risk. For example, a model may generate a well-structured clinical summary that omits a critical lab abnormality, suggest a treatment approach that conflicts with a patient's allergies or current medications, or provide recommendations that do not reflect the latest clinical guidelines or the organization's care protocols.

The most effective approach to delivering AI is to embed it directly within the enterprise applications where work already occurs. When AI is natively integrated into systems such as EHRs, revenue cycle platforms, supply chain systems, and HR systems, it operates with full access to real-time data and within established workflows. This allows AI to surface insights or recommendations precisely at the moment decisions are made, rather than requiring users to leave their workflow or interpret disconnected outputs. By operating within applications, AI can also trigger downstream agentic actions, such as scheduling follow-up care, initiating referrals, or updating care plans, ensuring that insights translate into execution. This improves adoption and ROI.

High impact AI starts with consistent and complete data

Even embedded AI is insufficient in isolation. While embedded AI improves workflow integration within a single system, each application typically operates on a limited subset of data, such as the data contained within a specific EHR or departmental system.

For example, an AI capability embedded within an EHR may have access to clinical data but lack visibility into payer authorization status, supply chain constraints, or post-acute care availability, while an AI tool within a revenue cycle system may submit a claim without key clinical data. Outcomes such as readmission risk are influenced not only by patient characteristics but also by local factors such as care pathways, discharge practices, availability of post-acute services, and population demographics. As a result, insights generated within individual applications may be locally optimized but globally incomplete, leading to gaps in care coordination, inefficiencies, or suboptimal decisions.

In addition, when AI capabilities are distributed across multiple applications without a unifying data and governance layer, organizations face challenges in maintaining consistency, ensuring data quality, and enforcing enterprisewide policies. This can result in conflicting, unnecessary costs, recommendations, duplicated logic, and increased operational complexity.

These limitations highlight a critical principle: embedding AI into applications is necessary but not sufficient. **To deliver meaningful, systemwide impact, embedded AI must be supported by a unified data platform that integrates data across systems, provides consistent context, and enables coordinated intelligence across the enterprise.** Only by combining embedded AI with a shared data foundation can organizations move from isolated pockets of automation to truly integrated, enterprise-scale intelligence.

Preparing enterprise data for AI

Successful AI strategies depend on a broader technological foundation and investment in modern data architectures capable of integrating, structuring, and governing data across the enterprise. This is a challenge because healthcare data is inherently fragmented and heterogeneous, spanning structured fields, unstructured clinical notes, imaging data, and genomic information, often stored across systems that use different standards and vocabularies.

A modern data platform addresses this challenge by transforming fragmented data into a unified, high-quality foundation that AI systems can reliably interpret. It connects foundation models to trusted enterprise data through mechanisms such as retrieval-augmented generation (RAG), which retrieves and incorporates relevant information at query time, as well as semantic layers or knowledge graphs, helping to ground outputs in current, context-specific information that can evolve to meet the needs of specific customers and workflows faster than the foundation models.

A semantic data layer maps raw data into standardized, normalized data, addressing variations in terminology such as "myocardial infarction" versus "heart attack" and ensuring lab values are interpreted consistently. Building on this foundation, healthcare knowledge graphs capture relationships among clinical and operational entities over time, representing patients, treatments, providers, events, and resources as interconnected nodes rather than isolated record (for example, linking a diagnosis to medications, lab results, providers, and outcomes across time). Together, the semantic layer and knowledge graph can provide structured, contextualized data that AI models can use both during training and during decision-making.

Legacy technology is not built for AI

The ability to prepare data for AI is difficult due to the underlying technology landscape, where legacy systems, including 1970s MUMPS-based data architectures, are optimized for transactional performance but not designed for modern AI workloads. Even when AI is embedded within applications, those applications typically lack the ability to integrate and contextualize data from other parts of the enterprise in real time.

Without modernization, organizations are often forced to maintain hybrid architectures in which legacy systems function as systems of record while separate platforms serve as systems of intelligence. This dual architecture introduces complexity, cost, and latency, and makes it more difficult to fully embed AI into workflows. By contrast, modern AI-enabled data platforms integrate structured and unstructured data, support semantic indexing, and enable real-time retrieval and reasoning within a unified environment, reducing architectural complexity and enabling deeper integration of AI into enterprise systems.

Data security

Another critical consideration is data residency and security. In many externally hosted AI models, prompts and data are transmitted to third-party infrastructure, introducing risks related to privacy, regulatory compliance, and exposure of proprietary information. For example, sending patient summaries or clinical notes to external model providers may raise HIPAA and data sovereignty concerns. Modern AI platforms can overcome this challenge by enabling organizations to deploy models within private, dedicated, or sovereign environments, allowing them to bring AI to the data rather than moving data to external systems.

Interoperability: the missing link between data and impact

Even the strongest data platform and AI strategy will underperform without interoperability, the ability for systems to exchange data reliably, securely, and in usable form across the care continuum. Interoperability is also a governance and operating-model issue. Organizations need clear policies for data access, patient matching/identity resolution, provenance (from where data came), consent and privacy controls, and auditing, especially when AI is involved.

Equally important is addressing non-technical barriers, including contractual constraints, inconsistent workflows, and “data blocking” behaviors that can limit exchange even when standards exist. Leaders should treat interoperability as a strategic capability, not an integration afterthought, because it directly determines whether AI can operate on a complete, trusted, and current picture of patient and system state.

Governance, validation, and trust

Successfully deploying AI in healthcare requires more than technical capability. It requires robust governance, validation, and ongoing oversight. Trustworthy AI requires human oversight, ensuring that AI systems support rather than replace clinical and operational decision-making, with qualified professionals retaining responsibility for final decisions (for example, clinicians validating AI-generated recommendations before acting on them.) Organizations must clearly define accountability, ensuring that AI supports, rather than replaces clinical judgment, and that appropriate safeguards are in place when outputs are uncertain or conflicting.

In addition, transparency is essential. Clinicians and administrators must be able to understand how models are inferenced with HER data, the purpose, data inputs, limitations, and appropriate use of AI-driven recommendations. This aligns with emerging regulatory expectations, including U.S. Food and Drug Administration (FDA) guidance on software as a medical device and Office of the National Coordinator for Health IT (ONC) requirements for decision support transparency within certified health IT.

Healthcare organizations must also evaluate AI systems across their full lifecycle, including initial validation, deployment, and continuous monitoring. This includes assessing performance across different patient populations, monitoring for model drift over time, and ensuring outputs remain clinically appropriate within evolving workflows.

The next era of healthcare

Healthcare AI will not fail because the technology lacks capability, but because organizations underestimate what it takes to operationalize it effectively. As this paper has shown, success depends on moving beyond fragmented, “bolt-on” solutions toward deeply embedded, enterprisewide intelligence supported by unified data, modern architecture, and rigorous governance. AI must be integrated into real clinical and operational workflows, grounded in high-quality, context-rich data, and aligned with the complex realities of healthcare delivery.

For healthcare leaders, the path forward is clear but demanding. It requires investing in foundational data platforms built for healthcare AI, interoperability at the data and AI layers, and governance frameworks that enable healthcare workers to use these models safely and effectively at scale.

It also requires maintaining human oversight, ensuring transparency, and building trust across clinicians and administrators. And to truly leverage AI’s transformational capabilities, healthcare leaders must also transform the way work is done in their health systems, redesigning workflows and business processes, investing in training and change management, and developing appropriate financial and non-financial performance metrics to ensure real value is being created through their AI investments.

Organizations that approach AI as a strategic capability, rather than a standalone tool, will be positioned to unlock its full potential, improving outcomes, efficiency, and patient experience. Those that do not risk falling short, not because AI cannot deliver value, but because it was never truly set up to succeed.

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com   facebook.com/oracle   twitter.com/oracle