Resultant

Healthcare AI:
What's Real, What Works, and What to Watch

CONTENTS

INTRODUCTION

As a healthcare data and analytics leader and former CDO for Indiana's largest health care provider, I experience the world of AI information overload daily. I also seem to be on everyone's mailing list. I opened my inbox this morning to 14 unread emails, every one of them promising to "revolutionize healthcare with AI." Whitepapers, webinar invites, breathless press releases, each one louder than the last. If you're anything like me, you've stopped reading most of them. Not because you don't care, but because the noise has become deafening.

What's missing in all that noise is a practical way to think about what AI in healthcare actually looks like today. Not the future promises, but the capabilities **already being used across the industry**, often quietly and often at less-than-ideal scale. One way to frame it is not as "Artificial Intelligence," but as **Assistive Intelligence**. These are real-world technologies that help providers, nurses, administrators, and back-office staff do their jobs more efficiently and consistently while improving patient experience, provider satisfaction, and health system performance.

Most healthcare organizations already use multiple forms of AI, even if they're not labeled as such. Clinical alert systems, imaging tools that flag abnormalities, predictive models that estimate readmission risk, and documentation assistants that convert conversations into notes all rely on different AI techniques. What often gets called "AI" in vendor marketing isn't a single technology but a collection of capabilities built for different types of problems.

WHY UNDERSTANDING AI TYPES MATTERS

This distinction matters because healthcare leaders are increasingly being asked to evaluate AI investments without a clear framework for understanding what they're buying. Two vendors may both claim to offer "AI solutions," yet one may be delivering a rules-based alerting system while another is deploying deep learning models that analyze imaging data. The technologies, risks, and operational implications can be very different.

AI's current capabilities span everything from simple pattern recognition to more advanced systems that support complex decisions and workflows. Understanding the different types of AI helps clarify what problems these tools can realistically solve in today's healthcare environment and where expectations should remain cautious. AI investments, whether single-point solutions, custom solutions, or a combination of both, should be tied to clearly defined clinical, operational, and financial goals.

In the sections that follow, I'll walk through several common types of AI used by health systems and providers today, along with the practical benefits they offer and the associated challenges organizations must be prepared to manage.

> **Each section also includes a brief question to help healthcare leaders evaluate how these capabilities currently appear in their own organizations.**

RULES-BASED AI

The pop-up warning that appears when a clinician prescribes two medications that may interact is a familiar example of rules-based AI at work.

What it is: Rules-based AI systems utilize predefined rules and logic trees (if-then logic or knowledge bases with inference engines) and are reliant upon explicitly programmed logic; they do not learn from data over time.

Use cases for healthcare providers:

- **Clinical decision support alerts** that warn clinicians about potential drug interactions, allergies, or dosage conflicts when prescribing medications in the EHR.
- **Emergency department triage protocols** that prioritize patients based on predefined symptom and risk criteria.
- **Medication safety monitoring** systems that automatically check prescriptions against patient history, lab results, and known contraindications to prevent adverse events.
- Clinical pathways that translate **evidence-based treatment guidelines** into recommended care steps.

Strengths:

- High transparency and auditability (decision logic can be traced)
- Reliable performance for narrow, well-defined problems

Challenges:

- Reliance on subject matter experts to manually create and maintain rules
- Limited ability to handle ambiguity or novel situations
- High volume of alerts can contribute to clinician alert fatigue
- Limited scalability as rules and workflows grow

What this means for healthcare leaders:

Rules-based AI is often the earliest form of AI deployed in healthcare environments. Many hospitals already rely on it within EHR systems and clinical decision support tools, but poorly tuned rules can contribute to alert fatigue and reduce their effectiveness.

Question for healthcare leaders:

Ask yourself: Do your clinical alerts meaningfully improve patient safety and assist care providers, or have they become routine pop-ups causing clinicians to experience alert fatigue?

MACHINE LEARNING

When a hospital forecasts tomorrow's patient census or flags a patient as high risk for readmission, machine learning models are often working behind the scenes.

What it is: Machine learning is a broad category of AI that includes both traditional statistical models and more advanced neural network approaches such as deep learning, AI systems where computers learn to perform tasks by analyzing data and identifying patterns.

In contrast with rule-based AI, these systems are not explicitly programmed with step-by-step logic and instructions. By identifying complex patterns within historical data, these models can generate predictions or insights when presented with new, unseen information.

These approaches power many of the predictive tools currently used across healthcare operations and clinical decision support. The most common include traditional machine learning models, deep learning techniques, and specialized applications such as computer vision.

Traditional Machine Learning

When hospitals use historical data to predict patient census, staffing needs, or readmission risk, they're usually using traditional machine learning models.

What it is: A method for analyzing structured data (like spreadsheets or databases) where **human expertise defines the relevant categories of information**, enabling the algorithm to identify which specific patterns lead to actionable insights.

Use cases for healthcare providers:

- **Predictive models for hospital operations**, such as forecasting daily patient census, estimating bed demand, and helping administrators plan staffing levels across departments.
- **Patient risk prediction models** that identify individuals at higher risk of complications, deterioration, or hospital readmission so care teams can intervene earlier.
- **Population health management tools** that analyze large patient datasets to identify trends in chronic disease, treatment outcomes, or gaps in care.
- **Remote patient monitoring programs** where machine learning analyzes data from wearable devices or home health monitors to detect concerning changes in vital signs.

Strengths:

- Rapid analysis of large healthcare datasets
- Strong pattern recognition in complex clinical data
- Predictive insights derived from historical patterns
- Ability to detect subtle correlations that may not be visible through rule-based logic

Challenges:

- Dependence on high-quality, representative training data
- Limited transferability across different health systems or patient populations
- Limited explainability of model predictions ("black box" challenge)
- Limited ability to incorporate human context such as patient preferences or cultural factors
- Model drift and performance decay as patient populations or clinical practices evolve

What this means for healthcare leaders:

Investing in traditional machine learning enables health systems to turn existing data into predictive insights that improve operations and support financial stability amid ongoing cost pressures. To sustain value, leaders must also invest in data quality, governance, and ongoing model monitoring to ensure performance remains reliable over time.

Question for healthcare leaders:

Ask yourself: As we look to continually improve operational efficiencies, are we investing at the appropriate level in this proven technology and do we have the data quality and monitoring processes in place to ensure our predictive models remain accurate and fair as patient populations and care practices evolve?

Deep Learning (Neural Networks)

When AI tools analyze X-rays, convert patient conversations into clinical notes, or extract insights from large volumes of clinical text, deep learning models often power those capabilities.

What it is: Deep learning is an advanced type of machine learning that uses layered neural networks to automatically discover complex patterns in raw information. Unlike traditional machine learning, these models do not require humans to pre-select which variables matter most. Instead, t**he system learns which signals are most important on its own**.

This approach is particularly effective for analyzing unstructured data, such as medical images, audio recordings of patient visits, and written clinical notes.

Use cases for healthcare providers:

- **Medical image analysis** that assists radiologists by detecting abnormalities in X-rays, CT scans, MRIs, retinal images, or pathology slides.
- **Precision medicine applications** that analyze genetic data alongside longitudinal patient records to identify which treatments or medications may be most effective for an individual patient.
- **Ambient clinical documentation systems** that capture provider–patient conversations and convert them into structured clinical notes within the EHR.
- **Extraction of insights from unstructured clinical data**, including physician notes, imaging reports, and other narrative documentation.

Strengths:

- Exceptional pattern recognition in complex clinical data
- Consistent analytical performance without fatigue
- Rapid analysis of large and complex data sets

Challenges:

- Dependence on large training datasets and significant computational resources
- Limited explainability of model predictions ("black box" challenge)
- Requirement for rigorous clinical validation and human oversight

What this means for healthcare leaders:

Deep learning unlocks the ability to analyze unstructured clinical data such as images, notes, and conversations. However, these models require substantial data infrastructure, validation processes, and clinical oversight to ensure safe and effective use.

Question for healthcare leaders:

Ask yourself: Do we have the data infrastructure, governance, and clinical oversight needed to safely deploy models that analyze unstructured data like images or clinical notes?

Computer Vision

When an AI tool highlights a suspicious spot on an X-ray or flags abnormal tissue in a pathology slide, computer vision technology is performing that analysis.

What it is: Computer vision is an application of machine learning, most often powered by deep learning, that enables computers to interpret visual information such as medical images or video. These models are trained on large datasets of labeled images so they can recognize patterns, objects, and abnormalities.

Use cases for healthcare providers:

Automated analysis of diagnostic imaging where AI highlights potential abnormalities in X-rays, CT scans, MRIs, or retinal images for radiologists to review.

Digital pathology analysis that scans tissue slides to detect cancerous cells, classify tumor types, and identify patterns that may indicate disease progression.

Video-based patient monitoring systems that detect falls, respiratory distress, or patient movement patterns in hospital rooms or long-term care settings.

Screening tools in underserved areas, where AI-assisted imaging helps identify conditions like diabetic retinopathy before referral to a specialist.

Strengths:

- Consistent analysis of visual data without fatigue
- Detection of subtle visual patterns that may be missed by the human eye
- Rapid image processing and case prioritization
- Extension of specialist diagnostic capabilities to underserved settings

Challenges:

- Difficulty translating pixel-level patterns into clinically meaningful interpretation
- Ongoing validation required as imaging technologies and clinical knowledge evolve
- Limited explainability of model outputs ("black box" challenge)
- Security and privacy risks when processing protected health data
- Dependence on diverse, representative training data to avoid biased results

What this means for healthcare leaders:

Computer vision is one of the most mature and impactful AI applications in healthcare today, particularly in radiology and pathology. However, successful adoption depends on careful validation, integration into clinical workflows, and ensuring clinicians retain final decision authority.

Question for healthcare leaders:

Ask yourself: Are we deploying AI imaging tools at the right scale to reduce clinician burden while maintaining clear human oversight and final clinical judgment?

NATURAL LANGUAGE PROCESSING (NLP)

When a doctor's conversation with a patient is automatically converted into a clinical note or when software scans physician notes to identify billing codes, natural language processing is doing the work behind the scenes.

What it is: Natural Language Processing (NLP) allows computers to interpret, extract, and organize human language in both written and spoken forms. By identifying patterns, context, and intent within text, NLP transforms "messy" human communication into structured, usable data. This field of AI leverages a combination of the machine learning and neural network techniques described earlier in this document to handle the complexities of medical terminology.

Use cases for healthcare providers:

- **Extraction of structured information from clinical conversations**, enabling ambient documentation tools that convert provider–patient interactions into draft clinical notes.
- **Structuring unstructured clinical documentation**, transforming physician notes and other narrative records into data that can be searched, analyzed, and reported on.
- **Automated coding assistance**, where NLP scans clinical notes to identify diagnoses and procedures that should be translated into standardized billing codes.
- **Clinical documentation search tools** that allow providers to quickly locate relevant information across large volumes of patient records.

Strengths:

- Significant reduction in clinical documentation burden
- Rapid search and analysis across large volumes of clinical text
- Reduced transcription and documentation errors
- Automation of administrative tasks such as coding and claims processing

Challenges:

- Dependence on high-quality, representative training data
- Privacy constraints that limit access to large clinical text datasets
- Complexity and variability of medical language and abbreviations
- Accuracy risks that require human verification of outputs
- High verification burden if AI-generated content requires extensive review
- Data normalization challenges across multiple healthcare coding systems and standards

What this means for healthcare leaders:

Natural language processing is becoming an increasingly visible AI capability in healthcare because it helps convert large volumes of clinical text into structured, usable data. When implemented thoughtfully, it can reduce documentation burden and unlock insights from physician notes, patient communications, and other unstructured records. However, successful adoption requires careful oversight to ensure accuracy, protect patient privacy, and avoid shifting the workload from documentation to verification.

Question for healthcare leaders:

Ask yourself: Are we using AI to support documentation and coding workflows, and if so, is it reducing workload or simply shifting it to reviewing and correcting AI outputs?

GENERATIVE AI

When software drafts a clinical note from a patient conversation, summarizes dozens of research papers in seconds, or suggests a response to a patient portal message, generative AI is producing that content.

What it is: Generative AI refers to models designed to generate new content such as text, summaries, or synthetic data rather than just classifying existing information. These models are trained on massive datasets to recognize deep patterns and relationships within language and data. Generative AI works in tandem with the technologies previously described. It uses Natural Language Processing (NLP) to interpret a user's prompt and neural networks to predict and generate the most relevant response.

Use cases for healthcare providers:

- **Automated drafting of clinical documentation**, including visit notes, discharge summaries, and referral letters based on provider–patient interactions.
- **Patient communication support**, generating draft responses to routine patient portal messages or creating personalized discharge instructions.
- **Summarization of medical research and literature**, helping clinicians quickly review findings across large volumes of studies and clinical publications.
- **Synthetic data generation for research**, creating realistic but anonymized datasets that allow organizations to train models without exposing protected health information.

Strengths:

- Significant reduction in clinical documentation workload
- Improved consistency in clinical documentation and communication
- Continuous availability for drafting, summarizing, and content generation

Challenges:

- Potential misalignment with specific clinical standards or patient populations
- Risk of bias from non-representative training data
- Limited understanding of specialized clinical context or local protocols
- Data privacy and security risks when interacting with third-party models
- Risk of hallucinated or inaccurate outputs requiring human verification
- High verification burden if AI-generated content requires extensive review
- Ongoing subscription and usage costs at scale
- Operational dependence on third-party vendors

What this means for healthcare leaders:

Generative AI is rapidly gaining attention in healthcare because it can produce new content like drafting clinical documentation, summarizing research, and assisting with patient communication. These capabilities offer meaningful opportunities to reduce administrative workload and improve information access. However, realizing those benefits depends on strong governance, clear policies for human review, and careful integration into clinical workflows to prevent errors or privacy risks.

Question for healthcare leaders:

Ask yourself: Do we have the right structures in place to encourage innovative use of generative AI while ensuring strong patient advocacy and human-in-the-loop oversight?

AGENTIC AI

When software detects that a patient's treatment order requires prior authorization, gathers the required documentation, and initiates the submission process automatically, an AI agent is coordinating that workflow.

What it is: Agentic AI represents a shift from reactive tools to proactive collaborators. Unlike standard AI that simply responds to prompts, Agentic AI works toward specific goals. It is most effective when applied to well-defined, finite tasks with known starting and ending points. Instead of just identifying risk factors or summarizing a patient history, an AI agent can assess that information and **take independent actions within a set boundary**. This includes triggering clinical workflows, coordinating multi-step care plans, or interacting with other software systems to complete specific tasks. These agents learn from the outcomes of their actions, allowing them to refine their strategies over time based on successes and failures.

Use cases for healthcare providers:

- **Diagnostic support systems**, that analyze longitudinal patient records including medical history, lab results, imaging studies, and claims data to identify patterns that may indicate potential diagnoses.
- **Continuous patient monitoring workflows**, where AI agents track data streams from wearables, bedside monitors, or EHR updates and trigger alerts when concerning trends appear.
- **Automated prior authorization processing**, where an AI agent gathers required clinical documentation, checks payor requirements, and prepares submission materials for insurer approval.
- **Operational workflow coordination**, where AI agents assist with scheduling operating rooms, coordinating care transitions, or managing staffing adjustments across departments.

Strengths:

- Automation of complex multi-step workflows
- Ability to trigger actions based on data analysis and clinical rules
- Reliable execution of repetitive operational tasks

Challenges:

- Need for strict clinical safety controls and human oversight
- Potential misalignment with local clinical standards or patient populations
- Risk of hallucinated or incorrect actions requiring verification
- Ongoing subscription and usage costs at scale
- Operational dependence on third-party vendors
- Limited explainability of automated decisions ("black box" challenge)
- Security and privacy risks when interacting with protected health data and connected systems

What this means for healthcare leaders:

Agentic AI has the potential to automate complex healthcare workflows that currently require significant administrative effort. However, because these systems take action rather than simply providing insights, they require strong governance, clear operational boundaries, and continuous human oversight to ensure patient safety and regulatory compliance.

Question for healthcare leaders:

Ask yourself: Are we deploying AI agents to assist and streamline defined workflows, or are we allowing them to operate in ways that could create risks if their actions are not carefully monitored?

Bringing
It All Together

Most healthcare providers don't use just one type of AI; they deploy multiple AI types layered together. For example:

- Natural Language Processing can extract key data from unstructured clinical notes.
- Machine Learning can predict a specific patient risk or forecast future staffing needs based on that extracted data.
- Generative AI can summarize these findings into a concise report for the provider.
- Agentic AI can coordinate the follow up actions, such as scheduling a lab test or sending a notification to the care team, within a well-defined and finite workflow.

This hybrid approach allows AI to support both clinical excellence and operational efficiency, while keeping human clinicians in control of life dependent decisions.

Much of the fear of applying AI in health care settings at the same rate it's being applied in other industries is born out of concerns that a wrong decision will be made or that privacy will be breached or that the underlying data isn't perfect. These concerns are all valid, but they've also been with us well in advance of AI.

When we use AI as "Assistive Intelligence" with a trained human still in the loop, we see successes. We see tools that remove or greatly reduce the manual work required to gather and organize information, allowing providers to focus on making better-informed decisions. As an industry, we certainly need to encourage investment, increased utilization, and faster adoption of these real capabilities. The benefits of doing so are anything but "Artificial."