BECKER'S HEALTHCARE aws

Generative AI + Healthcare Guidebook

6 questions — answered

While generative artificial intelligence (AI) has been around for years, 2023 marked an inflection point in the technology's evolution. In July, Amazon CEO, Andy Jassy called generative AI "one of the biggest technical transformations of our lifetimes."

Healthcare leaders have taken note. Sixty four percent of 1,000-plus hospital and health system leaders recently surveyed by Becker's Hospital Review said generative AI would have a "huge impact" on their industry in the next five years.

Many health systems are already integrating this technology, but the path forward remains unclear. Before healthcare organizations can put generative AI to good use, leaders must first know what exactly generative AI is and does. Strategic planning, integration, and iteration follow. As with any new technology implementation, getting started can be the hardest part.

This guidebook answers six questions to help organizations in the nascent stages of their generative AI journey.

Question 1: What is generative AI?

Generative AI is a type of AI that has the ability to create new content and ideas, such as conversations, stories, images, and videos. It is powered by foundation models (FMs) that are pretrained on vast amounts of data. This technology may have the power to improve productivity by automating administrative processes. Generative AI also has the potential to garner new insights about patients and populations to help predict outcomes and enable health systems to deliver proactive care and interventions.

FMs, including large language models, consist of millions to billions of parameters, and are pretrained on terabytes of text and image data to perform a wide range of tasks such as content summarization and text, image or video creation.

As these models are pretrained, they have the potential to lower training and infrastructure costs and enable customization for specific use cases.

The relationship between artificial intelligence (AI), machine learning (ML), deep learning (DL) and generative AI is shown in the figure below. Complex use cases can be addressed by combining multiple AI and ML technologies. For example, a healthcare wellness app helping a patient manage a long-term chronic condition, such as diabetes, requires a combination of different ML technologies to automate processes, develop insights, and engage patients.

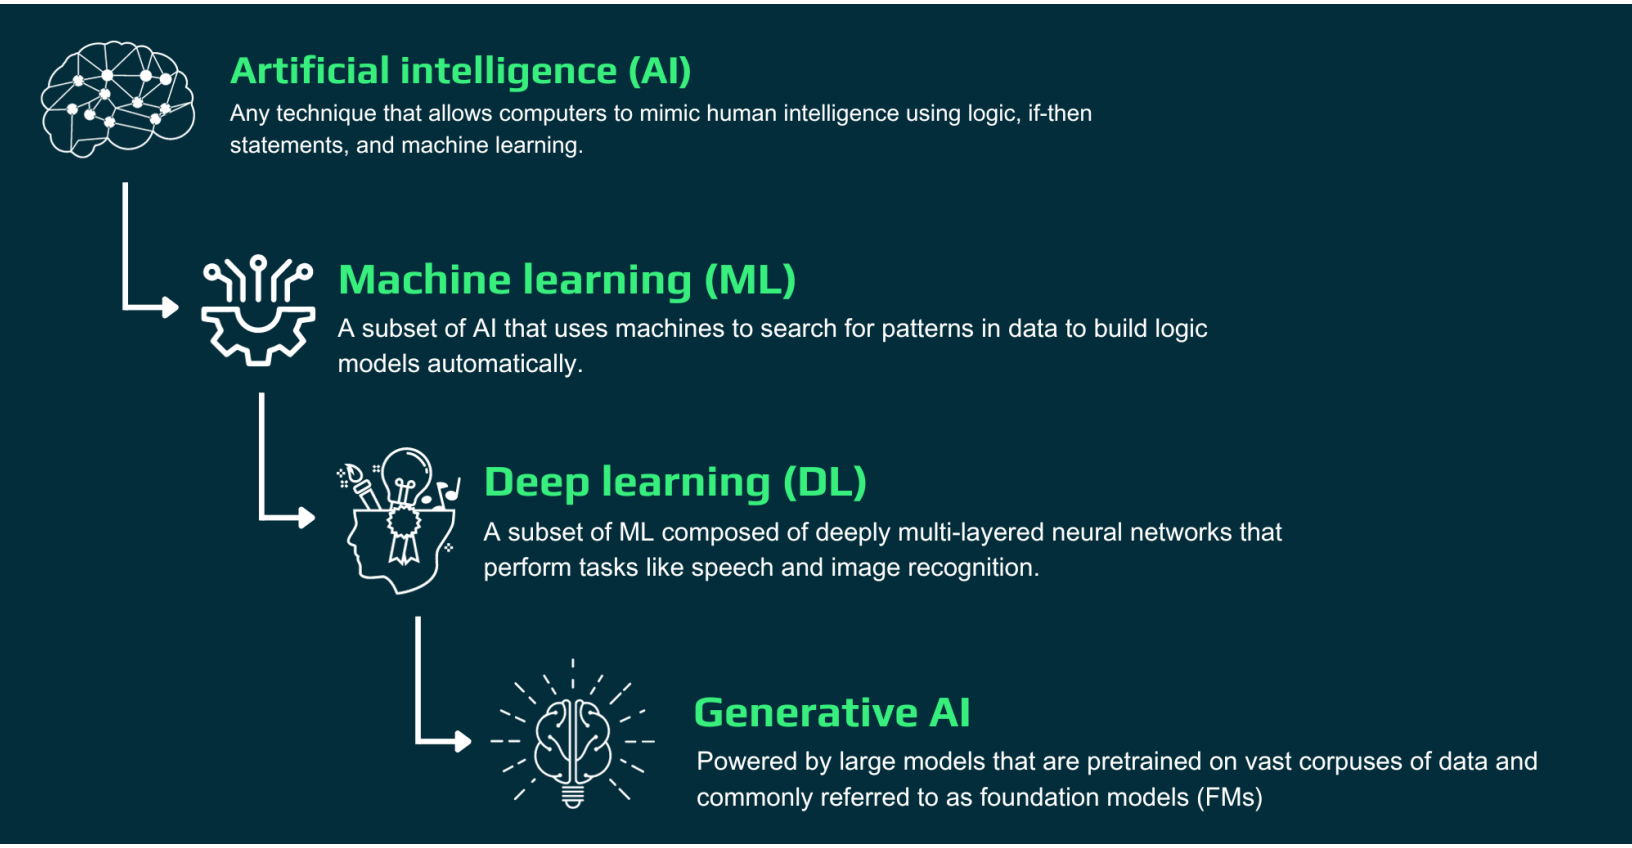

At Amazon, AI and ML research spans decades. This technology is used for personalization on the Amazon retail site, controlling robotics in Amazon fulfillment centers, and enhancing Alexa's intent recognition and speech synthesis. ML is an integral part of Amazon's DNA.

ML innovation is in Amazon's DNA and supports:

4,000 products sold per minute on Amazon.com

1.6 million packages handled every day

Billions of Alexa interactions each week

9 in 10 top pharmaceutical companies using AWS for ML

Question 2: How do FMs work?

FMs can learn from very large amounts of unlabeled, multi-modal data and can be rapidly adapted to new tasks with less labeled data. In healthcare, they can analyze large volumes of multi-modal data such as images, EHR data, and genomics. This information can support multiple AI-powered use cases that enhance clinical and operational efficiency, improve patient experiences, and reduce costs.

The evolution and proliferation of FMs in healthcare is ongoing and represents an exciting development for Rowland Illing, MD, chief medical officer and director of international public sector health at AWS. “In the future, there will be hundreds, if not thousands, of different foundation models,” Dr. Illing said. “AWS will not only democratize access to these different models, allowing a broader range of scientists to access them, but also provide secure customization, which is a real differentiator. I look forward to seeing how they are leveraged to improve health and wellness in the future.”

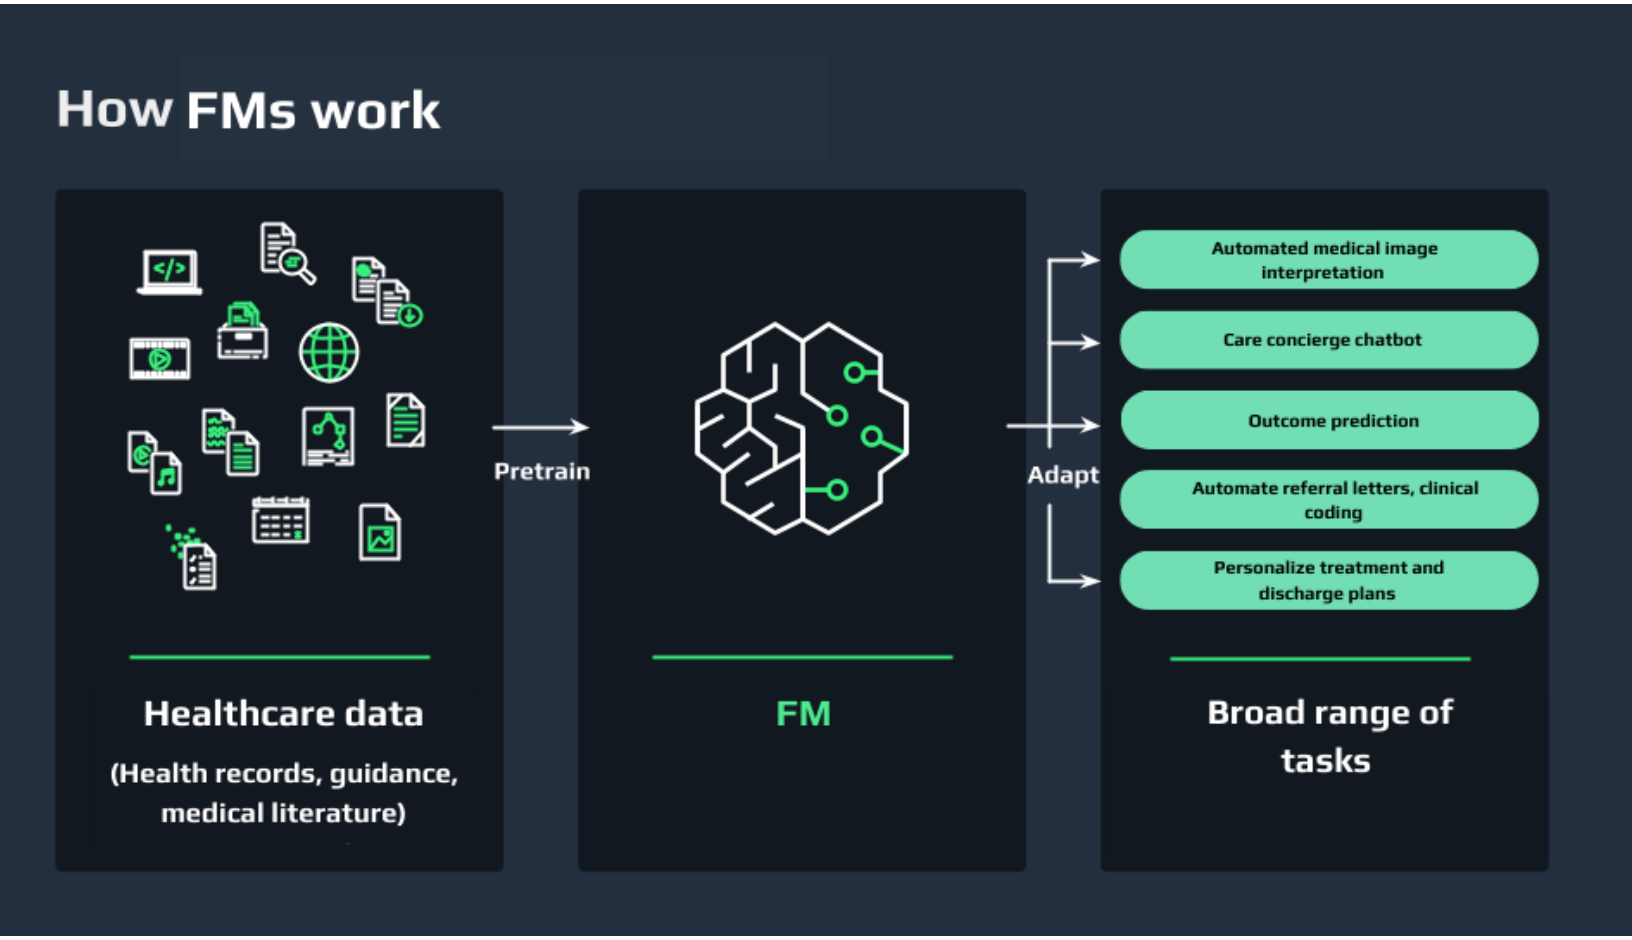

Question 3: What are the use cases for GenAI in healthcare?

The use cases for GenAI in healthcare will continue to evolve in the coming years, taking new shapes, and solving different challenges. Here are five current uses cases:

1. **Administrative automation:** Identifying waste due to manual workloads or duplicative processes to drive operational efficiencies, speed up reimbursements, and enhance value-based care reporting. Automating clinical documentation for example, summarizing clinical encounters, referral letters or discharge reports can reduce administrative workload.
2. **Analyzing and synthesizing images:** Rapidly develop AI tools to analyze medical images, such as X-rays and MRIs for the detection of diseases, or generate synthetic images to improve AI performance.
3. **Engaging patients through chatbots:** Conversational assistants and chatbots can reduce clinician workload, increase patient satisfaction, and help provide proactive healthcare to at-risk communities. These tools provide high-quality medical information and education, assist patients in managing long-term chronic conditions and even offer counseling support.

4. Supporting physicians with clinical decision support: Develop next-best-action systems based on massive, multi-modal clinical datasets that can help healthcare practitioners interpret unprecedented volumes of data and evidence. This can support patient care by accelerating diagnosis, identifying rare diseases, developing treatment plans, or identifying appropriate clinical trials.

5. Responding to population health crises: Generative models can be used to create synthetic patient and healthcare data, which can be useful for training AI models, simulating clinical trials or studying rare diseases without access to large real-world datasets, enabling faster and informed responses to population health crises.

Generative AI workloads in AWS

SUBSET OF USE CASES UNDER DEVELOPMENT AND/OR IN PRODUCTION

Medical Research

Patient to trial matching

Multi-modal data analysis

Clinical Efficiency

Longitudinal patient records for full patient picture

Automate medical image interpretation

Operational Efficiency

Auto-generate referral letters, clinical coding, and prior authorization

Intelligent document processing

Patient Experience

Patient outcome prediction

Personalize patient discharge instructions and treatment plans

Digital Health

Patient care concierge

Remote care management

Here are four challenges associated with the use of generative AI in healthcare today:

1. The risk of hallucination, a confident response by AI that does not seem to be justified by its training data.

2. The high computational/memory requirements and cost of developing FMs from scratch.

3. Reliance on third-party managed services, which have limited control and customizability, as well as limited visibility into how privacy and compliance are managed.

4. Slow response times when solutions are powered by third-party FMs.

In a recent article,"Responsible AI in the generative era," Michael Kearns, PhD, Amazon scholar and professor of computer and information science at the University of Pennsylvania, emphasized that addressing challenges related to generative AI will be a continuous endeavor.

"Generative AI has stoked both legitimate enthusiasm and legitimate fears … addressing responsible AI risks in the generative age will be an iterative process: there will be no 'getting it right' once and for all," Dr. Kearns wrote. "This landscape is sure to shift, with changes to both the technology and our attitudes toward it; the only constant will be the necessity of balancing the enthusiasm with practical and effective checks on the concerns."

Question 5: Why generative AI on AWS?

Developers are quickly creating solutions to the challenges associated with generative AI. Using techniques such as Retrieval Augmented Generation (RAG) supports generative AI responses based on appropriate data. RAG combines content curation and indexing technologies (like Amazon Kendra) with large language models and supports sophisticated identity and access management strategies to manage privacy and security. For example, a clinically accurate medical information chatbot could be developed using a RAG-based approach that combines an effective large language model with a validated, curated, and controlled set of clinical guidelines.

What does AWS offer for generative AI?

Access to FMs

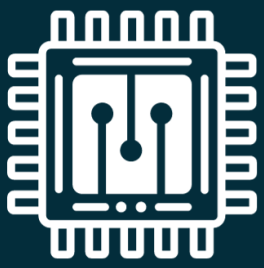

Evaluate and select from a set of foundation models based on use-case, price, and performance

Security, scale, and reliability

Domain train and tone with your data and protect your intellectual property

Optimized infrastructure

Most price performant cloud infrastructure including purpose-built ML accelerators and GPU instances

Solutions to speed up development

Deliver rich experiences and integrations to accelerate time to research and clinical impact

Unlocking the potential of generative AI

- The simplest way to build with FMs
- The most price-performant infrastructure
- Generative AI-powered applications
- Flexibility

AWS makes it simple for healthcare organizations to find the right models for any use case. It facilitates secure customization with the organization's specific data and enables large-scale deployment, integrating seamlessly into existing applications on the most secure, comprehensive global infrastructure. Whether an organization is looking to build a FM from scratch, customize an existing model or use an AWS service with generative AI built in, AWS provides the services healthcare organizations need to shape an individualized generative AI strategy.

Building on two decades of ML innovation to power healthcare breakthroughs

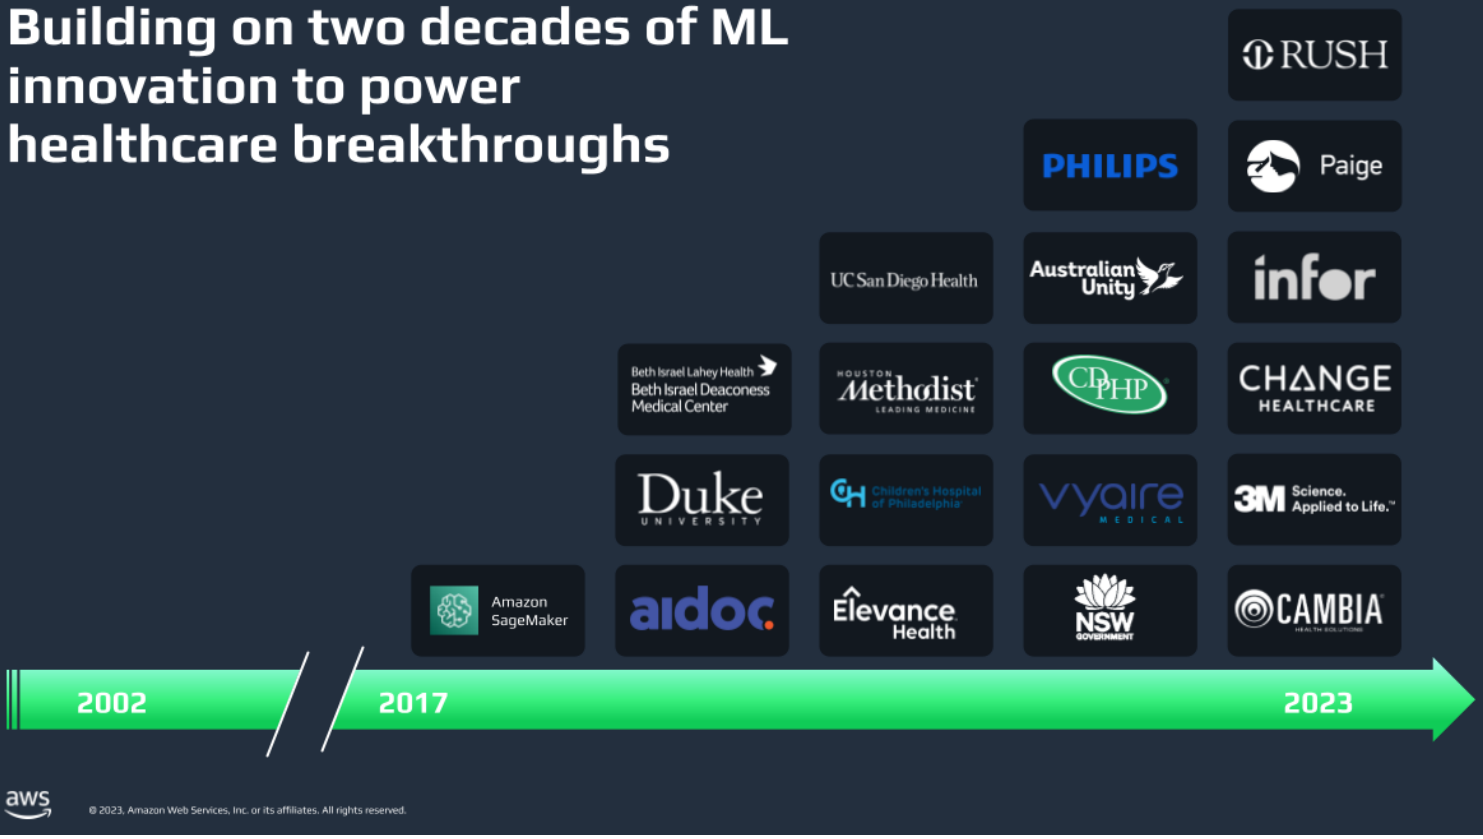

Here are three more opportunities for healthcare organizations to partner with AWS for generative AI offerings.

1. **Amazon Bedrock:** Amazon Bedrock is a new service that makes FMs from AI21 Labs, Anthropic, Stability AI, Cohere and Amazon accessible via an application programming interface (API). Bedrock offers the ability to access a range of powerful FMs for text and images, including Amazon Titan FMs, through a scalable, reliable and secure AWS managed service. Bedrock allows customers to privately customize FMs with their own healthcare data. Organizations can then easily integrate and deploy these models using AWS tools and capabilities without having to manage any infrastructure.
2. **Amazon SageMaker Jumpstart:** Amazon SageMaker JumpStart provides pretrained, open-source models for a wide range of problem types to help organizations get started with ML. Because these FMs are pretrained, they can help lower training and infrastructure costs and enable use case customization.
3. **AWS cloud adoption framework for AI, ML and generative AI:** The AWS cloud adoption framework is designed to help organizations generate business value from AI and ML. In this framework, AWS outlines a technology journey for organizations as their use of technology (and the technologies themselves) evolve. This framework also offers prescriptive guidance, as it provides an overview of the target state of these foundational capabilities and explains how to evolve them step by step to generate business value along the way.

AWS healthcare team is committed to supporting its partners at every step of their generative AI transformation journey. Healthcare companies and beyond have leveraged AWS offerings to design and operate reliable, secure, efficient, cost-effective, and sustainable technology programs.

Envision

Identify and transform opportunities in line with your strategic objectives

Align

Identify capability gaps and cross-organizational dependencies

Launch

Deliver pilots in production and demonstrate incremental business value

Scale

Expand pilots to desired scale and ensure that you can realize and sustain business benefits

Question 6: Why is now the time for healthcare organizations to start generative AI transformation?

Generative AI holds significant promise for healthcare, but as with all new AI-powered technologies, development needs to work backwards from the needs of patients, healthcare practitioners, providers, and systems to address complex use cases in accordance with regulatory, privacy, security, and ethical requirements.

“Generative AI and foundation models have enormous potential to unlock efficiencies and innovations across healthcare and accelerate a new era of medical progress and patient-centric care,” said Matthew Howard, PhD, head of healthcare AI, generative AI, and data sciences at AWS.

“This must be achieved with a deep commitment to data privacy, equitable AI decision-making, robust validation, and effective human-AI collaboration, ensuring we mitigate potential risks while maximizing the benefit to patients.”

The promise of generative AI in healthcare can’t be fully realized without organizations and leaders willing to do the work of adoption and continuous iteration. Organizations at the forefront of generative AI will help to write healthcare’s future. Join them by connecting with an AWS expert today.

AWS in healthcare

AWS's healthcare mission is to enable access and delivery of person-centered care in order to improve outcomes and lower costs by accelerating the digitization and utilization of healthcare data. Our aim is to help healthcare providers develop cloud strategies to achieve more with less, modernize technology, and digitally transform back office and clinical services.