

O'REILLY®

Compliments of  
OPTUM®

# State of Healthcare Technology

How Disrupting Technologies  
from AI to Genomics Will  
Transform the Industry

Kerrie L. Holley, Mike Jacobs,  
Sanji Fernando, Steve Graham,  
Ravi Kondadadi, Dan McCreary,  
Dima Rekesh, Kathy Tzeng  
& Matthew Versaggi

REPORT

HOW CAN WE MAKE THE PROMISE OF  
TECHNOLOGY A REALITY IN HEALTH CARE?

HOW

WE KNOW HOW. WE ARE THE HOW.

At Optum, we're not pursuing technology for its own sake. We're striving to make health care more human. We're a unifying force bringing together deep industry expertise with data, advanced analytics and emerging technology to move health care forward. Across millions of consumer touchpoints, Optum understands how the world of technology connects to the world of health.

[OPTUM.COM/TECHNOLOGY](https://www.optum.com/technology)

© 2020 OPTUM, INC. ALL RIGHTS RESERVED.



---

# State of Healthcare Technology

*How Disrupting Technologies from  
AI to Genomics Will Transform  
the Industry*

*Kerrie L. Holley, Mike Jacobs,  
Sanji Fernando, Steve Graham,  
Ravi Kondadadi, Dan McCreary,  
Dima Rekesh, Kathy Tzeng, and  
Matthew Versaggi*

Beijing • Boston • Farnham • Sebastopol • Tokyo

**O'REILLY®**

## State of Healthcare Technology

by Kerrie L. Holley, Mike Jacobs, Sanji Fernando, Ravi Kondadadi, Dan McCreary, Steve Graham, Dima Rekesh, Kathy Tzeng, and Matthew Versaggi

Copyright © 2020 O'Reilly Media, Inc. All rights reserved.

Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (<http://oreilly.com>). For more information, contact our corporate/institutional sales department: 800-998-9938 or [corporate@oreilly.com](mailto:corporate@oreilly.com).

**Acquisitions Editor:** Michelle Smith  
**Development Editor:** Corbin Collin  
**Production Editor:** Luan Chang  
**Copyeditor:** Rachel Monaghan

**Interior Designer:** David Futato  
**Cover Designer:** Karen Montgomery  
**Illustrator:** Rebecca Demarest

February 2020: First Edition

### Revision History for the First Edition

2020-02-05: First Release  
2020-06-12: Second Release

The O'Reilly logo is a registered trademark of O'Reilly Media, Inc. *State of Healthcare Technology*, the cover image, and related trade dress are trademarks of O'Reilly Media, Inc.

The views expressed in this work are those of the authors, and do not represent the publisher's views. While the publisher and the authors have used good faith efforts to ensure that the information and instructions contained in this work are accurate, the publisher and the authors disclaim all responsibility for errors or omissions, including without limitation responsibility for damages resulting from the use of or reliance on this work. Use of the information and instructions contained in this work is at your own risk. If any code samples or other technology this work contains or describes is subject to open source licenses or the intellectual property rights of others, it is your responsibility to ensure that your use thereof complies with such licenses and/or rights.

This work is part of a collaboration between O'Reilly and Optum. See our [statement of editorial independence](#).

978-1-492-08895-0

[LSI]

---

# Table of Contents

<b>The State of Healthcare Technology.....</b>	<b>1</b>
Healthcare Challenges and Technology	3
Artificial Intelligence	7
Deep Learning	11
Natural Language Processing	15
Graph Technologies	19
Genomics	25
Blockchain	27
The Internet of Things and Ambient Computing	32
Looking Ahead: Healthcare in 2025	36
Acknowledgments	39



---

# The State of Healthcare Technology

After lagging behind many other industries for years, the healthcare sector is now experiencing tremendous change. The industry is on the cusp of significant breakthroughs that will transform it and help make the healthcare system and healthcare work<sup>1</sup> better for everyone. While a full vision of healthcare technology is beyond its scope, this report will outline key technologies that are bringing about this revolution. It is organized into the following sections:

## *Healthcare challenges and technology*

Technology amplifies continuous change at an ever-increasing velocity. It drives digital transformation forcing healthcare companies to refresh their business models, reduce friction, and create engaging experiences. This section describes several of the key challenges and provides a preview of technologies driving continuous innovation.

## *Artificial intelligence*

AI is a general-purpose technology that encompasses several distinct technologies and architectures that will transform healthcare. This section explores what the emergence of AI means for healthcare and the healthcare system. Only two AI

---

<sup>1</sup> In this report we use the term *healthcare* to describe an industry or sector, a segment of the economy, and—in some contexts—a service provided by clinicians and/or the healthcare system.

technologies are explored in this paper: deep learning and deep natural language processing.

### *Deep learning*

Deep learning learns from data—and the more data, the better. Engineers feed a learning algorithm terabytes or more of data, train it using this data, and then let the computer figure out how to recognize sounds, words, objects, or language. This method has powerful benefits for healthcare.

### *Natural language processing*

Advances in computer vision are enabling machines to rapidly discover tumors or other medical conditions, meaning clinicians can provide faster and improved healthcare. Natural language processing (NLP), or **deep learning for NLP**, has arrived, and this section describes how its application is transforming healthcare.

### *Graph technologies*

Graph technologies redefine what's possible for connections, such as the intersection of patients, providers, clinicians, health-care companies, and more. Graph technologies allow IT to move at the speed of business, powered by AI. This section explores graph technology opportunities to transform healthcare.

### *Genomics*

Making precision and personalized medicine a reality is the promise of genomics. As this section describes, the human genome is the basis for developing a variety of solutions, diagnostic tools, and products to improve the identification, treatment, and prevention of poor health outcomes.

### *Blockchain*

Blockchain has the potential to do for resources and assets what the internet has done for information: decentralizing control and revolutionizing the way companies and people interact. This section details ways in which that is likely to work.

### *The Internet of Things and ambient computing*

Devices are becoming more connected and more intelligent. IoT is real and brings value to healthcare, especially when it comes to sensors and wearables. Ambient computing is computing that is invisible, contextually aware, and responsive to



interactions with the physical world. This section describes how both can improve healthcare.

### *Looking ahead: Healthcare in 2025*

The future of healthcare will be characterized by the interplay between people, physical things and environments, and the systems, applications, and tools created by software engineers and developers. Real-time healthcare is on the horizon and soon clinicians, consumers, and administrators will all demand that these technologies be part of their everyday use. This will require reimagining healthcare, and this section provides a few ideas for doing so.

## Healthcare Challenges and Technology

The rapid growth of computer technology touches every aspect of the healthcare sector due to its potential to address a daunting list of industry challenges. This section summarizes the current state of healthcare and begins the discussion of how new technologies are positioned to transform it.

Historically, healthcare has been slower to adopt new technologies compared to many other sectors of the economy, for several reasons:

- Complex ecosystem (providers, payers, clinicians, patients, and more), many of which are contractually bound
- Fragmented healthcare data stored in diverse sources and formats, making interoperability very challenging
- Complex healthcare and privacy regulations (state and federal)
- Physicians' desire that technology not intrude on caregiving
- Concerns regarding ability to maintain security of patient health data
- Reluctance of clinicians and providers to trust new technologies without full transparency in the reasoning and data behind the provided recommendation

As those traditional obstacles are overcome, we should expect more rapid growth of technology adoption in all areas of healthcare. New technologies are expected to help automate administrative, operational, and clinical tasks and processes; reduce administrative and medical costs; improve clinical outcomes; enable better-quality and

smarter experiences for consumers and providers; prevent fraud; and improve the effectiveness of value-based care. Computer technology tends to have a more transformative impact on healthcare when it's invisible to stakeholders in the healthcare ecosystem, especially clinicians. That is, it's better when computer applications and tools don't add to a clinician's work, and don't take time away from patient care. When the technology is seamless, clinicians become more productive and deliver better care.

It's important to understand that technology doesn't "do" healthcare. Technology providers and technologists must recognize that healthcare is what clinicians do—provide care for patients. Perhaps this may change in the future, but it holds true today.

## Challenges to Healthcare

Here are some of the current major challenges in healthcare that new computer technologies are poised to help address:

- Continued **rising costs** of healthcare must be slowed to make healthcare more accessible, in part by significantly reducing manual processes with technology automation.
- Healthcare data must be increasingly digitized to eliminate many paper-based processes and better address patient and clinician needs.
- Interoperability must improve to integrate new sources of data to enable the use of new technologies.
- Value-based care must be expanded to reward clinicians for keeping patients healthy rather than treating them after a diagnosis.
- The patient experience must be transformed to be closer aligned to that of a retail consumer experience, with price and quality transparency, modern and responsive customer service, and convenient at-home options for monitoring and treating ongoing healthcare concerns.
- Healthcare data analysis must shift from a historical or retrospective approach to one that is near real time, prospective, and predictive to help address issues before they become chronic conditions.

## Technology for Transformation

Conversational artificial intelligence, in which humans and AI systems conduct conversations using everyday language, has improved tremendously in recent years, and is expected to become the primary and preferred way that people interact with many future technologies. This kind of interface will be at the center of many healthcare applications, device apps, and systems.

AI includes many subfields, such as natural language processing (NLP), machine learning (ML), deep learning (DL), and knowledge representation (KR). **Figure 1** illustrates the relationships between various technologies (e.g., cloud, RPA, graph, IoT and blockchain) with AI as the engine, playing a role in the transformation of healthcare.

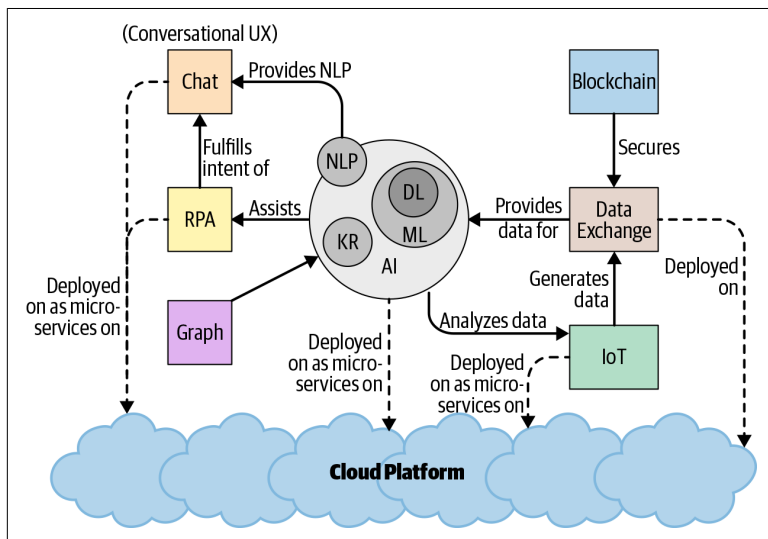


Figure 1. The many collaborative technologies at play for healthcare

AI will play a significant role in solving some of the challenges in healthcare and in driving a future vision of the patient, provider, and clinician experience. While this report organizes the topic of AI into different technological concepts, they are not distinct sets—there is significant overlap.

*The cloud* will be the dominant way in which healthcare organizations deliver value to internal and external customers. Virtual assistants and chats will become more useful as NLP feeds text-to-speech

and speech-to-text interactions to begin the pathway toward conversational AI.

Technologies like blockchain are set to usher in a new era of cooperation among competing entities toward a common goal. For example, inaccurate doctor directories are currently a barrier to accessing healthcare because each insurer, payer, and health plan maintains its own directories, often with outdated content. If they all used a shared and trusted source of information through the use of a blockchain, however, a single accurate provider directory could significantly reduce the friction in the healthcare system across many stakeholders.

*Robotic process automation* (RPA) technology and chats will improve the customer experience. RPA software using software bots can mimic human behavior by navigating existing systems the same way as an employee.

The *Internet of Things* (IoT), using sensors and other devices to monitor health and enhance the patient's experience, is on the rise with the healthcare system. IoT devices will be increasingly deployed and managed in combination with ML and DL models, giving rise to so-called *edge computing*, where computer technologies are deployed at the “edge,” such as homes, clinics, or hospitals. For example, a nurse engaging with a patient at home may use a mobile device that transparently uses NLP and ML/DL for quality note taking, enabling the nurse to maintain focus on the patient. Noninvasive sensors used to track movements or detect falls can be installed in homes and set to alert family and/or caregivers, helping to improve patient care.

Business capabilities such as benefit inquiry or ML-based disease prediction models can be built as *microservices*—Lego-like building blocks that get invoked or consumed as services by other systems. *Graph databases* are used for real-time queries and rapid flexible changes to the organization of the data. The ability to combine multiple data dimensions with graphs is ideal for many healthcare queries, such as a *patient journey*, which allows providers or payers (or anyone authorized) to see the steps a patient makes while interacting with the healthcare system.

Next we dive deeper into the application of these emerging technologies for healthcare.

# Artificial Intelligence

People use the term *artificial intelligence* to mean everything from things that computers can't really do yet, like carrying on a conversation, to things computers already know how to do, like contributing to decision support systems. With the hype surrounding AI, companies and vendors are eager to describe their products as AI even if they're just talking about well-established statistical methods or traditional databases.

If you're an engineer, executive, or business stakeholder, you have to dig a little deeper to determine which AI technologies are being used by a product to understand how they can solve your problem. With this understanding, you should be better equipped to decide if the product is using the most recent advances in AI or is a repackaged version of previous and much older generations of AI, some of which date back to the 1980s.

The last decade has changed everything for AI. It now seems clear that it is a *general-purpose technology* (GPT) that will change the way we live, work, and play—just as the computer and internet, to name a couple of other GPTs, have already done. AI will transform health-care by revolutionizing the tools and systems that guide healthcare providers and patients.

## What It Is

There's no consensus, but most would agree that artificial intelligence is a field of computer science heavily focused on building intelligent things: systems, apps, applications, and more. AI is or will be deployed in a variety of settings, including cars, homes, offices, factories, and hospitals.

As mentioned earlier, AI contains subfields including deep learning, machine learning, natural language understanding, natural language processing, and computer vision. It includes a lot of distinct classes of ML systems such as generative adversarial networks (GANs) and reinforcement learning. This report only covers a few classes of machine learning systems.

## How It Solves Healthcare Challenges

AI can and will solve healthcare challenges first and foremost as an aid or tool for clinicians. Adoption of AI improves decision making,

minimizes diagnostic errors, improves recommended care pathways, provides a holistic view of a patient, reduces friction and delays in the healthcare system, and, most importantly, enables clinicians to spend more time with patients. AI eliminates some back-office processes while accelerating others, and creates more meaningful, frictionless interactions among stakeholders in the healthcare ecosystem. AI can be used to analyze orders of magnitude more data than ever before, such as data coming from a variety of new sources like genomics, wearables, and intelligent things.

Deep learning is making a huge difference by improving results and outcomes when applying computer vision, natural language processing, natural language understanding and machine learning. DL networks have driven breakthroughs in self-driving cars and image recognition. Deep neural networks provide optimal results for computer vision and image classification, such as helping clinicians quickly identify patients with pneumonia. **Studies** have found that providers can diagnose conditions significantly faster using tools in which AI-embedded image recognition is used, enabling treatment to start sooner, which is of course vital for severely ill patients.

The **ImageNet competition** prompted a flurry of activity in image recognition and spurred innovation in ML and computer vision, leading to an AI renaissance. Healthcare is now having its own ImageNet moment for NLP. A few healthcare examples include advances in deep learning NLP to redefine manual work, medical billing and coding, and reviewing bills for accuracy. AI can scan enormous quantities of bills and find—and in some cases, correct—errors, which is an improvement over the current approach of sampling a small percentage of bills for inaccuracies. Deep learning neural networks “learn” from large and complex sets of data, not people, to correctly solve a wide variety of business problems.

Deep learning is well suited for healthcare because complex decisions often require skilled professionals like physicians and nurses. AI is being trained on previous decisions made by these professionals about policy changes, prior authorizations, and chart abstraction, and deep learning neural networks are helping to make new decisions. In all of these examples, AI helps the clinician be more productive.

## Use Cases

Another way to frame AI is to say that it enables machines to make judgments and infer meaning about data. Imagine all the areas in healthcare today in which humans need to make judgments, infer meaning, and make decisions. Here are just a few:

- Provider network management
- Prior authorization
- Claims processing
- Provider payments
- Clinical decision support

The goal is to create augmentation tools to improve clinician productivity, not to replace them. The rest of this section looks at specific cases in a little more detail.

### Remote patient monitoring

A bedbound state is dangerous for older patients. Even a few days of bed rest can cause muscles to deteriorate, bringing about long-term negative consequences. Too much bed rest can happen because of **caregivers not moving patients according to treatment plans**. Infrared sensors can measure the heat emitted by an object and detect motion. AI can use sensor data to detect when someone has fallen and is on the ground. Computer vision and ML are key AI technologies addressing this situation. AI detection may even go further to help with root-cause analysis and better decision making to prevent future falls.

AI will play a big part in filtering and organizing data into a consumable form. People could easily become overwhelmed by the sheer volume of data coming from sensors, cameras, scales, and other devices in their homes. AI will help correlate the sensor data with outcomes to make the information clinically useful. Noninvasive technology that monitors patients in their homes or outside of conventional clinical settings both is affordable and can improve access to care. Remote healthcare monitoring allows people to stay at home, providing an alternative to on-site clinical care monitoring.

## Care delivery processes

Clinician time is often spent on necessary but frequently inefficient or wasteful back-office tasks that don't directly affect patient care. Tasks such as operating room optimization, clinical appointment scheduling, and billing/coding are areas where AI could be implemented as an assistive technology.

## Conversational AI: Call centers

Patients often have questions about benefits or become upset when a claim is rejected or not fully paid. Time-critical prior authorization is needed now in certain urgent treatment situations such as cancer diagnoses, where a delay in treatment may result in a poor patient outcome, although if something is medically urgent then no prior authorization is required—immediate treatment is allowed. In the near future, patients will call their insurance company and talk to an AI-powered assistant. This assistant will listen in real time to the call and provide answers or support to a human call-center agent. The AI-powered assistant will be started when the patient calls and will authenticate the patient by voice—no more annoying prompts for personal information. It may even anticipate why the patient is calling. If the patient is likely to have a benefit question, the assistant may have the answer ready for the call center agent, perhaps even before the patient asks the question. If the issue is less complicated, like a prior authorization request for a routine covered procedure, it could even provide the authorization immediately.

## Where It's Headed

There is no doubt that computers can do some things better than people, as long as that thing is well defined, clearly explainable, and reproducible. Machines outperform people for any task where computation makes a difference. AI will be leveraged across the entire \$3.5 trillion US healthcare system to keep people well, predict and diagnose conditions, recommend treatment paths, reduce friction in the ecosystem, improve and eliminate manual processes, and more. It may even perform, or help perform, surgery.

Usage of AI can help reduce the almost **\$500 billion** spent every year on billing and insurance-related (BIR) costs. Automation using AI can help standardize and optimize care protocols and move the system toward preventative versus reactive treatment. Using AI with



the help of clinicians and researchers may help identify potential cures for various forms of disease.

Imagine a future in which AI diagnostic tools are available for physicians and consumers. AI-powered diagnostic tools could use speech, biometrics, keyboard dynamics, touchscreen patterns, face recognition or visual cues to detect behavioral disorders. Real-time monitoring of health issues becomes more common as wearables become ubiquitous.

Deep learning and natural language processing are subfields of AI. Deep learning, the topic of the next section, is a big factor in the emerging success of AI in healthcare.

## Deep Learning

Deep learning has sparked a renaissance in the use and capabilities of artificial intelligence. Key breakthroughs in recent years have demonstrated what's already possible in the application of AI. Technology companies are using DL not only to transform their businesses, but also to help them rethink whole industries.

You are observing the use of DL when you use social media services in which the app suggests you tag someone in a photo. The same is true when you sign-in to your mobile phone with FaceID or ask a voice assistant to set a reminder.

DL networks are transforming how we bridge understanding between people and software systems. In the past, programmers tried to do that by designing software systems to handle a range of situations, often by defining rules for each scenario. DL software systems can “learn” how to handle situations on their own; they don't need programmers to tell them exactly what to do.

## What It Is

### Neural networks

*Deep learning* is really shorthand for *deep learning neural networks*. Neural networks first gained popularity in the late 1980s and early 90s as programmers attempted to write software that mimicked brain neurons. Each “neuron” was represented as a nonlinear function that responded to a number of input signals. These neurons could then be connected together into a network such that they fed

each other and could respond to increasingly complex combinations of input signals. If you had an example of an input, you could pass your inputs into the network of equations, and it would return an answer.

A typical neural network learns by example. If your neural network is intended to recognize cats in photos, you need to provide example cat photos as input. The input is converted to a numerical representation and passed to a neural network (a network of equations). At first, the equations have randomly assigned values and will likely get the answer wrong. But through a feedback loop process called *back-propagation*, the network updates variables in the network of equations to “learn” from its mistakes. Subsequent examples are likewise learned until the network gets to be pretty good at classifying the inputs and returning the correct answer—that is, recognizing that a picture includes a cat. This process is called *training* your neural network.

### **Training by examples**

In the 1980s and 90s, researchers could link only a few layers of equations, which limited the types of use cases for which a neural network could be trained. Complicated image or text recognition was not possible because the process of training would overwhelm the computer processors that were available. Researchers also lacked the huge amount of example data necessary to train the neural networks.

A number of advances in the past 10 years have changed the landscape. The graphical processing units (GPUs) that make playing *Call of Duty* so realistic are also good at completing the kind of parallel math needed for efficient neural networks. These processors can handle very deep neural networks with hundreds of layers of networked equations. Another change was the *big data* revolution, because big data provides lots of examples to train on.

Now, instead of people programming complex sets of rules to classify or predict events (imagine trying to account for every kind of cat and every typical cat activity when trying to write a program to identify cats in a picture), DL neural networks can be trained on lots of previous examples. DL has many more uses than image recognition of cats.

## How It Solves Healthcare Challenges

The most well-publicized applications of DL in healthcare have been in classifying medical images. There are lots of examples of medical images that have been previously evaluated by clinicians (e.g., radiologists and dermatologists) to use to train DL neural networks. Initial seminal research studies focused on training a DL solution to recognize cancerous skin lesions.

Other types of medical diagnostic data are also well suited for classification with a DL model to assist clinical diagnosis. EKGs, MRIs, and genetic sequences have all been successfully classified by DL neural networks. These breakthroughs are likely to transform clinical care going forward.

Administrative use cases in healthcare, such as confirmation of benefits and reimbursement, are also well suited for automation with DL models. Other examples of administrative functions include prior authorization, third-party physician reviews, revenue cycle management, claims processing and adjudication, and risk adjustment.

These are complex processes and evaluations that ensure the appropriate reimbursement of care. The healthcare industry has developed very large sets of rules or guidelines in an attempt to automate processing. But these rules have grown substantially as the complexity of evaluating care and reimbursement has increased. Skilled professionals are relied upon to manually complete complex transactions and care as the rules break down. Physicians, nurses, and medical coders are all counted on to review and determine the appropriate level of care.

Records of claims and charts reviewed by skilled professionals, and the decisions they have made, represent valuable training data for DL neural networks. As policies, medical innovation, and regulations change, DL neural networks will need to be retrained. By training DL neural networks with the decisions of skilled professionals, the industry can build substantial, up-to-date, and collective knowledge and understanding.

## Use Cases

This section covers two use cases involving deep learning.

## Administrative Processes

There are many examples where nurses or doctors must manually review medical records. If every record is reviewed, this becomes a very time-consuming and costly process.

A DL neural network can be trained on hundreds of thousands of complex decisions using historical case reviews to determine cases where the complexity warrants a doctor review. The more past cases the model learns from, the more accurate it becomes, further reducing the time, cost, and number of cases requiring a manual review. Such a model can save hospitals and providers time and money, allowing physicians to focus on clinical care rather than administrative processing.

## Smart Prior authorization

*Prior authorization* is an administrative process which is used by insurance companies to determine if they will cover a medical procedure, prescription or service. Several companies are looking at accelerating prior authorization approvals with AI and automation. In such cases, an authorization determination engine is built using deep learning. AI provides an opportunity to automate the process and accelerate the time to get approval, while maintaining a human in the loop. By focusing the DL on automating approvals, the skilled professionals—nurses and clinicians—can focus on the most complex cases where there may be alternatives to consider.

## Where It's Headed

As noted, DL models have driven amazing breakthroughs in the use and adoption of artificial intelligence. With each breakthrough, the question arises: when will DL automate clinical diagnosis and treatment? The barriers and risks to automation in these areas are numerous. One of the biggest challenges is explaining why the model recommended a particular diagnosis or treatment.

The question of *why* is probably essential to ensure that appropriate care is recommended. Physicians will likely need to understand how the model arrives at a recommendation in explicit terms. Was it blood pressure, genetic sequence, or familial history that led an AI solution to determine a diagnosis and recommend treatment? Often this is referred to as *causality*—something not inherent to DL

models. For example, you can't point to the specific pixel or set of pixels in the image of a cat to say why the model classified it as such.

Some techniques, as seen in the first case study, can illuminate key parts of the input that may explain how the inputs are impacting the classification, but that isn't true causality. Work is under way to begin enabling the creation of causal models. *Invariants*, *causal graphs*, and *deep probabilistic programming* are all promising efforts intended to address these limitations. These and other breakthroughs will be essential to building true causal models that would be one key step toward automating clinical diagnosis and treatment.

## Natural Language Processing

Understanding and generating language are the most important aspects of human intelligence. The ability to communicate with others using complex language distinguishes humans from other species and has helped us become the dominant beings on the planet. Language comprehension is therefore one of the key aspects of an AI system.

Human language can be quite ambiguous. It's often hard even for a human to accurately interpret what another human is trying to communicate. Therefore, it should be no surprise that developing software that can understand human (or *natural*) language is quite tricky. Natural language processing is the area of research that attempts to make computers understand and generate human language, both written and spoken.

### What It Is

NLP is an interdisciplinary field of computer science and linguistics, and also a subfield of artificial intelligence, as [Figure 2](#) shows. NLP enables a computer to read unstructured data, analyze keywords and phrases, extract meaning from them based on experience, generate language that people use, and create actionable insights. ML and DL models have become critical components of many NLP systems.




Figure 2. How NLP fits into artificial intelligence

Since their advent in 2012, DL models have made significant advances in the fields of image processing and speech recognition, as discussed in the section “Deep Learning”. However, progress in the area of language understanding was relatively slow. Then, in 2018, one of the most path-breaking developments for NLP occurred: The creation of Bidirectional Encoder Representations from Transformers (BERT).

BERT, and models like it, can represent words and sentences in a way that captures underlying meanings and relationships better than past models. In NLP, *word embedding* models are used to generate vector representations of words to use in a neural network model. Traditional word embedding models like word2vec generate a single context-free representation for each word in a document’s vocabulary (called a *word embedding*). Word embeddings are context-free because only one representation is allowed per word regardless of the number of senses the word can have. But models like BERT can represent a word based on both the context before and after the word, and is therefore *bidirectional*. The embeddings these models generate are described as *context-aware*. BERT and similar language models have shown exceptional results in NLP.

Models like BERT facilitate the use of transfer learning techniques in NLP. *Transfer learning* allows for models trained on larger datasets to be used on tasks that have smaller datasets in a related domain. DL techniques are data hungry and therefore don’t perform well on smaller datasets. BERT-like models (also called *pretrained language models*) can be trained on much larger, unlabeled datasets, which can then be fine-tuned on smaller datasets for specific tasks like document classification, question answering, and machine translation.

BERT is an instance of a neural network architecture called a *transformer*. Transformers are better than other DL models like recurrent

neural networks (RNNs) at dealing with long-range dependencies (where words that are not close to each other in a text should be used together to extract the text's meaning), which are very critical to a few NLP applications like question answering.

Generative Pretrained Transformer (GPT) was the first transformer-based pretrained language model, where the model predicts the next word based on the prior context. BERT is the first bidirectional language model, but a surge of other transformer models for NLP have appeared in the past year, including XLNet and RoBERTa, each having more parameters than the one before. Nvidia's Megatron includes 8 billion parameters and is the largest transformer model as of this writing. These larger models provide better accuracy but take a long time to train, require more computational power, and take longer to run at processing time. Therefore, there has also been an influx of smaller transformer models, like ALBERT and DistillBERT. Though these models are much smaller, they retain much of the same computational accuracy. For example, DistillBERT is 40% smaller than BERT, but retains almost 97% of accuracy. ALBERT has 1/18 of BERT's parameters.

## How It Solves Healthcare Challenges

A large portion of healthcare data is unstructured. For example, most clinical information (such as history and physical notes, operative notes, procedure notes, and discharge summaries) is *narrative*—free-form writing that computers have never been good at interpreting. Deep NLP models (i.e., ones that use deep learning) can help identify important clinical elements and establish meaningful relationships between them from free-form clinical data in electronic health records (EHRs). This information can then be used in clinical decision support systems such as prior authorization, computer-assisted coding, clinical documentation improvement, and extraction of risk adjustment and hierarchical condition categories.

NLP can also be useful in extracting key elements from patient-physician interaction and autopopulating parts of EHRs, thereby reducing the time physicians need to spend on clinical documentation.

## Use Cases

Some cases in which NLP is being used in the healthcare field are described next. (Note that we use *deep NLP* to indicate the use of deep learning with NLP, which often provides outcomes superior to NLP without deep learning.)

### Patient experience

Although patients can access their health data from clinical visits, many have difficulty with understanding the content, which makes it challenging for patients to make health decisions. This can change with the application of machine learning/deep learning and natural language processing, which would increase patient health awareness. NLP algorithms could find their way into voice assistants or mobile apps and allow patients to use a voice interaction to make sense of their health data, bridging the gap between medical terms and patient's understanding of their health data.

### Speech recognition

The diagnostic potential of looking at speech patterns is growing. Voice data may tell a clinician a lot about neurocognitive conditions such as dementia, onset of Alzheimer's, and other psychological conditions—even cardiovascular diseases. Natural language processing and deep learning would be essential AI technologies for translating voice-to-text and processing it to a neural network.

## Where It's Headed

NLP investment for healthcare is projected to reach **\$4.3 billion** by 2024. New and innovative applications of NLP are likely in many areas, including the following:

- Improving customer experience
- Supporting care delivery
- Understanding clinical language

NLP will transform the healthcare continuum of care by extracting information from unstructured text to enhance the end-user experience. Virtual personal health assistants will create a better provider and patient experience using technologies like conversational AI. NLP will be the key technology to develop successful CAC



and CACDI products to help patient care centers improve clinical documentation and free up their staff to spend more time with patients.

## Graph Technologies

Healthcare data is complex. This complexity comes in many forms, but the primary factors are the many relationships and high variability within the data. For example, a patient chart has many subsections that describe a patient’s conditions, immunizations, allergies, prescriptions, medical tests, hospitalization admissions, and so on. Because of this complexity, traditional tabular data structures used by relational databases are not always a good fit for clinical data.

### What It Is

Bringing complex data together for display or analysis is a relationship-intensive query for a traditional relational database management system (RDBMS). We now know that these calculations can be run more efficiently using permanent in-memory pointers in a *graph database*. This is the key difference between the old and new ways to store clinical data.

They are called *graph* databases because they use the same structures as the well-established branch of mathematics called *graph theory*—a connected network of *nodes* and *edges*. Using in-memory *pointer jumping* is a sharp contrast to relational databases, which use JOIN operations to indicate relationships. Pointer jumping uses a memory location to store the memory address of a related object. Traversal across the data is fast because all the necessary information is directly in memory. No other data structures get in the way and slow down the relationship traversals.

**Table 1** shows some of the primary differences between relational and graph databases, along with why graph databases are preferred for many healthcare applications.

*Table 1. Comparing relationship and graph databases and their impact on clinical data management*

Relational	Graph	Business impact of using graph for clinical data
Optimized for row-by-row storage	Optimized for fast relationship traversal	Analysis of complex clinical data is three to five orders of magnitude faster.

Relational	Graph	Business impact of using graph for clinical data
Relationships calculated at query time	Relationships calculated at load time	Relationship-intensive clinical data is queried faster.
Fixed data model	Flexible data model	It is easier to add new clinical data to a graph at any time.
Limited analytical algorithm	Rich library of analytical algorithms such as pathfinding, PageRank, clustering, and similarity	It is easy to execute patient clustering, similarity, and care path recommendation engines.
Difficult to build custom hardware to optimize algorithms	Easy to build custom hardware to optimize algorithms	Hardware for optimizing algorithms is low cost.

Graph databases are only part of an emerging trend in healthcare as can be seen in [Figure 3](#), which shows a trendline in graph adoption since 2014. Graph analytics, graph visualization, and graph discovery are also core components to integrated healthcare solutions:

#### *Graph analytics*

The process of running complex queries over large amounts of data to find new insights. Because graph databases are so much more efficient at relationship traversal, graph developers can use many analytics algorithms to more efficiently find groupings of data, similar structures, and paths between structures.

#### *Graph visualization*

A collection of tools and processes for visualizing highly connected data. This includes using automatic layout tools to view thousands of nodes and their connections to see patterns within the data.

#### *Graph discovery*

The process of looking for unusual patterns in graphs that may not be evident by simple inspection. Graph discovery algorithms understand “normal” structures and highlight structures that might deviate from these norms.

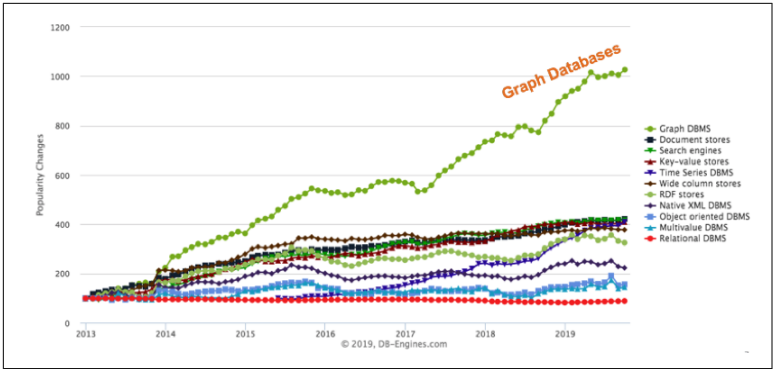


Figure 3. Popularity trendline for various database architectures (source: DB-Engines.com, used with permission)

## How It Solves Healthcare Challenges

Gathering all the related datasets for a patient in a central location is key to suggesting a care plan. The accuracy of care path recommendations improves as a *longitudinal health record* (LHR)—a comprehensive clinical summary of a patient-based clinical experience—becomes more complete.

Building the LHR data requires gathering data from disparate sources in many forms and loading it into a centralized graph structure, as seen in Figure 4, which shows an example of a graph model for an LHR. This data needs to be *normalized* (made consistent) so that all records can be queried with consistent results to be most effective. Graphs are superior to other forms of databases in that the retrieval of these complex LHR data structures requires only the simple and fast graph traversal of in-memory pointers.

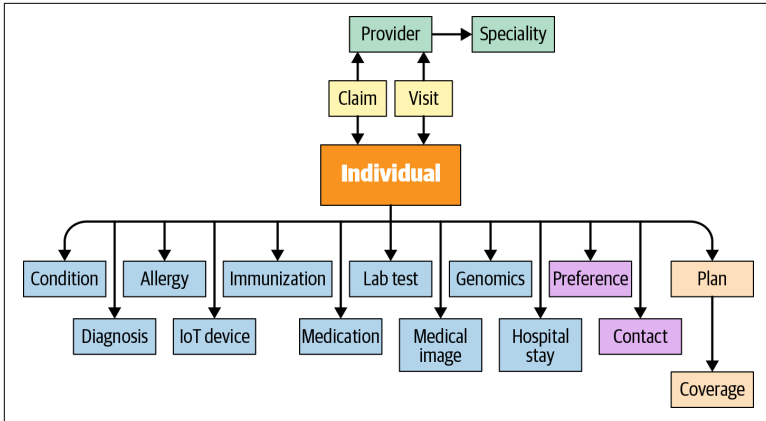


Figure 4. LHRs are characterized by many complex relationships, which makes them ideally suited for storage in graph databases

## Use Cases

This section discusses other areas of healthcare where graph technologies are already being used.

### Ontology-based data access

*Ontologies* are graphs that store terminology and rules about their use. Healthcare is notorious for generating complex jargon. For example, the Unified Medical Language System (UMLS) has over 3.5 million healthcare concepts in a single graph. These *concept graphs* can be used to normalize data into consistent formats. The process of passing data through ontology-driven rules systems is called *ontology-based data access* (OBDA).

### Single view of the patient

When a patient calls with questions or concerns about coverage, a call-center representative can generate a timeline view of all the patient's *touchpoints*. For example, each email, phone call, or survey can be considered a touchpoint. Even with hundreds of millions of patients, and thousands of touchpoints per patient, these timeline views can be generated in 1/10 of a second using graph database technologies. For example, consider a patient with multiple chronic conditions such as diabetes and heart disease. This patient may be on 10 medications and use a blood-glucose monitor. He may have weekly visits to clinics, and every day his updated healthcare data

needs to be integrated into his personal healthcare graphs from his internet-connected devices. Because all this data is already integrated, a next best action can then be suggested onscreen to the call-center representative.

### **Patient similarity and cohort building**

Every day new patients come into emergency rooms, urgent care centers, and clinics, each presenting clinicians with a unique set of symptoms. They may have specific conditions and be on specific drugs. Their genome dictates whether certain drugs are effective and at what dose these drugs should be taken. Key questions for the clinician are: “Have you seen similar patients like this?” and “What care paths had the best outcomes?”

One of the basic graph algorithms in graph technology is the *similarity calculation*. For example, *cosine* similarity uses a vector of features for each patient to quickly calculate a similarity score for any two patients. The number of features in these vectors can become very large. Algorithms such as the `graph2vec` function can use existing databases of patients to reduce vectors of tens of thousands of dimensions down to a more manageable 200 features. This means comparisons can be done in near real time during patient encounters. Physicians and researchers can also use these algorithms to take selected groups of patients and find similar patients to build cohorts for analysis and efficacy studies.

### **Real-time clinical decision support and recommendations**

You are likely familiar with ecommerce sites that recommend products and streaming entertainment systems that recommend movies and music that you might like. The algorithms that power most of these systems are graph-based. Similar graph-based algorithms can also suggest care paths for physicians and researchers.

### **Fraud detection**

Healthcare companies are increasingly overwhelmed by problems caused by sophisticated AI-powered fraud systems. Graph technologies are commonly used to build fraud-detection models that are more accurate than older systems. Graphs retain structural information that can reveal complex fraud rings. Once fraud incidents have been labeled, algorithms such as *graph convolutional neural networks* (GCNNs) can be used to find similar patterns in graphs.

Combined with machine learning, these algorithms avoid the need to create and maintain complex fraud-detection rules.

## Where It's Headed

Although graph technologies are already helping to lower healthcare costs, this is really only the beginning of the healthcare graph story. There are some particular areas to watch.

### Artificial intelligence

Graph databases are becoming the most common way to store knowledge used by intelligent agents. In his book *The Master Algorithm* (Basic Books, 2018), Pedro Domingos writes about five AI domains:

- Symbolic and logical reasoning
- Neural networks
- Bayesian probability graphs
- Evolutionary algorithms
- Associative algorithms

Graph technologies can help with many of these areas. Symbolic processing, long done with Resource Description Format (RDF) graphs, is used to store healthcare terminology. For example, a brand name drug (Advil) and its generic name, ibuprofen, can be stored in a graph indicating that they are the same or similar.

Bayesian graphs show the probabilistic and causality relationships between variables in a graph. For example, they can show that there is not just a probability that smoking and lung cancer are related, but also that there is causality.

Evolutionary algorithms are frequently simulated in graphs that emulate biological processes. Simulations can show how various care path recommendations for a group of patients might impact both their overall health and the cost of care.

### Improved explainability

Healthcare recommendations provided by typical DL systems often lack *explainability*—understandable reasons for the recommendations that people can process and comprehend. Graph technologies

will work with neural networks to associate text blocks with graph models. These text blocks can then be assembled to provide a written explanation of why specific recommendations have been made.

### Hardware graphs

Graph queries use only a narrow set of machine instructions, which today's computers weren't really built to handle. Current *complex instruction set computers* (CSICs) are not tuned for fast processing of graph algorithms. But in the next few years, a new generation of hardware will be optimized for graph traversal, making graph technologies even faster and more powerful.

## Genomics

Genomics provides a powerful approach to assess risk and diagnose and treat the underpinnings of human disease, creating a transformational opportunity to help make people healthier.

### What It Is

The human genome functions as a blueprint for the human body. DNA in the genome is located within chromosomes, comprises four building blocks, and represents the biological “plan” for each person. That plan tells each cell what to do and when to do it. A potential for disease occurs when a person's biological blueprint is flawed. These flaws are called *variants*. Matching variants with clinical signs and symptoms can establish the linkage with the disease. This matching process requires the ability to search millions of medical records and genetic data to establish a linkage between a disease and one or more variant.

The process of mapping a person's genome has become significantly cheaper. The first human genome, sequenced in 2003, cost \$3 billion and took 13 years to complete. The cost is now in the range of \$1,000, and the sequencing time is measured in hours. In addition, the data collected is more precise, and the insights obtained through DNA sequencing impact more conditions.

### How It Solves Healthcare Challenges

As more genomic data is integrated into healthcare, clinicians and providers envision using genomic information in conjunction with

other biomarkers to tailor a personalized approach to care known as *precision medicine*. Today, genomics is primarily used in oncology, rare diseases, infectious diseases, and pharmacology. Precision medicine is changing the way diseases are identified, prevented, and managed. It helps in personalizing treatments for cancer, diagnosing rare diseases, improving the success rate of organ transplants, and predicting disease outbreaks.

## Use Cases

There are numerous genomic use cases that are helping advance healthcare and enabling patients to live healthier lives.

### Precision medicine

Genomics provides new insights for precision and personalized medicine. Integrating genomics data with claim, EHR, and lab data will give a comprehensive view of patients, enabling both precision and personalized medicine and early discovery/translational research. Health records, lab data, claim data, and longitude information integrated with genetic data further improve the ability to predict disease, monitor health, and personalize treatment using disease prediction models. The combined data enables healthcare organizations to perform statistical analysis and develop models to identify new biomarkers associated with diseases, predisposition, and drug responses.

### Rare diseases

Genetic testing can help diagnose patients with previously undiagnosed diseases. Many rare diseases have been shown to be caused by mutations in genes. Genetic sequencing and improved data sharing result in faster discovery of disease genes and development of new diagnostics. Early diagnosis and intervention for patients with an undiagnosed disease, especially Neonatal Intensive Care Unit (NICU) patients, can often prevent the irreversible damage and reduce the length of NICU stays. The medical expense of unnecessary tests and treatments can also be significantly reduced.

### Pharmacogenomics

Pharmacogenomics is the low-hanging fruit for precision medicine and is one of the areas where genomics research is already mature enough for clinical use. With today's current standard of care, many



patients are prescribed drugs that don't work for them. In some cases, the drugs even worsen their health. The current trial-and-error treatment approach is costly and can lead to irreversible damage and adverse events. Genetic testing can accelerate personalized treatment by improving care and lowering costs by eliminating waste and getting to stable care sooner.

## Oncology

Cancer is a disease of the genome. Many genetically targeted therapies are currently available to cancer patients. Hereditary cancer tests are available to test for several cancer genes. Single blood test screening and liquid biopsy are available to detect different cancer types. Early-stage cancer is more easily treated than advanced-stage cancer, and treatment cost is significantly lower.

## Where It's Headed

As the pace of genetic sequencing accelerates and more genomic data is integrated into the health system, precision medicine will become a routine part of every person's medical care.

Right now, most physicians and patients aren't very familiar with genetic tests. In general, physicians don't know what tests are available, how and when to order them, whether they are covered by insurance plans, how to interpret the test results, or how to talk to patients about the tests. This will change as the adoption of genetic tests grows thanks to education in key disease areas, medical research centers publishing evidence-based guidelines and scientific papers, and genetic testing being included in EHR systems.

## Blockchain

The concept of a shared, distributed ledger was first described in a whitepaper published in 2008 by an unknown person or group writing under the name Satoshi Nakamoto. This report described the ledger as a *blockchain* dedicated to the exchange of a new virtual currency called Bitcoin. This idea has subsequently been extended into other implementations for use beyond virtual currencies, to include features needed by businesses.

## What It Is

A blockchain is a shared, distributed digital ledger on which transactions are chronologically recorded in a cooperative and tamper-resistant manner. Transactions, such as money transfers between two banks, are recorded on a list, or ledger. Each bank maintains its own ledger of all account debits and credits, along with when the transactions occurred. A transfer between banks may take several days while the two banks reconcile and confirm that the transfer amount is properly debited from the sender account and credited to the receiver account.

This latency on transfers could be shortened if the banks could share a single ledger, because then there would be no need to reconcile multiple ledgers. However, having a single shared ledger for all banks is impractical due to technical and business considerations. But if each bank had its own copy of the shared information that was automatically synchronized, that ledger would be a shared, distributed ledger that addressed the business considerations. Additional technologies can also be used to ensure the transactions are valid and audit-proof.

There's a lot of hype and misinformation about what blockchain can and can't do today. Businesses are experimenting with the latest version of blockchain technology to track, trade, find, collect, synchronize, and validate data. Key to these efforts is finding ways to reduce the costs associated with data management. Blockchain makes it possible to trace information back to its original source and view every detail in between. It simplifies commercial relationships by letting an organization track and trade something of value (such as data) without an intermediary.

Enterprises have learned that blockchain technology is most useful when independent or loosely coupled organizations, such as those that do business together or are otherwise linked, want to confidentially share and audit information and automate mutually beneficial processes. With built-in transparency and verifiability of transactions, control is decentralized, and all participating organizations are peers. These are important prerequisites if diverse stakeholders are to collaborate productively.

Blockchain technology is best suited to solve business problems involving the following:

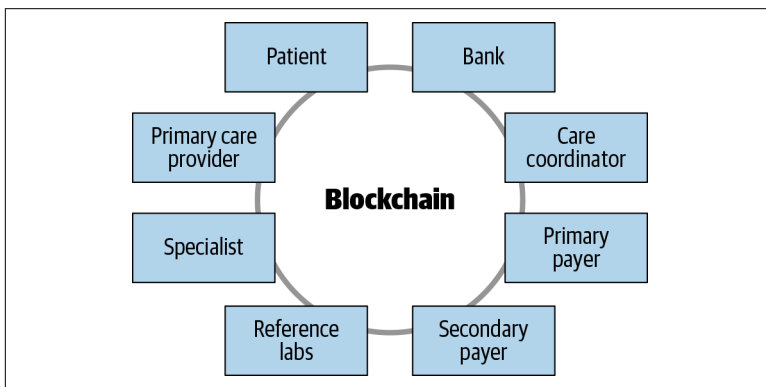
- Multi-enterprise islands of data
- High cost to reconcile inconsistent data
- Lack of complete trust between companies

Further, blockchain enables organizations to collaborate on problems that pose little or no risk to competitive positions. But even where there is willingness, incentives may need to be created to encourage participation.

The rise of blockchain has generated a lot of excitement, with many companies attracted to it as the “shiny new thing” and experimenting with it as a solution to a variety of business challenges. However, companies would do better to focus on whether blockchain technology *should* be used for a particular use case and let the business problem drive the technology selection. This approach is called *fit for purpose*.

## How It Solves Healthcare Challenges

As blockchain rapidly matures, many speculate that it will solve a lot of data problems in healthcare, such as improving data interoperability, helping prevent data breaches, and shifting ownership of personal health data from health systems to consumers. Healthcare organizations and big tech companies are experimenting with blockchain, seeking to understand what it is and how to effectively use it to derive real value. **Figure 5** illustrates some of the stakeholders that could participate in and benefit from a future healthcare blockchain.



*Figure 5. Many parties could participate in a permissioned healthcare blockchain*

## Use Cases

The use of blockchain is still in the early stages in healthcare. Due to the nature of the technology, blockchain requires *density of adoption*, also known as the *network effect*; that is, to be productive, a group of organizations must be willing to cooperate to test and leverage the efficacy of the technology. So, a good place to start with blockchain is on business problems shared by many organizations in an ecosystem. This section looks at a few examples.

### Healthcare provider data management

Healthcare provider *demographic* data—such as addresses, phone numbers, and whether or not a provider is accepting new patients—is a key building block of the US healthcare system. This data is essential for connecting patients with care providers and supporting consumers in making decisions related to their care. State and federal laws, regulations, and guidance require health plans, such as health insurance companies or *payers*, to regularly update provider directories. Plan members consult these directories when looking for a care provider. If the information is inaccurate, they may be inhibited from accessing care in a timely manner. The implications of inaccurate provider data are extensive and can negatively impact payers, providers, and consumers in many ways.

Currently health plans typically maintain their own provider datasets, which has created silos of data. On average, a provider is affiliated with 20 health plans and often has to update its information with each health plan individually. This large-scale duplication of effort results in redundant expenses. Across the US healthcare ecosystem, more than \$2.1 billion is spent annually by hospitals, doctors, and health insurers just to maintain provider data, according to conservative **estimates**. Despite this staggering investment, the system is still inefficient and reliant on manual processes.

A permissioned blockchain solution could enable participants to share some of the administrative burden and cost of data maintenance and reconciliation, improving data quality while substantially reducing the time and expense. In 2018 the Synaptic Health Alliance, along with other leading healthcare organizations, undertook a pilot project focused on provider data management. Through a provider data exchange (PDX), Alliance members actively share data with the aim of showing potential administrative cost savings

for payers and providers while improving data accuracy and consumers' healthcare experience. Early pilot results have shown that Alliance members are able to locate and update certain types of inaccurate information faster than they could on their own.

### **Coordination of benefits, longitudinal health records**

Blockchain could be a potential solution to other data-based problems in healthcare related to payment processing and sharing of patient-specific health data.

Blockchain technology has the potential to streamline *coordination of benefits* (COB), the costly process of coordinating proper payments across multiple health insurance companies. COB is how insurance companies determine the order of payments of claims (primary, secondary, and so forth) for a member who has coverage from more than one policy. Currently, providers and insurance companies do not have a cost-effective way to determine the correct order of payment, which leads to extensive overpayment, manual rework/recovery efforts, and lost profit across the insurance industry. A blockchain-based solution may enable providers and insurance companies to more immediately determine who should pay first and who should pay second, significantly alleviating the COB administrative burden of providers and insurance companies.

Blockchain technology may be a way to end the stalemate for generating LHRs. Blockchain could be used to reference islands of patient data in a single, secure view of health records accessed by providers who have permission. The goal of an LHR with blockchain technology is to provide a historical and real-time, whole-person view of health. This will improve care and give consumers control of their health data.

### **Where It's Headed**

In most cases, an alliance between organizations is necessary to effectively leverage blockchain. As noted earlier, blockchain works when organizations come together to share and automate common processes. Additionally, blockchain requires a different business mindset. All alliance members are responsible for the success of the alliance, but no one member owns the blockchain. Instead, every participant is working together to solve an issue for the industry.

As you build or join an alliance, find potential allies who have an appetite for early adoption of new technology. Look for candidates that are open and willing to learn, as these are often more risk-tolerant and flexible. Proceed deliberately and quickly to explore how blockchain may help advance your business. Blockchain has the potential to give you and your business partners improved velocity with higher-quality transactions and new business models.

## The Internet of Things and Ambient Computing

The term *Internet of Things* means connecting *things* (devices, cars, homes, appliances, sensors, etc.) to the internet and then using that network connection to gather data from them and even trigger various actions based on that data. *Ambient computing* refers to the idea of gathering data and performing computing tasks on that data in the background, with an example being voice assistants.

### What It Is

In healthcare, the most widespread form of IoT is *wearables*. These include fitness-monitoring devices (such as the Fitbit and Apple Watch) and health-monitoring devices (such as a continuous blood glucose monitor or a Bluetooth-connected blood pressure monitor). More and more fitness- and health-related devices are appearing all the time.

Ambient computing works a little differently, in that a human speaks into the environment, and the computing system responds by answering a question or performing some action. Another category of ambient computing is video camera-based systems used to detect movement, identify falls, accumulate information about how an individual moves through the environment, and so forth.

### How It Solves Healthcare Challenges

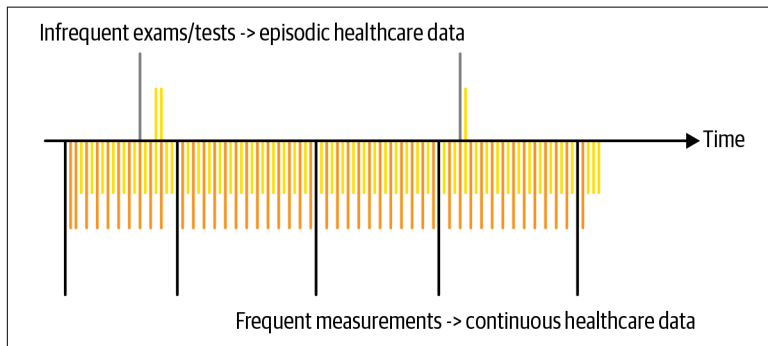
Ambient computing is a fundamentally new form of computing. Based on advances in artificial intelligence—specifically deep learning—voice and video data is now available for healthcare systems to process. For example, doctors can analyze video of a person's gait to detect signs of early-onset osteoarthritis or other conditions. Face recognition can be used to infer a person's emotions, and a

telemedicine caregiver can determine whether an explanation of a medical procedure was confusing or frightening to the patient. Analysis of a person's voice can even detect signs of depression.

New forms of interaction are available with ambient computing. *Conversational AI* uses ambient computing to provide a more natural and low-friction way for individuals to interact with healthcare systems. Gestures can be used to augment voice, and keyboard and touch technologies can enrich interaction with healthcare systems.

The healthcare industry uses the term *patient-generated health data* (PGHD) to refer to health-related data collected from IoT and ambient computing systems. PGHD becomes truly operational when it is integrated into the healthcare system. This involves merging the PGHD with each individual's EHR, using that additional data to generate insights about the individual's healthcare journey, and then assisting the individual's care team to tailor their treatment.

Healthcare will begin to transition from episodic care to continuous care as PGHD is incorporated on a near-continuous basis as shown in [Figure 6](#).



*Figure 6. Differences in the verity, velocity, and volume of healthcare data with continuous care*

In [Figure 6](#), the space above the timeline running through the middle depicts the current situation. A patient sees a family doctor once a year. A couple of labs are run during this visit, and the patient's EHR is updated. If all is well with the patient, there is no further update of the health system for a year or so. The space under the timeline indicates that the patient has one or more IoT devices, or maybe there is an in-home video camera. This continuous stream of

data can be analyzed much more frequently, allowing for a richer longitudinal picture of the patient's health.

## Use Cases

The use cases for IoT are as varied as the IoT devices themselves. Many devices help people with their fitness goals and with chronic care management.

### Using PGHD to inform better care plans

Most modern remote patient monitoring systems integrate PGHD collected from wearables as one kind of data used to understand a patient's current health situation. PGHD can be used, in conjunction with other healthcare data such as claims, labs and prescriptions, to tailor the way an individual engages with the healthcare system and learn what to do in order to properly manage their condition. It also helps the care team guide and monitor them on their healthcare journey and intervene when necessary.

### Using video to detect changes in movement habits

Stanford University's Departments of Medicine and Computer Science performed a **study of patients in an assisted living facility** using video to monitor *activities of daily living* (ADLs): sleeping, sitting, standing, walking, and so on. The study used DL models to combine depth-sensing and heat-sensing video feeds to form a privacy-preserving monitoring system, showing only the outlines of each individual. The study produced personalized heatmaps of the frequency and duration of the various ADLs that caregivers could use to quickly identify deviations from normal patterns and to plan interventions accordingly.

## Where It's Headed

As healthcare systems adapt to leverage the continuous stream of data, and as more IoT devices and ambient computing systems are deployed, healthcare companies could begin adopting what we'll call an *ambient healthcare flywheel*, as illustrated in **Figure 7**.



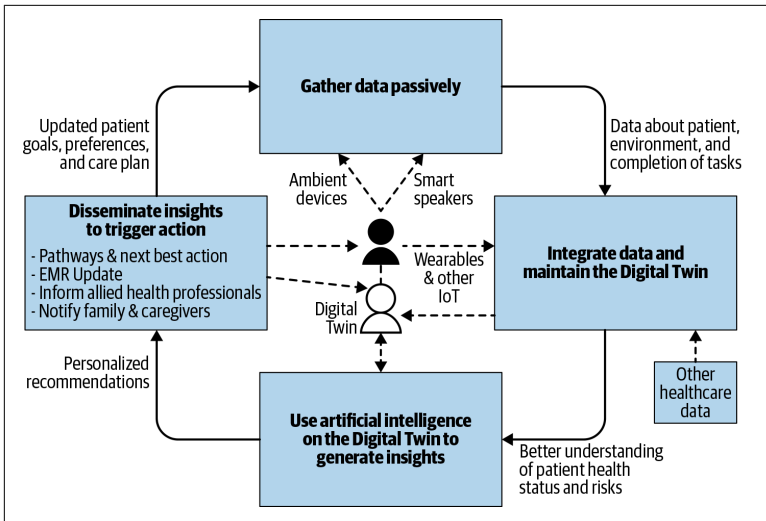


Figure 7. Ambient healthcare flywheel

The system would collect data from ambient systems and IoT devices and integrate that data with all the other healthcare data available about the individual (EHR data, labs data, prescription history, genomics data, claims data, etc.). The combined data would yield a *digital twin* of the individual—a model of the individual’s past and current health situation. Analytics, many of which are based on AI technologies, could be used to generate insights about the individual, addressing questions such as the following:

- Do the signals in the data suggest the onset of Parkinson’s?
- Does the data indicate that the individual is moving enough after his hip surgery?
- Is the individual able to do all the activities of daily living in order to reassure her loved ones that she is safe to continue living at home?

These insights may be disseminated to those in the healthcare system who need to know about and act upon them. If the individual slips and falls, this flywheel system can notify family members and 911. If the insight strongly suggests early onset or worsening of a disease condition, the healthcare team can be alerted.

With the flywheel, the actions performed by the healthcare system can be used in response to the insights derived from analytics run

on the digital twin to indicate what new kinds of healthcare data would be worthwhile to gather. For example, should new cameras be installed to be more precise about gait analysis? Should the patient have a wearable blood pressure monitor?

More data yields a richer digital twin, which allows for better analytics, generating more accurate insights and in turn more timely and effective healthcare actions. The iteratively improving system suggested by the ambient healthcare flywheel will deliver the continuous, anticipatory, and personalized form of healthcare called *continuous care*.

## Looking Ahead: Healthcare in 2025

Making confident predictions even five years into the future is difficult due to rapid changes in healthcare and technology. However, many scenarios that would have seemed far-fetched even a decade ago now look to be within reach. Consider the following:

- Your tablet notifies you of the pending onset of depression based on thousands of your facial images accumulated over a decade and a DL model.
- Your car detects a possible heart condition in your passenger through an ambient IoT device.
- Your doctor is alerted that you may have contracted a disease before symptoms appear thanks to hereditary characteristics and genomics monitoring.
- Your healthcare provider flags a possible Alzheimer's diagnosis by using accumulated voice data to screen your cognitive function and generate a mental state examination score.
- Caregivers use spatial heatmaps to see where their loved ones are spending their time or detect changes in their gait that suggest a possible future health issue.
- Noninvasive sensors in your parents' home alert you to potential problems based on their typical everyday activities.
- ML algorithms and virtual assistants learn the best time to remind you to take your medicine.
- The hypothesis that **aging is a treatable disease** is boosted by multiple studies using AI, triggering you to be more engaged with your health.

- Delivery of care is fully transparent to all stakeholders, and access to services can be coordinated with all parties—in real time—in compliance with all contracts.
- *Healthcare on the go* allows an aging population to travel with peace of mind as computing becomes part of the environment where we travel, work, and live.

Surgical robots; new medicines; improved treatment plans; more accurate diagnoses; reduced error rates in diagnosing and treatments; more interactive and personalized healthcare; AI-assisted healthcare in clinics, hospitals, homes, and assisted living facilities—all of these are poised to become reality in the next few years.

## Differing Worldviews

The opportunity for healthcare to reinvent itself through technology is abundantly clear, but exactly how and when it all will happen remains unclear. Robert M. Wachter summarizes this dilemma aptly in his book, *The Digital Doctor: Hope, Hype, and Harm at the Dawn of Medicine's Computer Age* (McGraw-Hill Education, 2015). While holding tremendous promise for improving healthcare, Wachter says, “technology sometimes overpromises, underdelivers, and can even cause harm.” With clinician oversight and keeping humans in the loop, we should be able to limit or avoid harm with these new technologies.

Vinod Khosla, a venture capitalist and cofounder of Sun Microsystems, **writes** that healthcare will see the “evolution from an entirely human-based healthcare system to an increasingly automated system that enhances human judgment.” Physician and author Abraham Verghese, in a **2011 TED Talk**, laments that the patient has become a mere icon for the “real” patient who lives in the computer, which he calls the *iPatient*. Instead of interacting and engaging with the patient, physicians are spending more time looking at EHRs, images, and data on the computer—treating the *iPatient*.

Dr. Verghese points out that no two patients are alike, and no two doctors are alike. Often the first signals of healthcare needs are received from one of these two sources, patient or clinician. This makes the science of which care pathway to use problematic. A lot of human judgment is often needed. How do you explain pain? Level of discomfort? The observations, eye contact, and person-to-person

contact with a doctor continue to be of paramount importance for improved patient care.

So, iPatient or human-centered healthcare? Wachter, in *The Digital Doctor*, predicts a future where technologies will “transform the work, the people who do it, and their relationships with each other and patients.” He sees a future of human-centered healthcare, dominant now and in the next five years.

## The Race Is On

The last five years have seen AI as a general-purpose technology powering innovations and new use cases for genomics, ambient computing, IoT, blockchain, graph, and more, as outlined in this report. Technology companies, startups, and consumers will continue to force change. New entrants into healthcare may, for example, lease a provider network in order to collect data about how networks work today, and then disrupt them by creating a true, real-time marketplace between consumers and providers, as has been the case with ecommerce in other sectors of the economy.

Deep learning will become the de facto approach for building intelligent learning systems rather than the existing rules-based systems. Forward-thinking companies will always choose machine learning over creation of a complex heuristic, reducing their technical debt and increasing their agility while building learning systems. In [his keynote speech at the 2019 Conference on Neural Information Processing Systems](#), Yoshua Bengio, Canadian computer scientist and winner of the 2018 ACM Turing Award for his work on deep learning, sees a world where DL is extended from its current state to a future state.

The race is on to provide better patient care, deal with the increasing flood of data, reduce the variations in practice that diminish the quality of care, and get patients to engage more actively in their health. Healthcare and technology companies, startups, and others will look to do all of the following:

- Reinvent healthcare
- Be the smartest at healthcare
- Excel at healthcare user experiences
- Empower players in the healthcare ecosystem to do more

Asymmetrical technology-driven assaults on the existing profit models will force change as compelling tools and systems are introduced. Healthcare will improve because of the technologies and methods discussed in this report, such as ML- and DL-powered diagnostic tools, improved aging at home, and healthcare on the go with wearables and computer vision.

Technology will help transform healthcare as more patients and consumers become more engaged. It will augment the work of clinicians, reducing their workload and allowing them to focus on providing actual healthcare and staying engaged with their patients. Diagnoses will become more accurate, and treatment will become more personalized. In summary, the technologies described in this report will play a significant role in making the healthcare system work better for everyone.

## Acknowledgments

The authors would like to gratefully acknowledge the following for their helpful reviews of this report:

- **Dr. Siupo (Cheung) Becker, MD** is a board-certified internist in Minnesota. She is vice president, Healthcare Strategies, at UnitedHealthcare.
- Dr. Vernon D. Smith, MD, is the chief health engineer at OptumHealth.
- **Dominick Dahlem, PhD**, is a distinguished engineer at Optum with an interest in the intersection of science, technology, and society. He studied at Trinity College in Dublin and MIT in Cambridge, did research at IBM Research in Ireland, and was involved in the startup Boxever.
- **Julie Zhu** is a distinguished engineer and data scientist at Optum with more than 20 years of experience in healthcare using advanced analytics in predictive health, with a particular focus in machine learning, deep learning, and natural language processing. She holds several innovation awards and patents, and is a sought after public speaker in AI and data science for healthcare.

## About the Authors

---

**Kerrie L. Holley** is a Technology Fellow and SVP at Optum. He focuses on an advanced technology agenda of artificial intelligence (AI), ambient computing, and graph. Kerrie is an IBM Fellow Emeritus, industry thought leader, and public speaker on AI and technology futures and trends.

**Mike Jacobs** is a Technology Fellow and SVP of Engineering at Optum. He leads the blockchain initiatives for UnitedHealth Group, and is a cofounder of the Synaptic Health Alliance, editorial board member of a healthcare blockchain journal, industry thought leader, and frequent conference speaker.

**Sanji Fernando** is a Senior Vice President at Optum, where he leads the AI and Analytics Platforms team. He is responsible for developing platforms that support the design and development of leading-edge AI models and analytic tools for the enterprise.

**Stephen (Steve) Graham** is a Senior Distinguished Engineer at Optum, working on Internet of Things (IoT) and ambient computing. He has a background in distributed computing, having worked on distributed computing projects at Dell/EMC, University of Chicago, and IBM.

**Ravi Kondadadi** is a Distinguished NLP Engineer with Optum. He has more than 18 years of experience in developing natural language processing (NLP) solutions. In the past, Ravi worked as an NLP researcher for Nuance Communications and Thomson Reuters. Ravi has authored over 20 papers that were published in refereed NLP conferences, which collectively have 500+ citations.

**Dan McCreary** is a Distinguished Engineer with Optum. Dan has worked in the areas of solution architecture, AI, graphs, semantics, metadata, hardware acceleration (FPGAs), graph query languages, and standards. He is a member of the DM32 GQL standards committee and coauthor of the book *Making Sense of NoSQL* (Manning, 2013).

**Dima Rekes** is a Senior Distinguished Engineer at Optum, focused on full-stack Deep Learning applications in healthcare. Prior to Optum, Dima spent 20 years with IBM working on large-scale web

platforms, high-performance computing, big data platforms, cloud computing, edge computing, and deep learning at the Edge.

**Kathy Tzeng, PhD**, is a Senior Distinguished Engineer and Research Fellow at UnitedHealth Group focused on genomics. Prior to joining Optum, she was IBM's worldwide Technical Team Lead for Healthcare and Life Science. She is a published author and public speaker.

**Matthew Versaggi** is a Distinguished Engineer and Senior Director of artificial intelligence in Optum's Advanced Technology division. Matt fills a multitude of roles including thought leader, evangelist, strategist, mentor, education SME, and delivery of advanced technology projects in the AI space.