PitchBook

# AI's Health Exam

## Analysis of artificial intelligence and machine learning in healthcare

PitchBook is a Morningstar company. Comprehensive, accurate and hard-to-find data for professionals doing business in the private markets.

## Credits & Contact


### Analyst

**CAMERON STANFILL** Analyst
cameron.stanfill@pitchbook.com

### Contact PitchBook

pitchbook.com

**RESEARCH**
reports@pitchbook.com

**EDITORIAL**
editorial@pitchbook.com


## Key Takeaways

- Artificial intelligence and machine learning (AI/ML) is beginning to resonate with VC investors in the healthcare sector, recording 48 deals in 2017 raising $718 million, representing 65% and 126% growth YoY, respectively.
- Within healthcare, we expect AI to be used mainly to augment workers' skills, as opposed to displacing them. Personal AI assistants for physicians/nurses can recall relevant medical research, compile population patient data, provide initial possible diagnoses, and optimize workflow. Time saved by these innovations represent countless hours that can be used more efficiently by industry professionals.
- Deep learning algorithms that leverage image recognition are already more accurate than panels of specialists in detecting and diagnosing pneumonia, TB, certain early stage cancers, and other diseases.

# Contents

## Introduction

In our initial coverage of the artificial intelligence and machine learning (AI/ML) vertical, we concluded that the technology has the potential to impact nearly every industry by automating a range of tasks. Due to this widespread potential, we want to provide deeper analyses into specific industries to highlight the rapidly growing and transformative use cases for AI/ML. The potential for AI/ML is beginning to resonate with venture investors particularly in healthcare, as 2017 deal count and value grew approximately 65% and 126% YoY, respectively.

### Healthcare AI deal activity by year

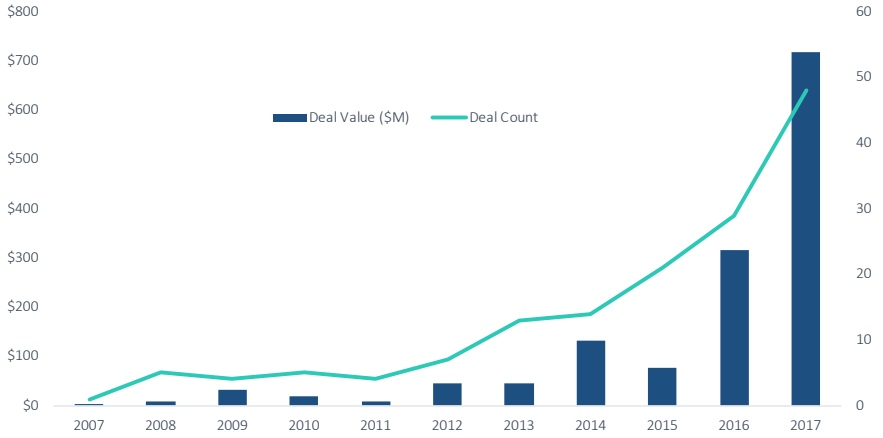


Source: PitchBook

The widespread data digitization in healthcare is the backdrop for all potential use cases in the space, but data structure and fragmentation (e.g. isolated within hospitals, insurance groups) serve as significant hurdles to widespread adoption.

The healthcare market is gigantic—representing 17.9% of US GDP in 2016[1]—and produces a massive amount of data, positioning the industry to receive significant benefits from further integration of AI/ML. Highlighting some of the overarching inefficiency, healthcare's percentage of GDP has doubled since 1980, but health outcomes haven't improved at the same rate. While this technology can increase business efficiency and profitability, AI/ML advancements in healthcare also aim to improve patient outcomes and enhance the overall patient experience. The widespread data digitization in healthcare is the backdrop for all potential use cases in the space, but data structure and fragmentation (e.g. isolated within hospitals, insurance groups) serve as significant hurdles to widespread adoption.

[1] Centers for Medicare and Medicaid Services

## Mapping the healthcare AI environment

Upon examining the deal flow data, we have segmented the current VC investment in AI healthcare market into five distinct subcategories:

1. **Automation & Assistance** – This segment represents 40% of deals and includes those companies focusing on assisting or augmenting day-to-day processes for medical professionals and patients.
2. **Computer Vision** – This segment represents 19% of deals and includes companies pursuing applications of computer vision of medical images to aid in diagnosis, preventative medicine, etc.
3. **Population Health & Predictive Analytics** – This segment represents 15% of deals and includes companies creating platforms focusing on improving population health using aggregated genome, patient, and treatment data, as well as businesses seeking to improve the healthcare system at a more macro level.
4. **Personalized Medicine** – This segment represents 18% of deals and includes companies working to create tailored healthcare analytics or devices to improve personal health outcomes. AI–enabled personal wearables and sensors are included in this subcategory.
5. **Drug Discovery** – This segment represents 8% of deals and includes companies using AI/ML to research and develop novel drugs.

**VC investment in healthcare-focused AI (#) by sector[2]**

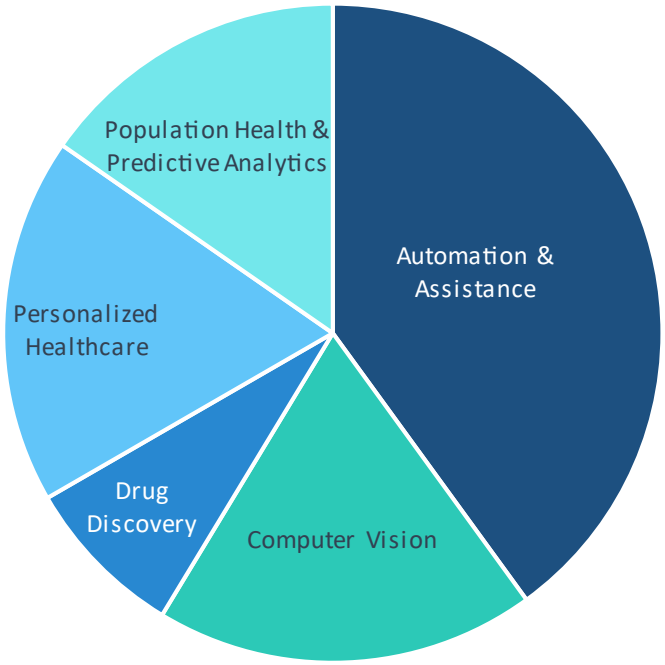


Source: PitchBook

[2] All data in this section is for the period from 2006–2017

## Large-cap outlook

While we can't know exactly what is mentioned behind closed doors, AI/ML hasn't quite broken into the public commentary from executives of public healthcare firms. Companies have instead opted for more broad classifiers like "novel technology", big data/analytics, personalization of treatment and commitments to "business development strategies". One explicit mention came from the insurer United Healthcare, which stated they were continuing to support "existing initiatives in AI" that could benefit both the insurance practice as well as their healthcare solutions subsidiary Optum. Given the nascence of the drug discovery space within healthcare AI/ML, the relative lack of penetration into the large-cap pharmaceutical market was fairly expected.

Some technology giants and AI pioneers also have ambitions within the healthcare sector. Most notably, one of the earliest movers is IBM's Watson Health, which according to the company is beginning to scale—reaching 115,000 patients and being used in 15 hospitals. Watson's breadth of applications has also increased over the past year; the program has now been trained to identify 13 types of cancer, up from four a year ago. Google's DeepMind also has a health initiative that has partnered with the NHS in the UK on an app called Streams that uses technology in mobile phones to alert healthcare staff when a patient's health deteriorates. As companies in the healthcare AI/ML field continue to grow and prove their commercial capabilities, we expect the attention from large corporates to increase over the next 10 years, including some notable acquisitions.

IBM's Watson Health, which according to the company is beginning to scale—reaching 115,000 patients and being used in 15 hospitals.

## Which segments are prime for near-term success?

### *Automation & Assistance*

**VC activity in automation & assistance-focused AI companies**

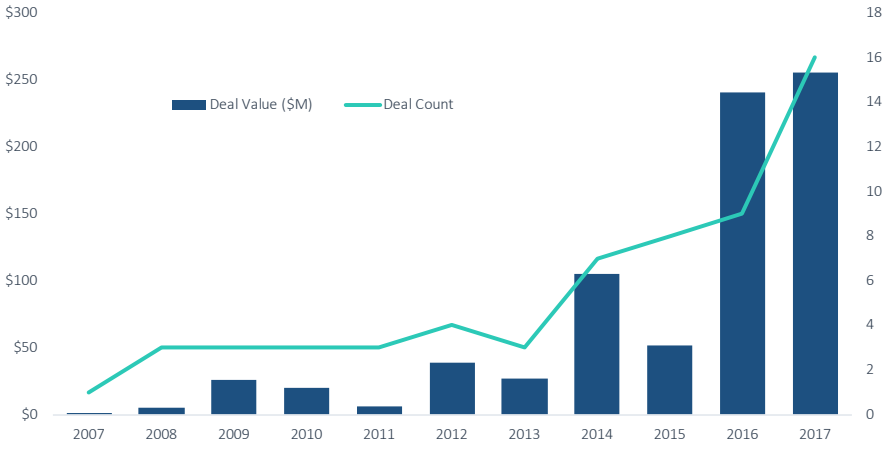


Source: PitchBook

The first broad application of AI/ML in healthcare in which we expect to see widespread adoption is the automation of menial or recurring tasks via personal assistance for both patients and practitioners. These will be focused on increasing workers' knowledge and improving day-to-day efficiency to allow more focus on value-additive tasks, with the ultimate goals of achieving better health outcomes and cost savings. With over $200 million of VC investment in the last two years, this is a large and growing niche of the healthcare AI space.

We believe the automation function is key because of the nature of labor in healthcare. AI/ML-enabled personal assistants for physicians/nurses could recall relevant medical research, compile patient data, provide initial possible diagnoses, and perform scribing duties, among others. These innovations should allow medical practices to reallocate countless worker hours toward more efficient tasks. One specific example is the automation of data entry into electronic health records (EHRs) via speech recognition technology.

Similar benefits can be realized for patients, as AI/ML applications can improve ease of access through initial contact points like chatbots. Digital preliminary diagnoses can also limit unnecessary clinic visits and re-admissions, decreasing costs for patients and freeing up capacity in the healthcare system. Once patients enter a medical facility, AI/ML can be incorporated into initial touch-points for patients to assess the severity of their condition to optimize physician scheduling.

Another costly and time-consuming task within the healthcare industry is the administration of clinical trials. AI/ML are aiming to help solve frictions in the drug approval process, from finding eligible patients to participate in trials, to aggregating and interpreting results. This is the niche that Clinithink, a UK-based startup with over $22 million in VC funding, aims to serve by working to transform unstructured clinical trial data into structured datasets to reduce risk and accelerate go-to-market time for drugs.

## *Computer vision*

### VC activity in computer vision-focused AI companies

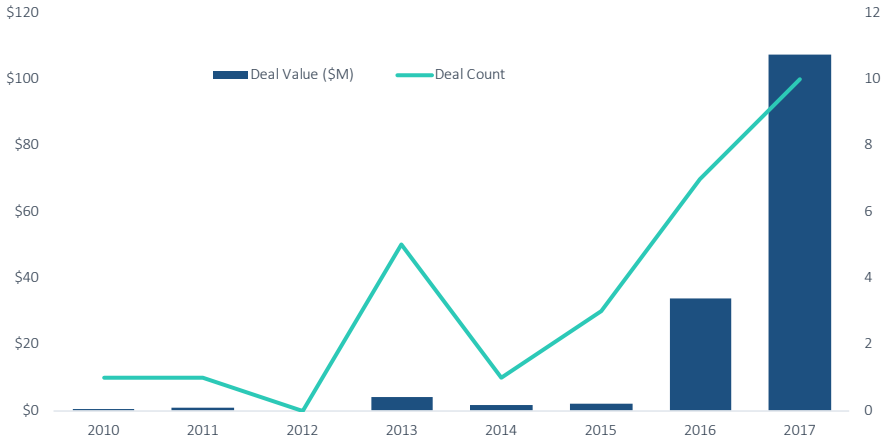


Source: PitchBook

Computer vision is the second area of healthcare AI in which we've seen extensive growth. A major use of AI/ML technology is applying it to medical imaging analysis, which has the potential to deliver staggering improvements in speed and accuracy. With the success of many general image recognition algorithms created by Google, Clarafai and others, we see opportunity in this space as nearing to commercial realization. The healthcare computer vision space saw a significant influx in VC investment over the last two years, tallying over $100 million in 2017 alone.

AI/ML algorithms built to analyze medical imaging have already exceeded panels of specialists in diagnosing pneumonia, TB, certain early stage cancers. AI implementation in diagnosis could drastically increase the amount of patient data that could be processed, expediting test results while increasing accuracy. Computer vision also applies to other visual-based diagnoses/assessments like dermatology. For example, Stanford's skin cancer classification algorithm matched the performance of dermatologists based on pictures of skin lesions, while a separate group of scientists developed another neural network created by that identifies onychomycosis (nail fungus).

These photo-based diagnosis applications could extend the health benefits beyond the clinic and serve as a preliminary point of contact via patients' mobile phones.

These photo-based diagnosis applications could extend the health benefits beyond the clinic and serve as a preliminary point of contact via patients' mobile phones. While early successes of AI/ML programs for medical imaging have fueled significant growth in financing, this market is still relatively nascent. As such, we expect this space will continue to receive more funding over the near term, and consequentially attract more new entrants.

## *Personalized healthcare*

AI/ML in the healthcare industry will also allow the opportunity for greater personalization of care. Armed with more detailed data on individuals, healthcare can be tailored based on personal history, genetics, demographics and real-time inputs. A key use for this is changing the way we prescribe medications. Currently, many treatments are given out in a standardized fashion, especially regarding the size of the dose. This can lead not only to less effective treatments, but also the potential for adverse reactions—the cause of thousands of deaths every year[3]. Chemotherapy treatments, for example, are largely thought of as a one-size-fits-all option; however, personalization in this instance has huge potential by determining the biomarkers found in the patient's tumor and tailoring molecular-based therapies or the most effective chemotherapeutic agents depending on that patient's specifics. Siris Medical is already providing personalized cancer treatments that should accelerate the treatment planning process to streamline workflow and strengthen health outcomes.

Reinforcement learning algorithms, such as contextual bandits, provide a promising solution to personalization interfaces. These programs use contextual and historical data to decide among the available options, then receives reward-based feedback based on decision it has made. The algorithm then learns to optimize its decision to the most favorable outcomes by repeating the process a multitude of times. This feedback process is not perfect and the bandit may discard a viable action during the course of its learning, so the algorithm must strike a balance between exploiting what it believes is current optimal decision while still exploring all available actions.

The reward function, especially when dealing with biology and health, is fairly abstract because there aren't clean, quantifiable data points or easily determinable right or wrong answers, especially not relayed in a timely fashion. This is where the integration of real-time data from wearables and other sensors begins to play a key role in proving the feasibility of many of these applications and, ultimately, making the analytics more accurate and valuable. For example, a wearable device could track blood sugar data to provide feedback to the bandit on the selected treatment, then determine a positive or negative reward in near real-time that serves to improve the program.

**Note: Contextual bandits** are, in essence, computerizations of a slot machine with multiple arms corresponding to the number of available options where every decision returns a reward to the program. Before the program (i.e., bandit) makes its next decision, it considers all the rewards received previously in conjunction with information about the current context (external data, multi-world testing, etc.). For example: programmatically generating personalized headlines on a news site to optimize clicks, the bandit reviews past decisions/rewards, headlines from other sources, what the user has clicked on before, etc. to come to a decision on which option to select.

[3] Shepherd, Greene & Mohorn, Phillip & Yacoub, Kristina & Williams May, Dianne. (2012). Adverse Drug Reaction Deaths Reported in United States Vital Statistics, 1999-2006. The Annals of pharmacotherapy. 46.

## Headwinds

Although deployment of AI/ML in the healthcare industry might deliver better health outcomes in a more efficient manner, many factors will impose limitations on the adoption and scalability of these technologies. The implementation of large-scale AI programs is likely to be mired by prevalent data fragmentation issues, as the healthcare sector produces large amounts of unstructured data that are handled differently by various institutions. As more data are collected (via wearables, IoT, etc.), more emphasis will be put on the collection and cleaning of data before it can be used to train algorithms. Additionally, companies that choose to operate in this space must consider the regulatory and security issues around healthcare data, a barrier that has kept many companies from entering the industry. Clean datasets and the ability to securely store and share this data are critical to expanding capabilities of the healthcare AI field in a sustainable manner.

Humans innately like to know the reasoning behind decision making, especially when it concerns our health. As AI/ML programs start assisting with diagnoses, more questions will arise about the rationale behind recommendations. The "black-box problem" – as this phenomenon is known – is an important consideration in the AI/ML field. However, as there is already a relative lack of interpretability of human intuition relating to complicated medical situations, we believe this fear in healthcare may be slightly overhyped. When an AI agent makes diagnoses, it could have the benefit of drawing on the complete patient history, aggregation of other relevant patients, relevant studies and other sources to come to an evidence-based conclusion. That said, it will still be important to continue to have humans make the final decisions as a sanity check and for patient confidence, because the importance of the physician's experience and direct contact with the patient cannot be denied. This is particularly pertinent when it comes to assigning culpability for misdiagnoses.

On a related note, we find another hurdle to widespread adoption of AI/ML in healthcare to be the preference for human relationships. Automating initial contact points with both physicians or administrative workers in clinics may not be well received by the broader patient population. However, this should be mitigated partially because in general we see AI/ML technology in this industry as working much more to assist workers rather than replace them.

## Looking forward

On the frontier of the AI/ML healthcare field is the possibility of new drug discovery by machines. The spontaneity and creativity required in novel discovery adds an extra layer of complexity to this use case of AI/ML. Bringing a new drug to market is a notoriously time-consuming journey that can take more than a decade and between $650 million and $2+ billion of investment[4]. Many times, large-cap pharma & biotechs try to avoid this commitment by acquiring smaller companies with promising drugs already in the clinical trial process. AI/ML programs may shine in this context due to their ability to model a multitude of outcomes which help predict side effects or lack of efficacy, thus reducing failure rates and rapidly shortening R&D timelines, while potentially allowing more efficient in-house drug development processes.

Though drug discovery with AI/ML is still in a nascent stage, algorithms have already been used to determine whether a drug molecule will bind with the target protein with up to 99% accuracy and are working to predict the outcomes of CRISPR-Cas9 treatments. These predictive modeling techniques can also be applied to genome data to find novel patterns and approaches to treating genetic diseases.

## Conclusion

Due to the healthcare industry's size and depth of data, it's clearly a huge market opportunity for AI/ML. As many applications remain in the proof-of-concept phase or single–clinic partnerships, the speed of adoption and total market penetration remains uncertain. Over the next several years, we expect the more developed areas like assistants, medical imaging, and diagnosis to increase reliability and drive further use of these platform services, with the winners enjoying positive network effects. As the industry receives more funding to continue to expand research efforts, we expect the pace of breakthroughs to accelerate. The prospect of improved health at the population level provided more efficiently and lower cost is an enticing outcome that AI/ML can provide and cannot be ignored.

[3] Tufts Center for the Study of Drug Development