Global Insights | Wellness

Smart Thinking on AI in Health: Part 1 – Overview

What's happened in AI in health, and what's coming next

28 Mar 2023 20:27:06 ET

Executive Summary

Adam Spielman
adam.spielman@citi.com

Wenyan Fei
wenyan.fei@citi.com

This is the first of five reports focusing on AI in healthcare. We believe that in time, AI will transform healthcare, gradually but profoundly. So far AI's impact has been modest, with interesting developments in radiology, BioPharma and in a handful of health-specific admin systems. In the next few years, however, AI is likely to automate a good deal of health administration, saving time and money. We also expect AI to revolutionize the BioPharma industry because the transformer technology that can understand ordinary languages, and which is driving innovations like ChatGPT, can also be used to analyze molecules like proteins and DNA. In the long term we believe AI will fundamentally change the relationship between doctors and patients. We believe diagnoses will become increasingly automated: over time, family doctors will become more like health coaches and specialists will become more productive, attempting more personalized and ambitious interventions.

"We always overestimate the change that will occur in the next two years and underestimate the change that will occur in the next ten."

– Bill Gates

"Healthcare is simultaneously one of the fastest evolving markets, and yet the one market most resistant to change. Global healthcare is so big – $8 trillion – that it can sustain this paradox."

– Po Bronson, General Partner SOSV and Managing Director of IndieBio

Contents

Smart Thinking on AI in Healthcare

> RELATED:
>
> The Transformation of Health:How Consumer Pressure and the Data Revolution Are Changing Healthcare for the Better

This is the first in a series of five reports on Artificial Intelligence in healthcare, and it aims to provide a broad overview, both on what's happening and the implications. We believe that in time, AI will transform healthcare, gradually but profoundly.

- **So far** AI's impact on healthcare has been relatively modest, with some interesting developments in radiology, BioPharma and in a handful of health-specific admin systems. *These products are mostly point solutions that address the needs of the pre-existing players in their pre-existing workflows.*
- **In the next few years** AI is likely to automate a good deal of health administration, relieving workers of mundane tasks, saving both money and time, and thereby improving clinical outcomes. *This sort of AI will integrate the different parts of the system, for example the payers and providers, much more tightly*.
- **We also expect AI to profoundly affect the BioPharma industry.** The transformer technology that can understand ordinary languages (like English), and which is driving innovations like ChatGPT, can also be used to analyze large molecules like DNA and proteins.
- **In the long term** we believe AI will fundamentally change the relationship between clinicians and patients. We believe diagnoses will become increasingly automated. We expect that family doctors will become more like health coaches and hospital specialists will become more productive, attempting more personalized and ambitious interventions. *We therefore think AI will trigger a significant change in the way healthcare is delivered.*

Mike Guarino is the global head of health technology at Citi's investment bank, and he believes AI will be a central driver in the future of health. "AI is going to be a critical component of healthcare moving forward," he says. "We believe it will accelerate trends that we have been seeing in the market, which include a focus on increased value through improved clinical outcomes and cost efficiencies, enhanced patient experience, and workflow simplification."

We have tried to keep this introductory report quite short, so we haven't included many case studies. However in the more detailed follow-up reports, we will include lots of examples of what companies are actually doing to bring the story of AI and healthcare to life.

Figure 1. Reports planned in the Smart Thinking on AI in Healthcare series

1	**Overview**	AI will progressively transform the health system
2	**Medical Devices**	AI devices are already becoming more common, especially in radiology
3	**Health Administration**	AI will help automate a lot of admin, taking over mundane tasks and saving time and cost
4	**BioPharma**	BioPharma will be deeply affected as transformers analyze large molecules
5	**Role of Doctors**	Gradually AI will automate diagnosis and prescription, letting doctors focus on higher tasks

Source: Citi Global Insights

Health is a critical area for AI as well

Not only is AI important for health, but health is important for AI.

"When you look at all the areas that AI could be used in, healthcare really stands out," says Sirisha Kadamalakalva, head of AI for Citi's investment bank. "It's partly because health is such a big industry, but it's also because there is just so much room for improvement. I think it's a long way behind financial services in terms of actually using AI, for example, and that's a real pity."

Last year, the two most frequentially cited academic papers involving AI both focused on applying AI to protein structures – and not (despite the hype) on generative AI.[1]

Figure 2. AI in Healthcare

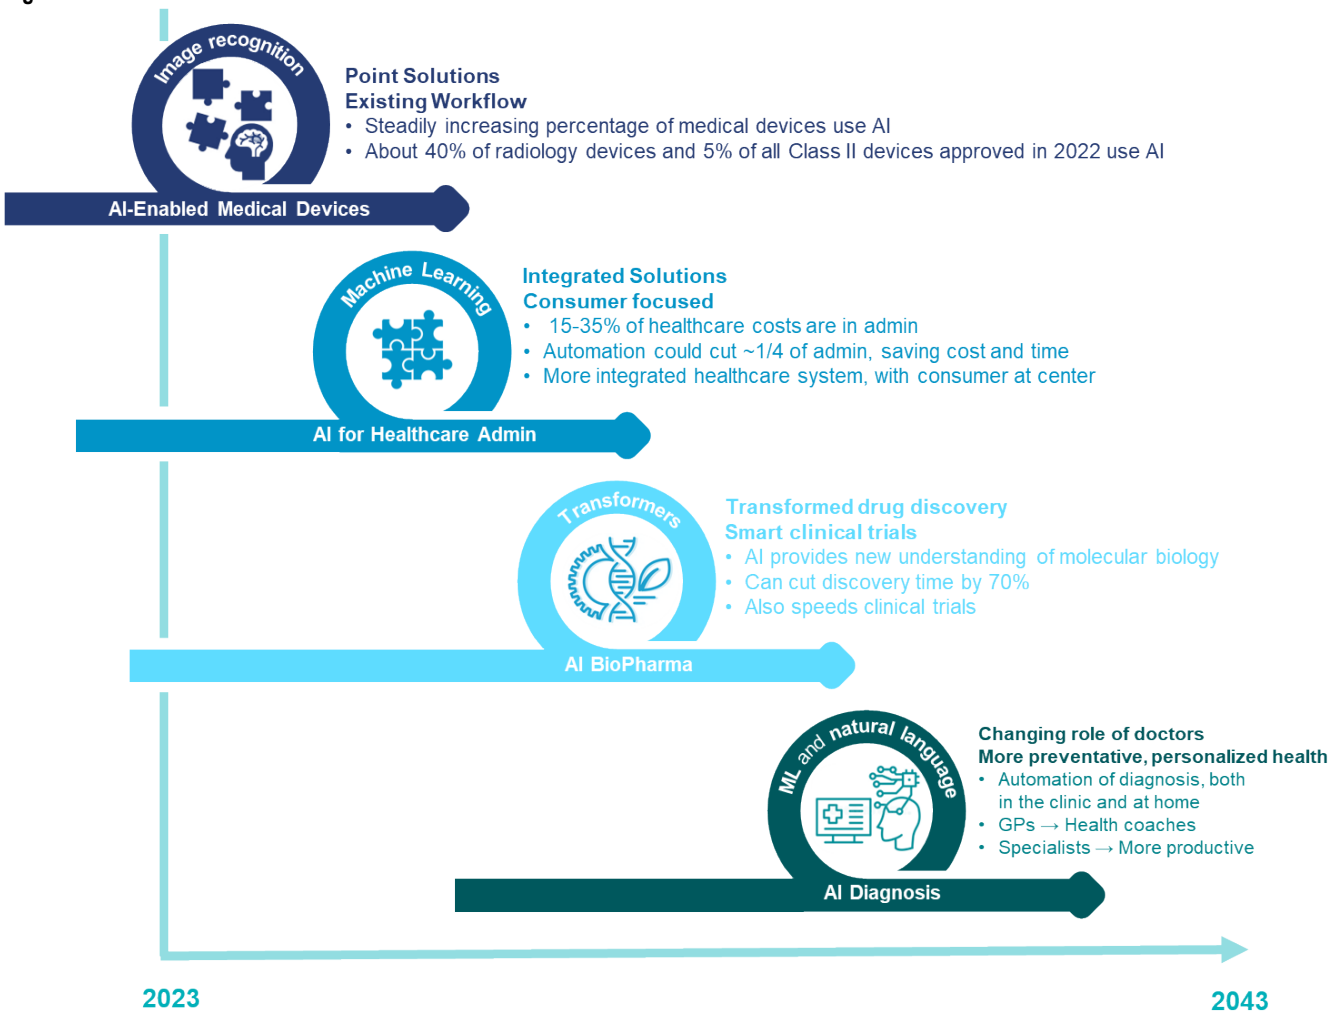

Source: Citi Global Insights

[1] AlphaFold Protein Structure Database (https://academic.oup.com/nar/article/50/D1/D439/6430488); ColabFold: Making Protein Folding Accessible (https://www.nature.com/articles/s41592-022-01488-1)

New types of AI are likely to lead to new types of use

AI comes in many forms, as we discussed in AI Time - 10 Ways Artificial Intelligence is Getting Real. Until now the main types of AI that have impacted health have revolved around image recognition – which is why the biggest impact has been in radiology – and machine learning (or ML).

Looking forward, however, we think transformer technology is going to be even more important, because transformers allow computers to understand both ordinary human languages like spoken English and the language of large molecules like DNA and proteins.

Image recognition was a key breakthrough, transforming radiology

AI is particularly suited to radiology because (1) medical images contain a huge amounts of data,[2] and (2) AI started outperforming humans in image recognition almost a decade ago.

ImageNet

One of the two most important breakthroughs in AI came back in 2012, with the ImageNet competition.[3] ImageNet was (and still is) a database that contains lots hand-labeled images – about 1.3 million at the time; about 16 million now. The breakthrough came when a program with a new type of neural network[4] won the competition to identify them accurately, and by 2015 neural networks were outperforming humans. Since then image recognition has gotten more and more accurate, and as a result, AI is now exceptionally well suited to categorizing images, which is the core of diagnostic radiology.

Machine learning is critical for health administration and diagnosis

Of course AI was never just about image recognition: there have been a series of improvements in other types of AI, principally around machine learning. ML allows computers to get gradually better at specific tasks – whether playing board games, administering medical payment claims, or diagnosing diseases.

The progress in machine learning was illustrated in 2016, when DeepMind achieved its celebrated victory in Go over Lee Sedol, using a program called AlphaGo. And then a year later, AlphaGo Zero – which was trained by playing against itself, with no reference to games that humans had previously played – beat AlphaGo convincingly, showing again how ML techniques were improving.

---

[2] In 2015, IBM researchers said 90% of data in healthcare comes from imaging. https://www.forbes.com/sites/matthewherper/2015/08/06/with-a-1-billion-deal-ibm-promises-to-teach-watson-to-see-can-it/

[3] The other really important breakthrough was the creation of transformers in 2017.

[4] A neural network is program that mimics elements of the biological networks in the human brain. The winning entry was called AlexNet, and it used what's known as a "convoluted neural network."

Transformers are the most important development of the past few years

The second great breakthrough in AI came in 2017, with the introduction of what are called “transformers”. We think this technology will have a much bigger impact on health than image recognition.

Transformers take any sequential information – for example the words in a text, or the amino acids in a protein – and analyze them by taking into account the relationships between the elements that come before and after, weighing them by their importance.

Transformer models don’t need to be trained on large labelled datasets. They are also good at altering their predictions based on context. As the CEO of Nvidia, Jensen Huang, put it, “transformers made self-supervised learning possible, and AI jumped to warp speed.”

Transformers will enable computers to interact with humans more and more effectively

Transformers are a breakthrough in understanding natural language, explains Jeffrey Zhu, a project manager at Microsoft. “Unlike previous [techniques] that processed words individually in order, transformers understand the context and relationship between each word and all the words around it.”

As a result it’s now reasonable to say that what are called Large Language Models, or LLMs, are getting better at *understanding* languages like English, Chinese and Spanish. AI can now extract the main points from an unstructured piece of text, and also write in an impressively human-sounding way, as ChatGPT and GPT-4 show.

AI is also getting much better at transcribing humans speaking. Working out what someone is saying requires the listener (whether human or computer) to have a sense of which word is likely to come next. Nuance, which Microsoft bought last year for $20 billion, has just launched a system that listens to the doctor-patient conversation and uses GPT-4 to instantly draft clinical notes.

But they will need to become more truthful

There are, however, limitations to the current generation of LLMs, including ChatGPT and GPT-4, that limit their use in medicine: they don’t really understand the concept of truth. As a result they can produce plausible sounding but misleading statements. That may not be a particular problem for a high school essay, but it is potentially deadly in healthcare.

This report assumes that in time, AI will be able to correct this, and that medical AI will stop generating plausible but erroneous statements. However it’s possible that this assumption turns out to be wrong, in which case the use of LLMs in healthcare will have to be very carefully monitored if they’re used at all.

Transformers also unlock the language of biology

Transformers were first used to analyze and predict natural language, but it turns out that the same techniques work equally well on biological structures – like DNA and RNA. AlphaFold, for example, can predict the 3D shape of proteins, and Nvidia’s AI-as-a-service aims to help researchers predict the function of biomolecules. These techniques have truly profound implications for BioPharma, which we will examine in the fourth report in the series.

Uses of natural language AI in health

It is too soon to talk of many successful companies using natural language AI in healthcare because developments in health take time, and the use of natural language understanding is nascent. However, we think it *is* instructive to analyze patents in the area.

We therefore did a search on small companies who have filed patents that combined "natural language" or "generative AI" and "healthcare" or "health."[5]

- One third of these companies are aiming to produce what we describe as natural language clinical assistants – in other words tools that use natural language inputs to collect data and provide medical advice either to clinicians or consumers. (It is significant, in our view, that often it isn't clear whether the business model is aimed at medical practitioners or consumers.) *We think that over time, AI's ability both to communicate fluently with people and (via ML) to diagnose and prescribe will change the role of many clinicians. We explore this in detail in the fifth report in our series.*
- A further big chunk, almost another 30%, aim to provide administrative support for health care professionals using natural language inputs, for example allowing spoken inputs for EHRs[6], or improving the process for dealing with pre-authorizations. *We discuss this in the third report, on AI and health admin.*
- For about a quarter of the companies with these sorts of patents, the prime focus is (still) medical imaging. *This is likely to accelerate the growth of AI-enabled radiology devices, which is the focus of the second report.*

Figure 3. Breakdown of small companies filing patents that referred to "natural language" or "generative AI" and health

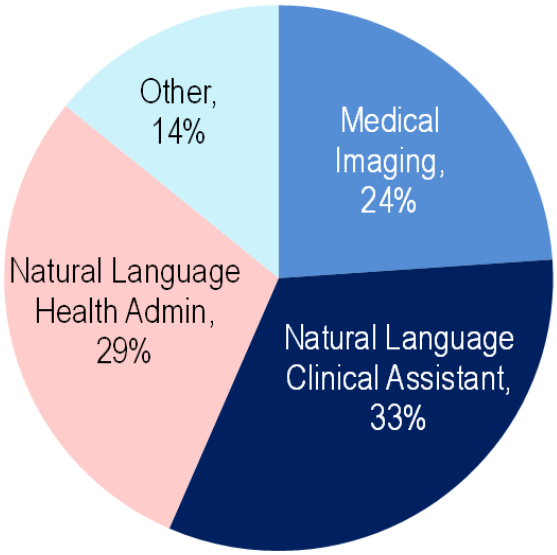

See Appendix for explanation of methodology

Source: Citi Global Insights

[5] See Appendix on page 29 for precise methodology.

[6] EHRs: Electronic Health Records. Also known as EMRs (Electronic Medical Records.)

Reasons for a measured view

As we've said, we think AI will *gradually* change health, over the course of decades. It's important not to get too carried away, however, because several factors make it particularly difficult to implement AI solutions in healthcare.

- **The human connection is vital in healthcare.** People want empathy as well as a prescription. A clear majority of Americans (60%) are uncomfortable with the idea of clinicians generally relying on AI, as Figure 4 shows, and two-thirds of women are uncomfortable.[7] The rejection rates increase when they think the issue is highly personal: 79% of the population don't want AI chatbots involved in mental health and 67% don't want AI used to determine pain medication. By contrast when the issue is determining whether there is skin cancer present, most Americans do want AI used.

Figure 4. American Attitudes toward AI and health

% of American adults comfortable with health providers relying on AI for clinical decisions, or wanting it used for particular treatments

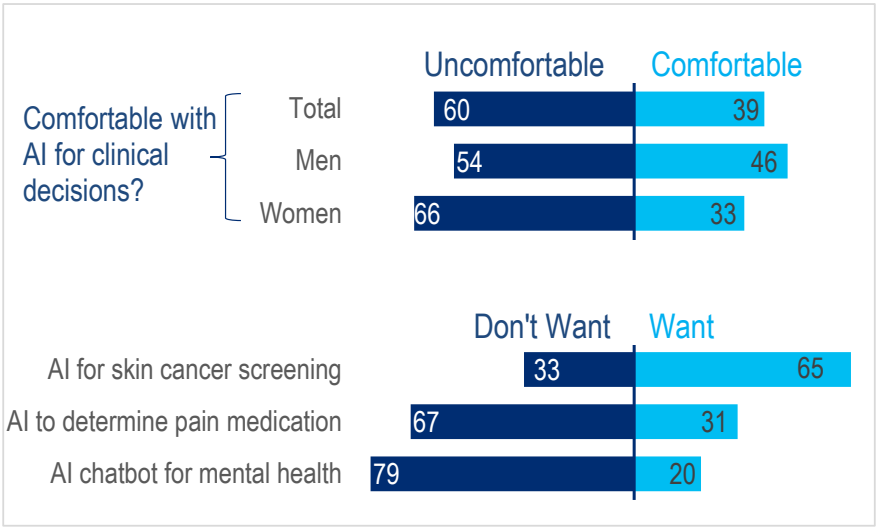

n = 11,004. Research conducted December 2022.
Source: Pew Research Center

- **"The stubborn persistence of the physical" -** Clinicians frequently need to see how people look, for example, or find out precisely where the pain actually is, or perform biopsies, and all these things need to be done by humans. In addition, many treatments go beyond prescribing pills. We don't expect robots to be drawing blood or injecting vaccines for many, many decades, for example.

- **Clinicians tend to be highly conservative -** The uptake of AI in diagnosis/ treatments is often slow because many clinicians are surprisingly hostile. There is a cultural mismatch between the medical profession – which is used to evidence based on multi-year RCTs[8] of specific products – and the AI, which uses constantly evolving algorithms applied to constantly evolving data.

[7] Older people and less educated people are also generally less comfortable with AI, according to the Pew survey.

[8] RCTs: Randomized Controlled Trials.

- **The reimbursement system often hinders adoption -** Executives find it very hard to predict whether there will be reimbursement for any particular digital health product whether in the U.S. or Europe (See Figure 16 on page 17.)

- **Legal liability is unclear -** This makes providers wary. Who bears the blame when things go wrong? The healthcare provider or the data scientists?

The cultural chasm between the health industry and AI

As we've mentioned, it seems to us there is a chasm in terms of cultures between the medical establishment and AI.

Whereas AI is constantly evolving, producing what seem like revolutionary new developments every few quarters, the medical establishment tends to be very risk averse.

Medicine assumes products should be tested for years or decades to show safety and effectiveness…

In medicine there is a general assumption that products are fixed and that it is reasonable for innovative products to be delayed for many years to gather sufficient high-quality evidence around safety and effectiveness.

One example is that the process of clinical trials for approving new drugs is so drawn out that it takes a median of 10½ years for a candidate molecule to go from the start of phase 1 clinical trials to authorization.[9]

… whereas AI products are constantly evolving

By contrast the AI community simply doesn't work like this:

- The underlying technology changes very rapidly: approaches developed two years ago now seem like commodity software.

- The products themselves learn as they come across more data. It's also possible for the underlying algorithms to change. As a consequence the interpretation of a given piece of evidence may well change over time.

- The tradition is for AI products – except in healthcare – to be released in open-source, beta trials for anyone to experiment with, partly because VCs typically fund startups on the assumption they can scale very rapidly.

… and many doctors don't want to use AI tools

The feedback we have from speaking to several CEOs is that many doctors don't want to use AI-powered clinical tools. In the second edition of this report series we will highlight the British Association of Dermatologists current position against an AI-driven triage tool, but there are other examples.

[9] https://go.bio.org/rs/490-EHZ-999/images/ClinicalDevelopmentSuccessRates2011_2020.pdf

Brian Pepin is the founder and CEO of Rune Labs, a software and data analytics company that works in neurology, in particular on Parkinson's. "Doctors don't want to be replaced," he told us. "Beyond self-interest, they feel strongly it is not in the patient's best interest."

Rune has been successful with clinicians, he says because it tries to be "a co-pilot not an autopilot. The difference is that an autopilot threatens to take over the physician's job, whereas a co-pilot helps them do their job better, and doesn't change their workflow too much."

…so many AI-first companies do not take the regulatory route

As a result many startups offering AI solutions do not seek regulatory approval, and instead aim their products at the "consumer wellness" market. This allows products to evolve rapidly in response both to technical advances and consumer feedback, but it also means that they are not allowed to offer medical-style readouts, for example a diagnosis.

However, this tendency to focus on "consumer wellness" can further reinforce the medical establishment's skepticism toward AI solutions.

Other issues around AI and health

Beyond consumers' desire for empathy and the cultural chasm between health and AI, there are several other issues that need to be treated carefully if AI is to become more common in healthcare, for example around privacy and the risk of bias.

Figure 5. American Attitudes toward AI in healthcare

% of American adults who think AI would make various aspects of healthcare better, worse or would make little difference

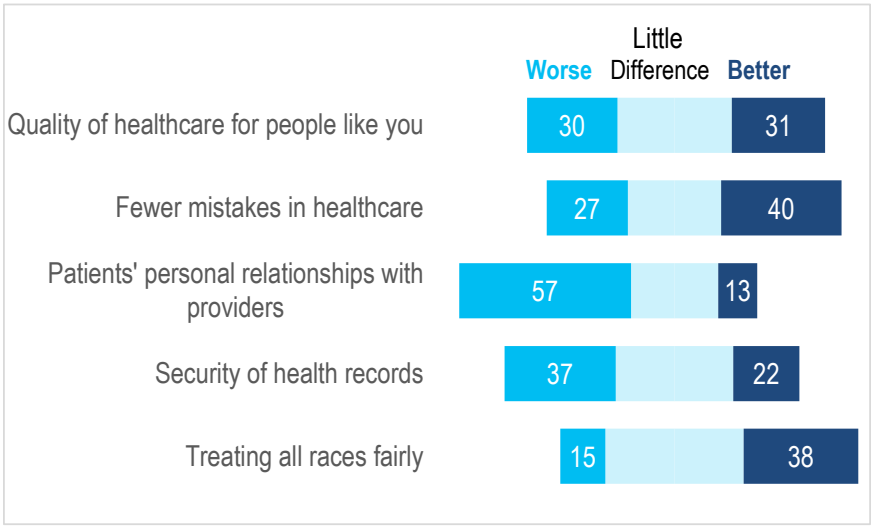

n = 11,004. Fieldwork in December 2022

Source: Pew Research Center

We find Figure 5 interesting. It shows Americans are split roughly equally between those who think AI will improve healthcare and those who think it will worsen it. On the positive side, more Americans believe it will reduce mistakes than think it will make more. On the other hand, the majority think it will worsen their personal relationships with clinicians, and unsurprisingly many think it will worsen data security. Most intriguingly of all, however, is that more Americans believe AI will reduce racial bias in the health system than increase it. And it's bias we turn to next.

Bias, health and AI

Many Americans, particularly African Americans, believe that racial bias is a problem in the healthcare system, as Figure 6 shows. What's important for this report is that of those people who believe that racial bias is either a major or minor problem, about half believe that AI is likely to make the situation better, whereas only 15% think it will make it worse.

The problem of bias in AI in healthcare isn't just about race of course. The important thing is that it should work well for all types of people - and ideally all types of people should feel comfortable using it.

Figure 6. % of Americans who think racial bias is a problem in healthcare

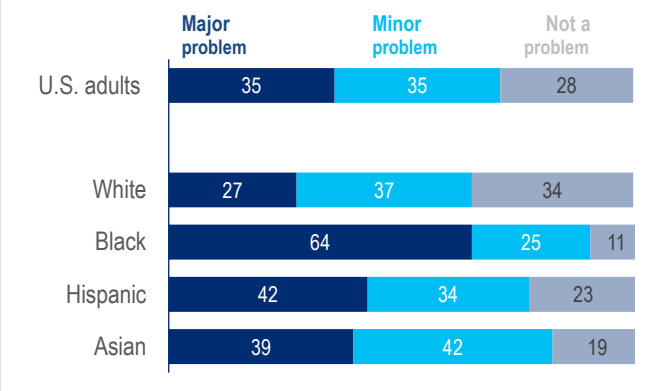

n = 11,004. Fieldwork in December 2022
Source: Pew Research Center

Figure 7. Of those who think racial bias is a problem, % who think AI will make it better or worse

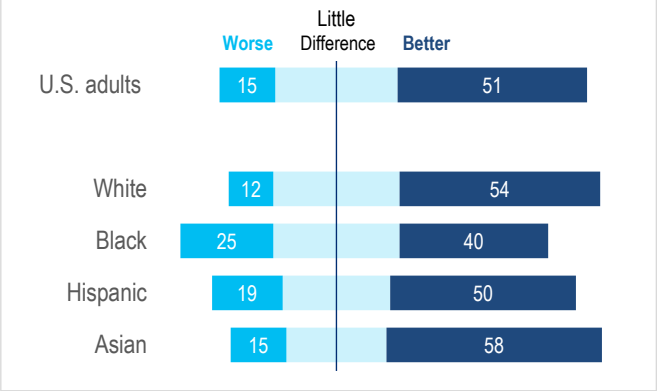

Source: Pew Research Center

People developing AI in health may find it easy to create products that they feel would help people like them. But these developers are not necessarily representative of the general populace. We assume that AI developers are more likely to be "trailblazers" in the language of Figure 8 (highly educated, in their early adult life and believe that technology can solve many problems) – but trailblazers represent only 16% of the U.S. population. By contrast 54% are considered either "bystanders" or "homesteaders" and both groups are much less interested in using tech to help with their health.

Morgann Carlon is a Health AI leader at Deloitte Consulting. "When we're staffing projects, it's really important to have a diverse team involved," she says. "The way farmers interact with the healthcare system is very different from someone from the city, and AI should understand that."

Figure 8. Distribution of Segments in Survey of U.S. Health Consumers, 2018

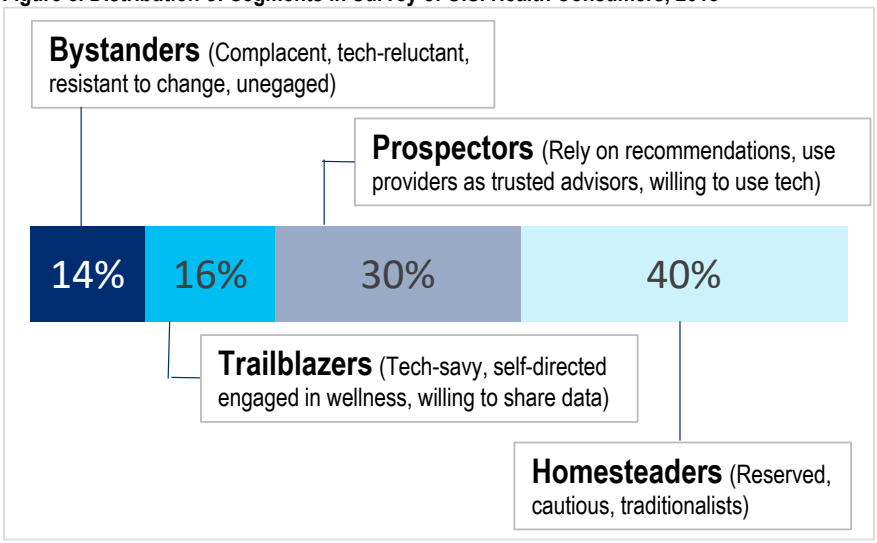

For more details, see Deloitte: Healthcare consumer segmentation

Source: Deloitte

Privacy is receiving more attention

Privacy is another important issue. It's not unique to AI because any product that can access people's health records raises issues around privacy. The issues are especially acute, however, with systems that combine lots of data – potentially from EHRs, genomic information and wearables.

Figure 9 shows that last year 70% of Americans were willing to share health data with their doctor, but only 8% with government agencies, down from 12% in 2020.

Figure 9. Percentage of U.S. adults willing to share medical data 2022 vs. 2020

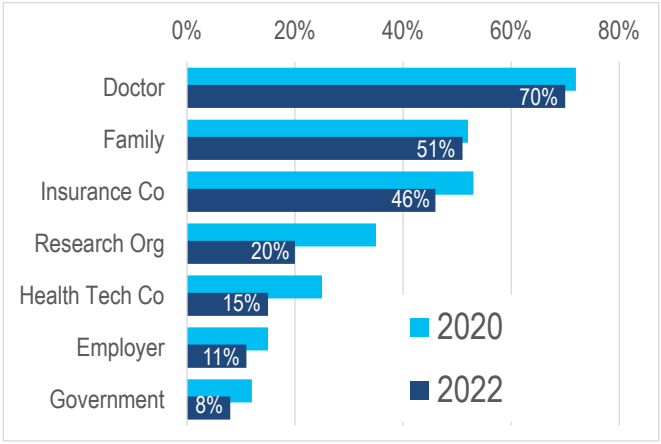

"Medical data" for this survey was defined as medical records, test results, prescription history, genetic information and physical activity.
n for 2022 = 8,014. n for 2020 = 7,980.

Source: Rock Health

Figure 10. Percentage of U.S. adults willing to share health device data obtained from wearables, 2018

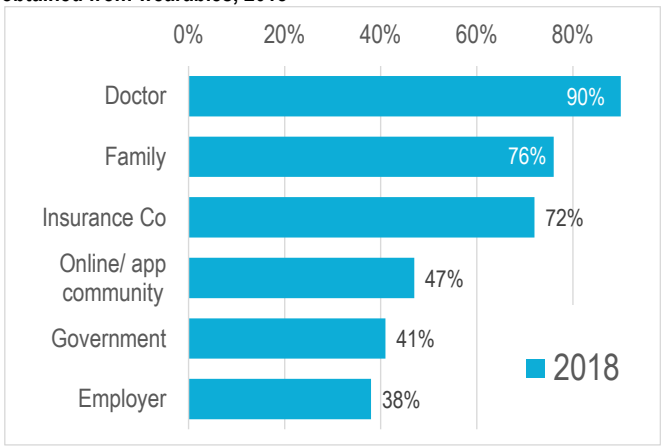

Source: Statista

We believe that Americans have become more guarded about their privacy during the past five years. Figure 10 isn't precisely comparable to Figure 9 as it focuses on wearable data but differences between the two are stark.

Morgann Carlon at Deloitte agrees that consumers are becoming more concerned about privacy. "People have become much more sensitive to this," she told us. "After the Roe v. Wade ruling, conversations about data ownership have skyrocketed, especially around healthcare apps."

She says that in her experience, however, the teams developing AI in healthcare settings are well aware of the need to protect privacy, and (in the U.S.) to obtain HIPAA clearance.[10]

[10] HIPAA: The Health Insurance Portability and Accountability Act of 1996. This sets out standards required to protect health information.

How we see AI progressing

Despite these issues, we expect that over time AI will become more and more embedded in health, via two different routes:

1. **Consumer push.** Disruption happens when there is both an unmet need and a technology to service that need. It is clear that the growth of "quantified self" technology (e.g. wearables) shows that many consumers want new types of health-focused products in their daily lives. We expect that over time this sort of technology will both deepen (e.g. with smartwatches that can advise consumers on their blood-sugar levels) and spread to an ever-larger percentage of the population. It's almost certain that AI will use this data to produce health insights – suggestions for both diagnosis and plans of action – and consumers will increasingly expect their doctors to assess these suggestions.

2. **Industry adoption.** At the moment there are only pockets on AI deployed in the health industry. However, we think they are likely to expand because we are optimistic both that doctors will see AI as helpful on the clinical side, and that AI will be able to reduce the administrative burden. We believe this will become ever-more evident in an industry where cost and other pressures are only increasing. We believe that as a result, more clinicians and managers will want to use more AI in more ways.

The convergence of health, tech and consumer will transform healthcare

All this fits in with our overall view that the future of health is being forged by the convergence of the health, tech and consumer industries, as set out in our Transformation of Health report, and summarized in Figure 11.

AI is a major pillar of what we call The Data Revolution, which is one of the three main drivers of the convergence. This revolution involves new ways of collecting, integrating, and analyzing data, and it's causing profound changes in many fields. Nonetheless it's particularly important in health, because data is often:

- Extremely **granular –** an image of a bone marrow biopsy, for example, can now be analyzed in minute detail by AI whereas previously pathologists could only give a generalized description[11];
- **Integrated** from many different sources **–** for example AI can synthesize data from wearables and from unstructured EHR records;
- **Communicated in real-time** to the individual consumer, their clinicians, and the public health authorities.

[11] The example of bone marrows comes from an Oxford research initiative on myeloproliferative neoplasms (or MPNs). These are a group of rare bone-marrow disorders, divided into five overlapping subtypes. They are diagnosed from blood counts and genetic mutations but the subtype is determined by bone marrow biopsy. Traditionally a human pathologist would assess about 10 different measures from the biopsies, by looking at about 200 fields under a microscope, trying to match the patterns with what's expected for the five MPN subtypes. The process is expensive and not especially accurate. However, AI can analyze 10,000s of data points on each bone marrow biopsy, mapping them into much more refined groups, greatly enriching our understanding of MPNs.

Figure 11. The Transformation of Health

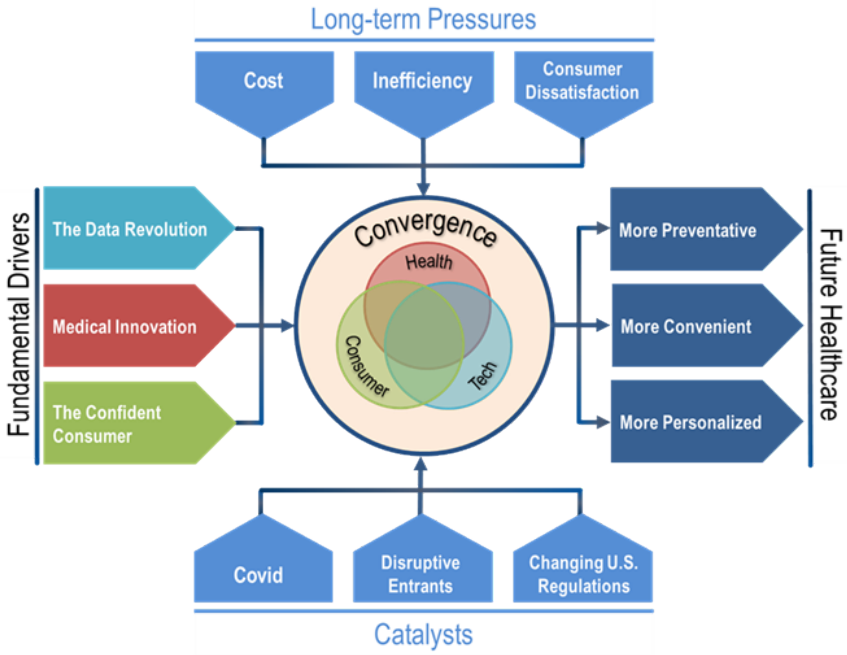

Source: Citi Global Insights

Over time, we think healthcare will become more preventative, more convenient, and more personalized. At the moment about two-thirds of health spend is entirely reactive, trying to sort out problems after they occur. Figure 12 shows that in 20 years' time, a much larger proportion of health spending will proactive, in other words trying to get ahead of problems. We think AI will drive a lot of that.

Figure 12. Expected Move from Reactive to Proactive Healthcare

Spending on healthcare in the U.S. -- Trillions of Dollars (inflation adjusted)

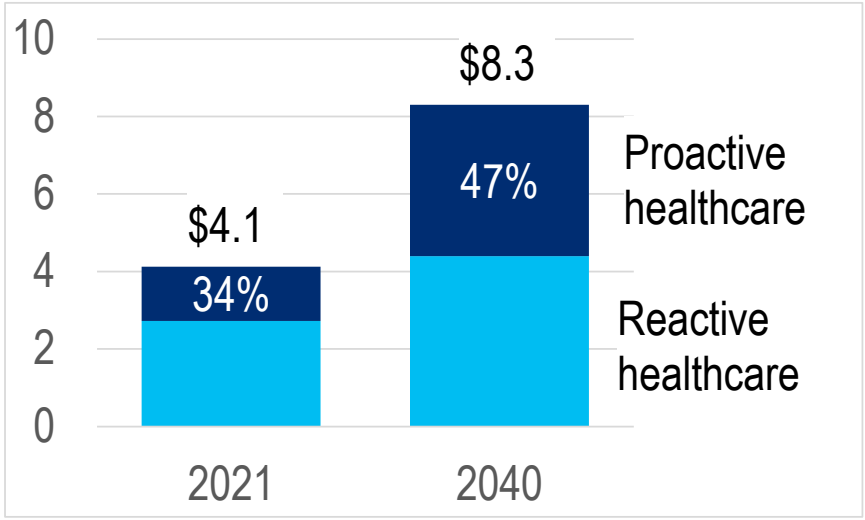

Source: Deloitte Consulting

AI and Medical Devices

***AI and Medical Devices* will be the focus of the second report in the *Smart thinking on AI and Healthcare* series**

It will include detailed case studies on several companies, and discuss how businesses can monetize AI in medical devices.

One part of healthcare where AI is already beginning to have a noticeable effect is in medical devices, particularly in radiology, in other words in medical imaging.

Peter Arduini is chief executive officer of GE Healthcare, the leader in radiology. "If you think of artificial intelligence, four years ago, we had hardly any products: no one did," he said at the investor day in December. "Today there really isn't a product coming out that's not enabled without some form of artificial intelligence."

In fact the number of AI-enabled radiology devices approved by the FDA has grown continuously, as Figure 13 shows, from close to zero in 2016 to about 40% of all Class II radiology devices approved.[12] About 5% of *all* Class II devices are now AI-powered. We expect these lines to continue upward.

Figure 13. AI-enabled device approvals as a % of all Class II radiology devices approved by the FDA

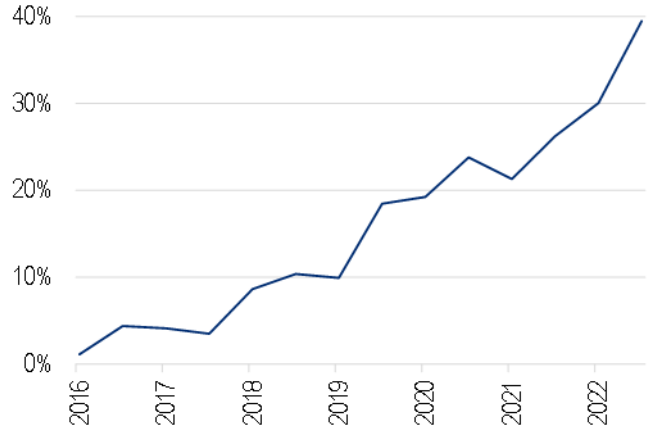

Note: FDA definition of AI
Source: FDA

Figure 14. AI-enabled devices as a % of all FDA Class II approvals

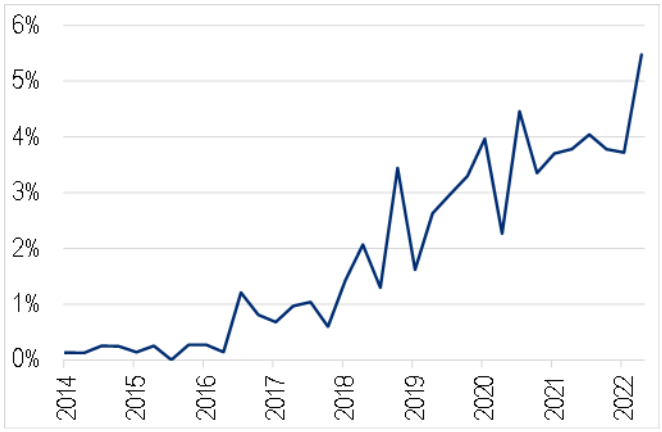

Note: FDA definition of AI
Source: FDA

Dollar sales are expected to grow fast

This means the sector is likely to grow rapidly. Of course, just because 40% of radiology devices that are *being approved* are AI-enabled does not mean that 40% of devices *in use* are AI-enabled. Clearly it will take many years for penetration to reach that level.

Overall the market – in dollar terms – for AI diagnostic devices is currently quite small: We looked at five independent estimates for total sales – combining hardware, software and services – and the average was under $1 billion for 2022. Nearly half (45%) of this comes from software.

The expected growth is very rapid. The average forecast for 2027 sales is $4 billion, implying a 33% 5-year CAGR. Last year we collated market forecasts for 100 growth areas in our Mapping Innovations report, and that 33% expected growth rate would put AI diagnostics in the fastest growing 10% of the areas we studied.

[12] Class III medical devices are the potentially most dangerous if they go wrong – often items inserted into the body, like pacemakers. Class I have very low risk, for example, an elastic bandage. Class II devices require an intermediate level of regulatory scrutiny, and include things like catheters and MRI scanners.

For start-ups and established players, that creates a big opportunity. In some areas, AI has now become a necessary cost of doing business; but it is possible to carve out a niche providing AI as a service, particularly when servicing biopharma

MedTech companies now prefer to launch AI enabled products in the U.S. than Europe

We find that medical device companies today overwhelmingly prefer to launch their products in the U.S., rather than in the EU. This is a complete turnaround from the early 2000s, when MedTech companies generally launched their products in Europe several years before they launched in the U.S.

In the past 15 years, the FDA has implemented numerous changes to make it easier to launch new devices, including setting up the Digital Health Center of Excellence. The EU by contrast, has implemented the Medical Devices Regulation, which is four times as long as the rule-book it replaced, and requires much more clinical data before new products can be authorized.

Figure 15 shows that 64% of U.S. respondents to a BCG/ UCL survey agree that the FDA has responded well to the growth of digital medical devices, whereas only 34% of European respondents would say the same about the EU. Almost a quarter of firms that already have EU approvals for products say that in future they will launch in China or Japan before they launch in Europe.

Figure 15. Assessment of whether US or EU regulators have responded well to digital technology

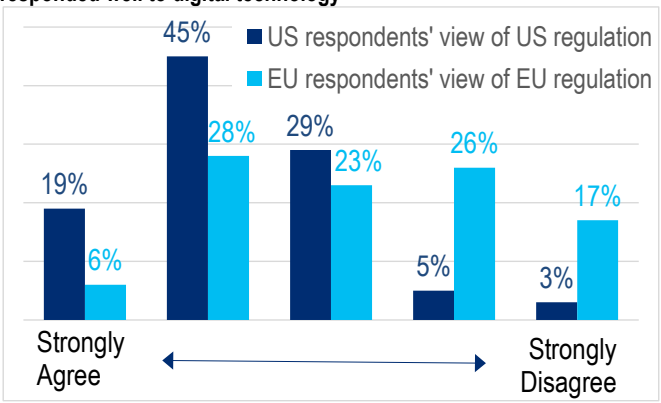

Source: BCG and UCLA Biobank

Figure 16. Predictability of Reimbursement for Digital Devices

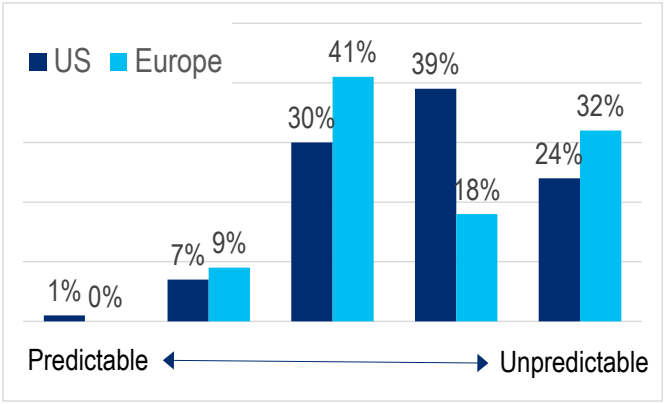

Source: BCG and UCLA Biobank

The limiting factor has become reimbursement

The FDA, or its equivalents overseas, may approve an AI-enabled device, but hospitals will buy it only if they believe payers will reimburse its use. Unfortunately companies find it extremely hard to predict whether hospitals will get reimbursed for new digital medical devices, both in the U.S. and in Europe, as Figure 16 shows.

In the U.S., reimbursement rather than approval now appears to have become the most important barrier to the use of new types of medical devices. More than 70% of the respondents to the survey said the primary barrier is now reimbursement; only 16% said it is FDA regulation. Many companies therefore say they want to see more coordination between the reimbursement authorities (e.g. the CMS) and the FDA.

AI and Health Administration

***AI and Health Administration* will be the focus of the third report in this "Smart thinking on AI and Healthcare" series**

We are optimistic that AI will also make a really significant impact on healthcare administration in the coming years.

This is important because (1) administration is a large part of healthcare that's often underappreciated by outsiders and (2) it's also very inefficient:

- **Excessive cost**: Very roughly ¼ of all spending on healthcare in the U.S. is on administrative – but automation could save perhaps 25-30% of this.[13]
- **Burden on clinicians**: More than 90% of clinicians say excessive admin is contributing to burnout.[14]
- **Industry fragmentation:** Friction in the flow of information means the healthcare system is currently split into distinct entities, e.g. the providers, payers, and government, with little integration between them, and mis-aligned incentives.

AI is already starting to help automate a great deal, and we believe the advent of really good understanding of ordinary language will accelerate this.

Dr Avi Mehra is a co-founder of Doctorpreneurs and CSO in IBM's healthtech team. "It's easy to think that sorting out health admin isn't that sexy," he told us, "But there will probably be more improvement in healthcare provision, for less effort, in applying AI to health administration, relative to the clinical side."

Figure 17. Admin Costs in U.S. Healthcare by Stakeholder (2019, $ in Blns)

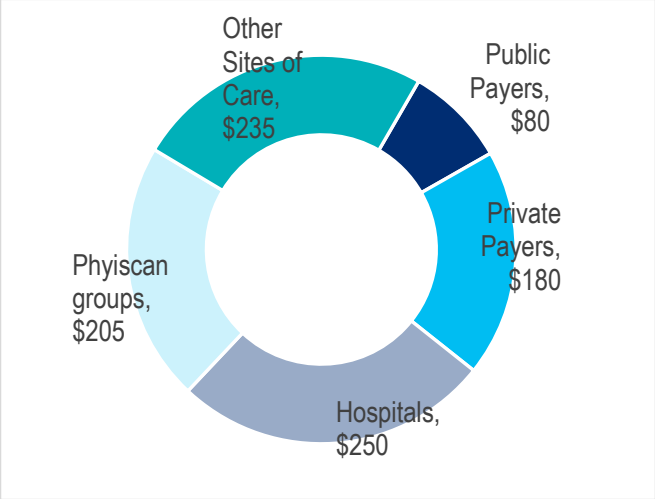

Other sites of care includes dental and nursing facilities and home healthcare
Source: McKinsey

Figure 18. Admin Costs in U.S. Healthcare by Functional Area (2019, $ in Blns)

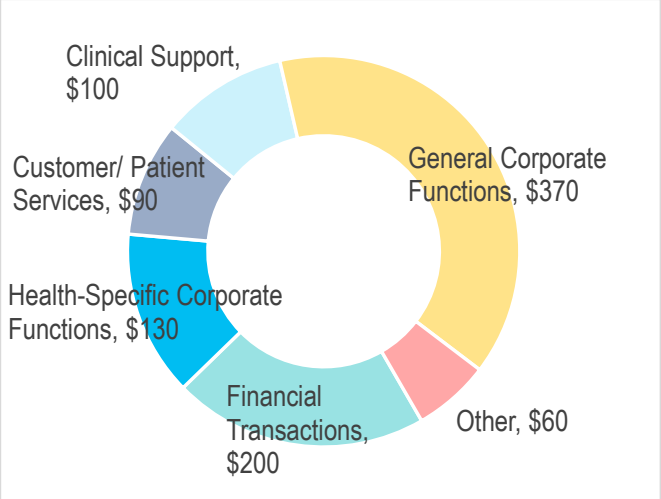

Source: McKinsey

[13] https://www.mckinsey.com/industries/healthcare/our-insights/administrative-simplification-how-to-save-a-quarter-trillion-dollars-in-us-healthcare#/

[14] Burnout is an increasing problem in many health systems globally, including in the U.S. https://discover.oliveai.com/rs/541-CSN-882/images/Report%20Findings%20Summary.pdf

AI likely to help integrate the currently fragmented system

Currently the healthcare system is split into distinct sub-industries, like providers, payers, and device manufacturers. In future, faster information flows between these companies mean that they will become more integrated, focusing more and more on the benefits of the consumer.

Neal Batra is responsible for the Future of Health at Deloitte Consulting, and he thinks there's a big opportunity to make the overall U.S. healthcare system more efficient. "Far too often at the moment, there's a big gulf between providers and payers and other parts of the healthcare system, like the government – and consumers get rather forgotten," he told us. "But AI will help integrate these things in the future -- and the consumers will be central in a way they just aren't now."

Figure 19. Likely evolution of industry fragmentation/ consolidation

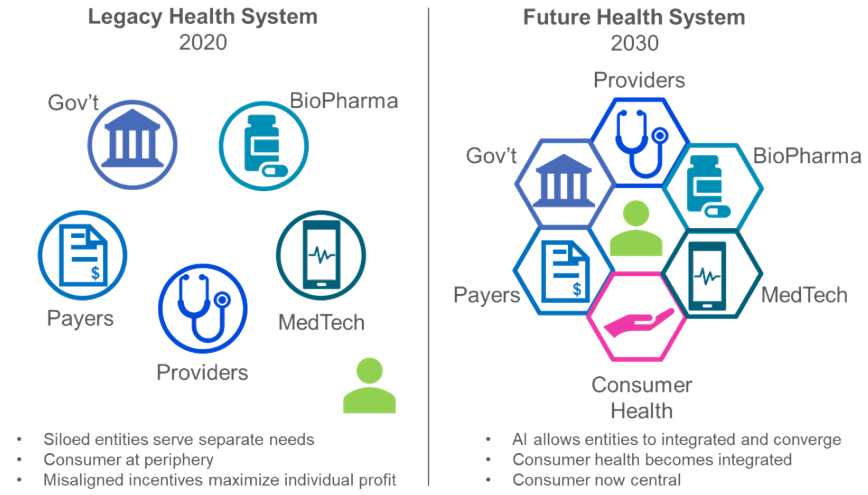

Source: Deloitte Consulting, Citi Global Insights

Natural language processing is likely to accelerate the change

As AI becomes better at processing natural languages, like English or Chinese, the burden of administration should fall further, and we think this will particularly benefit clinicians.

A number of companies – including Nuance (which is owned by Microsoft) and a startup called Abridge – are debuting products that offer ambient administration, in which a smart speaker listens in to the conversation between patient and clinician, and provides a draft of clinical notes for the physician to approve.

In future, we expect these sorts of products to:

- Extract the relevant information from the patient's EHR in real time and show it to the clinician;
- Recommend tests for the clinician to order; and
- Make the first draft of writing prescriptions, requesting prior authorizations, or adding whatever notes are necessary to the EHR, all subject to final approval from the human clinician.

AI and BioPharma

The fourth report in our "Smart thinking on AI and Healthcare" series will cover *AI and BioPharma*

We also think AI has the potential to transform almost every stage of the drug discovery process in the coming years.

"AI doesn't represent an incremental change; it's going to fundamentally change how drug discovery works, potentially reversing the decline in productivity seen across the industry," says Neal Batra at Deloitte. "And when quantum computing hits stride, the change will go even deeper."

AI-first companies are already making a big impact

The major BioPharma companies have struggled to develop new products in recent decades, despite spending more and more on R&D, as Figure 20 shows. On average it now takes 10 years and $2.6 billion to bring a drug to market.[15]

Figure 20. Selected Pharma Companies --- Combined R&D Expense in Billions of Dollars

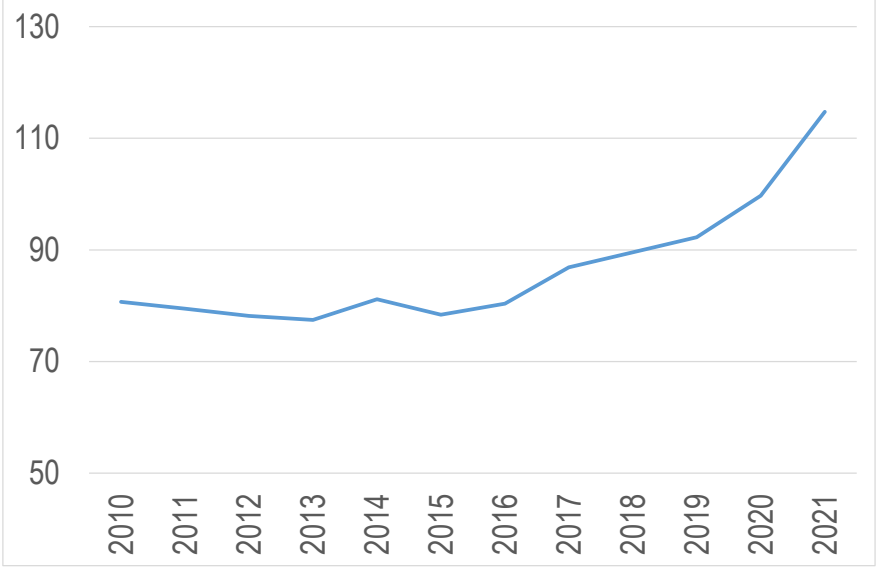

Companies included are: AstraZeneca, Bayer, Eli Lilly, GSK, Johnson & Johnson, Merck & Co. Merck KGaA, Novartis, Novo Nordisk, Pfizer, Roche, and Sanofi.
Source: Company Reports

Figure 21 shows that despite the increase in R&D spend, since 2010 the number of molecules developed by the top 20 companies that have reached the development/ preclinical stages has halved, from 705 to 333 (the number of assets in clinical trials has also dropped, albeit by a much smaller percentage, from about 1,100 in 2010 to about 1,000 now). Figure 22 shows that by contrast, the number of early stage assets generated by AI-first companies has surged in the last five years, and is now equivalent to about half the total generated by the top 20 pharma megacaps.

[15] http://phrma-docs.phrma.org/sites/default/files/pdf/rd_brochure_022307.pdf

Figure 21. Top 20 Pharma Companies -- No. of Assets in Discovery and Clinical Trials

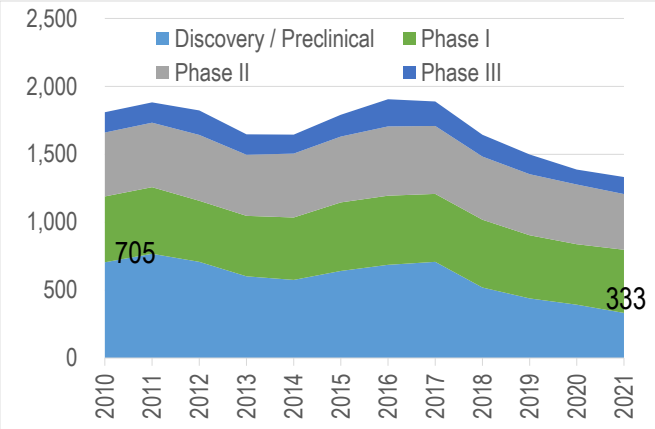

Source: Nature Reviews Drug Discovery via Recursion

Figure 22. AI-First Pharma Companies -- No .of Assets in Discovery and Clinical Trials

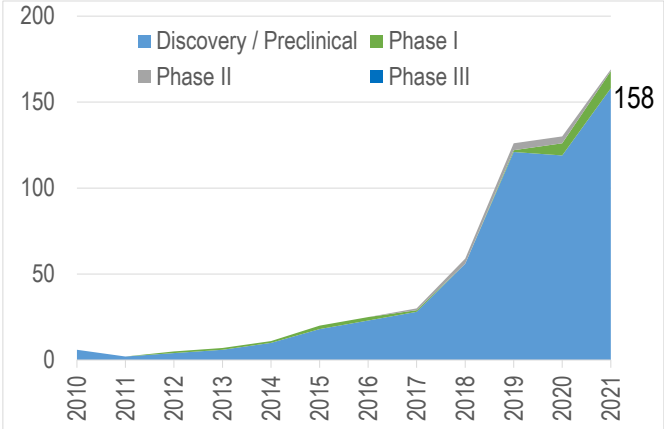

Source: Nature Reviews Drug Discovery via Recursion

One notable example of an AI-first BioPharma company is **Exscientia**, which has already taken three AI-designed drugs to clinical trials. It says its AI means it can:

- Reduce the discovery time[16] by 70%, relative to conventional processes;
- Reduce the number of compounds that need to be synthesized by 10x; and
- Reduce the capital required to go from finding a biological target to taking a molecule through preclinical toxicological studies by 80%.

Exscientia believes the greatest competitive advantage will come from designing the entire discovery process around AI, from soup to nuts. The company defines its mission as "encoding and automating every stage of drug design and development."

"The big question is who is going to be able to use AI really effectively first," Ben Taylor, the CFO and chief strategy officer of Exscientia, told us. "It's about integrating systems – every stage of the system has to be able to feed into all the others. If you get that right, and create a closed loop, then it's massively scalable. There will be a huge difference between those who can create a genuine learning system and those who can't."

AI can also help with gene therapy

Gene therapies are different from traditional medicines because they seek to change the multi-step process through which mutated DNA creates the proteins incorrectly. Scientists have developed a good knowledge of what many individual segments of DNA do. The trouble is that individual segments intertwine with each other and often act in combination, and that's much harder to analyze.

This is a bit like saying we know there's a series of books (the set of 23 pairs of chromosomes), and we understand many words and sentences – but that helps only a bit with understanding the books as a whole.

AI is much better at understanding the context of the DNA segments – and that opens up many new possibilities, as we will explore in the fourth report in the series.

[16] This is the time between finding a biological target for a potential therapy to identifying a candidate therapeutic molecule.

AI and Doctors

Report No5 in the "Smart thinking on AI and Healthcare" series will cover *AI and Doctors*

We believe that in the long term new forms of AI are likely to provide powerful tools for clinicians, gradually but profoundly changing their role. We expect diagnosing and prescribing will become largely automated, freeing up many family doctors to spend more time advising people in more wholistic ways. We also think AI will make clinical decision support systems much better, and allow specialists to become more productive, attempting more ambitious and personalized interventions.

In China, more than a third of doctors already use an AI-driven tool, called AskBob, to help them diagnose and treat patients. We think this reflects both the real commitment to AI shown by AskBob's owner (Ping An), but also the underdeveloped state of China's primary care, especially in rural areas. Part of the reason why there is no equivalent system in the West, is that human doctors do a better job (for now.) However, the success of AskBob does imply there is a big opportunity to improve primary care provision in geographies where the existing provision is weak.

We expect diagnosis tools to become much more common

In future AI is likely to become more user-friendly (both for patients and clinicians), and more accurate in more areas. We therefore think its role will increase globally.

AI will have better medical knowledge

ChatGPT – which isn't designed to be a medical specialist – is already nearly able to pass the USMLE exam[17], the gateway exam that everyone who wants to become a medical doctor in the U.S. is required to pass. It seems likely that in the relatively near future, generative AI systems with specific medical training *will* be able to pass.

And AI will have access to more data

AI will also have many advantages relative to humans in terms of medical knowledge, because it can analyze the entire corpus of medical literature, stay up-to-date on all medical advances, and track the potential interactions between all the various drugs on the market.

Furthermore, AI will be able to digest more data about individuals than human doctors. We expect that in coming years, AI will be able to synthesize: (1) Data from wearables – both around biometrics (e.g. around heart health and blood glucose) and on lifestyle issues, like exercise and sleep patterns; (2) Genomic data; and (3) EHRs, even if they are unstructured.

Accuracy is vital

As we've said, ChatGPT and GPT-4 are both liable to come out with plausible sounding but inaccurate statements. Our assumption is that fairly soon AI will be able to correct this, and that LLMs will not be prone to misleading their users. If this

[17] USMLE: U.S. Medical Licensing Examination. A preprint was recently published analyzing how ChatGPT would do on the USMLE. The authors concluded "ChatGPT approaches or exceeds the passing threshold for USMLE." Source: PLOS Digital Health Feb 2023

assumption turns out to be wrong – and it's not possible to make LLMs more accurate – they will have only a limited place in the health system.

The role of doctors is likely to change

Of course human doctors will remain vital, in part because people require a human connection and "the stubborn persistence of the physical." However we do think that AI is likely to have a big impact on the work of many doctors.

- **Family doctors**: we think that in 10 years' time, the "diagnose and prescribe" part of the job will become much less important because it will be largely automated. Rather, the role will be much more about being a health coach, discussing psychology and patients' health holistically. One of the implications of Figure 12 (on page 15) is that we think family doctors will concentrate more on preventative, proactive medicine, rather than on the reactive treatments, which is the focus today. And we think this will be a more satisfying, and value added role.

- **Hospital specialists**: AI is likely to make the systems that help physicians – the clinical decision support devices – much faster and more accurate. AI will provide better forecasts of likely outcomes of courses of action based on individual risk factors. We believe that as a result, specialists will (1) spend less time on diagnosis, and (2) use more personalized, and sometimes more ambitious, treatments. Overall we think AI will make specialist consultants more productive.

Likely path forward

Given the way the industry is regulated, two distinct sorts of diagnosis products are likely to evolve:

1. **"Medical" systems that work alongside doctors**, perhaps listening into to conversations, and helping in both in clinical terms (by suggesting potential questions to ask, possible diagnoses, and next steps) and with administration.

2. **"Consumer" systems,** which would work at home or on the go, and may or may not be FDA approved. We expect that over time consumers will use a growing number of wearables and other remote monitoring devices, and these will provide more accurate data on a growing number of biometrics and be linked by AI. We expect non-approved systems will make detailed lifestyle suggestions, while FDA-approved ones will eventually provide a sort of stripped-down AI doctor that would be available from smartphones, and that people could refer to several times a day at zero marginal cost.

In the very long-term, we think it is likely there will be a degree of convergence between the medical and consumer systems. For example we think that the annual check-up will be largely replaced by continuous monitoring, analysis and nudges.

Smart thinking on AI and healthcare

As we have said throughout this report, we think that in time, AI will gradually but profoundly change healthcare in many ways. It's already fundamentally altering radiology, and we think healthcare admin will be next. We believe drug discovery is ripe for disruption, and we think that eventually AI will fundamentally change the role of many doctors.

We will dig more into all of this in the future reports in our *Smart Thinking on AI in Healthcare* series, and bring the story to life with examples of what companies are actually doing. So please do keep reading.

Appendix: Analysis behind Figure 3

Citi Global Data Insights (CGDI) has established a thematic platform called the "Thematic Basket Incubator,"[18] which quantifies an entity's exposure to themes using alternative data such as job postings, patent applications, workforce profiles, news sentiment and others. The platform covers all types of entities including public and private companies as well as educational and other non-commercial entities.

The analysis behind Figure 3 (page 7) was limited to patents associated to generative AI and health. The data comes from over 80 offices across the globe, including China. Theme-relevant patents were identified using a combination of key word, proximity, fuzzy, and wildcard searches. The precise search terms are shown in Figure 23.

Figure 23. Patent queries used

"Natural language"~2 AND (healthcare OR wellness OR medical)'

'"Natural language "~2 AND ("early diagnosis"~2 OR "early detection"~2 OR "drug discovery"~2)'

'"Natural language"~2 AND ("Personalized medicine"~2 OR "treatment plans"~2)'

'("Generative Artificial Intelligence"~3 OR "Generative AI"~3) AND (healthcare OR wellness OR medical)'

'("Generative Artificial Intelligence"~3 OR "Generative AI"~3) AND ("early diagnosis"~2 OR "early detection"~2 OR "drug discovery"~2)'

'("Generative Artificial Intelligence"~3 OR "Generative AI"~3) AND ("Personalized medicine"~2 OR "treatment plans"~2)'

Source: Citi Global Insights

The analysis yielded 408 relevant entities. We then excluded ones where less than 10% of their activities revolve around AI and health (this effectively excludes large companies like IBM and Alphabet, because companies like these will inevitably be investigating many, many other areas). This left 111 companies for which AI and health is a major focus.

After the manual inspection, we filtered out a further 19 companies because we thought they were irrelevant. The remaining 92 companies fell under four categories: Medical Imaging, Natural Language Clinical Assistant, Natural Language Health Admin, and Other and the split is shown in Figure 3.

[18] Thematic Basket Incubator: Leveraging Alternative Data with Machine Learning and Knowledge Graph

If you are visually impaired and would like to speak to a Citi representative regarding the details of the graphics in this document, please call USA 1-888-800-5008 (TTY: 711), from outside the US +1-210-677-3788

Disclosures