


A GUIDE TO ARTIFICIAL INTELLIGENCE IN HEALTHCARE



DR. BERTALAN MESKÓ
THE MEDICAL FUTURIST



Authors:

Dr. Bertalan Meskó & Nóra Radó

Title:

A GUIDE
TO ARTIFICIAL INTELLIGENCE
IN HEALTHCARE

Copyright:

The Medical Futurist 2019



CONTENTS

Part I. THE BASICS OF ARTIFICIAL INTELLIGENCE

[Artificial intelligence: a reference point for innovation](#)

[Fears and expectations about A.I.](#)

[Let the quest for balanced views on A.I. begin](#)

[What is Artificial Intelligence?](#)

[Narrow, general, or super?](#)

[What do you need for developing A.I.?](#)

[Data Analytics, Machine Learning & Deep Learning – Methods of Teaching Algorithms](#)

[Data in healthcare](#)

[A brief history and the current state of electronic medical records](#)

[Why do we need help from A.I. when it comes to data?](#)

[Health data management](#)

[Treatment pathway design](#)

[Transforming diagnostics](#)

[Health assistance and administration](#)

[Patient management](#)

[Precision medicine](#)

[Supporting pharma: drug creation and clinical trials](#)

[FDA-approved Algorithms in Healthcare](#)

Part II. APPLYING ARTIFICIAL INTELLIGENCE IN HEALTHCARE

[Health data management](#)

[Treatment pathway design](#)

[Transforming diagnostics](#)

[Health assistance and administration](#)

[Patient management](#)

[Precision medicine](#)

[Supporting pharma: drug creation and clinical trials](#)

[FDA-approved Algorithms in Healthcare](#)

Part III. CHALLENGES OF ARTIFICIAL INTELLIGENCE

[Misconceptions and overhyping](#)

[Technological limitations of A.I.](#)

[Limitations of available medical data](#)

[The indispensable work of data annotators](#)

[Judgemental datasets and A.I. bias in healthcare](#)

[The need to regulate A.I.](#)

[The ethics of A.I.](#)

[Could you sue diagnostic algorithms or medical robots in the future?](#)

[Should algorithms mimic empathy?](#)

[Could A.I. Solve The Human Resources Crisis In Healthcare?](#)

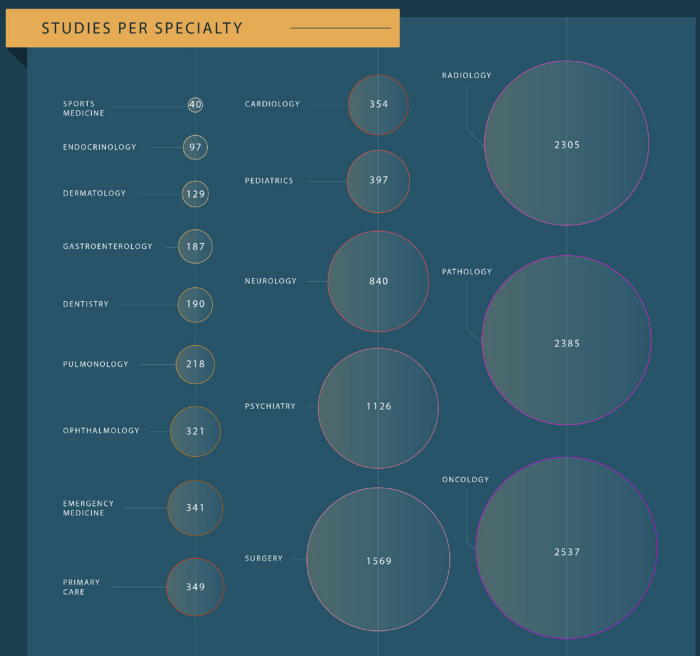
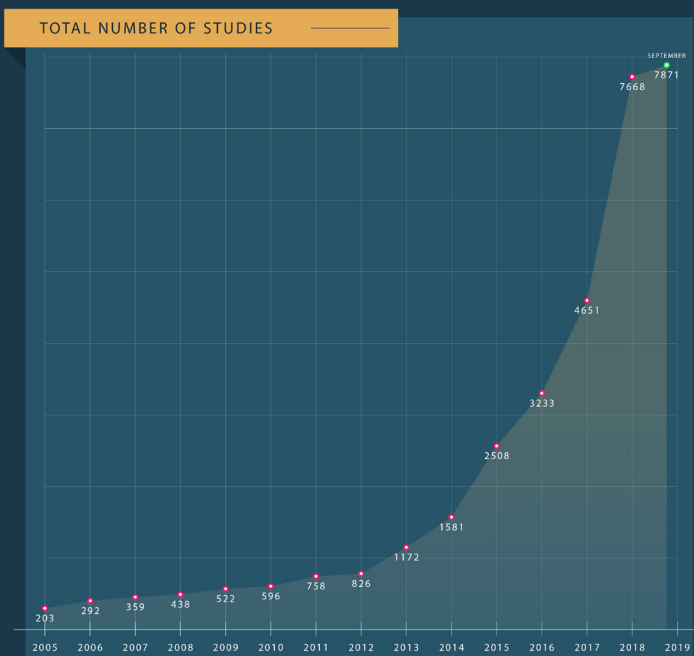
Part IV. MEDICAL PROFESSIONALS, A.I. AND THE ART OF MEDICINE

THE BASICS OF ARTIFICIAL INTELLIGENCE

Artificial intelligence: a reference point for innovation

In the last couple of years, artificial intelligence has evolved from a futuristic promise into an unavoidable reference point for innovation. Not only does it appear daily in news headlines, but the number of A.I.-related studies, research projects, university courses, and companies has grown exponentially, not to speak about the rapid improvement in the precision of the technology. The initiative called [A.I. Index](#) in its [2018 Annual Report](#) found that from January 2015 to January 2018, active A.I. startups increased by 113 percent, while all active startups only showed a moderate increase of 28 percent. Research into A.I. is similarly accelerating. Nothing presents that better as the number of studies published in the medical field. While in 2005, there were only 203 research papers on [Pubmed.com](#), one of the most prestigious databases for life sciences, the number rose to 7668 in 2018 and it's likely to hit a new record in 2019.

MACHINE AND DEEP LEARNING STUDIES ON PUBMED.COM

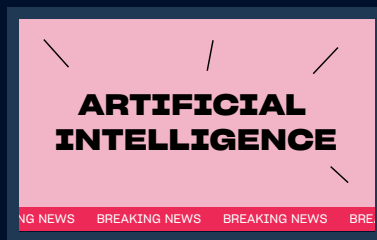


It was only a matter of time until the strategic importance of A.I. was recognized by fields other than tech industry, thus in recent years, we could follow closely how it has spilled over from the economy to politics as well. [A.I. Index noted](#) how the frequency of the expression's usage has skyrocketed in the last couple of months. In the U.S. Congress, the terms 'artificial intelligence' and 'machine learning' were used in 2018 more times than in the previous years in total. This is also reflected in the visible [A.I. race](#) between countries, especially China and the West, as well as political steps such as plans for the [International Panel on Artificial Intelligence](#) to assess the impact of the technology, including deeply divisive ethical considerations.

Thus, the expressions of deep learning, machine learning, smart algorithms, augmented intelligence, cognitive computing, and of course, artificial intelligence have all become part of our lives just as kale emerged as a necessary ingredient to a healthy diet or reusable mugs as means to reduce disposable coffee cups and the amount of trash. They definitely started to reshape the world around us.

"A.I. technologies will be the most disruptive class of technologies over the next 10 years due to radical computational power, near-endless amounts of data and unprecedented advances in deep neural networks," said Mike J. Walker, research director at Gartner, a leading research and advisory company providing information technology-related insight and market trend analysis already back in 2017. Since then, the abilities of self-driving cars, Siri, Cortana, [chatbots helping people appeal their parking tickets](#), [deep neural networks detecting sexual orientation from facial images](#), and recently the news about how deepfake videos – the synthesis of real footage and image sequences generated by artificial intelligence – [and audio recordings already as the source of crimes](#) have been all over the World Wide Web.

However, while A.I. definitely has the potential to fundamentally transform entire industries from transportation to healthcare, there has been too much hype around the capabilities of what A.I. can achieve today; and the term itself got diluted to the point where everything from big data analytics to pre-programmed robotic refrigerators could come under the term: artificial intelligence. Journalists overexcited about technology and click-bait news articles did not help either in offering clarification. The [story how Facebook shut down an A.I. experiment because chatbots developed their own language](#) after it was misrepresented by many news sites from [India](#) to [Hong Kong](#), is a brilliant example how aggravating fears about A.I. becoming conscious and aiming for destroying the human race spread around. And that's just one example out of a swarm of similar articles.



Fears and expectations about A.I.

A couple of years ago, when artificial intelligence appeared in the public narrative, it was often treated by two extremes: either as the source of ultimate evil aiming for the destruction of mankind or the means to solve every trouble on Earth. 'A.I. will cut out the possibility of making errors from a lot of industries'. 'Self-driving cars will liberate plenty of time for doing other activities than driving'. 'Intelligent robots will do monotonous work humans dislike such as administrative tasks, going underwater, and hacking in an underground mine'. These benefits sound amazing, but we have to acknowledge, A.I. in its current form is far from this level.

In addition, sentiments about A.I. sometimes swing in favor of the negative extreme nourished by public figures as well as dystopian sci-fi stories. Stephen Hawking said more than once that [the development of full artificial intelligence could spell the end of the human race](#). Elon Musk [told Bloomberg's Ashlee Vance](#), the author of the biography *Elon Musk*, that he was afraid that his friend Larry Page, a co-founder of Google and now the C.E.O. of its parent company, Alphabet, could have perfectly good intentions but still "produce something evil by accident"—including, possibly, "a fleet of artificial intelligence-enhanced robots capable of destroying mankind." In September 2017, the SpaceX and Tesla CEO even said that competition for A.I. superiority among countries such as China or Russia could lead to World War III. He explained that he was not just concerned about the prospect of a world leader starting the war, but also of [an overcautious A.I.](#) deciding "that a [pre-emptive] strike is [the] most probable path to victory". The possibility of developing autonomous, robotic, lethal weapons does not scare only Musk, however. In August, more than 100 robotics and A.I. pioneers, including Tesla CEO and Alphabet's Mustafa Suleyman, were [calling for a UN-led ban of lethal autonomous weapons](#).

Moreover, artistic interpretation tends to favor human-machine interaction as damaging: [the creation is pitted against its creators, aspiring ultimately to supplant them](#). Science fiction is full of robots-usurping-humans stories, sometimes entwined with a second strand of

anxiety: seduction. Machines are either out to eliminate us (Skynet from Terminator 2, Hal in 2001: A Space Odyssey), or to hoodwink us into a state of surrender (the simulated world of The Matrix, the pampered couch potatoes of WALL-E). On occasion, they do both.

However, the trend of representing the relationship between humans and intelligent machines got, perhaps, more sophisticated, and a pinch more positive lately: Interstellar presented how to work together with an A.I. system efficiently, Her showed what happens when humans fall in love with artificial intelligence, and Mother demonstrated what it might look like when an intelligent machine raises a human child.

Not only works of art, but articles about A.I. could also be characterized with a splash of positivity lately. The above-cited [A.I. Index report](#) shows that since early 2016, articles about the technology have become more positive, when articles went from 12 percent positive in January 2016 to 30 percent positive in July 2016, and the percentage of positive articles has hovered near 30 percent since then. The reasons might be on the one hand the recognition of the huge potential in A.I. to bring positive change in many industries, and on the other hand, the move from the abstract idea of A.I. deeper down the rabbit hole might have resulted in a more positive stance towards the technology - as with machine learning, deep learning, and co. people see how far we are from conscious A.I., but also how they can solve specific tasks much better than humans.


Let the quest for balanced views on A.I. begin

Thus, we believe that Elon Musk's fears about the development of A.I. being „[summoning the devil](#)“ are way too exaggerated. However, while the technology has much potential for transforming society, the latest positive attitudes might also be inflated.

As the [Gartner Hype Cycle](#), the most reliable indicator for hype around new technologies, has shown, artificial intelligence, machine learning, and deep learning have stationed for the last 3-4 years at the "[peak of inflated expectations](#)". So, next up should be a necessary recalibration of the realm - one that will separate the true meaning of A.I. from the hypes and fears, as well as the winning A.I.-driven companies from all the remaining noise.




HYPE CYCLE FOR EMERGING TECHNOLOGIES, 2018



We need balanced analyses about the nature of A.I., as well as its risks, its capability and its potential consequences so that we could prepare for the sweeping changes it brings for every industry, especially healthcare. Principally because public attitudes towards A.I. vary greatly depending on its application – and healthcare does not feature very high on the list of accepted uses. According to a [YouGov survey for the British Science Association](#) of more than 2,000 people, fully 70 percent of respondents are happy for intelligent machines to carry out jobs such as crop monitoring – but this falls to 49 percent once you start asking about household tasks, and to a miserly 23 percent when talking about medical operations in hospitals.

And why exactly should we set the terms of the human-machine interaction as a zero-sum game? What if we thought about A.I. rather as an innovative cooperative form between two players? In 1997, [IBM's supercomputer Deep Blue could beat Garry Kasparov](#), the reigning chess grandmaster. He said he could have performed better if he had had access to the same databases as Deep Blue. So later, freestyle matches were organized in which supercomputers could play against human chess players assisted by A.I. (they were called human/A.I. centaurs). As a result, in 2014 in a [Freestyle Battle](#), the A.I. chess players won 42 games, but centaurs won 53 games. The best potential pair is a human with technology. This is the only balance that can lead to a positive future with more and more disruptive innovations including ever-improving cognitive computing but also ever-improving human intelligence and wisdom. However, for arriving at a successful cooperation between these two players, we need to understand what we are dealing with. So, let the quest begin!



What is Artificial Intelligence?

Artificial intelligence or A.I., is a broad term „conjuring up“ everything from being the latest tool in salesforce through the conscious mind behind Scarlett Johansson's voice in the movie „Her“ until the program that trounces the best human players at Jeopardy! When you dissect it, you will find artificial narrow, general and super-intelligence, deep and machine learning, supercomputers and GPUs (graphic processing units - specialized electronic circuits designed to rapidly manipulate and alter memory to accelerate the creation of images) as well as the Turing test behind it.

The term was mentioned for the first time in 1956 by John McCarthy during the [Dartmouth Summer Research Project on Artificial Intelligence](#) (DSRPAI) where several scientists decided to meet to look at whether machines could reach intelligence. According to [the definition used at Stanford University](#), it is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but A.I. does not have to confine itself to methods that are biologically observable. In this context, intelligence means the computational part of the ability to achieve goals in the world. Varying kinds and degrees of intelligence occur in people, many animals and some machines.

A narrower definition signalling the most common understanding of A.I. [was coined by the European Parliament's Citizens' Rights and Constitutional Affairs Policy Department](#). According to their term, it is the capability of a computer program to perform tasks or reasoning processes that we usually associate to intelligence in a human being. It often has

to do with the ability to make a good decision even when there is uncertainty or vagueness, or too much information to handle. When we define A.I. as the latter, where intelligence means basically human-level intelligence, then we don't have artificial intelligence today yet.

There is also the question of how to know if we have reached the stage of human-intelligence level A.I. That's where the Turing test comes into play. [In 1950, British mathematician Alan Turing wrote a paper](#), titled "[Computing Machinery and Intelligence](#)", in which he examined the question, "Can machines think?" Turing's answer was yes, and he predicted that, by as early as the beginning of the 21st century, machines would be able to pass an intelligence test that he proposed, now famously known as the Turing test. A human interrogator is tasked with determining which of two chat-room participants is a computer, and which is a real human. The interrogator can say or ask anything, but all interaction is solely through typed text. If the interrogator cannot distinguish computer from human with better than 70 percent accuracy after 5 minutes of interacting with each participant, then the computer has passed the test. To date, there is [no computer which was able to pass](#). There is no consensus when we might reach that stage. Some people suggest that it'll happen somewhere [around 2029](#). Others think that this might happen somewhere closer to 2040. However, most people agree that it'll happen in our lifetime.

Narrow, general, or super?

The literature and experts break down A.I. into various types and stages according to its current capacity as well as its potential performance in the near or distant future. If we assume that the broader goal is to achieve a level of intelligence in machines similar to that of humans, first we have to see the howling difference between the way computers and humans work. When people see an object like a chair, they are able to recognize its distinctive features and apply it to exactly the same and many other different chairs as well. As Dave Gershgorn described it in Quartz [here](#), machines are very literal - a computer doesn't have a flexible concept of "similar." A goal of artificial intelligence is to make machines less literal, thus recognizing a swarm of chairs, the concept of a chair or the same chair photographed in different light.

This is called "[generalizing](#)" or forming an idea that's based on similarities in data, rather than just the images or text the A.I. has seen. "The goal is to reduce a complex human behaviour to a form that can be treated computationally," says Alex Rudnicky, a computer science professor

at Carnegie Mellon University. "This in turn allows us to build systems that can undertake complex activities that are useful to people."

To avoid confusion with the more general expression, [experts prefer to use the term Artificial General Intelligence \(AGI\)](#) to refer to that level of intelligence when a machine is capable of abstracting concepts from limited experience and transferring knowledge between domains. Thus, being able to pair the concept of a chair not only with pictures but also text, etc. AGI is also referred to as "Strong A.I." to differentiate against "Weak A.I." or "Narrow A.I.", which are systems designed for a specific task whose capabilities are not easily transferable to others. On the science-fiction sounding end of the scale is "Artificial Superintelligence", which refers to a form of intelligence which is more powerful than the human brain. It is very [tricky as scientists and experts have not even mapped what the human brain is capable of](#), thus we are very far from it in reality.

However, there is the field of Artificial Narrow Intelligence (ANI), which is developing at an incredible speed. These [narrowly intelligent programs defeat humans in specific tasks](#), such as IBM's supercomputer Deep Blue winning at chess but unlike human world champions, these algorithms are not capable of also driving cars or creating art. Solving those other tasks requires other narrow programs to be built, and it is an immense challenge for computers to connect the dots. Yet, there is incredible growth in computers' ability to understand images and video, a field called computer vision, as well as text in the frames of natural language processing. The former is used as a primary technology for self-driving cars, Google image search, automatic photo-tagging on Facebook, and it is extensively utilized now in healthcare, for example in the field of medical imaging.



What do you need for developing A.I.?

The theoretical background and ideas about A.I. have been around for 50-60 years already, but there were two decisive factors why A.I. research started to boom in the last decade. The first element is computational power, the second is the vast amount of data needed to be accumulated.

In the mid-2000s graphics processor unit (GPU) company [Nvidia](#) concluded that their chips were well-suited for running neural networks, and began making it easier to run A.I. on their hardware. Researchers [found that being able to work with faster and more complex neural networks](#) led to more improvement in accuracy. From that time on, there has been a fierce competition in the tech industry to build faster and stronger computers able to process and store more and more data. [IBM, Hewlett-Packard, Cray or Dell are all developing high performance computers, so-called supercomputers](#), which are able to process unimaginable amounts of information.

Thus, we arrived at the other key element for building artificial intelligence: data. Algorithms need to be fed with enormous amounts of data to be able to arrive at ANI. Currently, the amount of available digital data is growing at a mind-blowing speed, doubling every two years. In 2013, it encompassed 4.4 zettabytes, however by 2020 the digital universe – the data we create and copy annually – [will reach 44 zettabytes](#), or 44 trillion gigabytes (!). This vast amount of data not only enables, but already necessitates the presence of algorithms being able to make sense of big data.

Data Analytics, Machine Learning & Deep Learning – Methods of Teaching Algorithms

The methods for making sense of huge data sets and reaching ANI require IT and programming background. Here, we do not intend to go into the technical details, we only would like to familiarize interested readers with the broad concepts to be able to differentiate among them and select the most appropriate ones for their goals. Although more complex data analysing methods sound exciting and appealing, sometimes you can arrive at great results by using less advanced techniques. For example, [in a small Hungarian hospital, the pre-treatment waiting time for oncology patients dropped drastically from 54 to 21 days](#) only by optimizing patient management processes with the help of simple 'tricks' such as recording and following-up cases closely. The first step for data analytics is recording data appropriately – and then carrying out the necessary follow-up actions.

The second step is using various statistical methods, such as data mining [for collecting, analyzing, describing, visualizing and drawing inferences from data](#), for example from electronic health records or [the 'OMICS' universe](#). The focus is on discovering mathematical relationships and properties within datasets and quantifying uncertainty. Data mining usually means when [insights and patterns are extracted from large-scale databases](#).

However, this is only the ante-room of artificial intelligence. The next step should be the creation of symbolic systems where the most frequently used programs are so-called expert systems. The method was introduced by researchers at Stanford University, and represented the main line of artificial intelligence research from the mid-1950s until the late 1980s. Expert systems are comprised of a [knowledge base, an inference engine, which uses IF-THEN-ELSE rules to make sense of the dataset](#) and a user interface, where a user receives the required information. Such systems [are best applied to automate calculations and logical processes](#) where rules and outcomes are relatively clear.

Here, the program is only able to carry out the set of specific rules which a computer engineer programmed it to do. Machine learning goes far beyond that. It is the field of computer science that [enables computers to learn without being explicitly programmed](#) and builds on top of computational statistics and data mining. As with conventional statistics, [machine learning requires sufficient training datasets](#) (also known as sample size in traditional statistics) and the right algorithms to optimize its performance on the training dataset before testing. However, in contrast to the traditional methods, it is focused on building automated decision systems.

Machine learning has different types: it could be [supervised, unsupervised, semi-supervised or reinforcement learning](#).

The first one, supervised machine learning, is typically used for classification problems, e.g. pairing pictures with their labels. Thus, you have an input and output data – the image as well as the label –; and the algorithm learns general rules how to categorize. It is the most popular type of machine learning in medicine, and in a few years, it will be widely used in medical imaging in [radiology, pathology](#), and other image-intensive fields. However, it certainly has its limitations: it requires large data sets to become accurate enough, and the data has to be appropriately labelled. That's why the [work of data annotators is so crucial](#).

Nevertheless, supervised machine learning can also be effectively deployed to predict health events based on various input data, which can be classified in a linear way. For example, [from statistics on measles vaccination rates and disease outbreaks from the Centers for Disease Control and Prevention](#), as well as non-traditional health data, including social media and syndromic surveillance data generated by software that mines a huge range of medical records sources, an algorithm can provide a reliable map of future measles outbreak hotspots.

In the case of unsupervised machine learning, the computer program is asked to discover inherent structure and patterns that lie within the data. Unlike in the case of supervised machine learning, it has to make up its own groups, usually called 'clusters', and categories based on certain similarities in huge data sets. It is usually used [to predict unknown results and to determine how to discover hidden patterns](#). Unsupervised machine learning has subtypes: clustering algorithms and association rule-learning algorithms. It is widely used and has been implemented, for example, in self-driving vehicles and robots as well as being used in speech- and pattern-recognition applications. In medicine, for example, [tissues samples can be clustered based on similar gene expression values](#) using unsupervised learning techniques. As an example of association rule-learning algorithms, the [testing of novel drug-drug interactions](#) can be mentioned.

Semi-supervised learning lies between supervised and unsupervised learning, where the input-output pairs are incomplete – the labels are missing or the information is noisy. Its use is especially justified in healthcare, where a large body of data, especially looking at electronic medical records, lack information necessary to apply supervised machine learning. Thus, researchers are [looking at ways](#) to address the problem of making use of unlabeled or un-scored data, together with only a few supervised data, to improve the performance of analysis model for healthcare decision making.

The last category, reinforcement learning constitutes probably the most known type of machine learning: when the computer program learns from its mistakes and successes; and builds its experiences into its algorithm. The most famous example for it is [AlphaGo, the machine developed by Google that decisively beat the World Go Champion Lee Sedol in March 2016](#). Using a reward and penalty scheme, the model trained on millions of board positions in the supervised learning stage first, then played itself in the reinforcement learning stage to ultimately become good enough to triumph over the best human player.

However, the problem with applying reinforcement learning to healthcare, especially for optimizing treatment, is that [unlike with AlphaGo, we cannot play out a large number of scenarios where the agent makes interventions to learn the optimal policy](#) – as the lives of patients are at stake. Luckily, we already have examples where this issue can be mitigated. In a [study published by MIT researchers](#), the authors reported a successful formulation of clinical trial dosing as a reinforcement learning problem, where the algorithm taught itself the appropriate dosing regimens to reduce mean tumor diameters in patients undergoing chemo- and radiation therapy clinical trials.

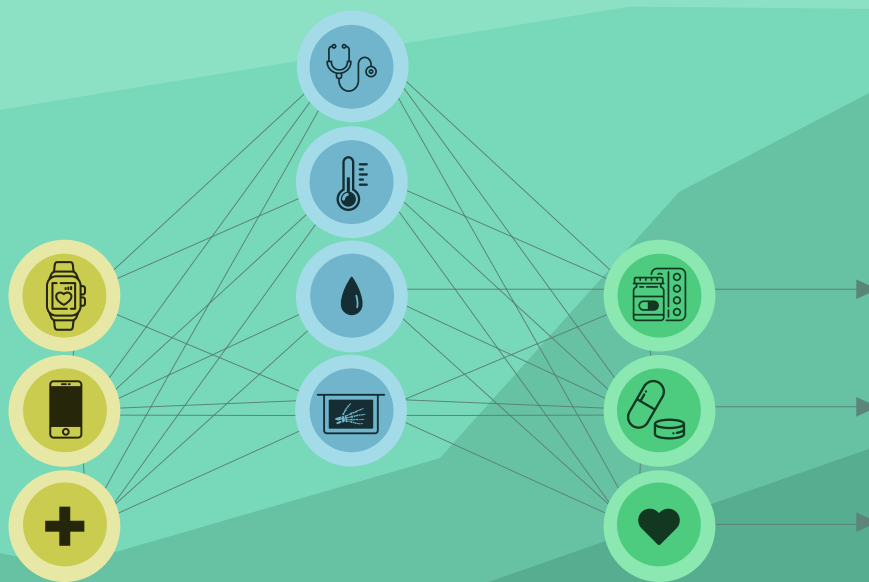
Deep learning is a subfield of machine learning where computers learn with the help of layered neural networks. What are neural networks? [Quartz formulated the explanation for them as the followings](#): algorithms which are roughly built to model the way the brain processes information, through webs of connected mathematical equations. Data given to a neural network is broken into smaller pieces and analyzed for underlying patterns thousands to millions of times depending on the complexity of the network. A deep neural network is when the output of one neural network is fed into the input of another one, chaining them together as layers. Typically, the layers of a deep neural network would analyze data on higher and higher levels of abstraction, meaning they each throw out data learned to be unnecessary until the simplest and most accurate representation of the data is left.

Deep learning has [different types used for reaching ANI in various areas](#). Convolutional neural networks (CNN) are typical for recognizing images, video, and audio data, due to their abilities to work with dense data. Recurrent neural networks (RNN) are used for natural language processing, while long short-term memory networks (LSTM) are variations of RNNs meant to retain structured information based on data. For instance, an RNN could recognize all the nouns and adjectives in a sentence and determine if they're used correctly, but an LSTM could remember the plot of a book.

Deep learning requires plenty of expertise, thus only a few technology companies managed to apply it to already marketed products. Google, Microsoft, IBM, Amazon, Baidu, or Tencent are the biggest developers of deep learning algorithms. Just think about Google's translation service, Google Translate, and you will be aware of the fact that you have [already used one of the most popular products powered by deep learning](#). As an example of deep learning in medicine, [researchers proposed an approach to deduce treatment policies for septic patients](#) by using continuous state-space models and deep reinforcement learning. In [another study](#),

experts attempt to solve the difficult problem of estimating polyp size using colonoscopy images or videos, which is crucial for making a diagnosis in colon cancer screening. Moreover, [unsupervised deep learning may facilitate the exploration of novel factors in score systems or highlights hidden risk factors](#) to existing models. It can also be used to classify novel genotypes and phenotypes from pulmonary hypertension, cardiomyopathy, and many other factors.

There is one more expression that is worth getting familiar with as it will be heard a lot more in the future: the concept of an open A.I. ecosystem. It was named as [one of the top 10 emerging technologies in 2016 by the World Economic Forum. An open A.I. ecosystem refers to the idea that](#) with an unprecedented amount of data available, combined with advances in natural language processing and social awareness algorithms, applications of A.I. will become connected. Devices, algorithms and humans will constitute one “living ecosystem” in which data will flow without barriers.



Data in healthcare

After getting to know the most important terms, definitions and methods used in the field of artificial intelligence, we need to be familiarized with the current situation of data and data analytics in healthcare in order to be able to find out where artificial intelligence can

come into play in healthcare; why and how we could make use of machine learning and smart algorithms – as the most important 'base material' for smart algorithms is data.

In healthcare, data can be related to health conditions, reproductive outcomes, quality of life, and many similar life events determining for an individual or an entire population. Health data includes clinical metrics along with environmental, socioeconomic, and behavioral information pertinent to health and wellness. A plurality of health data is collected and used when individuals interact with healthcare systems, and one of the most determining of these – and thus most widely used in A.I. research – is the electronic medical record.

A brief history and the current state of electronic medical records

A doctor from the beginning of the 20th century would truly be shocked by the changes in the medical system – only by looking at how current hospitals and healthcare systems treat and store medical records. Back then, there was the good old-fashioned pen and paper, where doctors usually noted the symptoms, the diagnosis of patients and the treatment they provided. So, how did it change throughout the decades? Let's look at the example of the United States, where one of the most extensive documentation for medical history is available.

The real history of electronic medical records [dates back to the late 1960s with problem-oriented medical records](#) – that is when third party facilities were able to verify the diagnosis independently. This development made it possible to transfer patients from one hospital to another for specific treatment, and determined also a general standardization for the whole industry.

The first electronic medical system in the US was developed in 1972, but it was only in the late 1980s, with the appearance of personal computers, that EMR was widely adopted by large hospitals and medical facilities. However, the technology did not spread around with smaller care providers. Not only medical records, but also general administrative paperwork such as billing or scheduling started to move to PCs. Later, with the appearance of the Internet, health information was accessible more easily than ever before – setting the stage for web-based EMRs. Already in 1991, the Institute of Medicine, [founded in 1970 to address the concerns of medicine and healthcare in the US](#), recommended that by the year 2000, every physician should be using computers to improve patient care.

In spite of the fast computerization and the appearance of internet-based and cloud-based services, the use of electronic medical records is not as wide-spread as average healthcare service users would assume. Not even in the US. [According to the Centers for Disease Control and Prevention](#), in 2013, 78 percent of office-based physicians used any type of electronic medical record system, up from 18 percent in 2001. And even healthcare facilities, which adopted an EMR system, report problems with its general functioning. [A survey even found that 37 percent of American physicians see EHR \(electronic health records\) as their number one challenge](#), while the same percentage list a financial issue as their primary concern. Another [research also showed that](#) even with all the advancements in EMR technology, 70 percent of physicians are unhappy with their current system. Specialists especially have difficulties adapting to their EMR software due to their different needs and areas of focus.

Thus, although patient medical records are more accessible than ever before in theory, meaning that data technology is on the verge of becoming fully portable and comprehensive, we need more efforts in advancing and streamlining EMR systems for both being able to handle ever-increasing medical data and being fit for the use of future technologies such as artificial intelligence. In addition, we only outlined here the situation in the United States briefly, but globally there are plenty of countries where medical records are still on paper and the process of making efficient EMR systems is even slower and more problematic.



Why do we need help from A.I. when it comes to data?

In spite of the fact that the good old-fashioned pen and paper seems to stick around for some more years in many countries, the accumulation of vast amounts of healthcare data

is unstoppable. Electronic medical records, clinical data, doctor's notes, lab results, medical images – or lately, data from health sensors and wearables are multiplying at an incredible speed. [According to the OECD](#), there has been tremendous growth in the range of information being collected, including clinical, genetic, behavioural, environmental, financial, and operational data.

If we only look at the market of wearables, in 2019 as many as 245 million wearable devices will have been sold, according to [CCS Insight's wearable tech report](#), Wearables Forecast Worldwide, 2015 - 2019. For wearables, this number is huge. In Q1 of 2019, the biggest sellers like Apple, Huawei, or Fitbit sold around 12.8 million, 5.0 million, or 2.9 million devices respectively. And that's just the first three months of the year! Imagine the amount of data that will be gathered through that vast amount of fitness trackers, sleep sensors and other wearables used on a regular basis.

And we haven't even spoken about the exponential growth of medical knowledge: medical and pharma research and development data, theoretical knowledge, etc. In 1950, [the time for medical knowledge to double was estimated to be about 50 years](#). In 1980 it was about 7 years, while in 2010, it only took about 3.5 years. It is estimated that by the year 2020, it will only take 73 days for the volume of medical knowledge to double. On [MEDLINE](#), between 1978 and 2001, a total of 8.1 million journal articles were published. On [Pubmed](#), there are 23 million papers. That is a lot for physicians, healthcare experts, healthcare lawmakers, and patients to try to read and make sense of.

Without the help of cognitive computing, these huge chunks of data cannot be dealt with. Big data analyzing tactics such as data mining or ANI methods such as machine learning are necessary for effectively carrying out many tasks in healthcare. The next chapter will show you how the application of artificial intelligence has started in various fields of healthcare.



APPLYING ARTIFICIAL INTELLIGENCE IN HEALTHCARE

Looking at the potential in artificial intelligence for information gathering and expertise sharing, analysing huge datasets and stumbling upon correlations or clusters that are invisible to the human eye, we can easily discern what types of tasks smart algorithms can carry out in medicine and healthcare.

1. Regarding information gathering and expertise sharing, A.I. can help medical teams collaborate and design programs that support patient care and treatment success.
2. After A.I. analyzes once invisible data, it can evaluate findings based on the leading medical literature, and update current medical knowledge and evidence-based guidelines, as well as patient care and treatment pathways.
3. Smart algorithms analysing huge patient datasets could help insurance providers manage preventive care for large populations.
4. In the case of pharma, A.I. can improve the quality, the time management and reach of clinical trials.
5. Regarding administration, smart algorithms can streamline administrative processes in the hospital, help in patient management and support the management of challenging electronic medical records systems.

Parallel to the increase of companies creating smart algorithms for medical purposes, policy-makers and regulators in healthcare also recognized the immense potential in A.I. to optimize the processes in medicine for healing patients in a more efficient way, as well as in pharma or in administration in general. The rise in the number of FDA-approved algorithms shows this change in the attitude of regulators. In 2014, only [AliveCor's algorithm for the detection of atrial fibrillation](#) was approved. Two years later, the FDA found four further solutions ready for clinical use, while in 2017, six new algorithms were approved by the US regulator. This exponential growth just accelerated last year, when the FDA endorsed 23 algorithms in medicine. As the first approvals in 2019 also show, we do not expect the trend to slow down. On the contrary, we will most likely see dozens of new medical A.I. solutions on the market. As a starting point, let's look at already applied smart algorithms, promising results on training datasets and future expectations in various medical fields.

1. Health data management

Not even the most acclaimed professors can match the abilities of cognitive computers. As the amount of information they accumulate grows exponentially, the assistance of computing solutions in medical decisions is imminent. A.I. will open new dimensions for doctors on the personal level as well as for hospitals and other medical institutions on a more structural level.

On the institutional level, the most obvious use of A.I. will be data management. Collecting, storing, normalizing, tracing its lineage – it is the first step in revolutionizing the existing healthcare systems. For example, the A.I. research branch of the search giant, Google, launched its [Google DeepMind Health](#) project, which is used to mine the data of medical records in order to provide better and faster health services. The project is in its initial phase, and at present, they are cooperating with Moorfields Eye Hospital NHS Foundation Trust to improve eye treatment. In June 2017, [DeepMind expanded its services – first of all, its data management app, Streams, to another UK hospital](#). This expansion came despite [controversy over the company's first NHS data-sharing agreement](#).

Data analytics can also streamline administrative processes in order to arrive at better care. 97 percent of healthcare invoices in the Netherlands are digital, containing data regarding the treatment, the doctor, and the hospital. These invoices could be retrieved easily. A local company, [Zorgprisma Publiek](#), analyses invoices and uses IBM Watson to mine the data. They can tell if a doctor, clinic or hospital makes mistakes repetitively in treating a certain type of condition in order to help them improve and avoid unnecessary hospitalizations of patients.



Dermatologists also recognized the huge potential of big data to bring lasting change to their specialty. The [American Academy of Dermatology introduced a clinical registry called DataDerm](#) in 2016. The database was created by dermatologists and it connects data on millions of patients from thousands of dermatologists throughout the US. It eases the pain of reporting and allows medical professionals to demonstrate the quality of care they provide to payers, policy makers, and the medical community. At the same time, it gives every member a private analysis of his or her practice's data against national averages – down to the patient level. It is great for setting standards in dermatology, measuring how each participant performs and ensuring the average quality of care.

2. Treatment pathway design


IBM Watson launched its special program [for oncologists](#) able to provide clinicians evidence-based treatment options. Watson for Oncology has an advanced ability to analyse the meaning and context of structured and unstructured data in clinical notes and reports that may be critical to selecting a treatment pathway. Then, by combining attributes from the patient's file with clinical expertise, external research, and data, the program identifies potential treatment plans for a patient.

IBM launched another algorithm called [Medical Sieve](#). It is an ambitious long-term exploratory project to build the next-generation "cognitive assistant" with analytical, reasoning capabilities and a wide range of clinical knowledge. Medical Sieve is qualified to assist in clinical decision making in radiology and cardiology. The "cognitive health assistant" is able to analyze radiology images to spot and detect problems faster and more reliably.

[Microsoft's research machine-learning project, dubbed Hanover](#), also aims to ingest all the papers and help predict which drugs and which combinations are the most effective for cancer cases. Moreover, [a series of start-ups](#), such as the Hungarian [Turbine](#) or [Cambridge Cancer Genomics](#) aim to build smart algorithms to make oncology better. They are building artificial intelligence solutions to design personalized treatments for any cancer type or patient faster than any traditional healthcare service.

[Robots enabled with artificial intelligence are increasingly assisting microsurgical procedures](#), too, to help reduce surgeon variations that could affect patient recovery. [Catherine Mohr](#), vice president of strategy at Intuitive Surgical and expert in the field of surgical robotics,

believes surgery will also take to the next level [with the combination of surgical robotics and artificial intelligence](#). She envisioned a tight partnership between humans and machines, with one making up for the weaknesses of the other. In fall of 2017, [Maastricht University Medical Center](#) in the Netherlands already used an A.I.-assisted surgery robot to suture small blood vessels – some no larger than .03 millimeters – and up to .08 millimeters across. In addition, A.I. is being used to provide analysis of a surgeon's technical abilities with products like [Caresyntax's qvident](#), a Web-based surgical risk and quality management tool, designed to eliminate the “black box” of surgical visibility.



3. Transforming diagnostics

One of the most advanced fields of ANI, computer vision, will have a huge impact on diagnostics through revolutionizing medical imaging. Smart algorithms will be able to analyse MRI, CT scans, X-rays and any other medical images in the future, moreover they could detect signs in the recordings which are not accessible for the human eye. Radiology, dermatology, but also ophthalmology, and other specialties could leverage on the vast potential of ANI.

Scientists at the University of Adelaide have been experimenting with an [A.I. system that is said to be able to tell if you are going to die](#). By analyzing CT scans from 48 patients, deep

learning algorithms could predict whether they'd die within five years with 69 percent accuracy. It is "broadly similar" to scores from human diagnosticians, the paper says. It is an impressive achievement. The deep learning system was trained to analyze over 16,000 image features that could indicate signs of disease in those organs. Researchers say that their goal is for the algorithm to measure overall health rather than spot a single disease.

But this is just the tip of the iceberg. There is also plenty of ongoing research to teach algorithms to detect various diseases. IBM's flagship A.I. analytics platform, Watson, is also utilized in the field of medical imaging. After the company purchased Merge Health in 2015, [Watson got access to millions of radiology studies and a vast amount of existing medical record data](#) to help train the A.I. in evaluating patient data and get better at reading imaging exams.


IBM decided to let also [dermatologists leverage on the results of Watson](#) in order to diagnose melanoma and other types of skin cancer faster, more accurately and preferably without the need for many biopsies. At the IBM T.J. Watson Research Center, experts found that their deep learning system was able to achieve a 76 percent accuracy at diagnosing melanoma cases based on dermatology images, while the average accuracy for the eight dermatologists on that data set was 70.5 percent.

Not only IBM but also giants like Philips, Agfa and Siemens have already started integrating A.I. into their medical imaging software systems. GE is developing predictive analytics software using elements of A.I. assessing the impact of sick leaves or the increase of patient volumes on imaging departments when someone calls in sick, or if patient volumes increase. Vital also has similar work-in-progress predictive analytics software for imaging equipment utilization. Not to speak [about the dozens of bigger and smaller start-ups trying to utilize the power of A.I.](#)

Yet, it is not just radiology. Researchers at Stanford University created [an artificially intelligent diagnosis algorithm for skin cancer](#) with the help of an algorithm developed by Google that was already trained to identify 1.28 million images from 1,000 object categories. Then, they made a database of nearly 130,000 skin disease images representing over 2,000 different diseases; and trained their algorithm to visually diagnose potential cancer. From the very first test, it performed with inspiring accuracy. It performed at least as well as dermatologists participating in the research, which is very impressive! Now, the team is considering to make the algorithm smartphone-compatible in the near future, bringing reliable skin cancer diagnoses to our fingertips.

In 2016, [Google developed an eye-scanning technique for looking at retinal images and detecting diabetic retinopathy](#) as well as a trained ophthalmologist. The disease is quite common among diabetes patients, and if it is not spotted early enough, it may cause blindness. The machine-learning algorithm uses Google's method for labelling millions of Web images as it examines photos of a patient's retina to spot tiny aneurysms indicating the early stages of diabetic retinopathy. A year later, [the search giant announced they had begun working on integrating the technology into a chain of eye hospitals in India.](#)

Google is not the only one working on A.I. solutions for eye care, though. A teenage girl from India, whose grandfather was diagnosed with diabetic retinopathy, [developed a smartphone app that can screen for the disease](#) with the help of a specially trained artificial intelligence program and a simple 3D-printed lens attachment. A truly disruptive innovation: smart, cheap and potentially life-changing!



4. Health assistance and administration

Dealing with patients requires plenty of administration, organization and paperwork. Looking at primary care, doctors and nurses often meet patients with minor issues that could be treated without the intervention from a medical professional, people who only want prescriptions or have organisational questions. Artificial narrow intelligence in the form of intelligent personal

assistants, a version of Siri or Alexa for healthcare, could definitely help out medical staff in this area. For example, these digital assistants with natural language processing programs converting voice to text could listen in patient-doctor visits and the 'conversations' between doctors and EHR systems, and provide a transcription without the doctor typing even one letter into his computer.

For example, in the first case, San Francisco-based [Augmedix](#) aims to [harness the power of Google Glass to make healthcare more patient-centric](#) and decrease the amount of paperwork. It provides a technology-enabled documentation service for doctors and health systems, so physicians do not have to check their computers during patient visits, while medical notes are still generated in real-time. [Voice assistants truly have the potential to free up the time doctors](#) spend on administration. Companies like [Nuance and M*Modal already provide software-based dictation services](#) to physicians. California-based company, [Notable](#), [launched](#) a wearable voice-powered assistant in May 2018 aimed at helping doctors capture data during interactions with patients.

Voice to text technologies mean a real alternative to medical administration done manually by doctors – and institutions agree. [Research firm Technavio published a report last year](#) projecting that hospitals will spend more than 72 billion by 2020 globally, representing 6 percent compound annual growth rate. However, recent voice recognition solutions do not eliminate transcription errors fully, thus the need for proofreading and human check-up will still take up the precious time of medical professionals. Even one letter recognized differently could mean a potentially life-threatening danger for patients. That's why we believe that an automated, artificial narrow intelligence-based voice recognition system with a similar, artificial narrow intelligence-based proofreading system might embody the final solution for medical administration. Only the problems flagged by the ANI system would be checked by a doctor. At least as the trends show, currently that's the objective that dozens of companies are working towards.



5. Patient management

Innovative technologies in patient management aim to enable patients to take their disease management in their own hands and take some burden off the shoulders of physicians.

In diabetes care, [Doug Kanter collected data about himself for a full year](#) – blood sugar readings, insulin doses, meals, sporting activities etc. His company, [Databetes](#), was born out of his own experiences with diabetes. It helps patients better manage their condition by providing a good way for logging and measuring data, as well as a revolutionary concept to analyze the big data behind one person's disease.

Supporting the visually impaired, [Horus](#), [OrCam](#), [BeMyEyes](#), and [Aira](#) all offer their solutions in order to open the possibility to live a more independent life. They are using various algorithms for describing the environment to the user, read out text, recognize faces and objects such as supermarket products or bank notes, or notify about obstacles. These algorithms are all able to learn over time. And that's just the beginning.


The company, [Kore.ai](#) offers smart bots for healthcare facilities. The digital assistant can connect patients to the right contacts directly, give appointment details or make any changes. It lets the patients easily refill prescriptions or pay bills. It delivers lab, test or procedure outcomes or recommended next steps. [Safedrugbot](#) embodies a chat messaging service that offers assistant-like support to health professionals, e.g. doctors who need appropriate information about the use of drugs during breastfeeding. [Izzy](#) helps women track their period and serves as a birth control pill reminder.

Bots like [HealthTap](#), [Your.Md](#), or [Ada Health](#) aim to help patients find a solution to the most common symptoms through A.I. However, a chatbot never replaces an experienced doctor. The bot itself exhorts the user to book an appointment with a doctor for a diagnosis, and eventually for the prescription of a therapy. The [AiCure app](#), supported by The National Institutes of Health, uses a smartphone's webcam and A.I. to autonomously confirm that patients are adhering to their prescriptions, or with better terms, supporting them to make sure they know how to manage their condition. This is very useful for people with serious medical conditions, for patients who tend to go against the doctor's advice and participants in clinical trials.

[Florence](#) might be a practical chatbot for older patients. It is basically a “personal nurse,” and “she” can remind patients to take their pills. The world’s first virtual nurse, Molly, was developed by the medical start-up [Sense.ly](#). It has a smiling, amiable face coupled with a pleasant voice and its exclusive goal is to help people monitor their condition and treatment. The interface uses machine learning to support patients with chronic conditions in-between doctor’s visits. It provides proven, customized monitoring and follow-up care, with a strong focus on chronic diseases.

Some institutions already recognize the potential in A.I.-based chatbots for patients and their services alike. Britain’s National Health Service (NHS) was reportedly scheduled to start to use a [chatbot app](#) for dispensing medical advice for a trial period in 2017, with the aim of reducing the burden on its “111” non-emergency helplines. The NHS developed the app with [Babylon Health](#), one of the new breed of paid, [doctor-on-demand services](#). It has already [launched an application](#) which offers medical A.I. consultation based on personal medical history and common medical knowledge. [Users report the symptoms of their illness to the app](#), which checks them against a database of diseases using speech recognition. After taking into account the patient’s history and circumstances, Babylon offers an appropriate course of action. The app will also remind patients to take their medication, and will follow them up to find out how they’re feeling. Through such solutions, the efficiency of diagnosing patients can increase by multiple times, while the waiting time in front of doctors’ examining rooms could drop significantly.

Unfortunately, the first experiences in the NHS system were mixed: [Patients participating in the trial indicated](#) that they would play the system to get an appointment with the doctor more quickly, so further pilots in North West London were dropped. This phenomenon indicates that patients do not yet have trust in health chatbots and need reassurance from a medical professional for their care.



6. Precision medicine

As the [National Institutes of Health \(NIH\) said](#), there is "an emerging approach for disease treatment and prevention that takes into account individual variability in genes, environment and lifestyle for each person." This approach will allow doctors and researchers to predict more accurately which treatment and prevention strategies for a particular disease will work in which groups of people.

Artificial intelligence will have a huge impact on precision medicine, including genetics and genomics as well. For example, [Deep Genomics](#) aims at identifying patterns in huge data sets of genetic information and medic

al records, looking for mutations and linkages to disease. They are inventing a new generation of computational technologies that can tell doctors what will happen within a cell when DNA is altered by genetic variation, whether natural or therapeutic.

At the same time, Craig Venter, one of the fathers of the Human Genome Project, [is working on an algorithm](#) that could design a patient's physical characteristics based on their DNA. With his latest enterprise, [Human Longevity](#), he offers his (mostly affluent) patients complete genome sequencing coupled with a full body scan and very detailed medical check-up. The whole process enables to spot cancer or vascular diseases in their very early stage.

7. Supporting pharma: drug creation and clinical trials

New drugs are approved through human clinical trials: rigorous, year-long procedures starting in animal trials and gradually moving to patients. They typically cost billions of dollars and take many years to complete, sometimes more than a decade. Plus, patients in trials are exposed to side effects that cannot be predicted or expected. And even if the trial is successful, it has to go through a regulatory approval: it may or may not receive the nod of the respective regulatory agency, e.g. the US Food and Drugs Administration (FDA).

Artificial intelligence can change the status quo for the better through lots of ways. It can help companies aggregate and synthesize a lot of information that's needed for clinical trials, thus shortening the drug development process. It can also support the understanding

of the mechanisms of the disease, establish biomarkers, generate data, models, or novel drug candidates, design or redesign drugs, run preclinical experiments, design and run clinical trials, and even analyze the real-world experience. [The number of already existing A.I. companies in drug development](#) reflects the manifold usage of the technology: they are many and increasing day by day.

For example, Insilico Medicine, working with researchers at the University of Toronto, made headlines in 2019 with the announcement that the process of developing a new drug candidate lasts just 46 days with the help of its smart algorithm. At first, it took [21 days for the team to create 30,000 designs for molecules](#) that target a protein linked with fibrosis (tissue scarring). They synthesized six of these molecules in the lab and then tested two in cells; the most promising one was tested in mice. The researchers concluded it was potent against the protein and showed “drug-like” qualities.

Another flagship company of A.I. drug discovery, [Atomwise](#), uses supercomputers that root out therapies from a database of molecular structures. In 2015, Atomwise launched a virtual search for safe, existing medicines that could be redesigned to treat the Ebola virus. They found two drugs predicted by the company's AI technology which may significantly reduce Ebola infectivity. This analysis, which typically would have taken months or years, was completed in less than one day. Imagine how efficient drug creation would become if such clinical trials could be run at the “ground zero” level of healthcare, namely in pharmacies.


For arriving at even more precise treatment solutions, some companies combine genomics and artificial intelligence. For example, [Cambridge Cancer Genomics](#) develops precision oncology solutions to transform the way cancer patients are treated. They believe that the more clinical and genomic data oncologists have, the smarter decisions they can make about drug usage in any given circumstance – and thus, they use machine learning and data analytics to give doctors this power. In another example, Brendan Frey's company, [Deep Genomics is leveraging AI, specifically deep learning to help decode the meaning of the genome](#). So far, the company has used its computational system to [develop a database that provides predictions](#) for more than 300 million genetic variations that could affect the genetic code. For this reason, their findings are used for genome-based therapeutic development, molecular diagnostics, targeting biomarker discovery, and assessing risks for genetic disorders.

Another company, Oxford-based [Row Analytics](#) operating since 2013, specializes in digital health, precision medicine, genomics, and semantic search. It is delivering a range of highly innovative data analytics platforms, for example, precisionlife for drug discovery. This platform combines A.I. methods and data analytics to look at multiple genetic variants in combinations across a range of diseases. As they are able to complete the process in weeks instead of months, even for large disease populations with tens of thousands of patients, this enables the rapid identification of novel drug candidates and potential drugs to repurpose.

Another way to [modernize the drug testing process](#) is applying technologies to the traditional framework, for example through online platforms to seek out participants. Various online services make it possible for more and more patients to participate in the process of drug creation. [TrialReach](#) or [Antidote](#) tries to bridge the gap between patients and researchers who are developing new drugs. If more patients have a chance to participate in trials, they might become more engaged with potential treatments or even be able to access new treatments before they become FDA approved and freely available. [TrialX](#) similarly matches clinical trials to patients according to their gender, age, location, and medical condition. The number of such services is growing to accommodate an increasing demand from patients.

Finally, we arrived at the technology that's currently considered the most 'science-fiction-like': [in silico clinical trials](#). In silico is [the term scientists use to describe the modeling, simulation, and visualization of biological and medical processes in computers](#). The emergence of in silico medicine is a result of the advancement in medical computer science, including artificial intelligence, over the last 20 years.

Modeling maps the elements of a biological system, visualization presents the predictions in a graphic form, and simulation attempts to realistically show how that system evolves over time under given stimuli. When looking at an in silico clinical trial specifically, it means an individualized computer simulation used in the development or regulatory evaluation of a medicinal product, device, or intervention. While completely simulated clinical trials are not feasible with current technology and understanding of biology, their development would be expected to have major benefits over current in vivo clinical trials, and [the FDA is already planning for a future](#) in which more than half of all clinical trial data will come from computer simulations.



When in silico clinical trials will be combined with artificial intelligence, the virtual patient models can provide answers to the usually asked question: "why does this treatment or drug work for this patient but not another one?" A.I.-building methods, namely [machine learning and deep learning have the potential to train a model](#) that will then be able to find patterns and clusters in otherwise unstructured data – and seek out information researchers don't expect to be there. In addition, neural networks can be used to predict adverse events or to anticipate the possible risk of patient dropout and compliance to the treatment.

FDA-approved algorithms in healthcare

While so far we have been looking at various fields in healthcare and pharma, as well as certain structures where artificial intelligence could make a difference, now we arrived at the smart algorithms that not only hold promise but have already proved themselves worthy for clinical use. In the previous chapter, we mainly enumerated examples which hold significant promise for the future, but the time has come for the artificial intelligence solutions that the regulators found worthy for clinical application.


We believe that's one of the most reliable indicators for a medical device, and the only yardstick for credible and accurate medical software – thus for evaluating the actual state of play when it comes to artificial intelligence in healthcare. Although in Europe, the [European Medicine Agency](#) has guidelines and statements about artificial intelligence, the FDA is the only regulator with efficient instruments in its toolkit to assess the credibility and accuracy of algorithms for medical purposes in detail. However, to understand fully what FDA approvals mean for the medical A.I. field, at first we have to understand what an FDA approval actually entails.

There's an entire scale of approvals [starting from 510\(k\) submission through de novo to premarket approval \(PMA\)](#). The first one refers to a premarket submission to demonstrate that a device aiming for market launch but not requiring premarket approval is as safe and effective as other similar instruments with PMA. The latter actually means an FDA process of scientific and regulatory review to evaluate the safety and effectiveness of medical devices supporting and/or sustaining human life and the most stringent of the device marketing applications. The de novo pathway for device marketing rights was added to address novel devices of low to moderate risk that do not have a valid predicate device – for example in the case of software solutions such as smart algorithms. Upon successful review of a de novo submission, the FDA creates a classification for the instrument, a regulation if necessary,

and identifies any special controls required for future premarket submissions of substantially equivalent devices.

As there was no database that contains all the smart algorithms with any type of FDA approvals worth applying to medical processes, we collected the data into one infographic.

FDA APPROVALS FOR ARTIFICIAL INTELLIGENCE-BASED ALGORITHMS IN MEDICINE



Looking at the infographic, [the distribution of smart algorithms in the various medical specialties](#) becomes apparent. Radiology and cardiology seem to be heavily populated by artificial intelligence-based solutions, there are already seven approved algorithms in cardiology, while 16 in radiology. However, geriatrics, orthopedics or pathology seem to be less prone to A.I. Certain medical specialties do not even appear on the list yet, such as pulmonology, dermatology, surgery, OB/Gyn or forensic medicine.

However, we should not draw definitive conclusions from this infographic only, as it constitutes a snapshot of the current state of play, not necessarily revealing anything about the trends. For example, in the case of pathology, while the number of FDA approved algorithms might be low at the moment, artificial intelligence is a promising technology in the field – although it might need the coming years to catch up with the number of solutions in radiology or cardiology.

These two fields represent the peak areas of artificial intelligence research due to several reasons. First and foremost, computer vision is one of the fastest growing fields in A.I. development, and medical imaging has both the data and the visuality that smart algorithms need to thrive. As a consequence, [researchers found that commercial software for automatically classifying breast density](#), and thus detecting breast cancer, can perform on par with human radiologists. What's more, in April 2018, [the FDA approved the first AI system](#) that can be used for medical diagnosis without the input of a human clinician.

The current list of FDA-approved algorithms in healthcare is as follows:

AliveCor supports the [early detection of atrial fibrillation](#), developed an ECG analytics platform – just as [PhysIQ Heart Rhythm Module](#), [Apple](#), and [Cardiologs](#) – and a six-lead smartphone ECG.

QbCheck helps with the [diagnosis and treatment of ADHD](#).

InPen [tracks insulin dosage](#).

One Drop Blood Glucose [quantifies blood glucose levels](#) and automatically sends the data to the paired app.

[Lumify](#) offers ultrasound image diagnosis.

Cantab Mobile acts as a tool for [memory problem assessment for the elderly](#).

EnsoSleep powers a [tool for recognizing sleep disorders](#).

AmCAD-US [evaluates thyroid nodules](#) and categorizes nodule characteristics.

[Lepu Medical](#) and [BioFlux](#) detect arrhythmias.

[Subtle Medical](#) offers a medical imaging platform.

[Bay Labs](#) offers echocardiogram analysis.

[Viz.AI](#) [detects stroke on CT scans](#) and helps clinicians win the race against time.

[Arterys'](#) algorithms are able to [spot cancerous lesions in liver and lungs](#) on CT and MR images, as well as [analyse cardiac MRI](#).

[Empatica](#) helps predict epileptic seizures.

[Cognoa's](#) algorithm built into an app [helps diagnose autism in kids](#).

[Medtronic](#) and [POGO](#) monitor and predict blood glucose changes.

[Idx](#) autonomously [detects diabetic retinopathy using retinal images](#).

[Icometrix](#) helps [neurologists interpret brain MR images](#).

[Imagen](#) [aids healthcare providers in identifying wrist fractures](#) with similar accuracy as human radiologists.

[NeuralBot](#) offers a solution for transcranial Doppler probe positioning.

[MindMotion Go](#) advances its algorithm for motion capture for the elderly.

[Dreamed](#) assists [healthcare professionals in the management of Type 1 diabetes](#).

[Zebra Medical Vision](#) detects and [quantifies coronary artery calcification](#), and analyses chest X-rays.

[Aidoc](#) is able to [flag brain bleeding in head CT images](#) and pulmonary embolism.

[iCAD](#) [classifies breast density and detects breast cancer](#) as accurately as radiologists.

[ScreenPoint Medical](#) assists radiologists with the reading of screening mammograms.

[Briefcase](#) triages and diagnoses time-sensitive patients.

[RightEye Vision System](#) [tracks eye movements for identifying visual tracking impairment](#).

[MaxQ](#) has been developing an acute intracranial hemorrhage triage algorithm.

[ProFound AI](#) detects and diagnoses suspicious lesions.

[ReSET-O](#) offers an adjuvant treatment of substance abuse disorder.

[Verily](#) developed an ECG feature on the [Study Watch](#).

[Paige.AI](#) provides a clinical-grade algorithm in pathology.

[FerriSmart](#) created a machine learning solution for the quantification of liver iron concentration.

[QuantX](#) offers the [first ever FDA-approved algorithmic tool for cancer diagnosis](#).

[Koios Medical](#) developed an A.I.-based decision support system for radiologists to aid breast cancer diagnostics.

[Canon Medical](#) works on CT noise reduction with its smart algorithm.

[CureMetrix](#) helps to triage and prioritize mammography workflow for radiologists.

SENDING DATA TO THE A.I.

ANALYSIS



CHALLENGES OF ARTIFICIAL INTELLIGENCE

Although the potential of artificial intelligence for making healthcare better is indisputable, and as the rundown of approved FDA algorithms shows, the number of operable software is exponentially growing, the successful integration of the technology into our healthcare systems is far from inevitable. For doing so, we have to overcome technical, medical limitations, as well as regulatory obstacles, soothe ethical concerns and mitigate the tendency to oversell the technology. In this chapter, we summarize the main challenges that concern the translation of artificial intelligence to everyday life, and ask the questions that are usually raised in the medical world when contemplating about A.I. and the years to come.

Misconceptions and overhyping

Overhyping the capabilities of A.I. through marketing tactics and oversimplified media representations does not help but destroy a healthy image about how A.I. could contribute to healthcare. It also adds to the fog of confusion and misconceptions which need to be cleared up when we want to implement the technology successfully into our healthcare systems. Definitions of machine learning, deep learning, smart algorithms, ANI, AGI or any other terms and concepts around A.I. need to be treated carefully. The same goes for its impact in healthcare.

That's the reason why anyone interested in the field [should read research papers and news about A.I. carefully](#). It matters in which scientific journal a study was published, and what kind of data the authors used. No algorithm can be trained without a good amount of quality data. The more images, text or any other source material the researchers have, the more precise the algorithms become. However, it is very difficult to get to a large amount of quality data, especially in medicine. Hospitals and other medical facilities, which have been accumulating health data for decades, are more often than not very reluctant to give their sensitive data to algorithmic experiments. Thus, some research groups do tricks on their dataset to make it bigger (e.g. invert images to double the database size). As a reader of an A.I. research paper, you should be aware of that.

Clinical collaborations also matter. If an algorithm performs well on a pre-selected dataset, that sounds alright, but clinical datasets make it excellent. For example, knowing that [DeepMind](#) has partnered with the NHS on several projects for years, e.g. [with the Moorfields Eye Hospital](#), makes a strong case for their algorithm.

In addition, if the authors only mention A.I. but don't describe the method how they are attempting to reach artificial intelligence, you should be skeptical and careful. A company or a research group should mention [machine learning or deep learning and be able to explain the method in detail](#) with which they are aiming to create A.I.

Regarding news pieces about smart algorithms, as online journalism centers around clicks, likes, and shares, many titles tend to be sensationalist and clickbait in nature. Of course, the word 'artificial intelligence' is much catchier than machine learning and saying that algorithms beat doctors will get a much wider audience than writing about percentages and comparisons in the title. That's why it is worth looking up the study itself when reading about the latest fantastic achievement of A.I. in an online article, and reading the conclusion of the research. In quality online magazines, the study is usually linked in the article, but if you want to 'play it safe' regarding its credibility, you can also search for the author and the medical journal itself on Google Scholar.

Technological limitations of A.I.

The term "artificial intelligence" itself might be misleading as due to the overuse of the expression, its meaning started to get inflated. Its meaning implies software with cognition and sentience, a far far more developed technology than it is standing at the moment, and a far far more likely used concept by marketing gurus to describe each and every software tool dealing with big data analytics.

At best, current technology – meaning various machine learning methods – is able to reach ANI in various fields. Yet, as we mentioned before, there is incredible growth in computer vision, as well as in natural language processing. Michelle Zhou, who spent over a decade and a half at I.B.M. Research and IBM Watson Group before leaving to become a co-founder of [Juji](#), a sentiment-analysis start-up, categorized ANI [for The New Yorker here](#) as recognition

intelligence, and as the first stage of A.I. It means what algorithms running on ever more powerful computers can currently do is recognizing patterns and gleaning topics from blocks of text or deriving the meaning of whole documents from a few sentences. Yet, we are nowhere close to AGI, that level of intelligence when a machine is capable of abstracting concepts from limited experience and transferring knowledge between domains – and that will remain the case most likely for decades to come.

Moreover, it turned out in several instances that artificial intelligence algorithms can be fooled through adversarial examples. Recently, a group of engineers from the university of KU Leuven in Belgium showed in a [study](#) in 2019 how simple printed patterns can fool an A.I. system that's designed to recognize people in images. If students printed out specifically designed patches and held the paper in front of them, the A.I. didn't recognize the students – as if they were cloaked in Harry Potter's Cloak of Invisibility. [As the researchers write](#): "We believe that, if we combine this technique with a sophisticated clothing simulation, we can design a T-shirt print that can make a person virtually invisible for automatic surveillance cameras." Computer scientists regularly test deep learning systems with so-called "adversarial examples" crafted to make the A.I.s misclassify them in order to find out the possible limitations of current deep learning methods. However, these adversarial examples could also be used to [fool self-driving cars](#) into reading a stop sign as a lamppost, for example, or they could trick medical A.I. vision systems that are [designed to identify diseases](#).

In that study, the researchers tested deep learning systems with adversarial examples on three popular medical imaging tasks—classifying [diabetic retinopathy](#) from retinal images, [pneumothorax](#) from chest X-rays, and [melanoma](#) from skin photos. In such attacks, pixels within images are modified in a way that might seem like a minimal amount of noise to humans, but could trick these systems into classifying these pictures incorrectly. The scientists noted that their attacks could make deep learning systems misclassify images up to 100 percent of the time, and that modified images were imperceptible from real ones to the human eye. They add that such attacks could work on any image, and could even be incorporated directly into the image-capture process. Now, that's a worrying phenomenon – as hackers could easily find a way to attack artificial intelligence-based medical software through such "adversarial examples" for medical fraud or for causing harm on purpose.



Limitations of available medical data

For building reliable algorithms, one of the most important components is having reliable datasets. However, health data is a difficult "beast". As it contains sensitive information, companies that are trying to create algorithms, often have difficulties to gain access to data other than publicly available datasets.

Moreover, health data itself was never crafted with artificial intelligence algorithms in mind, thus their streamlining and categorizing require tremendous effort and energy – even in the case of digitized data, such as electronic medical records. And in spite of the widespread notion about computerization, EMRs are not everywhere. As cited earlier, in 2013, only 78 percent of office-based physicians used any type of EMR system. And while EMRs have their specific troubles due to having been crafted rather according to the needs of providers than doctors or patients, paper-based medical records are the most useless from an A.I. point of view. As [researchers, who studied five large medical facilities where hospitals and clinics used conventional paper-based patient records, reported](#), 5-10 percent of patients were seen in clinics without an available record, while 5-20 percent of hospital records were entirely incomplete. Of the missing information, 75 percent consisted of laboratory test results and X-ray reports, and 25 percent of lost, incomplete, or illegible textual data. Now, try to craft a reliable algorithm from that.

The indispensable work of data annotators

As mentioned above, medical data was definitely not crafted with smart algorithms in mind. Let's say you want to have an algorithm to spot lung tumors in chest X-rays. For that, you will need to use tools for pattern recognition, such as supervised machine learning, - which will make the task very similar to spotting cats on Instagram.

It sounds easy, doesn't it? You say that if the image is of a furry animal with two eyes, four legs, that could be a cat. You describe its size, potential colors and how its cheek looks. Still, what if the animal is partly overshadowed by something? What if it's playing and it only looks

like a ball of fur? And ultimately, how do you even tell the computer all of this if it doesn't understand legs, eyes, animals, just pixels?

"You will need millions of photos where those photos that have a cat in them are appropriately labeled as having a cat. That way, a neural network and in many instances, a so-called multilayered deep neural network can be trained using supervised learning to recognize pictures with cats in them.", David Albert, M.D., Co-Founder and President of [AliveCor](#), the company that has been developing a medical-grade, pocket-sized device to measure EKG anywhere in less than 30 seconds, told The Medical Futurist. So, you won't tell the algorithm what a cat is, but you rather show it millions of examples to help it figure it out by itself. That's why data and data annotation is critical for building smart algorithms.

The task of annotating data is a time-consuming and tedious work without any of the flare promised by artificial intelligence associated with sci-fi-like thinking and talking computers or robots. In healthcare, the creation of algorithms is rather about utilizing existing databases which mainly encompass imaging files, CT or MR scans, samples used in pathology, etc. At the same time, data annotation will be drawing lines around tumors, pinpointing cells or designating ECG rhythm strips. Thousands of them. No magic, no self-conscious computers.

That's what Dr. Albert has been doing. He explained that "you must have accurately labeled and annotated data in order to develop these deep neural diagnostic solutions. For example, I may annotate or diagnose ten thousand ECGs over the course of several weeks, then another expert likely diagnoses the same ten thousand – and then we see where we disagree. After that, we have a third person, who is the adjudicator – who comes in and says, okay, regarding this five hundred where you disagree, this is what I think the answer is. And so it takes a minimum of three people to give you a reasonably confident answer that something is correct. Now, you can see it's an awful lot of work. Deep neural networks to perform correctly in order to take advantage of big data require a tremendous amount of annotation work".

Katharina von Loga works as a consultant pathologist at The Royal Marsden NHS Foundation Trust, and she's using software-based image analysis to monitor the changes of immune cells within cancerous tumors during therapy. The computer helps her count the cells after she designates carefully the set of cells she's looking for. "I have an image of a stain in front of me, where I can click on the specific set and annotate that that's a tumor cell. Then I click on another cell and say that's a subtype of an immune cell. It needs a minimum of all the

different types I specified, only after that can I apply it to the whole image. Then I look at the output to see if I agree with the ones that I didn't annotate but the computer classified. That's the process you can do indefinitely," she explained.

Sometimes data annotators not only need great skills for pattern recognition and medicine, but it is also beneficial if they are good at drawing. Felix Nensa works as a consultant radiologist at the University Hospital Essen, more specifically at the Institute of Interventional Radiology and Neuroradiology. He explained the hardships of data annotation through an example in a new medical subfield called [radiomics](#). "We identified a cohort of 100-200 patients with a certain type of tumor and we want to predict if the tumor responds to a certain therapy. In order to do it, you have to do a full segmentation of the tumor, a CT scan. If it's lung cancer than you have the full CT scan of the lung which includes slices of the lung in a particular slice thickness – let's say 5 mm. Then you have to draw a line around the tumor in each slide extremely precisely. It's really a lot of work to draw a smooth line around this shape because such a tumor can be really large and most often it's nothing like a ball."

Katharina von Loga said that although it sounds perfect in theory that you can train an algorithm within a period of time to support medical work in pathology, the practice is much more complicated. As medical data archives were (obviously) not created with mathematical algorithms in mind, it's gargantuan work to try to standardize existing sampling processes or to have enough "algorithm-adjusted" samples. In her field, for example, it matters how the sample was processed from getting the specimen from the patient until it's under the microscope. The staining method, the age of the sample, the department where the sample was produced – all factors to count in when it comes to making a decision about a sample for successful algorithmic teaching.

Beyond the problems of the massive variability in the samples, David Albert also mentioned the lack of experts for data annotation, as well as the difficulty to find databases of scale. Usually, the precision of an algorithm depends on the size of the sample – the bigger, the better. However, the leader of AliveCor mentioned how hospitals or medical centers, even really resourceful ones, don't have enough data or enough annotations. For this reason, he believes that it will take companies like Google, Amazon, Baidu or Tencent with unlimited financial resources and a global footprint to really derive the kind of scale that you need to develop accurate A.I. What is more, the human resources problem is aggravating. "There are only 30,000 cardiologists in the United States, all very-very busy. They don't have time to mark ECGs. On the same note, there are only 25-40,000 radiologists – they don't have time to

read more chest X-rays. So, we're gonna have to do something".

All three experts mentioned the option to employ medical students or pre-med students in university for simpler annotation tasks – to at least solve the human resources trouble. David Albert played with the thought of building online courses for training prospective annotators, who would afterward get some financial incentives for the annotation of millions of data points. Medical facilities could basically crowdsource data annotation through platforms such as [Amazon Mechanical Turk](#). The process could employ the "wisdom of the crowds". Imagine that you mix up a hundred chest X-rays and present them in batches of fifty to online annotators. If there are a thousand annotators, you will find out what the majority diagnosis on those hundred images was. Although you won't necessarily have experts of the field, statistics would say their responses will converge to the right answer, David Albert explained. He added that although the method was already ideated, it hasn't been applied yet.

Another option would be the employment of algorithms also for annotation tasks – so basically building A.I. for teaching another smart software. Felix Nensa mentioned that they are building deep learning-based tools that can do completely automatic annotation by themselves and then the user just has to correct where this automatic process did not work well. But certain easier tricks could also speed up the tedious annotation procedure. He mentioned how convenient it would be to "build an app for your iPad, where you can really take a pen, and draw the tumor, which is obviously much more convenient than clicking with a mouse connected to a desktop computer".

Katharina von Loga also mentioned how international and national committees are working on the standardization of the various sampling processes, which could really ease annotation



Judgemental datasets and A.I. bias in healthcare

Not only the accessibility to datasets and their formal appropriateness for becoming base material for an algorithm – whether or not the data is correctly labelled –, but also the quality of the content matters a lot when crafting an algorithm. In the last years, instead of the rational, logical algorithms expected to make reasoned and objective choices were to be found racially, gender-wise, ethnically or culturally biased.

We could hear from a lot of news outlets how facial recognition software favors white faces, but a [study](#) out of the MIT Media Lab published in February 2018 actually found that facial-recognition systems from companies like IBM and Microsoft were 11-19 percent more accurate on lighter-skinned individuals. They were particularly bad at identifying women of color. The smart algorithms were 34 percent less accurate at recognizing darker-skinned females compared to lighter-skinned males. [In another example](#), when A.I. was implemented in the U.S. criminal justice system to predict recidivism, it was found to disproportionately suggest black people were [more likely to commit future crimes](#), regardless of how minor their initial offense.

Regarding an instance of a gender-biased algorithm, [Amazon's HR department had to stop using their A.I.-based machine learning tool](#), which the company developed for sorting out the best job applicants, as it turned out that the smart algorithm favored men. As the tech scene is mainly dominated by men, and the data that the software was fed contained resumés from the previous 10 years, the program taught itself that women were less preferable candidates. While programmers tried to tweak the A.I., it still didn't bring the expected results, so in the end, they decided to scrap the program entirely. But what happened here? What went wrong with the algorithm? What's the difficulty with teaching A.I.?

There are basically three main reasons why a smart algorithm could turn out to be biased. Algorithms are trained on datasets, thus the quality of the data is crucial in the process. If the dataset is incomplete, not diverse enough, stems mainly from one area of study, the A.I. software could work flawlessly in the test environment, but come up with its inherent bias in the 'real world'.

Another, more complex issue is when the dataset is representative and diverse enough, but the algorithm still arrives at discriminative conclusions. The reason for that could be a social practice ingrained so deeply in society that will automatically be transferred into the judgment

process of the A.I. In a nutshell, that's the reason for Amazon's gender-biased HR algorithm: the program was fed with applications from the previous ten years, whose majority came from male candidates. As a consequence, the A.I. started to believe that the correlation between gender and qualifications in this area also meant causation – and a point of reference for selection.

And what if the programmer must choose or leave out some parameters to help the program learn? By describing factors, variables, elements in a certain way, they already include their hidden and frequently unconscious bias. When banks screen through loan applications with the help of algorithms, who decides who can get a loan? The programmer, the bank, or a human being? In such cases, the software developer can unconsciously include their own values and beliefs about the world into the code, and in an even more sensitive situation, perhaps with even riskier outcomes, the programmer can set some variables selecting specific characteristics for individuals or groups – which might have a biased outcome. Either way, individual choices can greatly influence how smart algorithms 'behave'.

Thus, the source, the quality, and the diversity of the data, the historical social practices ingrained into the data, meaning the bias of the deeper social structure, as well as the individual, conscious or unconscious preferences of individual programmers, determine whether and to what extent an A.I. will become biased. Now, let's look at some examples from healthcare where many could believe that as smart algorithms look at medical images, ECG strips or electronic medical records, the "bias factor" must be less prevalent.

Well, we shall bring some disillusionment. Even [comedian John Oliver said](#) that bias in medicine, in general, is a serious issue with consequences for American society. [Healthcare data is extremely male and extremely white](#), and that has real-world impacts. A 2014 study that [tracked cancer mortality over 20 years](#) pointed to a lack of diverse research subjects as a key reason why black Americans are significantly more likely to die from cancer than white Americans.

In another area of research, [meta-analysis](#) looking at 2,511 studies from around the world found that 81 percent of participants in genome-mapping studies were of European descent. This has severe real-world impacts: researchers who download publicly-available data to study disease are far more likely to use the genomic data of people of European descent than those of African, Asian, Hispanic, or Middle Eastern descent. And these distorted datasets would be the starting points for A.I. development.

Sometimes, ignorance of inherent bias in data could even jeopardize the applicability of an algorithm. [Winterlight Labs](#), a Toronto-based startup, which is building auditory tests for neurological diseases, [realized after a while that their technology only worked for English speakers of a particular Canadian dialect](#). That might be a serious problem for other companies, too, which are working with [voice-to-text technologies](#), [vocal biomarkers](#), or digital assistants such as Siri or Alexa for healthcare.

So, what should we do to eliminate these prejudices from programming smart algorithms? It is actually a very difficult task as human beings have their own bias in their thinking – and that has been a useful trait for thousands of years as it shortens the time needed for making snap decisions. It's also likely that human bias is here to stay, and technologies that are fed by information that is created in the real world [could fundamentally have the same outcome](#). So now the question is, how do you think about the situation when you're actually shifting a cognitive task completely into a machine, where you don't have the same kind of qualitative reaction that human beings will have?

The response might be twofold and still evolving. At first, we have to raise awareness of inherent bias in algorithms. Recently, [police officers have raised concerns about using "biased" artificial-intelligence tools](#), a report commissioned by one of the UK government's advisory bodies revealed. The report said policemen were worried about both data bias and becoming more reliant on automation. Another similar example was banning [facial recognition software from the streets of San Francisco](#). Activists and politicians, who pushed for the ordinance, cited studies that showed A.I.-based facial recognition technology is less accurate when distinguishing between individual women and people of color. That's an efficient move that might be followed by many others. As a second step, though, we might have to re-create these functions, such as facial recognition technology, to represent a more balanced attitude through minimizing bias. That's a tricky and a difficult process, especially because most A.I. algorithms are trained on biased datasets and researchers are just starting to bring them to the real-world.

Search

diversity m |



all




pictures




videos

matters



meaning



The Need To Regulate A.I.

As a huge step towards the future, the FDA approved the first cloud-based deep learning algorithm for cardiac imaging developed by [Arterys](#) in 2017. Ever since, the U.S. regulator has approved dozens of smart algorithms (see “The current list of FDA-approved algorithms in healthcare” above). However, regulations around artificial intelligence generally lag behind or are still literally non-existent. With the technology gaining ground and [appearing in hospitals within the next 5-10 years](#), decision-makers and high-level policy-makers cannot allow themselves not to tackle the issue.

They should rather step ahead of the technological waves and guide the process of implementing A.I. in healthcare along the principles and ethical standards they work out with other industry stakeholders. Moreover, they should push companies towards putting affordable A.I. solutions on the table and keeping the focus on the patient all the time. Governments and policy-makers should also help in setting up standards on A.I. usage as we need specific guidelines starting from the smallest units (medical professionals) until the most complex ones (national-level healthcare systems).



[CEO of Facebook had his hearing in the U.S. Congress last year](#), even though there were some members of the Senate as well as the House of Representatives who asked relevant questions and displayed their knowledge and interest, unfortunately, the main takeaway was that politicians [understand little of how the tech world works](#). Sometimes, [when questions like the one about how Facebook can sustain itself when it offers their services for free popped up online after the hearing](#), it was difficult to decide whether it is just a meme or someone indeed asked Zuckerberg that. After having followed the congressional hearing, we had a well-founded fear that if politicians and lawmakers have so much trouble with the data management issues of a social media site (and make no mistake here, we are aware that it is a complicated topic), how will they regulate artificial intelligence, robotics, virtual reality or healthcare apps? How will they handle topics where human lives are at stake?

The Ethics of A.I.

While the forecasting and predictive abilities of smart algorithms are anchored in previous cases and thus they might be quite useless in novel cases of drug side-effects or treatment resistance, these types of problems alongside the technological and data-related limitations will still be easier to overcome than ethical and legal issues.

ANI and at a certain point, AGI, should be implemented cautiously and gradually in order to give time and space for mapping the potential risks and downsides. Independent bioethical research groups, as well as medical watchdogs, should monitor the process closely. This is exactly what the [Open AI Foundation](#) does on a broader scale. It is a non-profit A.I. research company, discovering and enacting the path to safe artificial general intelligence. Their work is invaluable as they are doing long-term research, and may help in setting up ethical standards on how to use A.I. on micro and macro levels. Perhaps also in the healthcare sector.

Without the work of bioethicists, philosophers, futurists, and groups such as Open AI, who will answer questions like who is to blame if a smart algorithm makes a mistake and does not spot a cancerous nodule on a lung X-ray? To whom could someone turn when A.I. comes up with a false prediction? Who will build in safety features so A.I. will not turn on humans? What if algorithms will not only cognitively challenge humans but also on the level of feelings? What will be the rules and regulations to make a decision on safety?

As we are completely intrigued by these issues and dilemmas, we have chosen some of the most important ones and attempt to give you a response.

Could You Sue Diagnostic Algorithms or Medical Robots in the Future?

What if a deep learning algorithm misses a diagnosis, the doctor accepts the judgment and the patient dies? Who will be held liable in the future when A.I., acting autonomously, wrongs humans?

Let's imagine a future scenario. In 2031, Andrea went to Milan for a check-up to his GP because he felt nauseated all the time and noticed a strange pressure on the left side of his head. The doctor suggested to him that he should run a couple of tests and informed him about involving a diagnostic algorithm in the procedure. The machine learning algorithm was trained to identify brain tumors – one of the [first studies in the area dates back to March 2018](#) – with very high accuracy. In most cases, it diagnosed cancerous tissues far better than some trained histopathologists, but in Andrea's case, something went astray. The algorithm found something different than the diagnostician, and as the use of A.I. was already common practice, the histopathologist did not question the judgment. As a result, Andrea was mistreated: an unnecessary operation, ineffective medication cures and long-long weeks went by until someone discovered the algorithmic error. However, the patient's brain already suffered irreversible damages, and the family decided to sue.

Needless to say that we have highly theoretical reasoning here, without knowing the particulars which refine every case down to the point where there is a particular patient with a particular condition on a particular day in a particular place. Staying on the theoretical level, though, David Harlow, a US-based healthcare lawyer, consultant and [blogger focusing on digital health](#) told The Medical Futurist that it is still worth breaking down the case to key categories of concern: design flaws, implementation flaws, and user error. Thus, when we look at technology's encounter with the doctor and the patient, depending on the case, there might be a design flaw – in which case the company might be liable, an implementation flaw, in which case the doctor or the nurse might be responsible, and user error, which might go down to the patient.

Here, we are assuming that there was no user error, the patient could not have done anything differently, so our cases could either go down to design flaws or implementation flaws.

First of all, it is worth examining the differences in technologies when looking at the above-mentioned hypothetical case. In case of analog technologies – the first layer of technologies or traditional technologies – that provide data or let users access data without any algorithm (e.g., stethoscope); when a design flaw causes harm to patients, “the first step on the road to being able to hold the company liable, is often a trip to the FDA, seeking a recall of the medical device for failure to comply with the FDA approval”, says Harlow. In the case of digital technologies, constituting the second layer of technologies when looking at the level of advancement; which do have algorithms programmed into them without the code changing by itself (e.g., medical records software), the situation might be similar. As according to Harlow they might be considered a “black box,” i.e., a system that takes in some inputs and yields an output, without affording the clinician reading the output insight into the algorithm conducting the analysis, it is regulated as a medical device, the procedure might be similar to that of the analog technology.

The third category of technologies might be the most interesting and the most problematic to regulate and deal with – deep learning or machine learning algorithms for diagnostics in radiology and pathology. Here, Harlow asks the question of how we know whether the algorithm is progressing in the “right” direction. He says, in these circumstances, there are at least two intersecting bodies of thought that should affect the physician's decision-making when using A.I.: regulatory approval and standard of care. In the first case, assuming the FDA (or analog) can approve a machine learning tool that will change over time. Without regulatory approval of a device, the device may not be used in clinical practice. In the second case, the device needs to be permitted or required by the current professional consensus on the practice of medicine to be considered within the bounds of practice (i.e., it's not malpractice to use it).


Thus, legal practice might try to create a 'black box' out of machine learning algorithms as well, however, it might be more problematic here than in case of other digital technologies so, liability might go down to the company creating the algorithm.

Going further down the legal road, when there is no proven design flaw in the case, we have to examine whether the pathologist or phlebotomist or any other physician used the device as it was supposed to be used. Harlow says that in our case, the professional is open to liability if they used the tool in a situation outside the scope of its regulatory approval, or misused it, or applied it despite significant professional questioning of the validity of the evidence surrounding the tool, or with knowledge of the toolmaker obfuscating negative facts. In any other cases, the ball falls back on the creators and the companies.

And what if we go even further and imagine A.I.-powered robots in the future? What will we do with [Sophia-like creatures who already have citizenship in Saudi Arabia](#)? What will we do with fully autonomous machine learning algorithms making decisions based on their judgment as a result of considerations that might be out of human perception?

The European Union seems to experiment with a new legal status for the future. A European Parliament [report](#) from early 2017 suggests that self-learning robots could be granted "[electronic personalities](#)." Such a status would not mean that they could adopt kids or marry humans or other robots. It might only allow robots to be insured individually and be held liable for damages if they go rogue and start hurting people or damaging property. And how would the aggrieved parties receive any compensation? There is an idea for example by setting up a compulsory insurance scheme that could be fed by the wealth a robot is accumulating over the time of its "existence."

Although A.I. experts and researchers criticize the report for [allowing manufacturers to shake off their responsibilities](#), the idea might be a creative solution for a widening grey area in medical malpractice law. Other, similarly forward-looking legal notions and principles will be necessary for the near future, as Harlow estimated, the first scenarios with narrow artificial intelligence might arrive as early as this year at the medical malpractice law firms.



Should Algorithms Mimic Empathy?

Chatbots acting as life coaches sound astounding and terrifying at the same time. Extensive research has been going on lately in the field of applying human features, emotions, gestures, and reactions to digital technology; and it raises thousands of questions. Could not only smart, but emotional algorithms appear also in healthcare soon? Would there be a place or need for them? How would it impact the patient-doctor relationship or social interactions in general?

The human touch is the key part of practicing medicine. It is an integral part of the patient-doctor relationship, where patients feel that they are taken care of by a fellow human being, they are not alone in need. There is someone who not only understands their problem cognitively and offers a solution, but can easily "put themselves into the other person's shoes" in the first place. Research proves that this ability significantly boosts the healing process. For example, [diabetes patients who had compassionate physicians](#) had a lower rate of disease complications than their peers. People who caught a common cold [perceived their condition less severe](#) when they encountered an empathic medical professional. We are social beings; we need a caregiver to tell us everything is going to be fine. But then the question arises. Why are we building chatbots like [Woebot](#) or virtual assistants like [Nadia](#)?


At the dawn of modern healthcare, around the turn of the 18th century, medical professionals started to alienate themselves from the patients by not looking at them as persons, rather as symptom carrying medical cases to solve with the help of science. French philosopher, Michel Foucault even dedicated a book to explain what happened around [the birth of the modern clinic](#). It has been a dehumanizing process for both patients and doctors. People with problems are considered as just (statistical) numbers and symptoms in crowded waiting rooms, while doctors have only a few minutes for each patient on average and have to go on with their packed schedule as soon as possible. Thus, it is not surprising that patients experiencing grumpy doctors try to turn to empathy coming from someplace else, and digital technology tries to tap into that existing gap.

And why do medical professionals have so little time for patients? They are suffering from the burden of administration, hideous monotonous tasks and the lack of colleagues. Doctor shortages are global phenomena. The World Health Organization (WHO) estimates that there is a worldwide shortage of around 4.3 million physicians, nurses, and allied health workers. At the same time, the need for healthcare services is rising; illnesses are becoming easier to

catch, civilizational diseases such as diabetes and obesity are on the rise while aging societies need more and more care. Therefore, medical virtual assistants or healthcare chatbots with a pinch of empathy seize the moment and claim their places as new helpers of medical professionals.

Looking at the practical side and the harsh facts, it seems that digital technologies being able to reach out to patients through empathy and compassion should have a place in healthcare. However, from the perspective of the human-technology relationship as well as the interactions between humans themselves, it's more problematic. Would patients trust or accept AI-based chatbots as their companions during hard times?

And what is the psychological or individual reason why we want to program algorithms to emit human emotions? Is it a further step of estrangement and alienation in an already estranged world full of smartphones and televisions? Does it become so difficult to reach out to real people and engage in meaningful relations that the solution seems to be to build an echo of the human emotional spectrum? Is it a coincidence that research and development into empathetic and emotionally charged robots are the most advanced in Japan, where over [60 percent of unmarried people aged between 18 and 34 have no relationship](#) with a member of the opposite sex? Might future generations grow up with empathetic robots and emotional algorithms? Do we want to create technology so emotionally intelligent that they will not be able to tell the difference between a human and non-human apology anymore?



There are so many ethical, moral questions and possible outcomes regarding empowering technology with humanoid features. However, time is pressing for figuring out our possible responses and attitudes towards emotionally intelligent machines, as experiments for modeling human emotions with the help of machines is on-going and there are already amazing results. Mark Sagar and his team at Soul Machines is working on the [BabyX project](#), a virtual baby powered by A.I. modeled according to the already known workings of the human organism. When BabyX smiles, it's because her simulated brain has responded to stimuli by releasing a cocktail of virtual dopamine, endorphins, and serotonin into her system. This is part of Sagar's larger quest, using A.I. to reverse-engineer how humans work. Their ultimate aim is to build virtual assistants able to mimic humanoid features, as in the case of the above-mentioned Nadia project. Yet, if we look at the struggles and [challenges around Nadia](#), e.g. the lag time between a question being asked and Nadia delivering an answer, we have to recognize that there are still several years until technology reaches the point where mimicking empathy by machines becomes possible.

What is important for patients, medical professionals, as well as tech companies, is to realize that artificial emotions cannot replace human interaction, empathy, and compassion. But a coded gesture coming from a machine might reach its goal to temporarily offer comfort, especially when its limitations are fully acknowledged and accepted. So, if you don't expect the machine to act like a real human with unique reactions but rather as a programmed robot with guessable responses and gestures, you cannot get disappointed.

Could A.I. Solve The Human Resources Crisis In Healthcare?

Experts are pessimistic when it comes to the human resources situation in healthcare. There are worrying tendencies about the off-balance of supply and demand of medical professionals. According to [The Medical Futurist Institute](#), the needs-based shortage of healthcare workers globally is about 17.4 million. That's almost the entire population of Chile and twice as many as living in Austria or Israel.

The WHO's [Global Strategy on Human Resources for Health: Workforce 2030](#) reports that shortages can mount up to 9.9 million physicians, nurses and midwives globally by 2030. The U.S. could see a lack of up to 120,000 physicians by 2030, according to a [report](#) published in April 2018 by the Association of American Medical Colleges. In the next 10-12 years, [Southeast](#)

[Asia will need approximately 4.7 million more health workers](#) to achieve sufficient coverage, while the Western Pacific zone, which includes regional heavyweights such as China, South Korea, and Japan, will be missing roughly 1.4 million people. The situation is not an inch better in Europe: the continent may be facing a shortage of 1 million health professionals by 2020, according to a [European Union Joint Action on Health Workforce Planning](#) estimate.

At the same time, the number of the Earth's human inhabitants is increasing day by day. The [World Population Counter](#) says that there are three times more births than deaths on a given day, while more than 7.6 billion people populate the Earth at the moment. That's an insanely growing number. Parallely, life expectancy is widening, and populations are aging. According to data from [World Population Prospects: the 2017 Revision](#), the number of older persons — those aged 60 years or over — is expected to more than double by 2050 and to more than triple by 2100, rising from 962 million globally in 2017 to 2.1 billion in 2050 and 3.1 billion in 2100. As a consequence, the need for healthcare services is rising: illnesses are becoming easier to catch, civilizational diseases such as diabetes and obesity are on the rise while aging societies need more and more care.

These statistics not only shed light on the large amplitude of the problem but also on the widening gap between supply and demand. Thus, the healthcare workforce crisis is due to aging populations and a higher need for chronic care, the worldwide doctor shortage – and the aging and burnout of physicians themselves.

Some countries experiment with increasing the number of medical students and making their healthcare systems more attractive. For example, [Singapore accepted 110 more students in its medical training in 2015](#) than before. The government has also made an effort to attract Singaporeans who are midway through their medical studies abroad to return to work in the country; it offered to pay up to S\$50,000 a year for their final three years of education, and in return, they had to work in a public hospital for three to four years, which included their training year as housemen. However, in the long run, such incentives might not be enough to attract people to the medical profession as the amount of work will increase in parallel with the HR crisis. That's the point where digital technology, more specifically artificial intelligence, might come and save the day.

In its study, The Medical Futurist Institute argues that ANI has the highest chance of being used in the medical practice for analyzing large datasets, finding new correlations and generally supporting caregivers' jobs. A.I.-based services could facilitate more accurate diagnoses, administration, decision-making, big data analytics, post-graduate education, among others.

As highlighted earlier, smart algorithms could [assist medical professionals in designing treatment plans and finding the best-suited methods](#) for every patient. They can take over repetitive, monotonous tasks, so that physicians and nurses can concentrate on their actual jobs instead of, e.g. fighting the tread-wheel of bureaucracy. In the future, as in the brilliant movie, [Her](#), cognitive assistants could [prioritize emails in doctors' inboxes or keep them up-to-date](#) with the help of finding the latest and most relevant scientific studies in seconds. Moreover, its transformative power makes it [as essential as the stethoscope](#), the symbol of modern medicine, which appeared in the 19th century. By all means, we need to emphasize that practicing medicine is not a linear process. Not every single element and parameter can be translated into a programming language – but there are areas where ANI could definitely improve patient outcomes and ease the burden on medical staff.

Unfortunately, ANI is no wonder weapon, and many challenges will arise from technology. On the one hand, there is the financial burden: the cost of disruptive technologies might be too high for underdeveloped countries, pushing them further behind in improving healthcare. On the other hand, many technical and ethical questions pop up. What elements of physicians' repetitive tasks, such as note taking or administrative duties could A.I. ease, and which ones, such as diagnosis, treatment or monitoring, could it facilitate? Most doctors use online tools to help with research. Is there really a difference when it comes to using A.I.? Should A.I. be handled as another tool, such as a stethoscope or as an individual entity?

On the patients' side, will they stick to the human touch when shortages simply do not give them a chance to meet a physician in person for every medical issue? What if A.I. algorithms can mimic empathy either through an app or a chatbot? It's not yet known whether they will accept the use of A.I. in decision-making and learn its use in their care.

On the level of society, will it help shift focus from treatment to prevention? Will A.I. increase the cost of care? Might doctors and medical professionals become more efficient, because A.I. handles some of the time-consuming tasks? Will doctors provide better care in underdeveloped regions with the use of A.I.? And generally, how will it (if at all) change the current structures of insurance policies?

And most importantly, would artificial intelligence replace physicians?



MEDICAL PROFESSIONALS, A.I. AND THE ART OF MEDICINE

At the dawn of the [Fourth Industrial Revolution](#), automation and digitization of our worlds and workplace are continuing, changing the job market, the nature of many jobs and even the concept of what it means to be working. Many fear that robots and automation will take their jobs and will leave them without alternatives, not only physicians. The phenomenon is not new: in the 19th century, members of the Luddite movement – textile workers and weavers – destroyed weaving machinery in protest and fear that machines would take their place in their industry.

Lately, the same fears emerge in healthcare about [artificial intelligence taking the jobs of radiologists](#), [robots surpassing the skills of surgeons](#), or [taking jobs in pharma](#). A renowned voice in tech, Kai-Fu Lee, founder of venture capital firm Sinovation Ventures [told CNBC that](#) A.I. will be bigger than all other tech revolutions, and robots are likely to replace 50 percent of all jobs in the next decade.

As the fears of losing the battle against new technologies grow exponentially, alternatives on the individual and social level have already surfaced. The [most popular policy-level concept](#) is the introduction of the universal basic income, in which case the government would give everyone just enough money to live on while creating incentives for individuals to take risks, start businesses, change jobs, return to school or try a new career. Another idea is the negative income tax, where the state would give money to the poor the same way as in the case of taxing rich people; but [Bill Gates would tax robots](#) and some [economists think the solution lies at the heart of governments creating more jobs](#).

While these responses for the challenges of automation and digital technologies are only ideas at the moment – except for the [national-scale experiment of Finland with the universal basic income](#) –, it is natural that people are worried in the face of such fundamental changes.

[Huge waves are coming to healthcare to transform the job of physicians](#) into something distinctly different than before. Although some of their tasks will be taken over by A.I., they will have more time for others, for example, deal with patients with real care and patience.

[Vinod Khosla, the founder of Silicon Valley venture capitalist firm, Khosla Ventures, says](#) that it is inevitable that, in the future, the majority of physicians' diagnostic, prescription and monitoring, which may approach 80 percent of total doctors'/internists' time spent on medicine, will be replaced by smart hardware, software, and testing. This is not to say 80 percent of physicians will be replaced, but rather 80 percent of what they currently do might be replaced, so the roles doctors/internists play will likely be different and focused on the human aspects of medical practice such as empathy and ethical choices. He also emphasized how that 20 percent of their current tasks, the remaining 20 percent will be amplified in the future. Doctors and nurses will become more efficient and they can specialize better than they do today.

In spite of the promises, numerous experts voiced fears about A.I. taking their jobs and destroy the profession. They say the art of medicine, the creative process of understanding the uniqueness of each and every patient and tailoring treatments according to the arisen needs, as well as the highly value-added method of processing input, finding the right response and treating patients accordingly, [might disappear due to new technologies](#). Naturally, the biggest fears come from areas where deep learning is already present and produces incredible results: natural language processing and computer vision. As described above, both have mind-blowing achievements in medical imaging and diagnostics concerning accuracy and efficiency.

As a clear consequence, many radiologists went into panic mode. In his presentation at the [GPU Tech Conference in San Jose in May 2017](#), Curtis Langlotz, Professor of Radiology and Biomedical Informatics at Stanford University, mentioned how he had received an email from one of his students saying he was thinking about going into radiology, but does not know whether it is a viable profession anymore. But the assumption that the radiologist profession is dying, is just plain wrong.

Bradley Erickson, Director of the Radiology Informatics Lab at Mayo Clinic told me that some of the hype we hear from some of the machine learning and deep learning experts saying that A.I. would replace radiologists means that they are looking at radiologists just like looking at pictures. That would be me saying while I look at programmers, all they do is typing, so we can replace a programmer with a speech recognition system, he added. Langlotz compared the situation to that of the autopilot in aviation. The innovation did not replace real pilots, it augmented their tasks. On very long flights, it is handy to turn on the autopilot, but they are useless when you need rapid judgment. So, the combination of humans and machines is the winner solution. And it will be the same in healthcare.

Thus, I agree with Langlotz completely when he says that artificial intelligence will not replace radiologists. Yet, those radiologists who use A.I. will replace the ones who don't. Moreover, this enigmatic statement could also apply to ophthalmologists, neurologists, GPs, dentists, nurses or administrators. That's why I reframed the above sentence to articulate the core message of The Medical Futurist team as succinct as possible.

Artificial Intelligence will not replace physicians. Yet, medical professionals who use A.I. will replace those who don't.

In the next years, artificial intelligence will surely transform medicine as we know it. It will find new drugs, new treatments, and therapies through matching combinations that human physicians, pharma companies or medical innovators would never think of. As it won't be limited by the traditional pathways and thought patterns used for centuries in medicine, artificial intelligence may come up with entirely new solutions – without telling humans how it figured these out. Similarly to the supercomputer in Douglas Adams' *The Hitchhiker's Guide To The Galaxy*, where the answer to life, the universe and everything else is 42, smart algorithms might just spit out answers to questions without explanation. The real art of medicine will be the undertaking to figure out the logical path of how the A.I. arrived at a certain solution. That will definitely need the high levels of creativity, problem-solving and cognitive skills that the medical community possesses.

Thus, we are sure that A.I. is not going to replace us; it's going to be [the stethoscope of the 21st century](#). Digital health will give us more health data than ever before, and A.I. will help us analyze it to find new ways to treat diseases, to cut down on administrative tasks, to streamline medical practices, to optimize both physicians' and patients' schedules. However, we should never forget that they are going to be tools in the hands of physicians – and not the other way around. Compassionate care, empathy, creativity, problem-solving and profound human connection will never cease to be the terrain of physicians, moreover, it will be enhanced by artificial intelligence.

If we embrace it, the real art of medicine begins with the era of artificial intelligence.

Read more on medicalfuturist.com.

Get in touch with Dr. Bertalan Meskó, The Medical Futurist on [LinkedIn](#).

Follow The Medical Futurist on [Twitter](#), [Facebook](#) or [Instagram](#)!

Subscribe to [The Medical Futurist YouTube channel](#) to get access to all the videos about trends, technologies and devices that will shape the future of medicine!

Click [here](#) to sign up to The Medical Futurist newsletter for news and exclusive interviews about digital health.