Committee: World Health Organisation

Issue: The utilization of artificial intelligence in medicine

Student Officer: Andreas Hagis

Position: President

PERSONAL INTRODUCTION

Dear Delegates,

My name is Andreas Hagis and I will be acting as the president of the World Health Organisation during the 3rd American College of Greece Model United Nations conference. This will be my 11th conference overall and my 6th time acting as Chair, and I'm really looking forward to the debate and discussion around our 3 topics.

Throughout my time attending MUN conferences, I've noticed that most tend to distance themselves from topics that are heavily influenced by science and technology, and stick more to traditional topics that revolve around politics and diplomacy. A lot of the topics at ACG this year, however, do focus on some of the more pressing issues of today, due to science and technology; which is to me a welcome change. In WHO this year you will need to find solutions to 3 very interesting issues based on how societies' ethics, technologies and capabilities have progressed and interacted with each other in global medical communities.

In this study guide I'm aiming to provide you with a starting point for in depth research into our first topic, which looks at the role artificial intelligence does and can play within medicine whether it is in research, or actual patient interactions. I hope you find this topic as interesting and fascinating as I do. If you have any questions after reading the study guide, need any clarification or want to discuss something specific, feel free to email me at andreashagis@gmail.com prior to the conference, and I will be happy to respond. I look forward to meeting you all and seeing what new and innovative ideas you put forward.

Best regards,
Andreas Hagis

TOPIC INTRODUCTION

The fundamental goal of medicine across the years has always been to apply the safest and most modern technologies to improve people's health. Throughout the information age we've seen countless breakthroughs in both science and technology, a significant one being the dawn of artificial intelligence (AI) and refined machine learning.

AI has been a concept since Alan Turing published a paper in the 1950s that speculated about "Machines that can think". Just over sixty years later we have developed AI that can not only sift through and analyse practically infinite amounts of data, but can also make seemingly intuitive (The world's best player in GO losing to Google Deepmind) and ethical decisions (Self driving cars and the trolley problem, for example). AI has also shifted from hard-coded intelligence to machine learning from data sets and feedback. It is these characteristics of AI that have made it appealing to the world of medicine.

While technology has progressed, so has the science of medicine. We have uncovered enormous amounts of information about different kinds of medicine, their interactions with each other, and their interactions with the human body. We have come to understand the mechanisms that cause disease to spread and propagate and have managed to mathematically model potential outbreaks.

It is clear that bringing AI into medicine could be a significant step for the future of the world. It is impossible for a doctor to learn everything, but it isn't for AI; in fact, analyzing data, summarizing it and providing help is one of its simpler uses. AI can also be used to provide cheap and personalized healthcare which can predict illness based on symptoms and provide treatment recommendations potentially better than modern day doctors. Using AI in tandem with current medical practices and the scientific method could yield some extraordinary results.

DEFINITION OF KEY TERMS

Artificial Intelligence (AI)

Artificial Intelligence is a field in computer science that focuses on the development of machine that can carry out "Intelligent Processes" like decision making, visual recognition and pattern recognition. While being a very broad field its application in medicine is generally limited to more data-based approaches like machine learning.

Machine Learning

Machine learning is a data-based approach to Artificial Intelligence. It uses large data sets to make seemingly cognitive decisions, or predictions. It does this using techniques and methods like deep learning and neural networks which attempt to simulate the way the human brain processes information, but at a larger scale. This ultimately means that machine learning algorithms allow us make predictions and classify data incredibly accurately as it can identify patterns in huge data sets even when they are multidimensional.

Image:

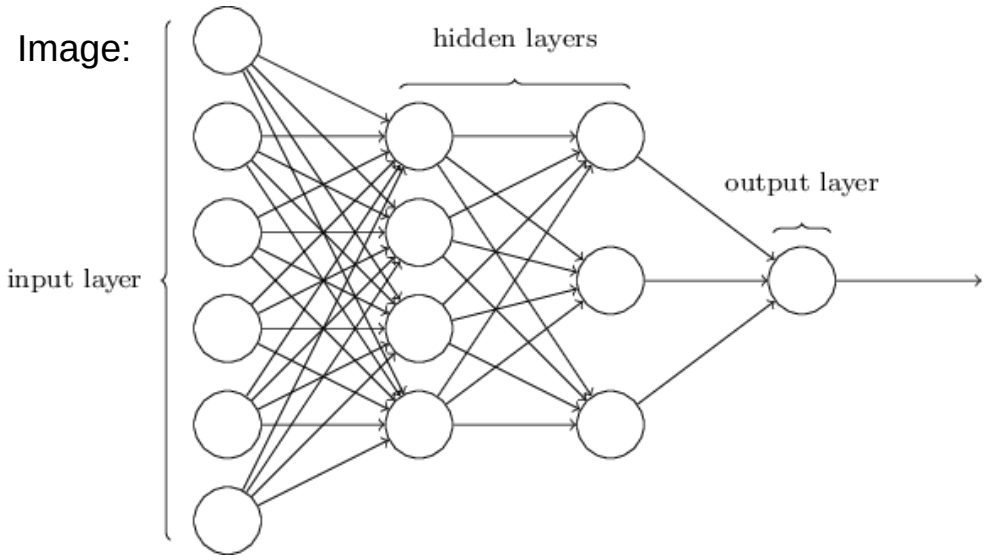

A simplified representation of how Neural Networks work. Data can be input, then analyzed and compared to other data from a data set or patterns the neural network has already identified in order to produce output. For example it may output a diagnosis from a input of a patient's symptoms based on a database symptoms.

Ethics and Morals

Ethics are essentially a set of rules put in place by a individual or even a society that dictate what the right, or just way to act or make a decision is based on a given situation. Making an ethically/morally correct decision is often very hard to quantify and as a result presents an issue for creating artificial intelligence that makes decisions that could affect peoples or patients' lives. As making these kinds of decisions is a fundamental aspect to practicing medicine, how this problem is dealt with is crucial to making the medical community accept many of AI's potential uses.

The trolley problem

The trolley problem is a thought experiment used to illustrate some of the difficulties in ethical decision making. In its simplest version the problem presents a train moving at full speed heading towards five workers on the tracks. You are given the option of diverting the train to take a different path with only one worker; we must assume that the train can't be stopped and that the workers don't realize it is coming and thus do not move out of the way of the train. The point of the problem is to present an ethical choice in whether to divert the train, saving five people's lives, or not, actively playing a role in sending the train towards one person. Naturally this ethical dilemma can be seen in other real life scenarios, such as which two patients should receive a transplant and so on. When we factor in these kinds of moral decisions it's easy to see how utilizing AI in medicine becomes difficult if there isn't a human to make a final call.

Big Data

Big data is a term used to describe gathering and storing huge amounts of data. To classify a data set as being a part of big data we generally refer to the so called "three V's" Volume, Variety and Velocity[1]. The data must arrive in huge amounts from different sources, in different formats and must be processed at a fast pace. We also consider variability and complexity, essentially data having unexpected spikes (for example a sudden surge in the amount of Google searches for a given topic) and all of the data coming from multiple sources with a different correlation and causation behind each source. It's important to note that this data is almost always classified and analyzed using machine learning techniques.

Electronic health records (EHRs)

Electronic health records are a collection of health and medical records on a population or group of people. Each person in that population will have their own electronic medical record (EMR) which provides details about an individual's medical history and other health information, such as their family's medical history or their genetic makeup. EHRs are the most frequent form of big data used in practiced medicine as they provide real-time and updated information that can inform a physician or healthcare providers' decisions for a given patient[2].

[1] "What Is Big Data? | SAS US." https://www.sas.com/en_us/insights/big-data/what-is-big-data.html.
[2] "EHR vs EMR | USF Health Online." https://www.usfhealthonline.com/resources/key-concepts/ehr-vs-emr/.

Black Box Medicine

Black Box Medicine is a term that describes the technologies that combines AI and big data like EHRs. It is a form of personalized health care that allows researchers to analyze huge amounts of data, for example using machine learning to find complex relationships and then suggest a course of action for a given patient. It's important to note that this technology often works by "training" neural networks, this means feeding data to an algorithm, allowing it to make a prediction, then confirming whether that prediction was correct or not. The algorithm then adjusts itself to yield a better result (as seen in Image 1, this would be altering the hidden layers between input and output). As a result, we often don't know what caused the machine to have a certain output without analyzing its hidden layers, something which most private corporations keep secret as it risks exposing their technology to competitors.[3]

Clinical Trials

Clinical trials are the medical communities' method for assessing whether a new drug, delivery mechanism or treatment is safe to be used and prescribed to patients. Clinical trials have four main stages, each testing for different effects on different numbers of people. The first stage is the most experimental, testing for the short-term safety of the drug on about 20-80 people. The second stage tests for effectiveness on about 100-300 people. The third stage tests for both safety and effectiveness on anywhere from a few hundred to a few thousand people. If the third stage is passed, the new drug or treatment will generally get approved for use by the general public while moving on to the 4th stage that monitors patients and the long-term effects. Clinical trials often take into account as much possible about the patients, even their quality of life, by gathering as much data as possible.

Prescription Drugs

Prescription drugs are pharmaceutical drugs that require a physician's recommendation and approval for use and cannot be legally obtained without a prescription.

[3] "Black-Box Medicine by W. Nicholson Price :: SSRN." 3 Apr. 2016, https://www.ssrn.com/abstract=2499885.

BACKGROUND INFORMATION

The Basics of AI: Technology, Industry and Capabilities

For an AI industry to develop in any nation, there needs to be an underlying foundation of already existing technologies and infrastructures to aid its development. Fundamentally large AI applications rely on both big data and large amounts of available computing power at any one point in time. Here we will be discussing how the AI technologies work, how the industry operates and what some of the potential capabilities of AI are.

At its core, AI is used to find patterns and then make educated predictions based on an input of data. Its ability to make accurate predictions depends on the quality and quantity of data provided to the system, the AI's efficiency in translating that data and the amount of computing power that can go into developing the AI. The systems used to develop AI also rely on training and refining the AI and its neural networks. The entire development process can be run several million times quickly if given enough data, time and computing power to do so. AI can also be used to check itself, for example researchers created an AI that generated images of fictional human faces and another that was trained in facial recognition which would determine the quality of the image. By the end of training the AI could produce faces that seemed so realistic that people couldn't tell the difference between fictional and real images. Other examples of training include providing AI-predicted diagnosis to doctors about their patients and then incorporating doctors' feedback into the program. The issue with this technology for utilization in medicine is two-pronged: firstly, we aren't always sure how the AI reached its conclusion or why it reached that conclusion; secondly, it could be fairly easy for false positives to slip past, resulting in constant errors of diagnosis or false causations between data.

The industry that developed these technologies largely started off as academic; around a decade ago, there was no sufficient technology or infrastructure for it. Most of the publicly known information about this technology, its studies and applications come from academic fields and white papers instead of the industry itself. This has occurred as the industry operates in a high degree of secrecy: major private companies and research firms avoid publishing the newest technologies or breakthroughs in order to maintain a competitive advantage. While this is beneficial in stimulating constant competition between private companies to develop the best and most innovative technology, there is certainly an argument to be made that the sharing of new technology and ideas generates more innovation.

This level of secrecy and the issue of the black box is a significant hindrance for utilizing artificial intelligence fully in medicine, particularly in healthcare, due to the desire for high levels of transparency when dealing with human lives. This ultimately raises two critical questions: what risks we are taking by entrusting aspects of healthcare and, by extension, human lives to AI whose creators don't fully understand why it made its decisions? And how do we push the industry to move towards transparency without reducing competition? It is important to note that a medical AI industry does exist and has several successful applications of the technologies as well as good proof of concepts, that are in fact more transparent, but then face issues of user privacy and rights. This presents the next question: what are we risking by giving so much personal and critical data about our health, and to what extent is this sacrifice worth it? It could be argued that the net benefit is greater than the potential risks, but member states are still concerned about data breaches or the AI being misappropriated to target individuals with certain medical conditions/history. Hence, it is important to implement necessary safeguards in order to minimize risk and prevent large scale errors. The answer to the above questions will likely be on a case-by-case basis for each member state, and each individual as we assess the potential risks.

Before we answer any of these questions, it is good to also be aware of some of the capabilities of these technologies. To break this down we can split the capabilities into two main sections: the front- and back-end Healthcare applications and the Research and Development (R&D) applications. The R&D applications rely on our ability to gather and compute data. A large section of the medical industry and intern medicine is R&D, and this ranges from developing new drugs to analyzing new bacteria and viruses, and even to simulating models of disease outbreak to plan a best response for a worst case scenario. All these aspects of R&D can be accelerated and improved upon through AI.

Developing new drugs and testing them involves two parts: synthesis and clinical trials. Utilization of AI here works in two ways: firstly, it can be used to find new chemical compositions to form new drugs by sifting through huge amounts of data, finding new chemical makeups, checking for any recorded interactions between the chemicals used, and testing them. If done efficiently, testing the drug initially could be automated with human oversight, if the drug is deemed to be appropriate for a clinical trial then one will proceed. The second way it could work is within the clinical trial itself: determining adverse effects of a drug could be done more easily if researchers had an AI to sift through medical data of the several thousand patients and try to find patterns or changes across the group. AI is also being used with animal trials: one of the most recent examples uses an AI called DeepSqueak originally developed for image recognition for self-driving cars. It detects ultrasonic sounds from rats from large samples of audio, and then give

researchers an understanding of the rat's condition or emotions in order to help with research into anxiety and depression in humans. From this we can see that, although originally intended for one application, AI can also be used extremely well in another.

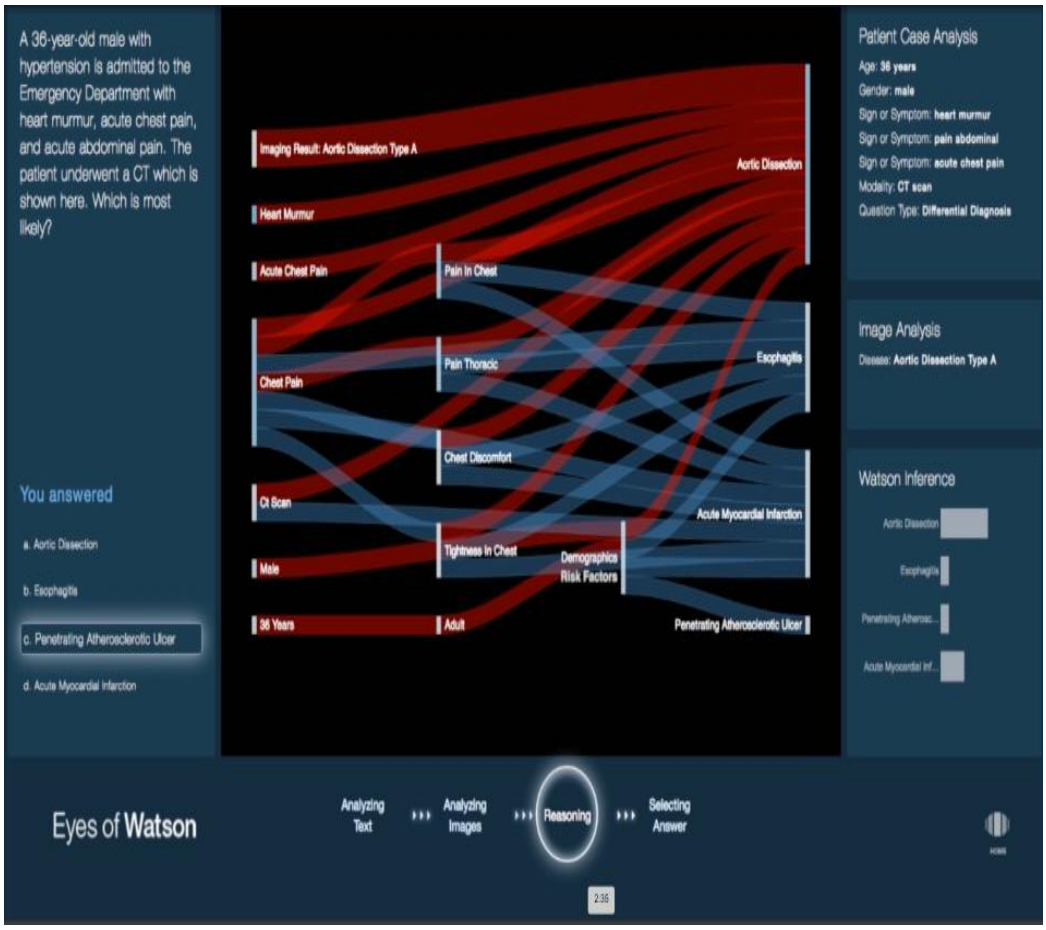

A current example of an AI called Watson developed by IBM being used to diagnose a patient with the help of a doctor or medical student. This is a fairly significant technology as it clearly indicates the AI's reasoning as well as taking input from real human healthcare

The second main aspect of AI in medicine as stated previously is dealing with doctors and patients directly (front-end) and dealing with the managerial work that comes with running a healthcare system like a hospital (backend). Depending on the level at which AI is adopted (something which will be discussed later in the guide), the image of what future healthcare looks like changes. For patients, AI could perform tasks such as analyzing data from personal health electronics (like Fitbits) or referring the patient to specialty doctors based on symptoms. For doctors, its uses are far more complex: AI is already used to help analyze data such as ultrasounds and retinal scans. One form is currently being developed to help in diagnosis and treatment by using databases of previous cases. This has interesting implications for areas of the world where healthcare professionals are in short supply. Taking

advantage of camera technology in smartphones and consumer technology of a similar availability can often provide enough data to inform a patient about whether they should see a medical professional. However, due to ethical and moral issues, this kind of technology will most likely remain in doctors' hands until either the status quo changes or it is further refined. Promoting innovation in this section of AI capabilities would be key to fully utilizing the technology.

Finally, we have the back-end capabilities, which mostly revolve around managerial work. This can be applied in two main ways: firstly, by dealing with administration in hospitals more efficiently, from filing patients' medical results to managing doctors' shifts; secondly, it could be applied at a far larger scale, for example managing the health records for any given country.

Ultimately the technology and industry behind AI directly impacts how and where we can use it. Understanding how both aspects function and relate to AI's capabilities and adoption is fundamental to forming appropriate legislation and frameworks that will not be overly restrictive or lenient regarding its applications.

The Ethical, Moral and Privacy Aspects

Ethics, morals and privacy have always been significant aspects of modern medicine and facing ethical and moral dilemmas is a significant aspect of a clinician's job. Decisions about euthanasia, using experimental treatments, or even taking advantage of new stem cell research all have ethical and moral questions attached to them. This section will explore how moral and ethical decisions impact the implementation of AI in medicine along with issues of providing large amounts of healthcare data.

Firstly, the ethical and moral aspect of AI in healthcare can be illustrated using the trolley problem described earlier. Fundamentally, the trolley problem is a question of picking the lesser of two evils. Incorporating AI into medicine can remove this responsibility that is placed upon clinicians. However, if AI were given this responsibility, who would be held accountable for any errors or oversights? The other aspect of this is accidental built-in biases that may unintentionally occur due to incorrect data being inputted, or an AI assuming causation when there is only correlation. The root issue is to what extent do we want to rely on and use AI to inform decisions about people's lives? On the one hand, due to the sheer amount of data the AI can learn from, technologies will generally make the correct decision once fully developed. They could be equivalent to or better than doctors at diagnosing and prescribing treatments. Therefore, the errors made by the AI are

negligible and a necessary compromise compared to the amount of correct choices it will make. On the other hand, it can be seen as unethical to leave healthcare in the hands of a machine when it is known that there are potential flaws. Ultimately again, the answer is very much based on an individual's ethics. Nonetheless, issues with AI still remain: should AI respond to a patient's time-sensitive condition before a doctor's approval in order to save the patient's life? Is it right to let a machine make decisions that could permanently impact someone's life? Establishing a set of ethical guidelines is therefore key to laying the groundwork for a wide implementation of AI technologies in medicine.

The other main ethical and moral issue is that of black box medicine: how much faith we can have in these machines when we don't know how exactly they work? This issue revolves around our inability to always understand how and why the relationships the AI has found exist, and our need to provide huge amounts of personal data to fully take advantage of this technology. Black box medicine allows us to shift from statistically based models for determining the effectiveness of a treatment or its side effects to far more biologically based models that account for the genetic makeup, lifestyle and medical background of the patient. However, an issue arises due to the possibility of the AI incorrectly relating two factors without us having a understand of why or how. This kind of error can occur simply due to having unrepresentative data. Additionally, patients often want a full account of what treatments to pursue and why, and black box medicine may not be able to provide this. It is important to remember that much of this technology can be useful, it is risky to developing a reliance on potentially unreliable technologies. Implementing legislation to increase transparency would be a significant step in the right direction for implementing AI.

Finally, we have the issue of patient privacy. To fully utilize much of these technologies, we need to take advantage of large amounts of data. Naturally this raises concerns about what happens with this data, how it is stored and who has access to it. Healthcare data is some of the most sensitive data that exists digitally and handing it over to private corporations to analyze has its issues. Firstly, it is practically inefficient as one would need every person to approve of their data for use; secondly, making this kind of data (which includes genetic predispositions and family medical history) available could lead to discrimination by employers or healthcare providers; thirdly, the data would need protection against third parties.

The 3 different adoption possibilities

When looking at the utilization of AI in medicine, we can classify the degree by which AI is utilized in Major, Medium and Minor adoption of the technology.

First off, we have the Major adoption case, where AI is an ordinary part of healthcare for the patient. The major class is one of the two extreme cases, and presents a scenario where AI is widely accepted into society and is used heavily in a lot of the front-end aspects of healthcare. At its most extreme a major adoption case is one where for the most part AI has replaced non-specialized civilians' jobs, and the convenience of getting a diagnosis from a phone an arm's length away trumps the comfort patients feel from dealing with real clinicians. While specialized doctors still exist, patients no longer really have a need to see a general practitioner for an initial diagnosis and instead refer to an AI that has records of their medical history and a significant database of information. The major adoption case hinges on AI's potential to outperform traditional healthcare providers, and when considering the amount of misdiagnosis each year, there is a fair amount of the industry that thinks this might be possible. IBM's Watson is an AI that currently exists that was created with the purposes of interpreting patients symptoms and referring them to the appropriate healthcare provider. Medical AI's efficiency here depends on its ability to understand every drug and analyze huge amounts of research, patient cases and even learn from every recorded medical mistake. Out of the three cases this is the second most likely due to push back from the medical industry, and the current state of the technology itself.

Next, we have the medium adoption case which has doctors cooperating with AI. This is where the current industry is heading and seems to be the most favored case by the medical industry. Here besides AI being used for research it will mainly assist healthcare providers in performing their jobs. For doctors, AI will help by analyzing data from medical technologies and may even play a role in diagnosis. Unlike the major case however, AI's role in diagnosis will be purely suggestive and will simply help advise doctors regarding the best course of action. This is currently where the healthcare industry is heading towards with AI, and it seems optimistic about implementing it with this level of adoption in mind. IBM is currently operating at this level working to help train medical students by analyzing hypothetical cases and displaying the logical steps that lead to conclusions in the case. AI here could also be used in personalized medicine, essentially taking tracked data from patients' smart watches and devices and analyzing it for any discrepancies or early signs of some medical condition and then regularly reporting back to the patient's doctor with a update of how well the patient is recovering.

Finally, we have the minor adoption case, which is essentially what would happen if the medical industry rejected the use of AI. Out of all three cases this is currently the most extreme and unlikely due to the trajectories of both the medical industry and the AI industry as well as the existing technologies. If the minor case occurred, for the most part AI would be used directly with patients and its use in research would be significantly limited. This can also occur if AI in medicine is over-

regulated to a point where innovation is hindered, and ultimately the growth of the industry, significantly limiting where it can be utilized.

Ultimately it would be unreasonable to say that we will definitely end up in one of the three cases and it would be more rational to think of these cases as landmarks across a spectrum of potential outcomes. While considering what legislation and frameworks to put in place we need to keep in mind that the use of this technology is almost inevitable, so the importance of establishing clear procedures for utilizing it is even more important. It's also important to keep in mind that utilizing AI in medicine largely depends on what the medical industry thinks of the technologies being presented to them. Furthermore at the end of the day this is a form of automation and will put white collar workers out of work, so finding a way to minimize that impact is also an aspect of this issue.

MAJOR COUNTRIES AND ORGANISATIONS INVOLVED

The United States:

While having no central policy on AI, the US has almost always been one of the leading technological powerhouses of Artificial Intelligence development and Industry and there is significant potential to begin integrating AI technologies into the medical system. On the one hand the government is investing for military purposes, which may also have some effect on healthcare and on the other hand the private sector involves promoting startups and tech giants such as Google's Deepmind Healthcare.

Despite all the development and clear interest in Medical AI technologies the US government has not taken much action to date in terms of regulation of the market. Ultimately, America's role in effectively utilizing AI in medicine comes down to its potential to regulate key member of the newly developing industry. The methods the US government takes ultimately will set a standard for most western countries with very well-established healthcare systems of how to integrate AI into said healthcare systems.

China

China is easily the second biggest player in the AI market. While as mentioned previously the US does have an overwhelming amount of private and academic support for these kinds of technologies, China has national programs in place to support the development of AI. In 2018 China revealed national plans to

become the world's leading country in AI by 2030, plans which have seemingly become one of China's leadership's top priorities.

The main thing to note would be Google and Microsoft starting up their own labs in China and while both these companies are US-based, having research and development done with in a country even if it is not done by a native company is beneficial due to expertise and training that it brings to employees living in the country (Employees who are likely to move on with their careers in more national companies).

Along with this China also has one of the most interesting use cases for AI. China's rural population has limited access to healthcare, so utilizing AI could greatly improve the living conditions for the affected population.

European Union (EU)

Ultimately the EU is the next big player in the development, integration and use of AI. This comes down to two main factors, the European Union's own push towards developing a European AI industry and the EU's member states push towards developing their own initiatives to fund AI.

The main EU member states who are investing in AI are France and The United Kingdom. Due to their inability to compete with USA and China, the EU is trying to specialize in specific areas and applications. The UK, having decided on focusing on the ethical aspect of AI has invested over $200 million in industry and education. The UK already has companies like Google's Deepmind working with the National Health Service (NHS) on AI applications in medicine as it is as an attempt to lighten the pressure on the NHS.

France is the next major EU member state to be investing in AI, pledging to invest 1.5 billion Euro by the end of 2022. France has chosen to focus on data and how it is used for AI research and applications following the announcement of the French strategy for AI "AI for Humanity". This strategy aims to utilize large centralized databases of France in conjunction with AI, while also focusing on privacy and ethics.

Finally, the EU itself has produced a coordinated plan for AI which follows suit with the directions for the French and British AI programs: developing a "human centric-approach" to AI and its applications. Since ethical use and privacy are two of the main barriers for full utilization of AI in medicine it wouldn't be surprising to see legislation being put in place over the next few years have a significant impact on the direction the industry takes as well as how quickly we move forward.

World Health Organization (WHO)

WHO fundamentally plays the regulatory role in the utilization of AI in medicine, and is tasked with laying out a framework for UN member states to begin to adopt or at least base their frameworks and legislation on. The direction that WHO takes when it comes to AI in medicine will ultimately dictate, not only how quickly AI is adopted but also by who and how. It would be fair to say that WHO has a global impact here, especially on countries that don't necessarily have the most developed AI programs and industries, as they will be following whatever global standard will have been set already, and WHO plays a role in that standard.

IBM Watson Health and Google Deepmind Healthcare

IBM Watson Health and Google Deepmind Healthcare are two of the largest players in the medical AI industry. They set a fundamental example as to how companies AI companies can operate within the medical industry. They're applications and solutions range from diagnosis to big data analysis. Deepmind's focus is diagnosing patients based off of big data patterns (they demonstrated in 2018 that they could detect over 50 eye disease from images better then doctors could), while one of Watson's main focuses is assisting doctors to diagnose and then prescribe treatments. Both companies have had their own issues with transparency, with Deepmind not notifying patients from the NHS that is was using their data, and Watson misdiagnosing hypothetical patients bringing to question its accuracy. This just goes to show that the industry still needs a clear standard in terms of what is ethical, and what is quality treatment and application.

TIMELINE OF EVENTS

Date	Description of event
1950	Alan Turing publishes the first paper to ever propose machine learning and genetic algorithms.
1951	The first Neural network (SNARC) is created based off of Turing's work by Marvin Minsky
1960-1970	Development of the first AI to be used in research, Dendral. Dendral was used to help organic chemists problem solve and analyze data. It essentially laid the foundations for later AI in research.

1994	American Medical Informatics Association released a paper about diagnostic decision and support systems for the past, present and future. Essentially acknowledging the potential of AI in medicine
2009	The Image net project is essentially a database of over 14 million images labeled by humans. it was created in order to help train AI in image recognition. This is significant as image recognition AI is utilized for some of the major applications of AI in medicine.
April 2015	IBM Watson launches Watson Health Cloud. A system that provides access to an enormous amount of computing power and powerful AI for hospitals and healthcare.
October 2015	Google Deepmind's AlphaGo AI beats world's best player at GO (a game considered to require human intuition) about 10 years earlier than predicted. Ultimately achieving a significant milestone for AI and demonstrating some very interesting use cases.
April 2018	EU announces its collective strategy for developing and supporting AI within Europe
May 2018	Apollo Hospitals (one of India's largest hospital companies) agreed to adopt IBM's Watson to help with oncology and Genomics

RELEVANT RESOLUTIONS, TREATIES AND EVENTS

UNHCR report on privacy in the digital age (August 2018):

A report created by the UNHCR that essentially outlines best practices for privacy in the digital age as well as elaborating on the responsibility of private companies' role in protecting privacy.

World Health Assembly resolution on health technologies (WHA60.29, May 27 2007):

A short resolution that was proposed to deal with new health technologies. In particular it looks at economic and technical challenges that introducing new health technology posed to member states.

First "Artificial Intelligence in Medicine" Conference (AIME) (1985):

This was the start of an annual conference that holds tutorials and panels about AI in medicine. It's a good example of how the industry is trying to adapt to or at least learn about this new technology

The Founding of "AI for Good" (2017):

The establishment of an organization committed to using AI to solve some of the world's largest problems and namely help tackle sustainable development. This is a UN organization (technically supported by the ITU), and it provides an example of the kind of NGO bodies that can be made to help deal with some of the issues around AI.

PREVIOUS ATTEMPTS TO SOLVE THE ISSUE

As we begin to look at what actions governments, organizations and companies have already taken to try and improve the utilization of AI in medicine, it's important to keep in mind that we are talking about recent technologies and fairly recent programs/legislation. As a result, a lot of the impact from these actions cannot be clearly seen nor determined due to the scale of the time frame in place (a lot of the actions being taken are part of programs meant to last until 2022 or even 2030).

When tackling this aspect of the issue, we have seen older legislation for other purposes come into play such as the American Health Insurance Portability and Accountability Act (HIPAA) and the European General Data Protection Regulation (GDPR). Both sets of legislation were not intended to deal directly with the privacy concerns of AI directly, but they have taken some effect in defending people's privacy in these sorts of cases. HIPAA at its core is legislation that safeguards medical information and therefore has indirect implications regarding AI. The GDPR is focused on the protection of individuals' data, and this therefore extends to a person's medical data and therefore to an extent to AI. Of course, legislation that does deal with some of the more intricate aspects of medical big data and black box medicine in terms of individual privacy and rights is key, and currently not a huge amount has been done to propose legislation that does.

So far, the closest we've gotten to legislation that deals with these intricacies is that of the UK, France and the EU mentioned before. Out of the three, France's

program is certainly the most robust and sets a good example of the kind of work we need for AI in medicine to be properly regulated so as to improve the public's acceptance of it; in fact the program includes direct references to dealing with black box medicine and a long term plan for improving the general public's comfort around AI. The EU's legislation on the other hand simply pushes for a wider adoption of AI across Europe and stipulates that every European country should have their own plan for AI by 2020's.

Along with this we have seen other steps being taken to help increase doctors' and patients' comfort and ease when using AI technologies for medical applications. One of these steps was the support for "user-based" design from the American Medical association discussed earlier. This was put forward in order to ensure AI companies make software that's easy to integrate into everyday practices is a simple way to remove some of the barrier to entry a lot of practicing clinicians might feel when then consider using this technology. Over the course of development working closely with practicing doctors and people in the medical fields seems to be fundamental to the way they develop the interfaces doctors interact with.

POSSIBLE SOLUTIONS

While we begin to look at solving this problem, we need to approach it by creating new innovative solutions while also building and reinforcing a lot of the already existing infrastructure we have in place. Since we are dealing with a problem of increasing utilization, we need to consider the stakeholder and their incentives to take advantage of this technology; ultimately appealing to them (doctors, patients, researches) would be the most effective way of utilizing the technology, as if they actively want to make use of the technology it will become an industry standard in healthcare far faster.

The first main issue we should look at is that of the ethics around black box medicine. Here we need to tackle the concerns of both patients and doctors. Throughout the study guide we have looked at how AI could make errors or have biases. One of the main ways to deal with concerns around these aspects of AI is by creating dedicated legislation and a dedicated framework for checking and testing AI. One example of this could be having each country create its own board to review medical AI and check whether it works alongside a set of global requirements/rules. Along with this, imposing global or even local legislation that establishes the roles of responsibility while using medical AI would be key to easing doctors' minds. Implementing processes that regularly check the state of an AI would also be a decent measure as it ensures a quality of care and eliminates risks for doctors,

patients and developers. Another option is increasing transparency regarding the AIs' workings, which would harm competitiveness but promote research.

Along with this, tackling the issue of privacy and security of medical information is key to making this technology more widely adopted. Finding ways to enforce encryption for data or setting up rules and regulations about what kind of data can be used by companies is one way to tackle this issue. It would be important to distinguish between companies that are using AI for consumer/patient medical purposes and those who are using it for research. So far companies like IBM have offered to pay for data like this to compensate patients for using their private data, which could also be an interesting route to take when tackling this aspect of the issue.

Besides addressing concerns of patients and doctors, setting up programs and initiatives to increase the rate at which this technology is adopted would be key to reducing barriers for utilization. The first way this could be done is through direct contact with doctors, training them on how to use the technology, introducing them to how it works and what its limitations are and so on. While putting programs like this in place would be useful, it is important to keep in mind that we don't want to build an overdependence on this technology as a whole.

Furthermore, we can also look at overall implementation of this technology on a far larger scale. Generally providing incentives for doctors, hospitals and healthcare providers would be one way to increase utilization. Both hospitals and insurance companies are likely to use these technologies to streamline their work flows, something which would ultimately lead to job loss. Dealing with the fallout from this scale of job loss due to automation as well as the backlash towards the technology is one of the fundamental aspects of ensuring that the utilization of this technology lasts. Also looking at how to incentivize AI to be adopted only for the sake of better care instead economic benefits would be important as at the end of the day we do want to improve general health care with the use of this technology instead of simply cutting corners.

Finally, it is worth mentioning that safeguards should be a part of implementation. Creating plans or procedures for when and if something goes wrong with an AI is fundamental to ensuring security and an overall acceptance of the new technology. We need to ensure that these technologies first of all, are not misused but more importantly don't malfunction (and release personal information, give patients the wrong treatment or misrepresent data) and so setting up guidelines for safeguards against this is important as we want to increase the acceptance of this technology and decrease the public's fear of it.

In conclusion, increasing the implementation and use of this technology needs to deal with the multifaceted issues of public opinion, ethical use, privacy and even protections.

BIBLIOGRAPHY

"AMA Passes First Policy Recommendations on Augmented Intelligence." *American Medical Association*, www.ama-assn.org/press-center/press-releases/ama-passes-first-policy-recommendations-augmented-intelligence

"Artificial Intelligence in Medicine." *Journal*, www.journals.elsevier.com/artificial-intelligence-in-medicine

Buch, Varun H, et al. "Artificial Intelligence in Medicine: Current Trends and Future Possibilities." *British Journal of General Practice*, British Journal of General Practice, 1 Mar. 2018, www.bjgp.org/content/68/668/143

Derrington, Dolores. *Artificial Intelligence for Health and Health Care*. Dec. 2017, www.healthit.gov/sites/default/files/jsr-17-task-002_aiforhealthandhealthcare12122017.pdf

European Union. "Coordinated Plan on Artificial Intelligence." *Digital Single Market - European Commission*, 7 Dec. 2018, www.ec.europa.eu/digital-single-market/en/news/coordinated-plan-artificial-intelligence

Ford, Allan, and W. Nicholson Price II. "Privacy and Accountability in Black-Box Medicine." *University of Michigan Law School Scholarship Repository*, 14 Feb. 2017, https://repository.law.umich.edu/mttlr/vol23/iss1/1/

French Government. “AI for Humanity.” *AI for Humanity*, 2018, www.aiforhumanity.fr/en/#s2

Gershgorn, Dave. “AI Is the New Space Race. Here's What the Biggest Countries Are Doing.” *Quartz*, Quartz, 1 May 2018, www.qz.com/1264673/ai-is-the-new-space-race-heres-what-the-biggest-countries-are-doing/

Hooijdonk, Richard van. “Artificial Intelligence in Medicine: Current Trends and Future Possibilities | Healthcare Tech.” *Clinician Today*, 17 Dec. 2018, www.cliniciantoday.com/artificial-intelligence-in-medicine-current-trends-and-future-possibilities/

“How Will A.I. Change Medicine?” *Futurity*, 14 Dec. 2018, www.futurity.org/artificial-intelligence-medicine-1932952/

“IBM Watson.” *IBM Watson*, IBM, www.ibm.com/watson

Ijaz, Rehan, et al. “10 Most Innovative AI Companies That Are Changing the Healthcare Market.” *Big Data Made Simple*, 30 Nov. 2018, www.bigdata-madesimple.com/10-most-innovative-ai-companies-that-are-changing-the-healthcare-market/

Mendeley Careers. “Artificial Intelligence in Medicine.” *Mendeley Careers*, Mendeley Careers, 2 July 2018, www.mendeley.com/careers/article/artificial-intelligence-in-medicine/

Nicholson, W. “Black-Box Medicine.” *SSRN*, 23 Sept. 2014, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2499885

Oxtoby, Kathy. “Is the Hippocratic Oath Still Relevant to Practising Doctors Today?” *The BMJ*, British Medical Journal Publishing Group, 14 Dec. 2016, www.bmj.com/content/355/bmj.i6629

Price II, W. Nicholson. “Black Box Medicine.” *The Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics at Harvard Law School*, 2015, https://petrieflom.law.harvard.edu/resources/article/black-box-medicine

Price, W. Nicholson. “Regulating Black-Box Medicine.” *Michigan Law Review*, U.S. National Library of Medicine, 2017, www.ncbi.nlm.nih.gov/pubmed/29240330

Shadoan, Rachel. “Reasoning About Opaque Algorithms.” *The Recompiler*, 23 Aug. 2016, https://recompilermag.com/issues/issue-2/reasoning-about-opaque-algorithms/

Teich, David A. “Artificial Intelligence (AI), Healthcare and Regulatory Compliance.” *Forbes*, Forbes Magazine, 2 July 2018, www.forbes.com/sites/tiriasresearch/2018/03/20/artificial-intelligence-ai-healthcare-and-regulatory-compliance/#5cf2dec74f66

“The Application Of AI To Augment Physicians And Reduce Burnout.” *The Application Of AI To Augment Physicians And Reduce Burnout | Health Affairs*, www.healthaffairs.org/do/10.1377/hblog20180914.711688/full/

“Top Artificial Intelligence Companies in Healthcare to Keep an Eye On.” *The Medical Futurist*, 20 Sept. 2018, https://medicalfuturist.com/top-artificial-intelligence-companies-in-healthcare

United Kingdom, House of Lords: *AI in the UK: Ready, Willing and Able?* 16 Apr. 2018, https://publications.parliament.uk/pa/ld201719/ldselect/ldai/100/100.pdf

Vincent, James. “DeepMind's AI Can Detect over 50 Eye Diseases as Accurately as a Doctor.” *The Verge*, The Verge, 13 Aug. 2018, www.theverge.com/2018/8/13/17670156/deepmind-ai-eye-disease-doctor-moorfields

"What Is Big Data and Why It Matters." *What Is Big Data? | SAS US*, www.sas.com/en_us/insights/big-data/what-is-big-data.html

Image 1:

Nielsen, Michael. "Neural Networks and Deep Learning." *Neural Networks and Deep Learning*, Determination Press, Oct. 2018, https://neuralnetworksanddeeplearning.com/chap1.html

Image 2:

Singh, Manprit. *IBM Watson The Most Comprehensive Enterprise AI Platform*. IBM, 2017.