"Checkmate – The Future is Here"

Discuss the role of AI in medicine. Is it possible to replace human doctors or any aspect of medical care with robots?

Apimaan Yogeswaran
Medical Student, University of East Anglia
Word Count (excluding title page and references = 1500)

The Unofficial Guide to Medicine Essay Competition

Introduction

Our NHS is struggling with funding cuts, rising demands, and limited resources. There is a shortage of staff, with a reported excess of 100,000 vacancies, across the NHS trusts[1]. Performance regularly falls below national standards – news headlines often report that the four-hour A&E waiting time target has not been met since July 2015[2]. Breaching commitments made to its patients in the NHS constitution[3], the future of the health service has come under scrutiny. However, while most discussion for saving our NHS focuses on improved funding and minimising waste of services, another solution has the potential to entirely turn health infrastructure on its head.

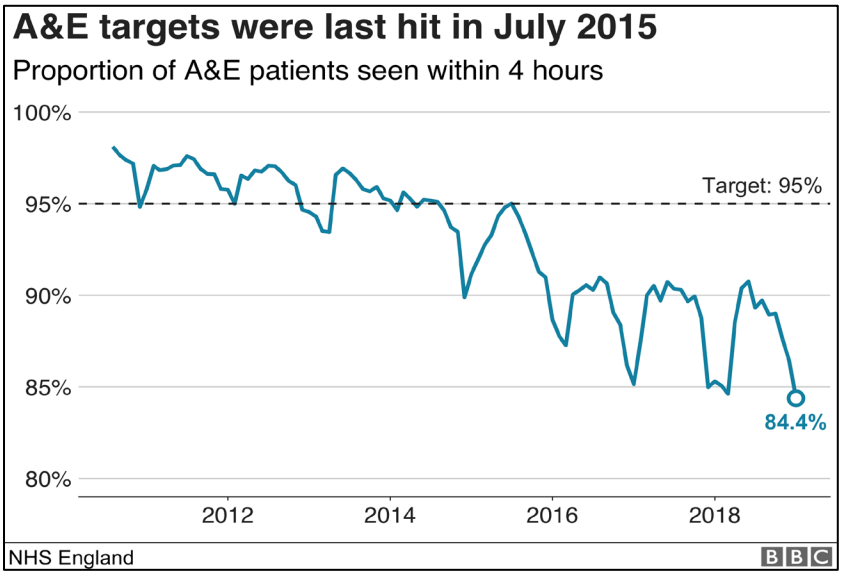

Figure 1

In 1950, the early pioneer of artificial intelligence (AI), Alan Turing, published his seminal paper detailing his famous 'Imitation Game', the first milestone for the question 'Can machines think?'[4]. Now In 2019, computers and AI seemingly still have not become sentient. Not yet. However, during the 69 years in between, software analysts and developers have found other more beneficial (and less intimidating) uses for AI, including in finance, e-commerce, and education[5]. Most dramatic is the rise in popularity of AI projects in healthcare, attracting more investment in the market than any other sector of the global economy, accruing $1.8B collectively since 2012[6,7]. From supporting diagnosis of individual cases to co-ordinating large-scale triage, this essay will discuss the current and potential applications of artificial intelligence in healthcare, and if at all, to what degree it could replace its human workforce.

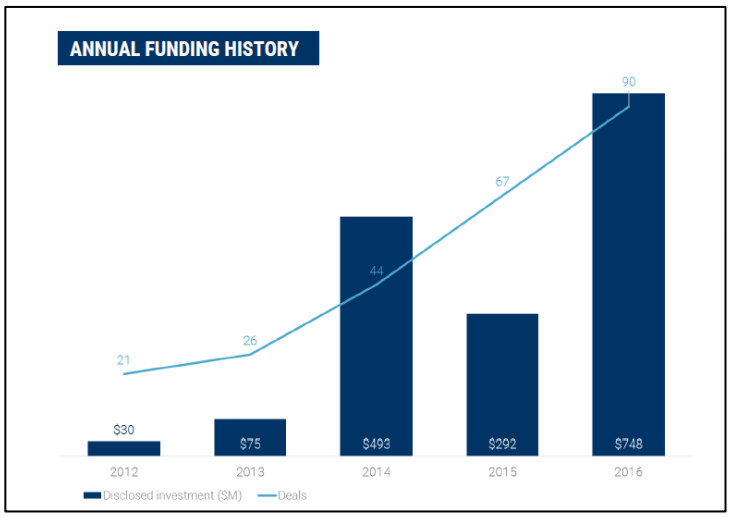

Figure 2

Beyond Human Limits

Checkmate. Before the turn of the 21st century, IBM had developed Deep Blue, a supercomputer powerful enough to defeat the world chess champion, Garry Kasparov, in a six-game match[8]. Though it is a ubiquitous norm in the modern age, it was unheard of for AI to be sophisticated enough to undertake complex tasks, such as driving cars, trading in the financial market, or even playing chess[9]. So, what is the connection between chess and healthcare? Quite a lot it seems. IBM (International Business Machines), creators of Deep Blue, have more recently created an AI programme called Watson, one of several advanced clinical-decision support systems (i.e. AI healthcare assistants)[9]. And the similarities do not end there. These programmes are centralised around the genius mechanism that first allowed AI to surpass human limits – 'machine learning' (ML).

Machine Learning

Since the technical aspects of ML are beyond the scope of this essay, a simplified analogy would be that Machine Learning is to AI, what evidence-based medicine is to a doctor[10,11]. The core principle behind evidence-based medicine is to incorporate past health data to inform current clinical practice. For a human doctor, this would perhaps be to assimilate all their knowledge of contemporary research, theory, and clinical experience and apply this to a consultation with a patient. However, this process is restricted by the individual doctor's own subjective experience. Machine Learning bypasses this human constraint entirely with its processing power. To add perspective to this advantage, in a game of chess, Deep Blue could process 200 million moves per second, while a human could at most process one move every 2-3 seconds[12].

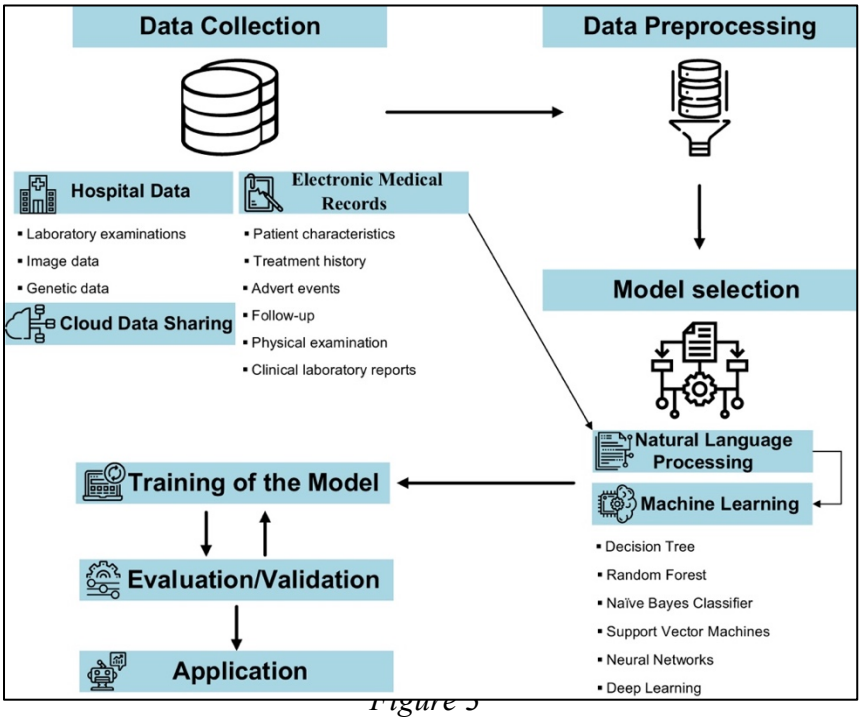

Figure 3

As the diagram above depicts, a system like Watson could have access to encyclopaedias worth of knowledge about illnesses, medicines and hospital data (laboratory results, CT scans etc.) to refine its learning algorithms - as part of its 'data collection'. For its 'model training & selection', the system could then be exposed to more cases (via previous patient data) than a doctor could in several lifetimes – permitting it to provide accurate diagnosis, prognosis, and recommended

treatment specific to a particular patient. All this, in only a matter of minutes[9]. Without the inherent biases that could throw a doctor off, nor the limits of a patient's memory, Watson would be able to bring prompt solutions to any clinical setting anywhere, available 24/7. Not only would AI be able to optimise and personalise management, as well as prevent all manner of clinical errors. But it could also be a key solution to ease the pressure of high caseloads and staff shortages on NHS staff– not through replacement but assisting them to perform more efficiently and expediently.

Data Security

ML brings with it a fundamental ethical concern. As Peter Eckersley (chief computer scientist of the Electronic Frontier Foundation) cautions, "most machine learning architectures require customers to hand huge amounts of sensitive data to a vendor in order to train classifiers"[13]. In order for deep learning algorithms to be effective in real-life scenarios, they require patient data. Lots of it. Raising concerns of data security and privacy, patients may find this unacceptable. In a recent survey, 23% of responders expressed concern specifically for sharing their data with an AI programme[14]. This comes as no surprise. After all, with sensitive patient information at the mercy of cyber-attacks, hacking, and system errors, could patients really trust their lives to be uploaded to the digital unknown?

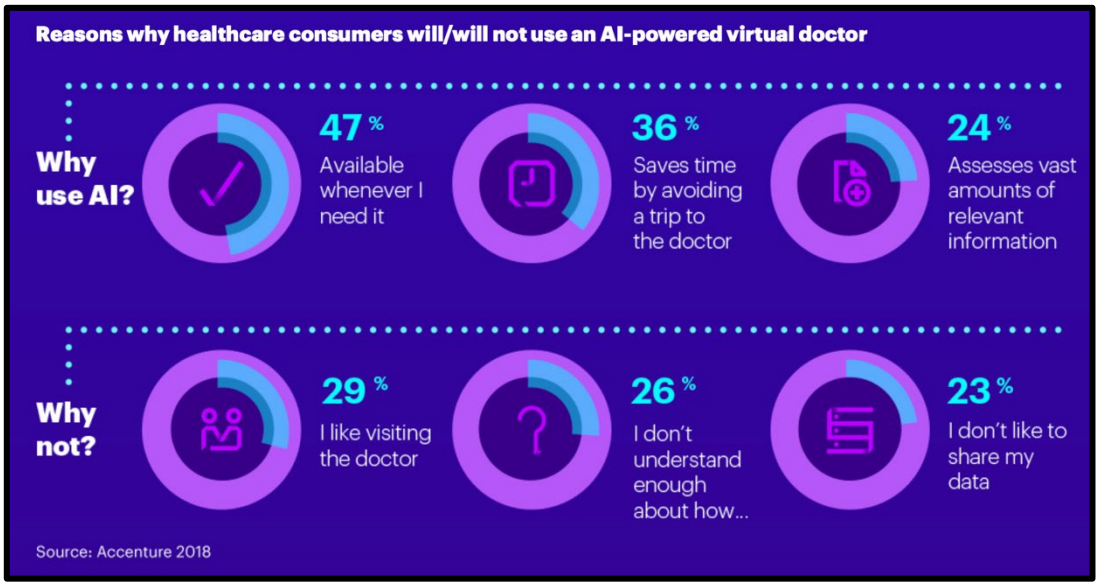

Figure 4

Human Touch

Another argument against a virtual doctor is that an AI could never wholly replace a human doctor. In the aforementioned survey, 29% of responders preferred to visit their doctor[14], which reflects how the face-to-face human interaction is as valued as the diagnosis or treatment that a doctor provides. A recent statistic also showed that at least 50% of physicians felt that this interaction is crucial for patient engagement[15], thus taking this dynamic away from the doctor-patient relationship could make AI implementation more of a detrimental barrier than an asset. Ask yourself this - would an AI be able to console you during grief as well as a doctor would? Would you be comfortable consulting an inorganic machine about a biological phenomenon, like pregnancy? Though software developers are looking into incorporating empathy into algorithms[16], we are realistically far from such sophisticated technology (a positive Turing Test, no less).

Automatization

Taking a step away from the individual patient-doctor consultation, this section looks at the wider impact AI may have on health infrastructure, as a whole.

Natural Language Processing

A system is only as strong as its individual working units, as is true in healthcare. Yet currently, there is a significant rate of burn-out among physicians ranging between 37-58% across all specialities, especially among those in general

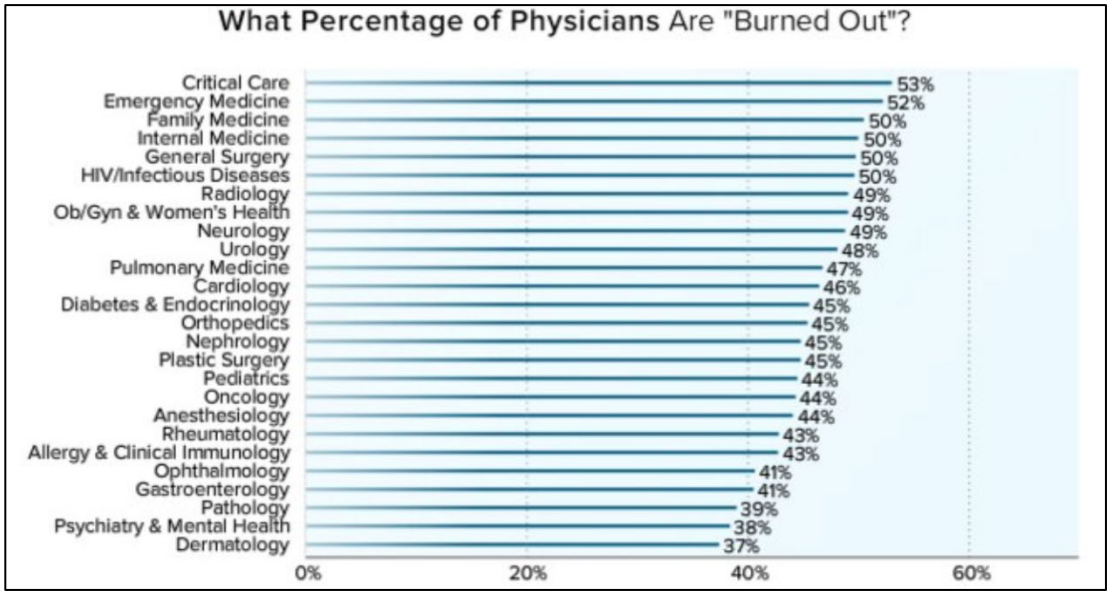

surgery and critical care[17]. One of the chief causes is the move to electronic medical records (EMRs). Ultimately, due to unintuitive user interfaces, clinicians must wade through data in order to obtain the information they need[18] – often staring at screens while conversing with patients. Natural Language Processing is a subset of AI that uses computer algorithms to cluster, extract and process information from this unstructured clinical data (EMRs, clinical laboratory reports etc.) to produce meaningful information for doctors to use. Streamlining health data in this way, like ML, could alleviate much demand on the individual doctor, and consequently the entire system too.

Specialised Tasks

Artificial Intelligence is found to excel in highly focused tasks, making it ideal for routine tasks such as "EMR documentation, administrative reporting, even triaging CXR scans"[19]. Through, the use of ML, there is evidence of efficacy among AI-based systems. Separate AI applications were capable of classifying suspicious skin[20], as well as pulmonary TB on chest X-Rays (with a sensitivity of 95%)[21] – both better than human specialists in the respective fields. This extends further to broader tasks, such as automated dispensing and large-scale triage. A fully automated robot pharmacy in San Francisco has not made a single error since its induction in 2011[22]. And c

Figure 5

ordina m

and for the wider community, artificial intelligence is competent in an extensive range of applications.

However, there is a risk of de-skilling that comes with this. As AI becomes more competent in a wider variety of tasks, while clinicians neglect their own clinical expertise, this leaves one fundamental question. Would there be any need for human doctors?

Conclusion

This brings us nicely back to the question; 'would AI be able to replace humans in the world of medicine?'. Researchers from Oxford and Yale predict that by 2053, fields such as surgery, pathology and radiology (thought to be the most automatable specialties) will be among the first to take the technological leap[24,25].Some may argue that with rising populations and limited resources, it's only natural that healthcare itself evolves in order to cope with these demands. This isn't hard fact though, and it's apparent that there are several ethical and practical obstacles before this future is realised.

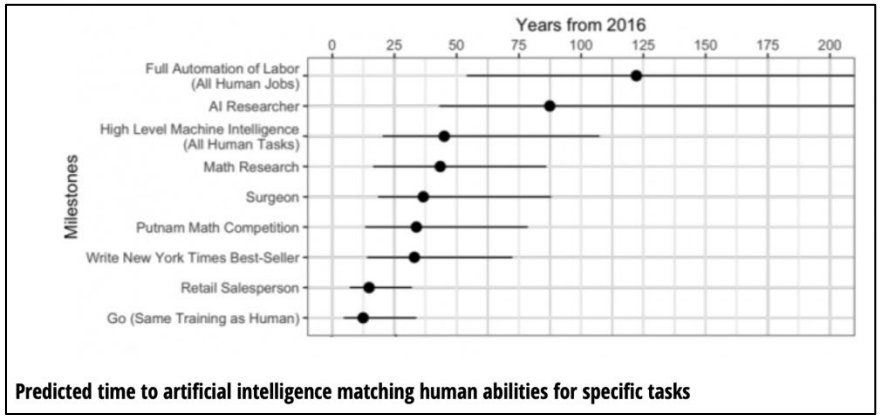

Figure 6

On a positive note however, AI has evolved more as an assistive technology (e.g. 'chess.com'), that enriches and trains their human counterparts. People seek the human touch in the game in the same way they do in healthcare. And so, it seems realistic that healthcare AI will better adopt a strong assistive rather than replace us. At least, in this lifetime.

****

References:

1. Murray R. The King's Fund responds to the latest NHS England hospital performance figures [Internet]. The King's Fund. 2019 [cited 31 December 2019]. Available from: https://www.kingsfund.org.uk/press/press-releases/kings-fund-responds-latest-nhs-england-hospital-performance-figures
2. NHS waiting times [Internet]. The King's Fund. 2019 [cited 31 December 2019]. Available from: https://www.kingsfund.org.uk/projects/positions/nhs-waiting-times
3. NHS Constitution for England [Internet]. GOV.UK. 2019 [cited 31 December 2019]. Available from: https://www.gov.uk/government/publications/the-nhs-constitution-for-england

4. TURING A. I.—COMPUTING MACHINERY AND INTELLIGENCE. Mind. 1950;LIX(236):433-460.
5. Application of AI - Javatpoint [Internet]. www.javatpoint.com. 2019 [cited 31 December 2019]. Available from: https://www.javatpoint.com/application-of-ai
6. Buch V, Ahmed I, Maruthappu M. Artificial intelligence in medicine: current trends and future possibilities. British Journal of General Practice. 2018;68(668):143-144.
7. Healthcare Remains The Hottest AI Category For Deals [Internet]. CB Insights Research. 2019 [cited 31 December 2019]. Available from: https://www.cbinsights.com/research/artificial-intelligence-healthcare-startups-investors/
8. Anderson M. Twenty years on from Deep Blue vs Kasparov: how a chess match started the big data revolution [Internet]. The Conversation. 2019 [cited 31 December 2019]. Available from: https://theconversation.com/twenty-years-on-from-deep-blue-vs-kasparov-how-a-chess-match-started-the-big-data-revolution-76882
9. Harari Y. Homo Deus: A Brief History of Tomorrow. Harvill Secker; 2015.
10. Simeone O. A Very Brief Introduction to Machine Learning With Applications to Communication Systems. IEEE Transactions on Cognitive Communications and Networking. 2018;4(4):648-664.
11. Wang R, Pan W, Jin L, Li Y, Geng Y, Gao C et al. Artificial intelligence in reproductive medicine. Reproduction. 2019;158(4):R139-R154.
12. Press G. The Brute Force Of IBM Deep Blue And Google DeepMind [Internet]. Forbes.com. 2019 [cited 31 December 2019]. Available from: https://www.forbes.com/sites/gilpress/2018/02/07/the-brute-force-of-deep-blue-and-deep-learning/#1cb9153249e3
13. Artificial Intelligence in Data Protection - Experts Predict future of AI to prevent data breaches [Internet]. Egress.com. 2019 [cited 31 December 2019]. Available from: https://www.egress.com/artificial-intelligence-for-data-protection
14. 2018 Consumer Survey on Digital Health | Accenture [Internet]. Accenture.com. 2019 [cited 31 December 2019]. Available from: https://www.accenture.com/us-en/insight-new-2018-consumer-survey-digital-health?utm_source=newsletter&utm_medium=email&utm_campaign=newsletter_axiosvitals&stream=top-stories
15. More Face-to-Face Interaction Is Key to Patient Engagement [Internet]. PatientEngagementHIT. 2019 [cited 31 December 2019]. Available from: https://patientengagementhit.com/news/more-face-to-face-interaction-is-key-to-patient-engagement

16. Empathy in Artificial Intelligence [Internet]. Forbes.com. 2019 [cited 31 December 2019]. Available from: https://www.forbes.com/sites/cognitiveworld/2019/12/17/empathy-in-artificial-intelligence/#839e5cd63270
17. 7. Burnout Rises Above 50% in Some Specialties, New Survey Shows [Internet]. Medscape. 2019 [cited 31 December 2019]. Available from: https://www.medscape.com/viewarticle/907834
18. Stress from using electronic health records is linked to physician burnout [Internet]. ScienceDaily. 2019 [cited 31 December 2019]. Available from: https://www.sciencedaily.com/releases/2018/12/181205134029.htm
19. Bresnick J. Arguing the Pros and Cons of Artificial Intelligence in Healthcare [Internet]. HealthITAnalytics. 2019 [cited 31 December 2019]. Available from: https://healthitanalytics.com/news/arguing-the-pros-and-cons-of-artificial-intelligence-in-healthcare
20. Esteva A, Kuprel B, Novoa R, Ko J, Swetter S, Blau H et al. Dermatologist-level classification of skin cancer with deep neural networks. Nature. 2017;542(7639):115-118.
21. Lakhani P, Sundaram B. Deep Learning at Chest Radiography: Automated Classification of Pulmonary Tuberculosis by Using Convolutional Neural Networks. Radiology. 2017;284(2):574-582.
22. UCSF Robotic Pharmacy, San Francisco - Pharmaceutical Technology [Internet]. Pharmaceutical-technology.com. 2019 [cited 31 December 2019]. Available from: https://www.pharmaceutical-technology.com/projects/ucsf-robotic-pharmacy-san-francisco/
23. Middleton K, Butt M, Hammerla N, Hamblin S, Mehta K, Parsa A. Sorting out symptoms: design and evaluation of the 'babylon check' automated triage system [Internet]. arXiv.org. 2016 [cited 31 December 2019]. Available from: https://arxiv.org/abs/1606.02041
24. Grace K, Salvatier J, Dafoe A, Zhang B, Evans O. Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts. Journal of Artificial Intelligence Research. 2018;62:729-754.
25. Bresnick J. Artificial Intelligence Could Take Over Surgical Jobs by 2053 [Internet]. HealthITAnalytics. 2019 [cited 31 December 2019]. Available from: https://healthitanalytics.com/news/artificial-intelligence-could-take-over-surgical-jobs-by-2053