STUDENT NAME (OPTIONAL)	Claire Thompson
STUDENT ID (REQUIRED):	51659440
COURSE CODE/NAME:	**ME33A ALTERNATIVE PROJECT**
COURSE TUTOR:	Dr Leeanne Bodkin
PROJECT SUPERVISOR:	DR HEATHER MORGAN
TYPE OF ASSESSMENT:	PROJECT
PROJECT TITLE:	BARRIERS TO THE GROWTH OF ARTIFICIAL INTELLIGENCE WITHIN THE NHS
	WORD COUNT: 3539

HAS AN EXTENSION BEEN GRANTED? ☐

If yes, please ensure grant of extension form is securely attached

PLAGIARISM STATEMENT

"Plagiarism is the use, without adequate acknowledgement, of the intellectual work of another person in work submitted for assessment. A student cannot be found to have committed plagiarism where it can be shown that the student has taken all reasonable care to avoid representing the work of others as his own."

I hereby acknowledge that I have read and understood the above definition of plagiarism. I declare that all material from other sources used in this piece of assessed work, whether directly quoted or paraphrased, has been clearly identified and attributed to the source from which it came.

SIGNATURE: *Cerni Thompson*

_____

DUE DATE:	DATE SUBMITTED (office use only)
1/12/17	

Barriers to the growth of artificial intelligence within the NHS

Introduction

Artificial intelligence (AI), regarded as the ability of systems to replicate human behaviour in an intelligent manner, shows great promise within the NHS. Increased automation of repetitive clinical and administrative tasks has the potential to save the NHS up to £12.5 billion a year worth of staff time (1) which could be invaluable to the NHS in addressing increasing pressures to meet financial and performance targets. Although artificial intelligence has been around for many years, advancements in AI methods in recent decades has resulted in increasingly sophisticated programs with the ability to revolutionise industries such as healthcare. Specifically, breakthroughs in machine learning, one of the most widely used branches of AI has led to algorithms capable of equalising the diagnostic skills of doctors, automating administrative tasks and assisting in complex treatment management to name a few. Within the NHS, some examples of the successful applications of AI can be seen such as the collaboration between British AI company DeepMind's and Moorfields eye hospital in accurately detecting and prioritising eye conditions (*appendix A)* (2) and Heart flow (3), an image analysis software using AI to construct detailed images of a patient's heart *(appendix B)*. From an administrative perspective, East Suffolk and North Essex NHS foundation trust is currently deploying a system using AI to directly tackle tasks such as the management of GP referrals and processing of blood and scan results. This system has already saved staff 570 hours of work in the space of three months and therefore shows the promise of automating repetitive admin tasks which currently demand substantial amounts of staff time (4). Although these, amongst others, are extremely promising examples of potential applications of AI within the NHS, it is important to note that the majority of these are trialled or used in practice within individual NHS boards rather than nationally. The UK government has stated that the NHS has the potential to become a world leader in AI (5) with advancements potentially transforming the prevention, early diagnosis and treatment of diseases. Despite this, the widespread use of artificial intelligence within the NHS evidently has not yet been achieved. The various factors which have been identified in preventing the growth of artificial intelligence within the NHS will be discussed within this project.

Technological advancements within AI

Broadly speaking, artificial intelligence is a field within computer science which aims to develop computer programs that demonstrate features of human intelligence. This includes the capacity to learn, plan, reason and self-correct. Although several branches of AI are currently used in technology today, one branch that has proven to be particularly prominent over the past decade is machine learning (ML), a method of analysing data which allows software to learn without being explicitly programmed *(appendix C)*. Programs initially learn from being exposed to large sets of data, identify features from this data and analyse this to make predictions about information it has not previously been exposed to. This recent growth in machine learning is partly due to the rise of the internet with large datasets being more readily available than ever before in addition to the recognition of the increased efficiency of the self-teaching nature of new ML methods such as deep learning. Machine learning systems are notably efficient in tasks involving the analysis of

large datasets of images as well as the automation of repetitive processes and therefore could be particularly advantageous in healthcare.

Despite its advantages, machine learning can pose multiple problems with regards to regulation, transparency and data usage. The most efficient and newest models of ML known as deep learning methods, are black box in nature meaning it is not always clear to the user or even the developer how exactly the program reaches a certain answer. This is acceptable in some industries, however, it can prove a considerable problem for regulation and accountability in safety critical industries. Additionally, the use of existing and 'real time' NHS data, which would be necessary to train algorithms specific for healthcare in the UK, raises ethical issues with regards to appropriate management, including collection, storage and distribution of the data. Despite these concerns, artificial intelligence and machine learning promise huge advantages within healthcare in terms of financial savings and improved clinical practice and so with careful navigation of some key barriers, it might be possible to develop appropriate models that sufficiently protect and support developers, NHS staff and the public in the increasing use of AI.

Current barriers preventing the growth of artificial intelligence within the NHS

The use of public data in training AI algorithms

The NHS collects and records vast amounts of data on a daily basis providing a valuable opportunity in AI where, as previously mentioned, methods such as machine learning require training by large datasets. As stated by Lord James O'Shaughnessy, NHS data "is an asset that no other country can bring together" (6) and utilising this appropriately is key to the successful adoption of artificial intelligence within the NHS.

Although the volume of data required for developing AI currently exists within the NHS, it is of equal importance that this data is of sufficient quality to ensure algorithms are accurately trained. According to a recent report carried out by NHS England and the AHSN network, however, respondents considered current NHS datasets patchy, dirty and often incorrect (7) and overall of an inadequate quality for use in AI training. Additionally, a research scientist at British AI company DeepMind stated that because of uncleaned, unrepresentative and disconnected NHS datasets it took several months to prepare relevant data before work could even begin on one of their diagnostic AI applications (8) proving to be a major inconvenience in the early stages of development. In addressing this, it would seem appropriate that it is necessary to establish a digital framework holding datasets appropriately prepared for AI use that can be easily accessed by approved developers. One of the first steps in achieving this, however, is the digitalisation of all NHS data which is a step that the NHS has struggled with across the board. Despite current targets claiming a paperless NHS will be achieved by 2020, this is now looking as overly ambitious as the previous target of 2018 with some experts predicting this cannot feasibly be achieved before 2027(9). Even after digitilisation, the requirement for data to be stored in an appropriate format ensuring data is machine readable, accurate and "clean" needs to be considered. The performance of ML algorithms is reliant on the information that is fed to them – if this information is incomplete, or contains inaccuracies, then this will be reflected in its output with potentially catastrophic consequences. Although the development of an appropriate digital infrastructure may be a challenge both in terms of labour and finances, the

significant time dedication for data organisation, which is currently required at the start of most AI developments, is currently proving to be a major impedance to developers and therefore is crucial to address.

In addition to the practical requirements mentioned above, the use of large sets of sensitive public information, carries the potential for inappropriate use and must be subject to relevant ethical and legal considerations by both developers and NHS trusts. This is particularly relevant for AI which, in some applications, use large datasets that even when anonymised contain significant amount of unique identifiers that could be linked to data subjects. As stated in a recent report carried out by KPMG, 51% of the public are currently concerned about their data privacy as the use of AI increases and this is particularly relevant for those with minimal knowledge of current AI technology (10). Although appropriate legal (e.g. The Data Protection Act 2018) and ethical frameworks (e.g. The Data Ethics Framework) are in place designed to safeguard patients, the growth of machine learning and its absolute need for real life data could present an ethical dilemma where public data is being used as an exploitable resource utilised for purposes outwith what it was originally collected for. An example of this has already occurred where Royal Free NHS trust failed to obtain appropriate consent for the use of 1.6 million patient's data in the development of one of its AI applications (11). Repeated episodes of this has the potential to decrease public trust in the NHS's handing of public data and with the recent introduction of National data opt-out programme for patient information (12), it is easier than ever for patients to simply refuse to hand over their valuable information. It is therefore crucial to ensure that with the growing AI movement, work is committed to ensuring the public are sufficiently educated and engaged in consenting to the use of their personal data for use in AI. Additionally, it is necessary to guarantee that patient information continues to be viewed as a valued asset subject to appropriate ethical and legal considerations as opposed to a "*free-for-all*" for use in training AI algorithms.

Regulation

Regulation of AI software is a particularly complex area with advancements exceeding the pace at which regulatory guidance is adapting. As observed by *Sherer* (13), the rise of AI has so far occurred in a regulatory vacuum with minimal guidelines existing specifically addressing the unique challenges posed by AI. Regulation of medical devices, including medical software, in the UK is currently enforced by the medicines and healthcare products regulatory agency (MHRA) which requires developers to obtain a CE marking thereby demonstrating products have conformed to EU legislation. Additionally, products using public data for uses such as AI must abide to the necessary Data Protection laws. Up until recent years these regulatory frameworks, although ensuring the safety of products, have failed to address key regulatory challenges within AI such as the associated transparency issues. Over the past year, however, regulatory updates have been made attempting to address these insufficiencies both within Europe through the updated General data protection regulation (2018) and within the UK by the government through the publication of a code of conduct for AI and data driven technology. (14).

The updated General data protection regulation (2018) is of particular interest to AI developers as its newly added clauses specifically address the use of public data within AI. Although these have been introduced to further safeguard patient privacy, it is arguable that some clauses have meant that new hurdles have been created for developers to overcome. For example, as part of the updated act developers must provide "*meaningful information*" about the logic involved within these tools. Although this attempts to address the transparency issue associated with black box coding

in AI and specifically ML, it can be difficult and sometimes impossible to demonstrate this logic. As stated by the Centre for Data Innovation, "Policymakers should understand that the GDPR will come at a significant cost in terms of innovation and productivity" and that failing to amend the GDPR will "consign Europe to second-tier status in the emerging AI economy" (15).

The government recognises and accepts that introducing regulation that sufficiently protects patients without stifling innovation within AI is central to creating a market that can lead on global adoption of AI technologies (16) and is therefore currently reviewing existing health and care regulatory landscape (14). It would seem appropriate that, when establishing the scope of AI regulation in the future, input from those with knowledge in AI technology is included in discussions to establish an appropriate regulatory solution that both considers both the public and developers' best interests.

Accountability

It is commonly accepted that nearly all AI software will fail to perform its intended purpose at some point during its lifetime. Although in some sectors AI failures can be trivial, in environments such as healthcare failures can have catastrophic consequences and in these situations it is vital to be able to hold the relevant party accountable. The assignment of accountability in AI and specifically machine learning, however, can be somewhat challenging primarily due to the lack of transparency used in many AI models. As stated by Dr Pavel Klimov (17), humans may no longer be in control of what decision is taken and may not even know or understand why a wrong decision has been taken, because the transparency is being lost. Currently, AI within the NHS is considered to be a tool supporting staff rather than acting as sole decision maker meaning medical professionals are always held accountable for decisions regardless of whether these decisions were influenced by AI technology or not. The fact that NHS professionals can be held accountable for decisions influenced by potentially inaccurate AI, which cannot even be proven in some situations due to lack of transparency, may deter them from embracing the AI movement.

If AI is to take on a greater role in supporting clinicians in decision making then it would seem appropriate that the issue of transparency is addressed. The development of transparent machine learning techniques would not only benefit the current accountability dilemma in AI but also address the previously mentioned regulatory issues and therefore is an important area for research to focus on. Tackling transparency in artificial intelligence and specifically machine learning can be somewhat challenging however current research suggests that it could be possible. For example, MIT scientists have recently developed methods for developing transparent deep learning neural networks (18) and while developing its AI eye detection technology, DeepMind specifically tackled this issue by creating a solution that displays selected information regarding how the program arrived at its recommendation (2). Despite these promising examples, technological solutions to making machine learning transparent is still in its youth and would require developers with expertise in this new area which may not be an available resource. Another approach is to develop auditing systems that explain the inner workings of these algorithms. The realising Accountable Intelligent Systems (RAInS) project, currently being developed by the collaboration of the university of Aberdeen, Cambridge and Oxford aims to do just that by developing prototype solutions that generate "secure, tamper-proof records of intelligent systems' characteristics

and behaviours" and can be provided to the necessary bodies in the case of an adverse event in order to ensure AI systems are transparent and accountable (19).

Public trust

Obtaining public trust and support in AI is crucial to its successful deployment within the NHS. Current reports, however, suggest that the NHS may currently be in a less than desirable position with one recent poll finding that only 20% of respondents support the use of AI in healthcare with 48% opposing its use (20). Although there are numerous AI success stories in recent years, these have arguably been outshined by catastrophic AI failures, which have inevitably dented the publics already limited trust in the use of AI within daily life. For example, in March 2018 an AI powered Uber vehicle failed to detect a female pedestrian crossing the road ultimately resulting in the first pedestrian death from an autonomous car (21). These high-profile incidents in combination with the unfamiliarity and lack of understanding a significant portion of the public currently associate with AI presents as a significant problem for the NHS. Gaining public trust in the use of AI, however, is essential for its growth and measures must be carried out by relevant parties including the government to achieve this. As stated by Norman Lamb with regards to AI in healthcare, "trust is going to be of central importance – if we lose the trust of people then we won't be able to realise the great opportunities ahead of us." (22)

If the NHS is to gain public trust in the use of AI within its daily practice, then empowering patients with appropriate information is crucial. To combat high profile AI failures, successful developments in AI use within the NHS should be communicated to the public. If people are aware of the benefit AI is regularly bringing to healthcare then failures may be seen as an anomaly rather than a common occurrence. Efforts also need to be made not only to ensure that patients are sufficiently educated on current technology but also reassured this technology is safe and developed according to the relevant technical and ethical standards. In the recent AI sector deal, the government established various organisations designed to oversee the safe and ethical development of AI such as the centre for data ethics and innovation, an AI council and a central office for AI (5). Additionally, the government has stated it will aim to introduce a trusted approval scheme for digitial and healthcare products, similar to the NHS health apps library which currently identifies trusted digital apps for use by the public and organisations (14). The existence of these schemes and institutions, in combination with relevant regulatory bodies, should hopefully provide an appropriate framework to reassure the public that development of AI is subject to sufficient monitoring, ethical considerations and appropriate regulation.

Education

Education of NHS staff is a hugely important factor in aiding the growth of AI within the NHS. Ultimately, even if trusts have access to the latest AI technology, this is of little use if staff are not equipped with the necessary skillsets in order to utilise this technology. Commonly, technological advances are introduced to businesses with insufficient support to staff in adapting to these new technologies which inevitably leads to failure. As stated by Matt Hancock, with regards to AI integration in the NHS "only 10% of the challenge is the tech, 90% is the culture" and this is particularly applicable to NHS staff where AI technology may appear at present to be an enigma (23). The Government has identified

the need for this and has currently commissioned a review *(The Topol Review)* into how NHS staff should best be trained in order to support use of new digital technologies. In the report, it states that all clinicians should be educated in the ethical standards and good practice of working with AI including the ability to interpretate clinical statistics and understand recommendations generated by AI systems (24). Specific details regarding how this will be addressed and deployed is not currently clear however will be important to address in the near future to best aid NHS staff.

Within the interim report, the importance of relevant teaching within undergraduate education is also highlighted. Current research suggests that medicine is moving from an information age into an age of artificial intelligence and medical education reforms need to acknowledge and embrace this movement (25). Despite this need being recognised, medical school curriculums and graduate medical education programmes across the world are failing to address this and educate students appropriately on the use of AI and other relevant technology within medicine. Within Boston University Medical school, this requirement was recognised by one professor who responded by establishing an introductory course to machine learning for medical professionals. As part of the course, students are educated on real life ML applications in medicine and taught the necessary skillsets to analyse data using machine learning techniques (26). By adopting a similar approach within UK universities, not only would students be more confident and comfortable using AI technology at later stages of their career but also the burden of educating staff within the NHS could be lessened if this teaching has already been delivered at an earlier stage.

Outlook on the future of AI within the NHS

Ultimately, the use of artificial intelligence within the NHS has the potential to significantly improve patient care if appropriate measures are put in place to address current barriers the NHS is facing. The recent development of more accurate AI methods, specifically deep learning, has the potential to benefit a variety of areas including medical imaging analysis, administration and diagnostic support. If automation is sufficiently adopted, hundreds of hours of staff time per year could be freed, greatly impacting the current healthcare landscape and addressing financial and performance pressures on the NHS.

To reap the eventual benefits of such sophisticated technology, there are several barriers which first need to be overcome. Notably, current known inadequacies associated with NHS data sets continue to pose problems to developers in implementing AI applications throughout the development process. The establishment of an improved framework would ensure accurate machine-readable data sets are easily accessible to approved developers and perhaps combat this issue. Additionally, a multitude of issues surrounding current deep learning capabilities come to light, with research into transparent methods of machine learning vital if AI is to become incorporated into daily healthcare practice. Addressing this could quell the aforementioned regulatory and accountability concerns by ensuring the logic behind algorithms is available and clear to the user. Expanding appropriate education to the public and medical professionals is also crucial to the successful integration of AI within healthcare as it is clear without such measures the value of this technology will be limited or even hindered. While the *Topol Report* indicates that education for NHS staff is a key priority for the near future, concrete plans on how education at an undergraduate level will be addressed are yet to be seen. It is against this backdrop we can conclude that although it may be some years before the NHS is fully able to utilise the undeniable potential of AI, such time allows for necessary advanced education and preparation for the inevitable AI movement.

References

1. The Lord Darzi review of health and care (better health and care for all - a 10-point plan for the 2020s). London: Institute for public policy research; 2018 p. 24.

2. Moorfields Eye Hospital NHS Foundation Trust | DeepMind [Internet]. DeepMind. 2018 [cited 29 November 2018]. Available from: https://deepmind.com/applied/deepmind-health/working-partners/health-research-tomorrow/moorfields-eye-hospital-nhs-foundation-trust/

3. HeartFlow | Non-Invasive Coronary Artery Disease Detection [Internet]. Heartflow.com. 2018 [cited 26 November 2018]. Available from:
https://www.heartflow.com/?gclid=EAIaIQobChMIu_uNzqTg3gIVD853Ch2IigxnEAAYASAAEgIjL_D_BwE

4. East Suffolk and North Essex NHS Foundation Trust employs Thoughtonomy Virtual Workers alongside human teams to improve patient experience and boost employee satisfaction - Thoughtonomy [Internet]. Thoughtonomy. 2018 [cited 29 November 2018]. Available from: https://thoughtonomy.com/esneft-nhs-patient-employee-satisfaction/

5. Department for Business, Energy and industrial Strategy/ Department for digital, Culture, Media and sport. AI Sector Deal. UK Government; 2018.

6. UK can be 'world leader in AI' if NHS data can be tapped [Internet]. Digital Health. 2018 [cited 29 November 2018]. Available from: https://www.digitalhealth.net/2018/04/uk-world-leader-ai-nhs-data/

7. Hancock M, O'Shaughnessy L. Accelerating Artificial Intelligence in health and care: results from a state of the nation survey [Internet]. The AHSN Network; 2018. Available from: https://www.kssahsn.net/what-we-do/our-news/news/Documents/AI-Strategy.pdf

8. Algorithms in decision-making - Science and Technology Committee - House of Commons [Internet]. Publications.parliament.uk. 2018 [cited 29 November 2018]. Available from:
https://publications.parliament.uk/pa/cm201719/cmselect/cmsctech/351/35105.htm

9. NHS hospitals won't be paperless before 2027 | Digital Health [Internet]. Digital Health. 2018 [cited 29 November 2018]. Available from: https://www.digitalhealth.net/2017/04/nhs-will-not-be-paperless-before-2027/

10. How the UK can win the AI race [Internet]. KPMG; 2018 p. 7. Available from:
https://assets.kpmg.com/content/dam/kpmg/uk/pdf/2018/09/how-the-uk-can-win-the-artificial-intelligence-ai-race.pdf

11. Hern A. Royal Free breached UK data law in 1.6m patient deal with Google's DeepMind [Internet]. the Guardian. 2018 [cited 29 November 2018]. Available from: https://www.theguardian.com/technology/2017/jul/03/google-deepmind-16m-patient-royal-free-deal-data-protection-act

12. National data opt-out programme - NHS Digital [Internet]. NHS Digital. 2018 [cited 29 November 2018]. Available from: https://digital.nhs.uk/services/national-data-opt-out-programme

13. Scherer M. REGULATING ARTIFICIAL INTELLIGENCE SYSTEMS: RISKS, CHALLENGES, COMPETENCIES, AND STRATEGIES. Harvard Journal of Law & Technology [Internet]. 2016 [cited 29 November 2018];29. Available from: http://euro.ecom.cmu.edu/program/law/08-732/AI/Scherer.pdf

14. Initial code of conduct for data-driven health and care technology [Internet]. GOV.UK. 2018 [cited 29 November 2018]. Available from: https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology

15. Wallace N, Castro D. The Impact of the EU's New Data Protection Regulation on AI [Internet]. Www2.datainnovation.org. 2018 [cited 29 November 2018]. Available from: http://www2.datainnovation.org/2018-impact-gdpr-ai.pdf

16. UK Government. Government response to House of Lords Artificial Intelligence Select Committee's Report on AI in the UK: Ready, Willing and Able?. 2018.

17.. House of Commons. Algorithms in decision- making. Science and Technology Committee; 2018 p. 27.

18. Mascharka D, Tran P, Soklaski R, Majumdar A. Transparency by Design: Closing the Gap Between Performance and Interpretability in Visual Reasoning. [Internet]. 2018 [cited 29 November 2018];. Available from: https://arxiv.org/abs/1803.05268

19. Aberdeen U. Universities receive £1.1m for AI accountability project | News | The University of Aberdeen [Internet]. Abdn.ac.uk. 2018 [cited 29 November 2018]. Available from: https://www.abdn.ac.uk/news/11851/

20. Trendall S. Lack of public support for AI risks 'concerted backlash', study suggests [Internet]. PublicTechnology.net. 2018 [cited 29 November 2018]. Available from: https://www.publictechnology.net/articles/news/lack-public-support-ai-risks-'concerted-backlash'-study-suggests

21. Levin S, Wong J. Self-driving Uber kills Arizona woman in first fatal crash involving pedestrian [Internet]. The Guardian. 2018 [cited 29 November 2018]. Available from: https://www.theguardian.com/technology/2018/mar/19/uber-self-driving-car-kills-woman-arizona-tempe

22. Sharman J. AI could have 'immense' benefits for NHS, says tech committee chair [Internet]. The Independent. 2018 [cited 29 November 2018]. Available from: https://www.independent.co.uk/news/health/ai-nhs-benefits-artificial-intelligence-chair-norman-lamb-reform-a8152651.html

23. Hancock M. My vision for a more tech-driven NHS [Internet]. GOV.UK. 2018 [cited 29 November 2018]. Available from: https://www.gov.uk/government/speeches/my-vision-for-a-more-tech-driven-nhs

24. Topol E. Preparing the healthcare workforce to deliver the digital future - Interim Report [Internet]. Health Education England; 2018 p. 30. Available from: https://www.hee.nhs.uk/sites/default/files/documents/Topol%20Review%20interim%20report_0.pdf

25. Wartman S, Combs C. Medical Education Must Move From the Information Age to the Age of Artificial Intelligence. Academic Medicine. 2018;93(8):1107-1109.

26. Current Students New Course: Machine Learning for Biomedical Applications | Graduate Medical Sciences [Internet]. Bumc.bu.edu. 2018 [cited 29 November 2018]. Available from: https://www.bumc.bu.edu/gms/2018/07/06/current-studnets-new-course-machine-learning-for-biomedical-applications/

Appendix A

DeepMind and Moorfields Eye hospital

The collaboration between DeepMind and Moorfields eye hospital began in 2016 with the aim of investigating whether Artificial intelligence could be used to aid clinicians in diagnosing and treating a range of eye conditions. The technology uses deep learning methods to analyse and classify OCT (optical coherence tomography) scans which can be difficult to interpret from human sight alone and currently require analysis from experts. As stated by DeepMind, the time required to analyse these scans in addition to the vast number of scans produced each day can produce lengthy delays which could be significantly reduced with widespread use of this software. Additionally, the system can prioritise patients based on urgent need of care. The initial results of the project were recently published (1) demonstrating that the new technology can correctly identify categories of eye disease 94.5% of the time, the same accuracy if not better than leading consultant ophthalmologists at MoorFields Eye Hospital. Currently, the technology is set to be used within clinical trials with the ultimate goal of integration into daily clinical practice.

(2)- https://deepmind.com/blog/moorfields-major-milestone/#image-31930
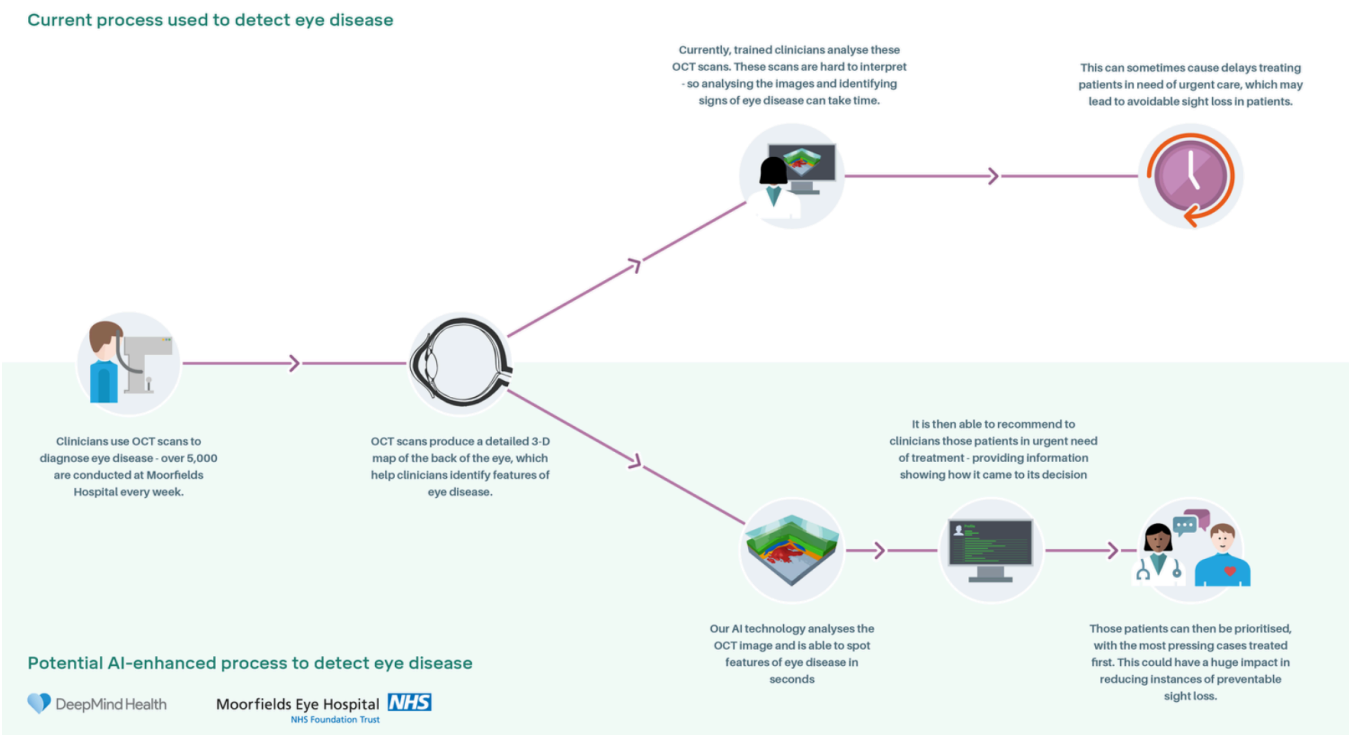
Appendix B

Heartflow software uses deep learning technology to analyse data from CT scans in order to create a three-dimensional model of coronary arteries analysing blood flow and blockages within the heart. Estimated by the National institute for health and care excellence (NICE) to save the NHS almost £9.1 million a year (3), Heartflow aims to reduce the time taken to diagnose and establish a treatment plan for coronary heart disease if used widely within the NHS. As part of the innovation and technology payment programme (ITP) which was established to aid appropriate technology in being adopted faster within the NHS, Heartflow (4) was selected to be fast-tracked into the NHS with the aim to distribute the software within 35 NHS hospital by 2019.
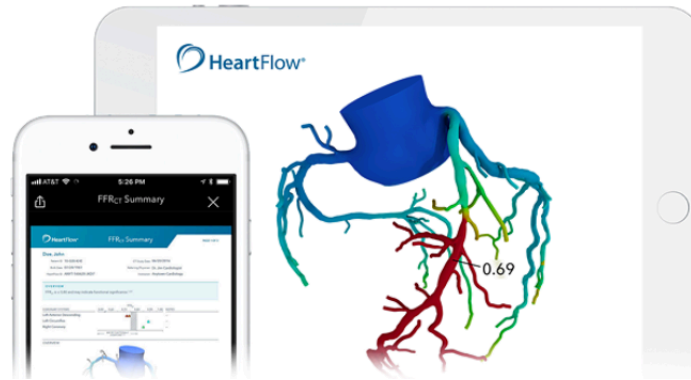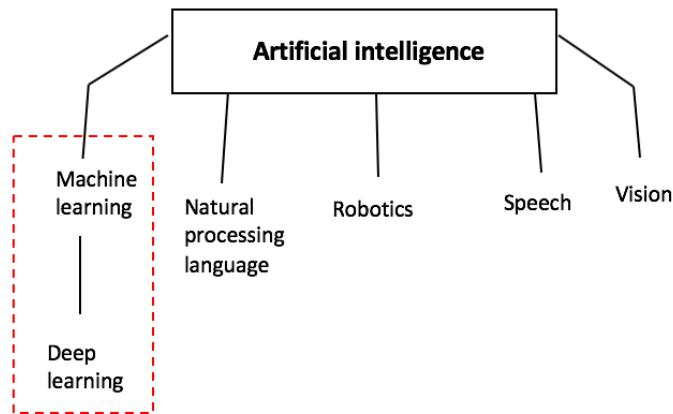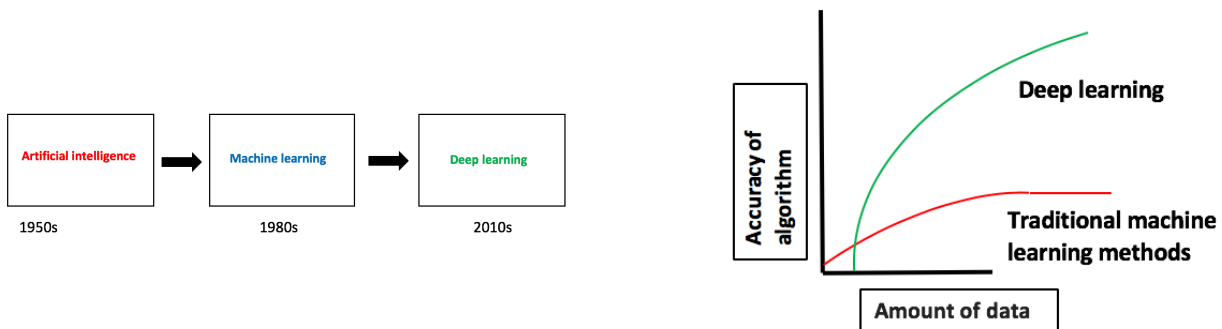

(5) -
https://www.heartflow.com/?gclid=EAIaIQobChMIo9W_uLPg3gIVGOR3Ch0lxwvSEAAYASAAEgIauvD_BwE

Appendix C

Artificial Intelligence – Subdivisions


Machine learning

Machine learning:
Machine learning is a subset of artificial intelligence that allows programs to learn without being explicitly programmed. In machine learning, systems learn from datasets, identify patterns within these datasets and use this information to make predictions about data it has not been previously exposed to.

Deep learning:
Deep learning is a relatively new (*see graph 1*) subset of machine learning which uses algorithms that replicate the human brain (neural networks) in structure and function. In machine learning, the more data the algorithm can utilise, the greater the performance of the system. In older styles of machine learning, algorithms eventually reach a plateau in performance regardless of how much additional data is used to train the algorithm (*see graph 2*). Deep learning, however, can utilise greater quantities of data leading to higher levels of accuracy in algorithms that previously wasn't achievable – this is one of the main reasons deep learning has revolutionised the AI field today.

Despite the advantages with regards to increased performance, one of the major drawbacks of deep learning is the lack of transparency in the operation of the algorithm (known as a black box system). In deep learning, it is possible to view the input and output of algorithms but difficult, even for developers, to determine how the algorithm produced its output. This proves significantly problematic in disciplines such as healthcare where justification for decisions reached is crucial to ensure maximum safety for patients.
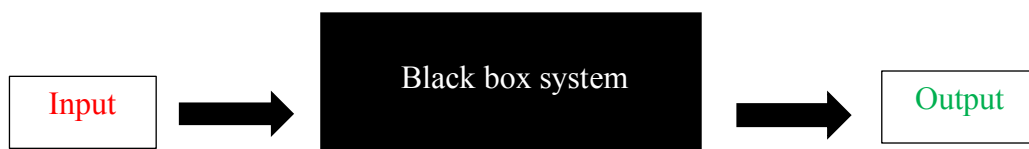
Artificial intelligence within the NHS

Benefits of the use of AI within the NHS

£12.5 billion

WORTH OF STAFF TIME SAVED PER YEAR

Through increased automation of repetitive clinical and administrative tasks

Experts have stated AI can be used to help prevent **22,000 cancer deaths** a year in addition to contributing to the prevention of **diabetes**, **heart disease** and **dementia**.

Improved efficiency and potential applications within:

Treatment

Medical imaging

Rehabilitation   Diagnosis   Drug discovery

AI chatbots   Surgery   Administration

Personalised medicine

Timeline in AI developments

Artificial intelligence

Machine learning
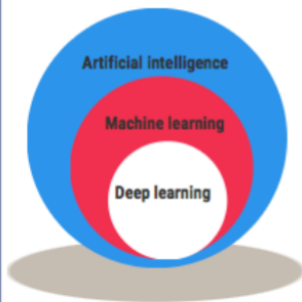
Deep learning

1950s- Artificial intelligence

1990s - Machine learning (subset of AI)

2010s: present - Deep learning (subset of ML)
Development of algorithms capable of utilising greater amounts of data than previously possible resulting in improved algorithm performance and major breakthroughs in the AI field.

Accuracy of algorithm

Deep learning

Traditional machine learning methods

Amount of data
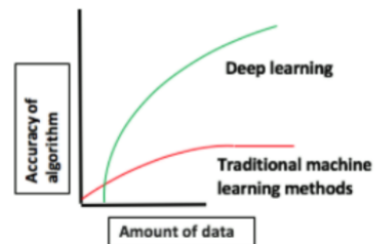
AI within the NHS: case studies

Administrative - Thoughtonomy My virtual workforce

- **Technology**: Using AI technology, "virtual workers" manage incoming GP referrals, extract relevant data from these referrals and create a single pdf document summarising this information.

- **Results**: Time taken to process GP referrals cut from **25 to 5 minutes** freeing up **500** hours of staff time within the first 3 months of deployment

Detection of eye disease

- **Technology**: Using machine learning technology, British AI company DeepMind developed a system able to analyse OCT scans, identify signs of eye disease and recommend appropriate follow up plans.

- **Results**: Initial results show accuracy levels of **94.5%**, matching the accuracy of consultant opthalmogists at the hospital.

Barriers to the growth of AI within the NHS

REGULATION

AI is advancing at a greater pace than regulatory guidance - current regulation does not appropriately address the unique challenges associated with AI such as the black box nature of deep learning algorithms.

USE OF PUBLIC DATA

Currently, much of the data held within the NHS is not in an appropriate format for training AI. A digital infrastructure storing complete, accurate and machine readable data needs to be established to provide easily accessible data sets to approved developers.

ACCOUNTABILITY

At present, AI is considered a support tool rather than a sole decision maker meaning staff will always be held accountable for decisions regardless of whether this was influenced by AI. Given some AI is unexplainable, staff may be uncomfortable using this technology unless sufficient explanations for decisions made is given

EDUCATION

Technological advancements are commonly introduced to organisations with insufficient staff support which inevitably leads to failure. The introduction of relevant staff education is therefore crucial in successfuly integrating AI into the NHS.

PUBLIC TRUST

Public trust is currently limited due to a lack of understanding of AI technology in addition to negative AI incidents making high profile news. Educating the public on the benefits AI can bring to healthcare is therefore crucial in gaining public support.

Appendix References

1.De Fauw J, Ledsam J, Romera-Paredes B, Nikolov S, Tomasev N, Blackwell S et al. Clinically applicable deep learning for diagnosis and referral in retinal disease. Nature Medicine. 2018;24(9):1342-1350.

2.A major milestone in the treatment of eye disease [Internet]. 2018 [cited 29 November 2018]. Available from: https://deepmind.com/blog/moorfields-major-milestone/#image-31930

3. HeartFlow FFRCT for estimating fractional flow reserve from coronary CT angiography | Guidance and guidelines | NICE [Internet]. Nice.org.uk. 2018 [cited 29 November 2018]. Available from: https://www.nice.org.uk/guidance/mtg32/chapter/4-nhs-considerations

4. England N. NHS England » Heart patients among those to benefit as NHS England backs innovation [Internet]. England.nhs.uk. 2018 [cited 29 November 2018]. Available from: https://www.england.nhs.uk/2018/04/heart-patients-among-those-to-benefit-as-nhs-england-backs-innovation/

5. Heartflow software [Internet]. 2018 [cited 29 November 2018]. Available from: https://www.heartflow.com/?gclid=EAIaIQobChMIo9W_uLPg3gIVGOR3Ch0lxwvSEAAYASAAEgIauvD_BwE