SUPPLEMENTARY DIGITAL CONTENT

Getting started with artificial intelligence healthcare projects: ten tips from the frontline

APPENDIX 1: Author Biographies

Dr Anthony Wilson (AW)
Anthony is a consultant in anaesthesia and critical care medicine at Manchester University NHS Foundation Trust. He is the research informatics lead for the deployment of the Trust's new, comprehensive electronic patient record. As a PhD scholar at the University of Manchester, he is working to link disparate sources of routinely collected ICU data and to apply AI techniques that utilise this data for better patient care. He is currently exploring the factors that make data-derived recommendations acceptable to clinicians at different points of care.

Mr Haroon Saeed (HS)
Haroon is an ear nose and throat specialist trainee and University of Manchester PhD scholar. His research aims to create a clinical prediction model for congenital progressive hearing loss using machine learning techniques. He has successfully built a collaborative science team who have optimised clinical data collection, navigated the process of research ethics committee approval and created a data processing agreement encompassing multiple NHS organisations. From the outset Haroon has worked to ensure the prediction model has maximum chance of being used on the clinical frontline by engaging with end users throughout development and acting early to ensure methods for model reproducibility and generalisability.

Miss Catherine Pringle (CP)
Catherine is a neurosurgery specialist trainee and University of Manchester PhD scholar. Her PhD focuses on the potential roles of machine learning approaches in outcome prediction and risk stratification of paediatric brain tumours. Paediatric brain tumours present a challenging patient cohort due to the relatively small overall numbers which generate a high dimensional data problem; a deep, multi-faceted, variable rich data set generated from a small number of patients to which traditional survival outcome statistical analysis is often difficult to apply. Through the application of machine learning techniques to radiological, biological and clinical data that accompanies these patients, Catherine is aiming to identify potential biomarkers of tumour diagnosis, risk-stratification and prognosis for children's brain tumours.

Dr Iliada Eleftheriou (IE)
Iliada is a lecturer in healthcare sciences at the University of Manchester. Her expertise lies in mapping complex data landscapes in healthcare settings to identify and address socio-technical challenges stemming from disparate information systems and data formats. She is also an academic consultant at The Christie NHS Foundation Trust where she investigates the feasibility of embedding artificial intelligence models in existing pathways to improve patient outcomes. She is leading a project to automate the validation of chemotherapy prescription regimes in the cancer care setting. The validation process is based on well-defined protocols and pharmacists can use AI-powered decision support tools to accommodate increasing numbers of patients undergoing chemotherapy.

Dr Paul A. Bromiley (PB)
Paul is a Lecturer in health data sciences at the University of Manchester. His research involves the development of machine learning based computer aided diagnostic systems for use in radiology, with a particular focus on musculoskeletal radiology and neuroradiology. He developed the machine learning software used in the ASPIRE™ teleradiology service marketed by Optasia Medical Ltd. (Manchester, UK; www.optasiamedical.com). ASPIRE™ automates the diagnosis of vertebral fragility fractures visualised incidentally in CT images, allowing hospitals to improve diagnostic rates for osteoporosis without increasing radiology department workloads.

Prof Andy Brass (AB)
Andy has over 25 years' experience in building capacity for informatics in healthcare, including developing the world's first Masters in Bioinformatics. His current focus is on working in partnership with Health Education England and the National School of Healthcare Sciences to develop career pathways in bioinformatics and data science for the NHS. As a part of this activity he has worked with NHS trusts across the UK to explore the challenges faced in embedding these new professions and methodologies within the existing workforce and working practices.

APPENDIX 2: GLOSSARY OF TERMS

The following glossary of terms may be useful for health care professionals in collaborative science teams as they seek to develop a shared vocabulary with their data scientist colleagues. The terms are divided into two parts, general healthcare data science terms and AI-specific terms. The former provides a minimum shared vocabulary for all project team members, whilst the latter may be of use to HCPs wishing to gain a deeper understanding of data science techniques.

Term	Definition
Healthcare Data Science Terms	
Algorithm	A formula or set of rules for performing a task. In Artificial Intelligence, the algorithm tells the machine how to go about finding answers to a question or solutions to a problem. Link to article by AI Glossary – AI Trends
Artificial Intelligence (AI)	Artificial Intelligence (AI) has numerous definitions which is due to how the interpretation of the concept of intelligence and also how the field has evolved over time. In this paper, we refer to AI as the following: "Refers to a broad field of science encompassing not only computer science but also psychology, philosophy, linguistics and other areas. AI is concerned with getting computers to do tasks that would normally require human intelligence." (Stefan van Duin and Naser Bakshi, 2017). Link to article "Part 1: Artificial Intelligence Defined" on Deloitte's website "the science of making machines do things that would require intelligence if done by people" (McCarthy and Minsky,1955) Link to article by John McCarthy on Stanford University
Big Data	Represents large amounts of data that are unmanageable using traditional software or internet-based platforms. It surpasses the traditionally used amount of storage, processing and analytical power. Link to article by Dash, S., Shakyawar, S.K., Sharma, M. et al. Big data in healthcare: management, analysis and future prospects. J Big Data 6, 54 (2019)
Cloud Computing	The practice of using a network of remote servers hosted on the Internet to store, manage, and process data. Link to article by Matt Dryfhout and Scott Hewer on Scout Technology Guides (2019)

Decision Trees	Decision tree learning is commonly used in data mining, aiming to create a model that predicts the value of a target variable based on several input variables. A classification tree models discrete target variables, where leaves represent final class labels and branches represent conjunctions of features that lead to class assignment, such as tumour classification. Regression tree models have continuous value target variables, for example length of stay in hospital. Link to book by Rokach, Lior; Maimon, O. (2008). Data mining with decision trees: theory and applications. World Scientific Pub Co Inc. ISBN 978-9812771711.
Deep Learning	A subfield of machine learning concerned with algorithms inspired by the structure and function of the brain called artificial neural networks. Link to article by Jason Brownlee on Machine Learning Mastery (2019)
General AI	Also referred to as strong AI or deep AI, is the concept of a machine with general intelligence that mimics human intelligence and/or behaviours, with the ability to learn and apply its intelligence to solve any problem. General AI can think, understand, and act in a way that is indistinguishable from that of a human in any given situation. AI researchers and scientists have not yet achieved strong AI. To succeed, they would need to find a way to make machines conscious, programming a full set of cognitive abilities. Link to article by Serena Reece on Codebots (2020)
Machine Learning	The study of computer algorithms that improve automatically through experience. It is seen as a subset of artificial intelligence. Machine learning algorithms build a mathematical model based on sample data, known as "training data", in order to make predictions or decisions without being explicitly programmed to do so. Examples of machine learning algorithms are: Random Forest, Nearest Neighbour, Support Vector Machine (SVM), Deep Learning, etc. Link to book by Sammut, C., & Webb, G. I. *Encyclopedia of machine learning and data mining*. Springer Publishing Company, Incorporated. (2017).
Narrow AI	Also known as weak AI, is a specific type of artificial intelligence in which a technology outperforms humans in some very narrowly defined task. Unlike general artificial intelligence, narrow AI focuses on a single subset of cognitive abilities and advances in that spectrum. Link to article by Techopedia (2020).

Natural Language Processing (NLP)	A branch of artificial intelligence that deals with the interaction between computers and humans using natural language. The ultimate objective of NLP is to read, decipher, understand, and make sense of the human languages, in both speech and written form, in a manner that is valuable. Most NLP techniques rely on machine learning to derive meaning from human languages. Link to article by Jason Browniee (2019)
Neural Network	Artificial Neural Networks (ANNs) consist of a layered network of nodes between input and output layers, with weighted relationship connections between each node. The weight of connection increases or decreases the strength of a signal at a connection, and auto-adjusts as the model learns Link to article on Wikipedia
Robotics	Robotics – is an interdisciplinary branch of engineering and science that includes mechanical engineering, electronic engineering, information engineering, computer science, and others. Robotics deals with the design, construction, operation, and use of robots, as well as computer systems for their control, sensory feedback, and information processing. Link to Wikipedia Robotics page
Semi Structured Data	Semi-structured data does not have the same level of organisation and predictability of structured data associated with databases or data tables, but does contain searchable and retrievable tags or labels. These tags, labels or codes separate semantic elements from the data and enforce hierarchies of records and fields within the data, and is also known as a self-describing structure. Items within a semi-structured class may have different attributes despite being grouped together. Link to article by Buneman, Peter. "Semi structured data." Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems. 1997.
Structured Data	Structured data describes the information that can be stored and viewed in a consistent, organised and reproducible manner, typically in a standardized format such as tables with headings in columns and values in rows. This type of data can be easily read by a machine, analysed over time, and validated against expected values or biologically plausible thresholds. Link to book by Cormen, Thomas H.; Leiserson, Charles E.; Rivest, Ronald L.; Stein, Clifford (2009). Introduction to Algorithms, Third Edition (3rd ed.). The MIT Press.

Unstructured data	This type of data doesn't come pre-defined or with a pre-defined data model, therefore it lacks the organisation and precision of structured data. As a consequence, this type of data cannot be analysed easily by a machine. In order to be used as input to an algorithm, the data needs to be transformed into a structured form and often needs to be manually interpreted and analysed. Unstructured data is often text heavy, any may also contain dates, numerical values and factual information, making the application of traditional programs difficult when compared to structured data. Link to Wikipedia page on Unstructured data

AI - technical terms	
Error-driven learning	This is a sub-field of machine learning concerned with how an agent ought to take actions in an environment so as to minimize some error feedback. It is a type of reinforcement learning. Link to error-driven learning page on Wikipedia
Generative Adversarial Network	Algorithmic architectures that use two neural networks, pitting one against the other (thus the "adversarial") in order to generate new, synthetic instances of data that can pass for real data. They are used widely in image generation, video generation and voice generation. Link to article by Ian Goodfellow "NIPS 2016 tutorial: Generative adversarial networks." arXiv preprint arXiv:1701.00160 (2016).
Hidden Layer	In neural networks, a hidden layer is located between the input and output of the algorithm, in which the function applies weights to the inputs and directs them through an activation function as the output. In short, the hidden layers perform nonlinear transformations of the inputs entered into the network. Hidden layers vary depending on the function of the neural network, and similarly, the layers may vary depending on their associated weights. Link to article on DeepAI.
Junction tree algorithm	Also known as 'Clique Tree' is a method used in machine learning to extract marginalization in general graphs. In essence, it entails performing belief propagation on a modified graph called a junction tree. The graph is called a tree because it branches into different sections of data; nodes of variables are the branches. Link to article on Wikipedia

KNN (Nearest Neighbour)	Nearest neighbour algorithms classify a test example by finding its closest neighbours in a multidimensional feature space populated by known examples from a reference (training) data set. The class prediction is estimated to be that of the nearest neighbour, or by a weighted average of the classes of the k nearest neighbours. Link to article by Nisbet, R., Elder, J., & Miner, G. Handbook of statistical analysis and data mining applications. Academic Press. (2009)
Logistic Regression	A supervised learning classification algorithm used to predict the probability of a target variable. The nature of target or dependent variable is dichotomous, which means there would be only two possible classes, but there can be two more categories of target variables that can be predicted by it: multinomial and ordinal. Logistic regression can be used for various classification problems such as spam detection, diabetes prediction, cancer detection etc. Link to article by TutorialsPoint.
Multiple Imputation	Imputation is an approach that deals with missing data within a dataset by imputing values based on a statistical method. Single imputation of missing values often fails to account for the uncertainty about the missing values. Multiple imputation, allows for the uncertainty about the missing data by creating several different plausible imputed data sets and appropriately combining results obtained from each of them. Link to article by Sterne, J. A., White, I. R., Carlin, et al. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. Bmj, 338, b2393. (2009).

Overfitting	Overfitting occurs when a supervised learning algorithm becomes over familiar with a small training set of data, and may therefore fail to fit and analyse newly presented data reliably. The contrary position to this is Underfitting. Link to article by Tetko IV, Livingston DJ and Luik AI. Neural Network studies.1. Comparison of overfitting and overtraining. *J. Chem. Inf. Comput. Sci.* 1995, 35, 5, 826–833.
Principal component analysis (PCA)	Principal component analysis (PCA) simplifies the complexity in high-dimensional data while retaining trends and patterns. It does this by transforming the data into fewer dimensions, which act as summaries of features. PCA projects the data onto a few principal component (PC) directions, without losing too much information about the subjects.

	Link to article by Lever, J., Krzywinski, M. & Altman, N. Principal component analysis. *Nat Methods* **14,** 641–642 (2017).
Random Forests	Random forests use multiple decision trees to enhance the classification and regression outputs at time of training, with the mode of the classes (classification) or mean prediction (regression) of the contributing trees generating the final output. Random decision forests correct for decision trees' habit of overfitting to their training data set. Link to book Hastie, Trevor; Tibshirani, Robert; Friedman, Jerome (2008). The Elements of Statistical Learning (2nded.). *Springer*. ISBN 0-387-95284-5.
Reinforcement Learning	The process of teaching a machine to make specific decisions using trial and error in order to optimize output. Reinforcement learning is one of three basic machine learning paradigms, alongside supervised learning and unsupervised learning. However, it differs as it does not require labelled input and output data nor correction of poor outcomes. Instead it aims to balance exploration of uncharted data and exploitation (of current knowledge). Link to paper: Kaebling L,; Littman ML.; Moore AW. (1996). "Reinforcement Learning: A Survey". Journal of Artificial Intelligence Research. :4: 237-285.

Statistical Shape modelling	Statistical shape models (SSMs) are often used in medical imaging analysis, based upon the idea that a statistical shape model represents the normal shape variation of a class of shapes, which is then used as prior knowledge in an algorithm when encountering new data. SSMs describe the shape of an object through the application of PCA to reference points, and work on the assumption that newly encountered shapes are a deformed version of the reference shape. Link to article by Lüthi, Marcel, et al. "Shape modeling using gaussian process morphable models." Statistical Shape and Deformation Analysis. Academic Press, 2017. 165-191.
Supervised Learning	Supervised learning approaches infer a function from labelled training data consisting of a set of training examples. These learnt inferences can then be applied to unseen data and correctly determine class Link to article by Mehryar Mohri, A.R., Ameet Talwalkar The MIT Press ISBN 9780262018258., Foundations of Machine Learning, . The MIT Press 2012.

Support Vector Machines (SVM)	Support Vector Machines analyse and group labelled input data into classes separated by the widest plane (support vector), often in a non-linear relationship where data is not easily distinguishable. The distance of all data points from the separation plane is calculated, and the two closet data points on each side of the plane are used as reference points to classify new data against. It is used in clinical research, for identifying imaging biomarkers, to diagnose cancer or neurological diseases. Link to article by Erickson, B.J., et al., Machine Learning for Medical Imaging. Radiographics, 2017. 37(2): p. 505-515.
Underfitting	Underfitting occurs when a model or algorithm has not learned enough from the training data set, and cannot capture the trend or patterns within a dataset, resulting in low generalization and unreliable predictions. Link to article by Anas Al-Masri on Towards Data Science (2019)
Unsupervised Learning	Unsupervised algorithms process large amounts of data to identify unknown patterns through clustering techniques to identify previously unseen patterns and associations. In contrast to supervised learning approaches, the supplied dataset does not have any labels or known outcomes. Link to article by Hinton, Geoffrey; Sejnowski, Terrence (1999). Unsupervised Learning: Foundations of Neural Computation. MIT Press. ISBN 978-0262581684.