An Open Science Approach to Artificial Intelligence in Healthcare

A Contribution from the International Medical Informatics Association Open Source Working Group

Chris Paton^{1,3}, Shinji Kobayashi^{2,3}

¹ University of Oxford, UK

² Kyoto University, Japan

³ IMIA Open Source Working Group

Summary

Objectives: Artificial Intelligence (AI) offers significant potential for improving healthcare. This paper discusses how an “open science” approach to AI tool development, data sharing, education, and research can support the clinical adoption of AI systems.

Method: In response to the call for participation for the 2019 International Medical Informatics Association (IMIA) Yearbook theme issue on AI in healthcare, the IMIA Open Source Working Group conducted a rapid review of recent literature relating to open science and AI in healthcare and discussed how an open science approach could help overcome concerns about the adoption of new AI technology in healthcare settings.

Results: The recent literature reveals that open science approaches to AI system development are well established. The ecosystem of software development, data sharing, education, and research in the AI community has, in general, adopted an open science ethos that has driven much of the recent innovation and adoption of new AI techniques. However, within the healthcare domain, adoption may be inhibited by the use of “black-box” AI systems, where only the inputs and outputs of those systems are understood, and clinical effectiveness and implementation studies are missing.

Conclusions: As AI-based data analysis and clinical decision support systems begin to be implemented in healthcare systems around the world, further openness of clinical effectiveness and mechanisms of action may be required by safety-conscious healthcare policy-makers to ensure they are clinically effective in real world use.

Keywords

Artificial Intelligence; machine learning; expert systems; open science; open source

Yearb Med Inform 2019;47:51

<http://dx.doi.org/10.1055/s-0039-1677898>

Introduction

Artificial intelligence (AI) is a broad term that encompasses a range of technologies (many of which have been under development for several decades) that aim to use human-like intelligence for solving problems [1]. The earliest AI systems used symbolic logic (such as “if-then” sequences) to create so-called “expert systems”. In the healthcare domain, these expert systems have been designed by clinicians (providing the clinical knowledge) working with programmers (translating their expertise into symbolic logic). For example, healthcare workers triaging patients could be guided by an expert system through a series of questions asked to their patients. They would then enter patient responses into the system which would be used to determine which questions to ask next and, ultimately, the system could support the management of patients by suggesting a triage decision. These types of expert systems are now widely used by services such as the NHS 111 telephone triage service (currently being beta-tested as a mobile app in the UK [2]) and other commercial symptom-checking apps [3].

Another approach to AI is the use of machine learning (ML) techniques including artificial neural networks (ANNs). Using an ANN approach, computer programmes create decision-making networks of artificial “neurons” that operate in ways similar to biological nervous systems [4]. The main difference between a ML AI and an expert system AI is that the former is not explicitly specified by experts but rather is created by a process of automated iterative

improvements (supervised or unsupervised by human experts). By using “training” datasets that match patterns of inputs (such as symptoms, medical images, or biomedical signals) with specified outputs (such as medical diagnoses), the ML programme can iteratively learn which network will match the most patterns correctly. The programme will aim to maximise the “area under the curve” (AUC) when the false negative rate is graphed against the true positive rate [5]. A predictive ML algorithm will have a low false negative rate and a high true positive rate resulting in a high AUC on a scale of 0-1.

Although ML has been used for decades, it is only recently that the combination of sufficient processing power and large enough training datasets has enabled the creation of ML algorithms that can compete with or out-perform symbolic approaches. The performance of ML has now been established through a number of high-profile achievements such as Google’s AlphaGo and its successors winning a series of Go championships from 2015-2017 [6]. Commercial applications of ML include improving search engines (especially for images and voice-powered searching), image processing, and robotic control systems (such as to increase automobile safety systems using lane-keeping and automatic braking and acceleration technology). In healthcare, ML-powered clinical decision support systems (CDSSs) have been proposed for image analysis and interpretation for radiology [7], dermatology [8, 9], and pathology [10], and for improving the scope and accuracy of biomedical signal interpretation [11–13].

Against this recent history of rapid technological developments, there are several concerns as AI begins to be applied to the healthcare context. Healthcare is a safety critical industry. In contrast to the low-risk technology demonstrations where systems such as AlphaGo and IBM Watson have won high-profile game shows or competitions, the healthcare community will expect a higher level of openness and validation before AI-based tools are adopted. Through a rapid review of the recent literature and an IMIA Open Source Working Group discussion, we have examined how the principles of Open Science [14] (Figure 1) could inform the adoption of AI in healthcare.

Open Data

New ML systems require large datasets to improve their accuracy and predictive capabilities. Open data is a concept whereby

data is freely available under an open license [15]. It offers significant promise within healthcare (as previously reported by the IMIA Open Source Working Group [16]). There is now a large range of open biomedical datasets available for training new ML algorithms developed by governments, medical societies, and international research collaborations (Table 1).

The major open datasets outlined above are generally for bioinformatics purposes rather than for clinical informatics. However, new open datasets for ML derived from clinical records such as the CheXpert dataset of chest X-rays [26], are now becoming available. When dealing with sensitive healthcare data from electronic health record (EHR) systems and other clinical systems privacy is a major issue. Even after de-identification, sensitive healthcare data often cannot be released as open data [16] and secure environments need to be used to conduct ML analyses on sensitive data sources [27].

Open Research

There is a large and growing body of research on AI technologies and techniques with many studies focusing on the predictive abilities of new algorithms. There is less research, however, on the mechanisms of action (what happens inside the “black box”) of AI-based algorithms or on the clinical effectiveness of AI-based clinical decision support systems in real world applications. The “black box” concept in AI research usually refers to neural networks that are so complex that it can be very difficult for even skilled programmers to understand how decisions are made. For some uses in healthcare, a balance will need to be struck between the advantages that the “black box” approach offers (in terms of enabling a very large number of factors to be weighed in a particular decision) and the problems caused by the obscurity of the decision-making process. For digital health companies, however, the “black box” may offer a commercial

Fig. 1 Open Science Taxonomy (Attribution: Petr Knoth and Nancy Pontika CC BY 3.0).

Table 1 Examples of Open Data for Biomedical Research.

Organisation	Open Datasets
US National Center of Biotechnology Information (NCBI) and National Cancer Institute (NCI)	GEO (Gene Expression Omnibus), GenBank, and PubMed and -omics datasets [17, 18].
European Bioinformatics Institute (EBI)	Gene-Wide Association Study (GWAS), UniProt [19] and EBI-Pfam (Protein, DNA, RNA, and X family), and other catalogues [20].
Research Collaboratory for Structural Bioinformatics: Rutgers and UCSD/SDSC	The Protein Data Bank (RCSB-PDB) 3D structured protein data bank [21].
Gene Ontology Consortium (GOC)	Gene Ontology, a development framework of computational representation of genes and their biological functions [22].
SIB - Swiss Institute of Bioinformatics CPR - NNF Center for Protein Research EMBL - European Molecular Biology Laboratory	STRING database of protein–protein interactions and outputs data [23].
Center for Genomics and Personalized Medicine at Stanford University	RegulomeDB is also an open database of SNPs with known regulatory elements combined with GEO and other open datasets [24].
US National Cancer Institute	LIDC-LDI is a lung image database of 1,018 cases [25].

advantage and regulators and policy-makers will need to strike a careful balance between fostering innovation and insisting on openness that may harm commercial interests.

Implementation research shows that even if new medical technologies work “in the lab”, they might not be readily transferable to real world clinical practice [27]. AI systems may have a high AUC with strong predictive accuracy but run into problems when used in clinical practice. Technologies need to be aligned with clinical workflows and be acceptable to clinicians without introducing new problems [28]. For example, in order to be able to use AI-based CDSSs in a clinical setting, users may need an accurate mental model of how the algorithms are using the various data sources to make decisions they can trust [28]. With “black box” systems, this may not be possible and could limit their clinical effectiveness. If healthcare workers give up on understanding how decisions are made, it could lead to technology over-dependence with consequences on the degradation of clinical decision-making skills leading to worse healthcare outcomes for patients [29].

As highlighted by regulatory agencies such as the National Institute for Health and Care Excellence (NICE) [30, 31] and the US Food and Drug Administration (FDA) [32, 33], AI-driven CDSSs will usually require

a high standard of clinical evidence before they can be safely used. However, there are now a large number of digital health start-ups operating in “stealth mode”, with large investments and research and development programmes but without published evidence of mechanism of action or clinical effectiveness [34]. This conflict is likely to be resolved as new regulations are implemented, but vigilance will be required by clinicians, researchers, and policy-makers to ensure that open research approaches are adopted so that evaluations of AI systems are clinically relevant and reproducible, and can enable informed decisions about whether or not they should be adopted.

Open Access

The AI community has benefited immensely from the current trend towards open access publishing. Pre-print archives such as ArXiv.org, open access proceedings from large AI conferences, and researcher-created GitHub code repositories mean that tools, techniques, and results of countless experiments are open at no cost to researchers around the world. To underline the field’s commitment to open access, a recent campaign by the

Journal of Machine Learning Research (JMLR), one of the field’s open-access journals, advocated for researchers to state when their papers describing ML research were rejected by closed-access journals and garnered the support of more than 3,000 researchers [35].

Although major medical publications have not traditionally embraced the open access publishing movement to the same degree as AI research organisations, there has recently been a significant shift to the open access publishing model, driven by new open access requirements from research funders and pressure from academics and clinicians for greater openness of research.

Open Educational Resources (OER)

One of the first and largest “massive open online courses” (MOOCs) was on the topic of Machine Learning [36]. The prevalence of open courseware, now numbering hundreds of courses with millions of students from around the world, is supporting a growing international cohort of data scientists and also provides healthcare workers from a wide range of backgrounds the opportunity to become conversant in AI technologies. As AI technologies are applied in healthcare, it will be critical that healthcare workers and managers develop an overall understanding of how AI will impact the future of healthcare provision. Developers of AI systems will likewise need a good understanding of how healthcare differs from other contexts in order to ensure systems are safe and clinically effective. There are now a number of open health informatics courses for both healthcare workers and software developers [37–40] that cross this interdisciplinary boundary. These courses will need to keep up to date with new developments in AI as they emerge.

Open Tools

The field of artificial intelligence has a long history of using open source tools for creating AI systems and managing the large

datasets needed to train neural networks. Table 2 shows examples of major open source ML software libraries, many of which have been released by large Internet companies such as Amazon and Google. These organisations may not be able to generate fees from software licences but can still generate significant income for the use of their cloud systems based on open source frameworks. This business model supports the further development of open source frameworks. These tools are routinely used by healthcare AI research groups with, for example, Google's TensorFlow library and open data being used for classifying lung cancer and predicting mutations [41].

Table 3 shows statistical software together with other software packages used for ML applications in healthcare. GNU R and CRAN (Comprehensive R Archive Network) are statistical languages/environments with packages used for ML [49]. WEKA is an open-source data mining and ML tool [50], and TopHat is a collection of bioinformatics tools utilized for gene ontology and mapping [51].

Conclusions

The field of AI has adopted open science approaches to education, data sharing, research, and software development. These approaches have also been adopted for AI research and development in the healthcare domain. However, as AI-based data analysis and CDSSs begin to be implemented in healthcare systems around the world, further openness of clinical effectiveness and mechanisms of action will be required. The use of "black-box" systems or the introduction of systems that have not demonstrated clinical effectiveness may not be acceptable approaches for the safety-critical healthcare context.

References

1. Russel, SJ, Norvig P. Artificial intelligence: a modern approach. Malaysia: Pearson Education Limited; 2016.
2. Armstrong S. The apps attempting to transfer NHS 111 online. *BMJ* 2018 360: k156.
3. Fraser H, Coiera E, Wong D. Safety of patient-facing digital symptom checkers. *Lancet* 2018;

- 392:2263–4.
4. Michalski RS, Carbonell JG, Mitchell TM. Machine Learning: An Artificial Intelligence Approach. Springer Science & Business Media; 2013.
5. Bradley AP. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit* 1997;30:1145–59.
6. Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, et al. Mastering the game of Go with deep neural networks and tree search. *Nature* 2016;529:484–9.
7. Hosny A, Parmar C, Quackenbush J, Schwartz LH, Aerts HJWL. Artificial intelligence in radiology. *Nat Rev Cancer* 2018;18:500–10.
8. Fogel AL, Kvedar JC. Artificial intelligence powers digital medicine. *NPJ Digit Med* 2018;1:5.
9. Zakhem GA, Motosko CC, Ho RS. How Should Artificial Intelligence Screen for Skin Cancer and Deliver Diagnostic Predictions to Patients? *JAMA Dermatol* 2018;154:1383–4.
10. Salto-Tellez M, Maxwell P, Hamilton P. Artificial Intelligence-The Third Revolution in Pathology. *Histopathology* 2018; Available from: <https://onlinelibrary.wiley.com/doi/abs/10.1111/his.13760>.
11. Orphanidou C, Bonnici T, Charlton P, Clifton D, Vallance D, Tarassenko L. Signal-quality indices for the electrocardiogram and photoplethysmogram: derivation and applications to wireless monitoring. *IEEE J Biomed Health Inform* 2015;19:832–8.
12. Triantafyllidis A, Velardo C, Shah SA, Tarassenko L, Chantler T, Paton C, et al. Supporting heart failure patients through personalized mobile health monitoring. In: *Wireless Mobile Communication and Healthcare (MobiHealth), 2014 EAI 4th International Conference on*. IEEE, 2014:287–90.
13. Triantafyllidis A, Velardo C, Chantler T, Shah, SA., Paton C, Khorshidi R, et al; SUPPORT-HF Investigators. A personalised mobile-based home monitoring system for heart failure: The SUPPORT-HF Study. *Int J Med Inform* 2015;84:743–53.
14. Pontika N, Knoth P, Cancellieri M, Pearce S. Fostering Open Science to Research Using a Taxonomy and an eLearning Portal. In: *Proceedings of the 15th International Conference on Knowledge Technologies and Data-driven Business*. New York, NY, USA: ACM, 2015:11:1–11:8.
15. Open Knowledge Open Definition Group. Open Definition 2.1 - Open Definition - Defining Open in Open Data, Open Content and Open Knowledge. Available from: <http://opendefinition.org/od/2.1/en/>.
16. Kobayashi S, Kane TB, Paton C. The Privacy and Security Implications of Open Data in Healthcare. *Yearb Med Inform* 2018;27:41–7.
17. Edgar R, Domrachev M, Lash AE. Gene Expression Omnibus: NCBI gene expression and hybridization array data repository. *Nucleic Acids Res* 2002;30:207–10.
18. Benson DA, Karsch-Mizrachi I, Lipman DJ, Ostell J, Rapp BA, Wheeler DL. GenBank. *Nucleic Acids Res* 2000;28:15–8.
19. UniProt Consortium T. UniProt: the universal protein knowledgebase. *Nucleic Acids Res* 2018;46:2699.
20. Finn RD, Mistry J, Schuster-Böckler B, Grif-

Table 2 Major open-source machine-learning libraries and their management bodies.

Framework	Organization	License
TensorFlow [42]	Google	Apache 2.0
Torch [43]	Facebook	3-Clause BSD license
Microsoft Cognitive Toolkit [44]	Microsoft	MIT license
Chainer [45]	Preferred Networks Inc.	MIT license
Caffe2 [46]	Facebook	3-Clause BSD license
Scikit-learn [47]	Community base	Apache 2.0
DSSTNE [48]	Amazon	Apache 2.0

Table 3 Software packages frequently used for AI in healthcare.

Name	Description	License
GNU R	Statistical language/environment	GPL3
WEKA	Data mining and machine learning tools	GPL3
TopHat	A spliced read mapper for RNA-Seq	Boost Software License 1.0

- fiths-Jones S, Hollich V, Lassmann T, et al. Pfam: clans, web tools and services. *Nucleic Acids Res* 2006;34:D247–51.
21. Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, et al. The Protein Data Bank. *Nucleic Acids Res* 2000;28:235–42.
 22. Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, et al. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat Genet* 2000;25:25–9.
 23. Szklarczyk D, Franceschini A, Wyder S, Forslund K, Heller D, Huerta-Cepas J, et al. STRING v10: protein-protein interaction networks, integrated over the tree of life. *Nucleic Acids Res* 2015;43:D447–52.
 24. Boyle AP, Hong EL, Hariharan M, Cheng Y, Schaub MA, Kasowski M, et al. Annotation of functional variation in personal genomes using RegulomeDB. *Genome Res* 2012;22:1790–7.
 25. Armato SG. 3rd, McLennan G, Bidaut L, McNitt-Gray MF, Meyer CR, Reeves AP, et al. The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): a completed reference database of lung nodules on CT scans. *Med Phys* 2011;38:915–31.
 26. Irvin J, Rajpurkar P, Ko M, Yu Y, Ciurea-Ilcus S, Chute C, et al. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. *arXiv [cs.CV]* 2019; Available from: <http://arxiv.org/abs/1901.07031>.
 27. Greenhalgh T, Hinder S, Stramer K, Bratan T, Russell J. Adoption, non-adoption, and abandonment of a personal electronic health record: case study of HealthSpace. *BMJ* 2010;341:c5814.
 28. Amershi S, Weld D, Vorvoreanu M, Fourney A, Nushi B, Collisson P, et al. Guidelines for Human-AI Interaction. 2019; Available from: <http://dx.doi.org/10.1145/3290605.3300233>.
 29. Artificial Intelligence in healthcare - Academy of Medical Royal Colleges. Academy of Medical Royal Colleges 2019; Available from: <https://www.aomrc.org.uk/reports-guidance/artificial-intelligence-in-healthcare/>.
 30. Evidence standards framework for digital health technologies. National Institute for Health and Care Excellence Available from: <https://www.nice.org.uk/about/what-we-do/our-programmes/evidence-standards-framework-for-digital-health-technologies>.
 31. Greaves F, Joshi I, Campbell M, Roberts S, Patel N, Powell J. What is an appropriate level of evidence for a digital health intervention? *Lancet* 2018;392:2665–7.
 32. Cortez NG, Cohen IG, Kesselheim AS. FDA regulation of mobile health technologies. *N Engl J Med* 2014;371:372–9.
 33. Food and Drug Administration. Mobile medical applications: guidance for industry and Food and Drug Administration staff. *FDA Med Bull* 2013.
 34. Cristea IA, Cahan EM, Ioannidis JPA. Stealth research: lack of peer-reviewed evidence from healthcare unicorns. *Eur J Clin Invest* 2019;e13072.
 35. Statement on Nature Machine Intelligence. Oregon State University. Available from: <https://openaccess.engineering.oregonstate.edu/>.
 36. Year in Review: 10 Most Popular Courses in 2017. Coursera Blog 2017; Available from: <https://blog.coursera.org/year-review-10-popular-courses-2017/>.
 37. Paton C. Massive open online course for health informatics education. *Healthc Inform Res* 2014;20:81–7.
 38. Paton C, Malik M, Househ M. Disseminating health informatics research and Informatics in Developing Countries and the Health Informatics Forum MOOC. *Australas Epidemiol* 2014;21:29–32.
 39. Konstantinidis ST, Hansen M, Bamidis PD, Paton C. Health Education in the Era of Social Media, the Semantic Web and MOOCs. *Stud Health Technol and Inform* 2013(192) :MEDINFO :2013. IOS Press; 2013.
 40. Koch S, Häggglund M. Mutual Learning and Exchange of Health Informatics Experiences from Around the World - Evaluation of a Massive Open Online Course in eHealth. *Stud Health Technol Inform* 2017;245:753–7.
 41. Coudray N, Ocampo PS, Sakellaropoulos T, Narula N. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat Med* 2018;24(10):1559–67.
 42. TensorFlow. TensorFlow Available from: <https://www.tensorflow.org/>.
 43. Scientific computing for LuaJIT. Available from: <http://torch.ch/>.
 44. Microsoft Cognitive Toolkit. Microsoft Cognitive Toolkit Available from: <https://www.microsoft.com/en-us/cognitive-toolkit/>.
 45. Chainer. A flexible framework for neural networks. Chainer Available from: <https://chainer.org/>.
 46. Caffe2. Caffe2 Available from: <https://caffe2.ai/>.
 47. scikit-learn: machine learning in Python — scikit-learn 0.20.2 documentation. Available from: <https://scikit-learn.org/stable/>.
 48. amazon-ds2tne. Github Available from: <https://github.com/amzn/amazon-ds2tne>.
 49. Team RC. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, 2013; 2014.
 50. Frank, E., Hall, M.A., Witten, IH. The WEKA Workbench. Morgan Kaufmann; 2016.
 51. Trapnell, C., Pachter, L., Salzberg, S.L. TopHat: discovering splice junctions with RNA-Seq. *Bioinformatics* 2009; 25: 1105–11.

Correspondence to:

Dr Chris Paton

Centre for Tropical Medicine and Global Health

Nuffield Department of Medicine

University of Oxford

Oxford, OX1 3SY, UK

E-mail: chris.paton@ndm.ox.ac.uk