S Afr Gen Pract. 2024;5(3):108-111
https://doi.org/10.36303/SAGP.0954

S Afr Gen Pract
ISSN 2706-9613 EISSN 2706-9621

EDITORIAL

The GP's guide to Artificial Intelligence (AI) in medicine

K Outhoff, **K Trathen**

University of Pretoria, South Africa
Corresponding author, email: *kim.outhoff@up.ac.za*

Artificial Intelligence (AI) in medicine involves using computer programs, algorithms, and machine learning to analyse complex medical data and support clinical decision-making. In general practice, AI has the ability to transform patient care through predictive analytics, enhancing diagnostic accuracy, personalising treatment plans, and forecasting chronic disease risks, leading to proactive, tailored health care, and improved patient outcomes. Uptake of AI technologies in the medical workspace will be accelerated by increasing our awareness and understanding, and by heeding ethical and regulatory constraints.

Introduction

Artificial Intelligence (AI) aims to create machines and software that can think, learn, and make decisions like humans. In simple terms, AI is about making computers smart enough to handle tasks autonomously by learning from data, just as we learn from experience. Thus, AI is used in virtual assistants like Siri or Alexa, or in recommendation systems like those used by Netflix® or Amazon®, which suggest movies or products based on our preferences.

In general practice, AI-powered systems can enhance patient care by offering time-saving assistance to general practitioners (GPs), who manage many patients with diverse conditions within limited timeframes and often under significant uncertainty. GPs may have a basic understanding of AI, particularly as it relates to its applications in healthcare.[1] They might be familiar with AI technologies like diagnostic tools, predictive analytics, or electronic health records (EHR) systems: the digital versions of a patient's comprehensive medical history and information.[2] EHRs typically include patient demographics, medical and family history, medications, allergies, immunisations, laboratory data, progress notes, treatment plans, medical imaging, and care coordination. That is, a lot of data! Some might be very knowledgeable about AI and actively use it in their practice, while others may have a more general awareness of its potential benefits and limitations.[1] There may be a fair deal of scepticism and reluctance to engage with AI too.

A decade ago, predictions suggested that AI could replace radiologists and pathologists, though the impact on primary care was less definitive. Some researchers believed AI would enhance practices and reduce routine visits, while others warned of significant disruption and the potential for unemployment among doctors. A survey of 720 United Kingdom (UK) GPs in 2019 revealed that most viewed AI's potential as limited.[3]

Medicine's adoption of AI has been slow, but its role is rapidly expanding, with the potential to democratise advanced medical care globally.[4] However, many are unaware of AI's implications, limitations, and ethical challenges.[5] AI could become more accessible and/or relevant if doctors better understood its underlying principles and potential applications. A basic AI guide for primary care doctors that includes a primer on basic tech jargon could help bridge this knowledge gap. This editorial sets out to aid GPs with an improved understanding of AI and its implementation in the medical workspace, so that its adoption may be more closely considered.

AI categories and abilities

Bostrom's framework categorises AI into levels based on cognitive capabilities: Artificial Narrow Intelligence (ANI), Artificial General Intelligence (AGI), and Artificial Superintelligence (ASI). Current AI technology is at the ANI level, excelling at specific tasks but lacking broader cognitive abilities.[6] AI algorithms and computational power may never fully replace human wisdom and prudence, i.e. reason, which is envisaged for AGI and ASI.[7] Yet its applications in healthcare are growing.

ANI may be better understood through its methodological evolution from Good Old-Fashioned Artificial Intelligence

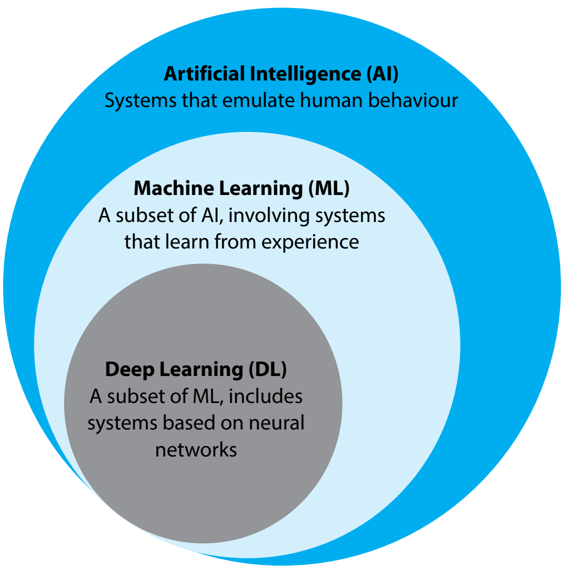

Figure 1: A schematic of the hierarchical relationship among AI, ML and DL

(GOFAI) to more advanced techniques like machine learning and deep learning, each representing a distinct approach and level of sophistication (Figure 1). This progression and its applications in medicine are outlined here.

Classic AI

Classic AI, often referred to as "Good Old-Fashioned AI" (GOFAI), involves rule-based systems where behaviour is guided by explicitly programmed instructions and logical reasoning rather than learning from data.[8] The many EHR alerts and clinical decision support tools are examples of this classic AI iteration, i.e. the process where a computer is programmed to do a task which is limited in scope, functioning only within the narrow domains it is designed for, without the ability to learn from data. Examples include EHR alerts for drug-drug interactions, allergy, duplicate testing, condition-specific care reminders (provide reminders for best practices, such as the need for screening tests in patients with chronic conditions like diabetes or heart disease), and abnormal laboratory results.

Classic AI is also used in industrial and surgical robots which execute specific movements directed by a surgeon. These systems often focus on precision and dexterity. While classic robotics relies on fixed algorithms and does not typically adapt or learn from data, many modern robots incorporate machine learning and other forms of AI that go beyond classic AI, allowing them to adapt, learn, and respond to new situations in real time.

Machine learning

Machine learning, a (significant) step up from classic AI, involves algorithms that enable computers to learn from data rather than relying on programming.[6,8] An algorithm is a set of well-defined, step-by-step instructions or rules designed to solve a specific problem or perform a task. In computing, algorithms take input data, process it through a sequence of operations, and produce one or more outputs. Essentially, a machine is trained to perform tasks typically done by humans while recalling and leaning on all training material, and the key is that it may continuously improve its performance as it learns from more examples, i.e. continuous learning.

Predictive models in healthcare – such as forecasting chronic illnesses like diabetes, hypertension and heart disease, or tailoring care plans based on genetic makeup, medical history, and other factors – are based on machine learning, thus minimising adverse drug reactions and improving treatment outcomes.

Natural language processing (NLP)

Natural language processing (NLP) tools, driven by machine learning, can assist GPs by transcribing, interpreting, and summarising patient visits. NLP tools may extract relevant information from unstructured data in large volumes of unstructured text, such as electronic patient records, clinical notes, and research articles, saving time and improving the accuracy of patient care. NLP is also used to analyse social media posts, news articles, and health reports to monitor outbreaks and track the spread of diseases, providing data for public health interventions.[9]

Generative AI tools are designed to produce new content – such as text, images, or music – based on their trained parameters.[10] Open AI's ChatGPT, for instance, uses techniques in NLP and machine learning to understand and generate human-like text based on the input and prompts it receives, forming a large language model (LLM). LLMs are trained on vast amounts of text data and use patterns and information learned during training to respond to queries, carry out conversations, and perform various language-based tasks.

NLP-powered chatbots and virtual assistants may interact with patients, answer their questions, provide reminders for medication, and even assist with scheduling appointments, improving patient engagement and adherence to treatment. Chatbots can be integrated with wearable devices or mobile apps to monitor patients' health metrics, such as blood pressure or glucose levels. They can offer personalised recommendations based on real-time data by exploiting continuous learning, reminding patients to take medications or encouraging healthy lifestyle choices. They can also be used as a triage tool.

Deep learning

A particularly advanced form of machine learning is deep learning, which uses algorithms modelled after the human brain.[6] These are known as artificial neural networks comprising multiple interconnected layers that operate on mathematical models to analyse and interpret complex data.[8] Each layer of the neural network extracts increasingly abstract features from the input data, enabling the model to learn patterns and make decisions with high accuracy. These algorithms significantly enhance a computer's learning capabilities, even enabling it to learn autonomously without explicit programming.

The range of deep learning using artificial neural networks is extensive in medicine, enabling the management of data in fields like genomics and molecular biology.[8] Artificial neural networks are most commonly used for analysing visual images, though, where data can be gathered quickly through image scanners, digital cameras, remote sensors, electronic devices, or the Internet of Things (IoT).[4] Thus, many diagnostic tools utilise deep learning including in medical imaging to detect and diagnose conditions such as tumours or fractures from X-rays, MRIs, and CT scans, to identify circulating tumour cells in whole blood samples,[11] and to analyse pathology slides, dermatology photos, or even a patient's physical movements to generate differential diagnoses.[5]

Deep learning models also aid drug discovery including by predicting how new compounds interact with biological targets, potentially speeding up the development of new treatments. This is possible, even in South Africa! The University of Cape Town's (UCT's) Holistic Drug Discovery and Development (H3D) Center has successfully collaborated in tuberculosis and malaria drug discovery with a technology non-profit that disseminates data science tools in low-resource settings.[12] Additionally, deep learning is employed in personalised medicine to analyse genetic data and predict how individual patients might respond to specific treatments, leading to more tailored and effective healthcare interventions.

Progress in computing, internet connectivity, advanced statistics, machine learning, and neural networks, along with the proliferation of handheld and wearable devices like smartphones and smartwatches, have led to revolutionary changes in healthcare driven by AI technology.[13] This leap is partly because AI relies heavily on data. In fact, data is the foundation that AI systems use to learn, make decisions, and improve over time. And healthcare generates high volumes of data.

Big Data

Big Data refers to extremely large and complex datasets that traditional data processing tools cannot handle efficiently.[7] In AI, Big Data provides the vast amounts of information necessary for training and refining AI models. AI algorithms, particularly those in machine learning and deep learning, rely on Big Data to learn patterns, make predictions, and improve their performance over time.

Characteristics of Big Data include volume, variety, velocity (the speed of data generation, processing, and analysis), veracity (the quality and reliability of the data, which can vary and affect the outcomes of AI models), and value (the potential insights and benefits that can be derived from analysing the data). The quality, quantity, and relevance of the data directly influence the accuracy and effectiveness of an AI system. Without sufficient high-quality data, AI systems would struggle to function properly or generate reliable, statistically significant outcomes. The greater the data, the better the learning – thus Big Data, which is abundant in healthcare and includes EHRs, "omic" data (such as genomics, metabolomics, and proteomics), along with sociodemographic and lifestyle-related information captured on wearables, lends itself to advancing AI processes.

Big Data has significantly impacted medicine, with patient health care data reaching genomic scale in both volume and complexity. The systematic use of this data holds the potential to personalise treatments and reduce healthcare costs by improving resource allocation. However, the anticipated benefits have yet to fully materialise, as healthcare professionals often lack the tools and time needed to integrate Big Data into everyday clinical practice.[14]

Black box AI systems

The most advanced AI systems are often the least transparent, leading to the "black box" problem, where users, patients, and even developers struggle to understand how results were obtained.[15,16] We can observe the data input and output, and understand the general workings of AI, often through deep neural networks. However, it can be challenging to determine why an AI made a specific decision, diagnosis, or action in a particular case. This opacity arises from the immense complexity of these systems, which may rely on millions of parameters. Additionally, AI systems can autonomously alter their algorithms without human oversight.[14]

Black box AI systems, which operate with non-transparent decision-making processes, present several risks. For instance, they are vulnerable to cyber-attacks, as the lack of transparency makes it difficult to detect and defend against malicious tampering, which could lead to compromised outcomes or data breaches. Systematic bias is also a concern, as the AI's decision-making is often influenced by underlying training data that may contain hidden biases, leading to unfair or discriminatory outcomes, especially in sensitive areas like medical diagnosis. Lastly, there is the risk of a mismatch between AI decisions and real-world needs or values, as black box systems may produce results that are technically correct but misaligned with human expectations, ethical standards, or specific contextual requirements, leading to potentially harmful or inappropriate actions.[14]

When AI systems become highly complex and enter the realm of black-box medicine, the fundamental principles of patient-centred care are challenged.[17] It is asserted that black-box medicine does not effectively support informed decision-making, which relies on shared information, collaborative deliberation, and a mutual understanding between practitioner and patient. Because of the scepticism and fear associated with opacity, strides are being made to explain the inexplicable in AI.[16] Some argue that there is an unrealistic expectation for explainable AI (xAI), suggesting that current methods of explainability are unlikely to meet the needs for patient-level decision support.[18] That being said, to what degree can we even trust our own human reasoning? High-stake medical decisions are often made without consciously connecting all the neural dots.

Accuracy limitations

On average, AI may outperform humans in certain areas of medical diagnosis and drug development; in the administration of treatment and surgery, AI may already have surpassed trained medical professionals.[14] However, with innovation comes limitations. Concerns have emerged about the impact of algorithmic clinical decision support on health equity, especially due to the use of datasets that lack representation from minority or previously disadvantaged populations and the potential for algorithms to perpetuate and amplify existing biases. Additionally, the risks to data security, confidentiality and privacy are increasingly apparent.[8,19]

Addressing coded bias requires diverse data, inclusive development teams, and ongoing monitoring to ensure equitable care.[20] Key principles steer these efforts, namely promoting equity throughout all phases of the algorithm life cycle; ensuring transparency and explainability of healthcare algorithms; engaging patients and communities authentically to build trust; identifying and addressing algorithmic fairness issues and trade-offs; and establishing accountability for equitable outcomes from healthcare algorithms. Ensuring data quality and origin is essential to building patient trust in AI systems and preventing unethical practices. However, achieving this level of data integrity is costly.[4]

Patient concerns and ethical, legal, and regulatory considerations

The primary goal of AI in healthcare is to enhance health outcomes and patient experiences,[21] but patients have concerns about safety, threats to their autonomy, potential rises in healthcare costs, and data security. Patient acceptance of AI depends on addressing and alleviating these potential risks.[22]

Legal responsibility is also a critical concern especially as AI is a relatively new field with few or no established legal precedents. Key questions include who will be held accountable for actions resulting from AI-based decisions and who is responsible for errors made by AI programs: doctors, institutions, or developers?[23] Legal responsibility is further complicated by unclear regulations around sensitive data processing, consent, transparency, and storage. Ultimately, the human element will remain crucial in integrating AI into medical practice, where human oversight may still be necessary even in fully automated processes.[4]

AI's potential to transform clinical practice is significant, but realising its benefits requires focused efforts on design, implementation, and evaluation. Regulatory approval is essential, and AI developers must address issues like confirmatory bias and alert fatigue. The FDA has recently introduced a pilot certification programme that evaluates both AI developers and their products, a move aimed at ensuring public trust in new medical AI applications.[4]

Doctors must be aware of AI's risks and opportunities, advocating for ethical and equitable systems while preserving patient trust. However, AI also has the potential to address some of medicine's current challenges with bias and inequality. Physicians should be aware of both risks and opportunities, advocating for the development and implementation of ethical and equitable systems. Furthermore, they must act as responsible stewards of patient data to preserve the trust between providers and patients.[19]

AI has generated significant excitement in healthcare due to advances in machine learning models. However, realising AI's potential for scalable and sustained patient care in real-world settings remains a challenge. To effectively integrate AI into healthcare, a focused effort is needed not only on developing AI technologies but also on their implementation. This integration involves applying design thinking, process improvement, and implementation science to create and deploy AI-enabled systems. Understanding how AI can enhance healthcare requires careful design, implementation, and evaluation of these systems.[24]

As AI technologies continue to advance, their integration into general practice will likely improve healthcare delivery and patient outcomes. Embracing these innovations while addressing ethical and regulatory considerations will be key to realising the full potential of AI in medicine.

ORCID

K Outhoff https://orcid.org/0000-0002-0851-4802

References

1. Buck C, Doctor E, Hennrich J, Jöhnk J, Eymann T. General practitioners' attitudes toward artificial intelligence-enabled systems: interview study. J Med Internet Res. 2022;24(1):e28916. https://doi.org/10.2196/28916.
2. Katurura MC, Cilliers L. Electronic health record system in the public health care sector of South Africa: A systematic literature review. African Journal of Primary Health Care & Family Medicine. 2018;10(1):1-8. https://doi.org/10.4102/phcfm.v10i1.1746.
3. Blease C, Kaptchuk TJ, Bernstein MH, et al. Artificial intelligence and the future of primary care: exploratory qualitative study of UK general practitioners' views. J Med Internet Res. 2019;21(3):e12802. https://doi.org/10.2196/12802.
4. Ahmad Z, Rahim S, Zubair M, Abdul-Ghafar J. Artificial intelligence (AI) in medicine, current applications and future role with special emphasis on its potential and promise in pathology: present and future impact, obstacles including costs and acceptance among pathologists, practical and philosophical considerations. A comprehensive review. Diagn Pathol. 2021;16:1-16. https://doi.org/10.1186/s13000-021-01085-4.
5. Lin S. A clinician's guide to artificial intelligence (AI): why and how primary care should lead the health care AI revolution. The Journal of the American Board of Family Medicine. 2022;35(1):175-84. https://doi.org/10.3122/jabfm.2022.01.210226.
6. Meskó B, Görög M. A short guide for medical professionals in the era of artificial intelligence. NPJ Digital Medicine. 2020;3(1):126. https://doi.org/10.1038/s41746-020-00333-z.
7. Fjelland R. Why general artificial intelligence will not be realized. Humanities and Social Sciences Communications. 2020;7(1):1-9. https://doi.org/10.1057/s41599-020-0494-4.
8. Coppola F, Faggioni L, Gabelloni M, et al. Human, all too human? An all-around appraisal of the "artificial intelligence revolution" in medical imaging. Front Psychol. 2021;12:710982. https://doi.org/10.3389/fpsyg.2021.710982.
9. Kostkova P, editor. A roadmap to integrated digital public health surveillance: the vision and the challenges. Proceedings of the 22nd international conference on World Wide Web; 2013. https://doi.org/10.1145/2487788.2488024.
10. Boscardin CK, Gin B, Golde PB, Hauer KE. ChatGPT and generative artificial intelligence for medical education: potential impact and opportunity. Acad Med. 2024;99(1):22-7. https://doi.org/10.1097/ACM.0000000000005439.
11. Li Y, Mahjoubfar A, Chen CL, et al. Deep cytometry: deep learning with real-time inference in cell sorting and flow cytometry. Sci Rep. 2019;9(1):11088. https://doi.org/10.1038/s41598-019-47193-6.
12. Hlozek J, Chibale K, Woodland JG. Ongoing implementation and prospective validation of artificial intelligence/machine learning tools at an African Drug Discovery Center. ACS Med Chem Lett. 2024;15(7):989-93. https://doi.org/10.1021/acsmedchemlett.4c00243.
13. Guo J, Li B. The application of medical artificial intelligence technology in rural areas of developing countries. Health Equity. 2018;2(1):174-81. https://doi.org/10.1089/heq.2018.0037.
14. Kiener M. Artificial intelligence in medicine and the disclosure of risks. AI & Society. 2021;36(3):705-13. https://doi.org/10.1007/s00146-020-01085-w.
15. Cadario R, Longoni C, Morewedge CK. Understanding, explaining, and utilizing medical artificial intelligence. Nature Human Behaviour. 2021;5(12):1636-42. https://doi.org/10.1038/s41562-021-01146-0.
16. Combi C, Amico B, Bellazzi R, et al. A manifesto on explainability for artificial intelligence in medicine. Artif Intell Med. 2022;133:102423. https://doi.org/10.1016/j.artmed.2022.102423.
17. Bjerring JC, Busch J. Artificial intelligence and patient-centered decision-making. Philosophy & Technology. 2021;34:349-71. https://doi.org/10.1007/s13347-019-00391-6.
18. Ghassemi M, Oakden-Rayner L, Beam AL. The false hope of current approaches to explainable artificial intelligence in health care. The Lancet Digital Health. 2021;3(11):e745-e50. https://doi.org/10.1016/S2589-7500(21)00208-9.
19. McCoy LG, Nagaraj S, Morgado F, et al. What do medical students actually need to know about artificial intelligence? NPJ Digital Medicine. 2020;3(1):86. https://doi.org/10.1038/s41746-020-0294-7.
20. Chin MH, Afsar-Manesh N, Bierman AS, et al. Guiding principles to address the impact of algorithm bias on racial and ethnic disparities in health and health care. JAMA Network Open. 2023;6(12):e2345050-e. https://doi.org/10.1001/jamanetworkopen.2023.45050.
21. Bindra S, Jain R. Artificial intelligence in medical science: a review. Irish Journal of Medical Science (1971-). 2024;193(3):1419-29. https://doi.org/10.1007/s11845-023-03570-9.
22. Richardson JP, Smith C, Curtis S, et al. Patient apprehensions about the use of artificial intelligence in healthcare. NPJ Digital Medicine. 2021;4(1):140. https://doi.org/10.1038/s41746-021-00509-1.
23. Price II WN, Gerke S, Cohen IG. Liability for use of artificial intelligence in medicine. Research Handbook on Health, AI and the Law: Edward Elgar Publishing; 2024. p. 150-66. https://doi.org/10.4337/9781802205657.00016.
24. Li RC, Asch SM, Shah NH. Developing a delivery science for artificial intelligence in healthcare. NPJ Digital Medicine. 2020;3(1):107. https://doi.org/10.1038/s41746-020-00318-y.