

Time to Regenerate: The Doctor in the Age of AI

Now is the time to prepare the next generation of doctors to work in the AI-enabled health system, bringing their humanity to the machine: patient interface.

Xiaoxuan Liu^{1,2,3}, Pearse A Keane^{1,4}, Alastair K Denniston^{1,2,3,4}

1. Department of Ophthalmology, University Hospitals Birmingham NHS Foundation Trust, Birmingham, United Kingdom
2. Academic Unit of Ophthalmology, Institute of Inflammation & Ageing, University of Birmingham, United Kingdom
3. Centre for Patient Reported Outcome Research, Institute of Applied Health Research, University of Birmingham, United Kingdom
4. NIHR Biomedical Research Centre for Ophthalmology, Moorfields Eye Hospital NHS Foundation Trust and UCL Institute of Ophthalmology, London, United Kingdom

Corresponding author: Prof Alastair Denniston, PhD MRCP FRCOphth, Academic Unit of Ophthalmology, Institute of Inflammation & Ageing, University of Birmingham, United Kingdom; Tel: +44 (0) 121 371 6905; Email: a.denniston@bham.ac.uk

Contributors: XL, PK and AD have special interests in retinal imaging, the development of objective endpoints and the use of artificial intelligence to improve patient care. XL is a doctoral researcher working in automated analysis of ophthalmic imaging. PK leads the Moorfields Eye Hospital NHS Foundation Trust/DeepMind collaboration evaluating the potential for machine learning algorithms to diagnose sight-threatening retinal diseases from standard ophthalmic imaging.

Contribution: AD conceived the article; all authors contributed to the manuscript and reviewed the final draft.

Conflicts of interest: PK is a consultant for DeepMind, a company specialising in artificial intelligence; AD and XL have no conflicts of interest.

Acknowledgements: PK is supported by a National Institute for Health Research (NIHR) Clinician Scientist Award (NIHR-CS--2014-14-023). The views expressed in the publication are those of the author and not necessarily those of the Department of Health. AD and XL receive a proportion of their funding from the Wellcome Trust Health Improvement Challenge Fund.

Time to Regenerate: The Doctor in the Age of AI

Now is the time to be preparing the next generation of doctors to work in the AI-enabled health system, bringing their humanity to the machine: patient interface.

Introduction

We are experiencing a rapid expansion of new technologies which are fusing the digital and biological worlds. New digital technologies such as artificial intelligence (AI), electronic health records and Big Data, telemedicine, “wearables” for home monitoring, and virtual/augmented realities, are shaping the future of medicine to become more efficient, more accurate, and more sustainable.[1] Digital systems from industry leaders such as DeepMind and IBM Watson are already being tested for use in healthcare in the United Kingdom and the United States.

Faced by machines that outperform us in many areas, some clinicians fear that AI will render them redundant. This is, however, to underestimate the role and value of the doctor to the patient and society. Yes, AI has the potential to precipitate one of the greatest changes in the role of the doctor to date, but this is not something to be feared. Change is inevitable, but we believe the core values that characterise a good doctor will remain unchanged.

In contemporary culture there is a timeless hero known simply as ‘The Doctor’ who regularly regenerates to meet the demands of a new generation of humans. Despite the anticipation - and even anxiety - of each regeneration, we quickly get used to each new incarnation of ‘The Doctor’ because the core principles hold firm. [2]

So too, in the age of AI. Doctors will need to adapt: to let go of old roles, and to find where they can be most relevant, and make most impact. In most reviews of AI, it is the computer algorithm that takes centre-stage: what can it do better than humans. In this article we have put the spotlight back on the human doctor: their role and their unique gift - to be human, when a perfect algorithm is not enough.

The Doctor as the Human:AI Interface in Diagnosis

Clinical diagnosis relies on a doctor’s interpretation of signs, symptoms, physical examination and relevant investigations. It is vulnerable to a doctor’s fallible memory, incomplete knowledge and cognitive bias.[3,4] AI can potentially provide highly accurate data on probable diagnosis, investigation pathway, and recommended treatment, based on objective evaluation of all data and medical evidence available at that point in time: comprehensive, and up-to-the-minute.

AI will require accurate data input in order to generate a correct diagnosis. Human experiences of symptoms do not always translate perfectly into medical terminology, and 'taking a good history' will remain a key skill in clinical diagnosis. This is also the time that trust is earned. A doctor who listens well and shows empathy, is more likely to be given the 'hidden data'. Sometimes this is direct – 'I've not wanted to tell anybody but...' or via an invitation to ask more – 'I didn't come last week because things have been difficult...'. A human doctor knows that a wrong move at this point can close down any future engagement, whereas the right cue may enable the patient to seek and receive the help they need.

There is also the issue of our human tendency to report inaccurate or irrelevant information. Including embellishment, exaggerations, or even lies. This is more readily recognised by other humans than machines. The physician will have an important role as the interface between the 'human' perception of illness and the 'accurate' data input into the machine.

The fundamental question though may not be 'Can this machine understand me?', but 'Do I want this machine to understand me?'. At some point AI will almost certainly be able to simulate empathic listening, and evaluate the veracity of what it is being told. Chatbots are on the rise, and AI interpretation of body language is progressing. But do I, as a patient, want to share this information with a machine? And do I want that machine to be the first 'person' to tell me about my cancer, however 'emotionally appropriate' their simulated responses may be?

Effective communication of a serious diagnosis requires considerate assessment of a patient's hopes, fears and expectations. Much of this is non-verbal. A skilful physician 'reads between the lines'. These channels of communication are instinctual and they influence the doctor's consulting behaviour on a minute-by-minute basis, often without them realising. The most sophisticated dimensions of this human-to-human interaction happen at an innate level, and cannot be replicated by an algorithm.

Sometimes AI-algorithms may simply fail due to a lack of appropriate data. For example in rare diseases there may be insufficient 'training' data to support AI. A vital new skill for the doctor in this era will be to know the limits of AI, and how to make diagnostic decisions in these situations. Similarly in cases of multi-morbidity, treatment decisions become more complex, and potentially more nuanced as some decisions may improve one disease at the expense of another, and again the data may be less certain. It is likely that this will be a

particular area where human value judgements are helpful. Another challenge will be in cases of diagnostic equivalence, i.e. where multiple AI-derived diagnoses or solutions are proposed which may be of equal likelihood. The physician will need to manage this uncertainty and communicate with the patient.

Triaging the Doctor to the High-Risk Patient

The current health system depends on triage done by humans, sometimes according to protocol, sometimes according to knowledge and experience. Such protocols are based on few variables are therefore fairly blunt tools.

Triage by AI could be faster, more accurate and more sensitive, based on many more variables than are currently possible. Input variables would include both symptoms elicited by phone and live tracking of clinical measurements elicited by wearable or implantable technology. Triage need no longer be simplistically divided into broad categories (such as red, amber and green) but can be adjusted on a continuous scale of risk and need for rapid intervention. The continuous datastream could provide early triggers to the emergency services, such that your ambulance - driverless but with human paramedic - may be pulling up at your door almost before you know you need it.

The role of the doctor in the emergency room is likely to be as team-leader, knowledge-handler, and communicator. The coordination of the rapidly evolving estimates of diagnostic certainty, recommended treatment pathways, and, where possible, discussion of potential benefit and risk with patient and relatives is going to be key. These roles do not need to be restricted to a doctor, but they do require a human.

Directing the Doctor to the Complex, the Nuanced and the 'Doesn't Fit'

Many milder ailments may be handled almost entirely by AI-informed virtual health interfaces. In cases where there is diagnostic certainty and well-established, effective and safe treatments it may not be necessary to have a human interface at all or it could be provided as an option, similar to many telephone banking systems or electronic help desks.

If most routine low-risk ailments are being dealt with by AI, it is anticipated that there would be more time for clinicians to focus on those patients who specifically need an experienced human clinician. These might be patients with more complex conditions (rare disease, or multi-morbidity) or where the diagnostic certainty is low.

Alternatively it may be the patient's needs, rather than their condition, that is complex. Healthcare decisions for a patient with learning difficulties, dementia, addiction or significant social deprivation are likely to require more human support than for other patients.

The Doctor as Patient Educator and Advisor

Traditionally the doctor has been gatekeeper for medical knowledge, and making medical decisions for the patient. In the AI era, the medical knowledge which forms the basis of decision making will be as accessible to the patient as to the doctor. Humans are however notoriously poor at comprehending probability and evaluating risk, and especially so when it pertains to their own health or the health of a loved one. Therefore, for most patients, one of the most important roles for the doctor will be to understand risk and communicate this to the patient, whether around diagnostic certainty, the safety of an intervention or the efficacy of a treatment. The doctor will also need to be able to explain the way AI has 'formulated' a treatment plan. It is worth noting that this does not require a detailed knowledge of machine learning techniques. Just as doctors use other investigations such as a magnetic resonance imaging (MRI) scans without a detailed knowledge of its mechanics, so it should be possible to communicate the value of AI in informing clinical decisions without deep computational knowledge.

The Doctor as Patient Advocate

The doctor has a unique experience of the front line of healthcare, being privileged to listen to patients' on a daily basis, often caring for the same individuals over many years, and possessing a critical appreciation of both the possibilities and limitations of medicine. From this vantage point they are well-placed to listen and respond to the needs and priorities of individual patients and the collective patient group. This advocacy role is particularly important where there are competing interests – for example limited resources to be allocated between patients or patient groups. Such issues may be complex, and emotionally charged, but are at least reasonably transparent. Not everybody may agree with the final decision but the process leading to that decision is open to scrutiny.

In the AI era there is a risk that stakeholders can influence patient care through embedding 'hidden' values within the algorithms. As Paul Hodgkin notes, "But what happens when different values conflict? A drug firm funding a machine learning system might want to increase sales, whereas a healthcare system might want to hold down costs, and patients might prioritise safety." [5] All of us – whether patients, public or doctors - will need to engage with this process, and to hold the 'rules' of the algorithm to account. We need to move from

a 'black-box' to a 'glass-box' mentality. In this debate, the key contribution of the doctor will be their understanding of the two overlapping domains - the experience of patients in the 'real world' and their specialist knowledge of both the potential and risks of medicine.

The Doctor at the End of Life

In his 'zeroth' law of robotics, Issac Asimov proposed that the most basic tenet of AI was to 'not injure a human or through inaction allow a human being to come to harm'.^[6] This principle works in most situations, but may fail spectacularly when it comes to end of life decisions. In contrast, a human physician recognises that some human decisions are not simply a matter of survival-based logic. Despite the similarity of Asimov's law to Hippocratic concepts, humans have a more nuanced understanding of beneficence and non-maleficence which incorporates not just length of life but also quality of life. The limitations of AI in this area cannot be overcome by simply inserting a quality-adjusted-life-year threshold at which life is no longer deemed worth preserving. One patient with a terminal disease may choose palliation; another will opt for further chemotherapy. The patient may make this decision based on many variables that would also be available to the AI-algorithm, and yet, the final decision needs to be the patient's alone. Such decisions must always remain outside an algorithm.

Conclusions

The advent of AI will be a seismic shift in healthcare, and the doctor's role will need to evolve. In this article, we have considered just some of these aspects, highlighting areas of particular opportunity or challenge. Being a good physician in the era of AI will require a refocussing of our skillset and an even bigger shift in mindset. The rate of progress in AI is such that medical schools and postgraduate training schemes need to be engaging with this revolution now. We need to ensure that our new doctors are equipped to handle the AI-deconstructed world that they will be living in. In this new world, AI will seamlessly transcribe every patient and every clinical presentation into input data from which it will generate output probabilities of disease, treatment efficacy, adverse events and death. In most cases it will do this faster, more reliably and cheaper than a human. Some will see this as a threat; others as an opportunity.

This article is not about AI - it is about the new Doctor, and how they find their place in the AI-enabled health system. The need will be for the human: AI interface, the knowledge-handler, the empathic communicator. We believe that now is the time to start preparing this

next generation of doctors to work alongside AI, knowing both its value and its limitations - and, in so doing, to discover their own irreplaceable value to patients and society.

Bibliography

- [1] Topol E. Creative Destruction of Medicine: How the Digital Revolution Will Create Better Health Care. 2013. Basic Books. ISBN: 978-0465061839.
- [2] "It's time to meet the thirteenth doctor" at <http://www.bbc.co.uk/doctorwho/> [Last accessed 08/10/17]
- [3] Simpkin AL, Vyas JM, Armstrong KA. Diagnostic Reasoning: An Endangered Competency in Internal Medicine Training. *Ann Intern Med*. 2017 Oct 3;167(7):507-508. doi: 10.7326/M17-0163.
- [4] Hussain A, Oestreicher J. Clinical decision-making: Heuristics and cognitive biases for the ophthalmologist. *Surv Ophthalmol*. 2017 Aug 30. pii: S0039-6257(17)30115-7. doi: 10.1016/j.survophthal.2017.08.007
- [5] Hodgkin P. The computer may be assessing you now, but who decided its values? *BMJ* 2016;355:i6169
- [6] Asimov I. Robots and Empire. 1985. Doubleday Books. ISBN 0-385-19092-1