The Price of Artificial Intelligence

Enrico Coiera

Australian Institute of Health Innovation, Macquarie University, Sydney, NSW, Australia

Summary

Introduction: Whilst general artificial intelligence (AI) is yet to appear, today's narrow AI is already good enough to transform much of healthcare over the next two decades.

Objective: There is much discussion of the potential benefits of AI in healthcare and this paper reviews the cost that may need to be paid for these benefits, including changes in the way healthcare is practiced, patients are engaged, medical records are created, and work is reimbursed.

Results: Whilst AI will be applied to classic pattern recognition tasks like diagnosis or treatment recommendation, it is likely to be as disruptive to clinical work as it is to care delivery. Digital scribe systems that use AI to automatically create electronic health records promise great efficiency for clinicians but may lead to potentially very different types of clinical records and workflows. In disciplines like radiology, AI is likely to see image interpretation become an automated process with diminishing human engagement. Primary care is also being disrupted by AI-enabled services that automate triage, along with services such as tele-medical consultations. This altered future may necessarily see an economic change where clinicians are increasingly reimbursed for value, and AI is reimbursed at a much lower cost for volume.

Conclusion: AI is likely to be associated with some of the biggest changes we will see in healthcare in our lifetime. To fully engage with this change brings promise of the greatest reward. To not engage is to pay the highest price.

Keywords

Artificial intelligence, electronic health record, radiology, primary care, value-based care

Yarb Med Inform 2019;14-5

<http://dx.doi.org/10.1055/s-0039-1677892>

We are not ready for what is about to come.

It is not that healthcare will be soon run by a web of artificial intelligences (AIs) that are smarter than humans. Such general AI does not appear anywhere near the horizon. Rather, the narrow AI that we already have, with all its flaws and limitations, is already good enough to transform much of what we do, if applied carefully.

Amara's Law tells us that we tend to overestimate the impact of a technology in the short run, but underestimate its impact in the long [1]. There is no doubt that AI has gone through another boom cycle of inflated expectations, and that some will be disappointed that promised breakthroughs have not materialized. Yet, despite this, the next decade will see a steadily growing stream of AI applications across healthcare. Many of these applications may initially be niche, but eventually they will become mainstream. Eventually they will lead to substantial change in the business of healthcare. In twenty years time, there is every prospect the changes we find will be transformational.

Such transformation however comes with a price. For all the benefits that will come through improved efficiency, safety, and clinical outcomes, there will be costs [2]. The nature of change is that it often seems to appear suddenly. While we are all daily distracted trying to make our unyielding health system bend to our needs using traditional approaches, disruptive change surprises because it comes from places we least expected, and in ways we never quite imagined.

In linguistics, the Whorf hypothesis says that we can only imagine what we can speak of [3]. Our cognition is limited by the concepts we have words for. It is much the same in the world of health informatics. We have developed strict conceptual structures that corral AI into solving classic pattern recognition tasks like diagnosis or treatment recommendation. We think of AI automating image interpreta-

tion, or sifting electronic health record data for personalized treatment recommendations. Most don't often think about AI automating foundational business processes. Yet AI is likely to be more disruptive to clinical work in the short run than it will be to care delivery.

Digital scribes, for example, will steadily take on more of the clinical documentation task [4]. Scribes are digital assistants that listen to clinical talk such as patient consultations. They may undertake a range of tasks from simple transcription through to the summarization of key speech elements into the electronic record, as well as providing information retrieval and question-answering services. The promise of digital scribes is a reduction in human documentation burden. The price for this help will be a re-engineering of the clinical encounter. The technology to recognize and interpret clinical speech from multiple speakers, and to transform that speech into accurate clinical summaries is not yet here. However, if humans are willing to change how they speak, for example by giving an AI commands and hints, then much can be done today. It is easier for a human to say "Scribe, I'd like to prescribe some medication" than for the AI to be trained to accurately recognize whether the speech it is listening to is past history, present history, or prescription talk.

The price for using a scribe might also be an even more obvious intrusion of technology between patient and clinician, and new risks to patient privacy because speech data contains even more private information than clinician-generated records. Clinicians might simply replace today's effort in creating records, where they have control over content, to new work in reviewing and editing automated records, where content reflects the design of the AI. There are also subtler risks. Automation bias might mean that many clinicians cease to worry about what should go into a clinical document, and simply accept whatever a machine has generated

[5]. Given the widespread use of copy and paste in current day electronic records [6], such an outcome seems a distinct possibility.

At this moment, narrow AI, predominately in the form of deep learning, is making great inroads into pattern recognition tasks such as diagnostic radiological image interpretation [7]. The sheer volume of training data now available, along with access to cheap computational resources, has allowed previously impractical neural network architectures to come into their own. When a price for deep learning is discussed, it is often in terms of the end of clinical professions such as radiology or dermatology [8]. Human expertise is to be rendered redundant by super-human automation.

The reality is much more nuanced. Firstly, there remain great challenges to generalizing narrow AI methods. A well-trained deep network typically does better on data sets that resemble its training population [9]. The appearance of unexpected new edge cases, or implicit learning of features such as clinical workflow or image quality [10], can all degrade performance. One remedy for this limitation is transfer learning [11], retraining an algorithm on new data taken from the local context in which it will operate. So, just as we have seen with electronic records, the prospect of cheap and generalizable technology might be a fantasy, and expensive system localization and optimization may become the lived AI reality.

Secondly, the radiological community has reacted early, and proactively, to these challenges. Rather than resisting change, there is strong evidence not just that AI is being actively embraced within the world of radiology, but also that there is an understanding that change brings not just risks, but opportunities. In the future, radiologists might be freed from working in darkened reading rooms, and emerge to become highly visible participants to clinical care. Indeed, in the future, the idea of being an expert in just a single modality such as image interpretation may seem quaint, as radiologists transform into diagnostic experts, integrating data from multiple modalities from the genetic through to the radiologic.

The highly interconnected nature of healthcare means that changes in one part of the system will require different changes

elsewhere. Radiologists in many parts of the world are paid for each image they read. With the arrival of cheap bulk AI image interpretation, that payment model must change. The price of reading must surely drop, and expert humans must instead be paid for the value they create, not the volume they process.

The same kind of business pressure is being felt in other clinical specialties. In primary care, for example, the arrival of new, sometimes aggressive, players who base their business model on AI patient triage and telemedicine is already problematic [12, 13]. Patients might love the convenience of such services, especially when they are technologically literate, young, and in good health, but they may not always be so well served if they are older, or have complex comorbidities [14]. Thus, AI-based primary care services might end up caring for profitable low-cost and low-risk patients, and leave the remainder to be managed by a financially diminished existing primary care system. One remedy to such a risk is again to move away from reimbursement for volume, to reimbursement for value. Indeed, value-based healthcare might arrive not as the product of government policy, but as a necessary side effect of AI automation.

There are thus early lessons in the different reactions to AI between primary care and radiology. One sector is being caught by surprise and playing catch up to new commercial realities that have come more quickly than expected; the other has begun to reimagine itself in anticipation of becoming the ones that craft the new reality. The price each sector pays is different. Proactive preparation requires investment in reshaping workforce, and actively engaging with industry, consumers, and government. It requires serious consideration of new safety and ethical risks [15]. In contrast, reactive resistance takes a toll on clinical professionals who rightly wish to defend their patients' interests, as much as their own right to have a stake in them. Unexpected change may end up eroding or even destroying important parts of the existing health system before there is a chance to modernize them.

So, the fate of medicine, and indeed for all of healthcare, is to change [15]. As change makers go, AI is likely to be among the biggest we will see in our time. Its tendrils will touch everything from basic biomedical discovery science through the way we each

make our daily personal health decisions. For such change we must expect to pay a price. What is paid, by whom, and who benefits, all depend very much on how we engage with this profound act of reinvention. To fully engage brings promise of the greatest reward. To not engage is to pay the highest price.

References

1. Roy Amara 1925–2007, American futurologist. In: Ratcliffe S, editor. *Oxford Essential Quotations*. 4th ed; 2016.
2. Schwartz WB. Medicine and the Computer. The Promise and Problems of Change. *N Engl J Med* 1970;283(23):1257-64.
3. Kay P, Kempton W. What is the Sapir-Whorf hypothesis? *Am Anthropol* 1984;86(1):65-79.
4. Coiera E, Kocaballi B, Halamka J, Laranjo L. The digital scribe. *NPJ Digit Med* 2018;1:58.
5. Lyell D, Coiera E. Automation bias and verification complexity: a systematic review. *J Am Med Inform Assoc* 2017;24(2):423-31.
6. Siegler EL, Adelman R. Copy and paste: a remediable hazard of electronic health records. *Am J Med* 2009 Jun;122(6):495-96.
7. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. *Med Image Anal* 2017 Dec;42:60-88.
8. Darcy AM, Louie AK, Roberts LW. Machine learning and the profession of medicine. *JAMA* 2016;315(6):551-2.
9. Chen JH, Asch SM. Machine Learning and Prediction in Medicine - Beyond the Peak of Inflated Expectations. *New Engl J Med* 2017;376(26):2507-09.
10. Zech JR, Badgeley MA, Liu M, Costa AB, Titano JJ, Oermann EK. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS Med* 2018 Nov 6;15(11):e1002683.
11. Pan SJ, Yang Q. A survey on transfer learning. *IEEE Trans Knowl Data Eng* 2010;22(10):1345-59.
12. McCartney M. General practice can't just exclude sick people. *BMJ* 2017;359:j5190.
13. Fraser H, Coiera E, Wong D. Safety of patient-facing digital symptom checkers. *Lancet* 2018 Nov 24;392(10161):2263-4.
14. Marshall M, Shah R, Stokes-Lampard H. Online consulting in general practice: making the move from disruptive innovation to mainstream service. *BMJ* 2018 Mar 26;360:k1195.
15. Coiera E. The fate of medicine in the time of AI. *Lancet* 2018;392(10162):2331-2.

Correspondence to:

Enrico Coiera
Australian Institute of Health Innovation
Macquarie University
Level 6 75 Talavera Rd
Sydney, NSW 2109, Australia
E-mail: enrico.coiera@mq.edu.au