

High hopes for ‘Deep Medicine’? AI, economics, and the future of care

Robert Sparrow, Monash University

Joshua Hatherley, Monash University

NOT FOR PUBLICATION WITHOUT PERMISSION.

This is the “pre-press” version of a paper that was subsequently published as:

Sparrow, R. and Hatherley, J. 2020. High hopes for ‘Deep Medicine’? AI, economics, and the future of care. *Hastings Center Report* 50(1): 14-17. DOI: 10.1002/hast.1079.

Please cite that version.

* * * * *

In *Deep Medicine*, Eric Topol argues that the development of artificial intelligence (AI) for healthcare will lead to a dramatic shift in the culture and practice of medicine.¹ In the next several decades, he suggests, AI will become sophisticated enough for us to delegate many of the everyday tasks of physicians to it. According to Topol,

The promise of artificial intelligence in medicine is to provide composite, panoramic views of individuals' medical data; to improve decision-making; to avoid errors such as misdiagnosis and unnecessary procedures; to help in the ordering and interpretation of appropriate tests; and to recommend treatment. (p. 9)²

However, rather than replacing physicians, Topol suggests, AI could function alongside of them in order to allow them to devote more of their time to face-to-face patient care. Thus:

The greatest opportunity offered by AI is not reducing errors or workloads, or even curing cancer: it is the opportunity restore the precious and time-honoured connection and trust — the human touch — between patients and doctors. Not only would we have more time to come together, enabling far deeper communication and compassion, but we would also be able to revamp how we select and train doctors... Eventually, doctors will adopt AI and algorithms as their work partners. (p. 18)

Topol is perhaps the most articulate advocate of the benefits of AI in medicine, but he is hardly alone in spruiking its potential to allow physicians to dedicate more of their time and attention to providing empathetic care for their patients in the future.³ Unfortunately, these high hopes for AI-enhanced medicine fail to appreciate a number of factors that, we believe, suggest a radically different picture for the future of healthcare. Far from facilitating a return to “the golden age of doctoring”,⁴ the use of medical AI seems likely to further erode therapeutic relationships and threaten professional and patient satisfaction.

The fundamental problem with Topol’s optimistic vision for the future of medicine after AI is that substitutes fantasies about what AI might make possible for a realistic account of what it is likely to bring about.⁵ In particular, like many pundits who focus on technology rather than society when they think about the future, Topol neglects the role of economic and institutional considerations in determining how AI is likely to be used.

The economics of healthcare, especially where it is provided in a for-profit context, will dictate that any time savings made possible by a reduction in the administrative burdens on physicians in the course of patient consultations will be used to move more patients through the system rather than to allow practitioners to spend more time talking with, and caring for, their patients. Even in the public sector, the institutional drive to cost savings and efficiency prompted by concerns about the rising costs of healthcare,⁶ as well as concerns about social justice in access to healthcare, are likely to mean that AI is likely to be used to improve access to healthcare, by increasing the number of people that a given service can treat per day, rather than to increase the amount of time spent with each patient.

Another powerful institutional dynamic may be expected to parallel and reinforce this economic imperative. Organizations tend to concentrate — one might even say fixate — on things that they can measure rather than the more subtle and intangible aspects of their operations.⁷ Time per patient (or per procedure) is easily measured and optimized, whereas “care” is subtle and hard to measure. For this reason alone, there will be a tendency for institutions to use AI to treat more patients rather than devote more time to each patient.

Topol is conscious that institutions might adopt AI in ways that exacerbate rather than mitigate the dynamics that currently work to prevent physicians spending quality time with their patients. He writes,

The increased efficiency and workflow could either be used to squeeze clinicians more, or the gift of time could be turned back to patients — to use the future to bring back the past. The latter objective will require human activism, especially among clinicians, to stand up for the best interest of patients. (p. 21)

Topol hopes that physicians will mobilize politically to defend their interests – and the interests of patients — in longer conversations about care. We hope so too... but it is vital to acknowledge that this is a hope rather than a prediction. Moreover, there are a number of reasons to believe that is naïve hope.

Political action requires a confident and empowered group of people who share goals and (usually) an identity. Defeats make political action harder, whilst victories make it more likely. Unfortunately, the introduction of AI is likely to demoralise, fragment, and disempower the medical profession at the very point at which Topol expects doctors to rise up and demand better working conditions and outcomes for their patients. One thing that everyone agrees on in discussions of AI in medicine is that its introduction is likely to be highly disruptive to existing practices and institutions.⁸ Such disruption tends to be unsettling for those who work in the disrupted settings. Even if AI is unlikely to replace physicians entirely,⁹ it is likely to render redundant skills that the current generation of physicians spent years learning and have placed at the heart of their professional self-conception.¹⁰ Especially if combined with advances in robotics, AI may also break down complex tasks in healthcare into a number of different tasks that can be performed by people with smaller skill sets, as well as reduce the number of people required to be employed to complete various procedures. More generally, as with previous generations of information and computing technology, the introduction of AI into hospitals and healthcare settings is likely to lead to a shift in power and authority away from frontline practitioners to those who manage and design the IT systems.¹¹ Finally, research is already being directed toward using AI to monitor physician performance,¹² suggesting that physician surveillance will be one of the first uses of AI in the health sector. Physicians who are demoralized, disempowered, concerned for their jobs, and feel themselves to be under surveillance are ill-placed to win political victories.

It must also be observed that the historical record doesn't inspire much confidence here. Doctors in the US have as yet been unable to motivate US governments to adopt universal basic healthcare or even get the US public to endorse it, despite the fact that universal

healthcare would be in the interests of all Americans.¹³ They were unable to resist the rise of managed-care in the 1990s or the destructive impacts of the introduction of electronic medical records in the 2000's.¹⁴ Topol himself notes that

“it was, after all, doctors themselves who allowed the invasion of grossly inadequate electronic health records into the clinic, never standing up to companies like Epic, which has, in its contracts with hospitals and doctors, a gag clause that prohibits them from disparaging electronic health records or even publishing EHR screenshots.” (p. 288)

This history of failure provides little grounds for confidence that the medical profession will be able to resist the same economic, political, and institutional dynamics when it comes to the adoption of AI. Conversely, if one is concerned about care in medicine, there is little need to await the coming of AI to begin campaigning to defend it.

There are also a number of reasons to think that AI may reduce rather than increase the amount of time that healthcare practitioners have to spend talking with patients.

Most obviously, the fact that the lifeblood of AI is big data suggests that, as AI is introduced, the demand for things to be measured and recorded in medical settings will only increase.¹⁵ That is to say, healthcare workers may be expected to spend more time rather than less staring at screens and filling in forms on computers when they would rather be talking to patients. Again, the lesson of previous generations of technological change, which for the most part have shifted – or even increased – administrative burdens rather than relieved them is relevant here. For instance, when the introduction of computers and electronic health records into hospitals made it easier to record data, the result was that more data was demanded rather than that the same amount of data was recorded more swiftly. Importantly, the operations of AI are themselves likely to generate even more data both about the internal functionings of the systems and about their performance.

As we observed at the outset, Topol's hope is that AIs will record and manage all this data by themselves and thus not further burden healthcare workers. There may be some settings where this is the case. However, any optimism here should be tempered by the recognition that one of the lessons of AI research over the last six decades is that “sensing” turns out to be a much harder problem than calculating, planning, or analyzing. Despite remarkable

progress in natural language processing in recent years, extracting the meaning of interactions with patients in the clinic in real world conditions, which may require taking into account both patient and physician's accent, colloquialisms, body language, and social context, remains a formidable challenge. While patients may report to their healthcare providers with more and more data generated by their online-behavior, by apps, and by wearables, working out which datasets are relevant and integrating them with the patient's medical records often requires human judgement. Until we are prepared to rely entirely on AI for diagnosis, every new scan or test will demand that a clinician looks over the results. In the short-to-medium term, then, AI is likely to require human beings to provide the data that it needs to function.

It's also possible that the use of AI to gather and record data in some contexts may itself work to the detriment of care. Sometimes physicians gather information by examining, or asking questions of, the patient and this process is also an opportunity for the conversation to roam more widely and thus an opportunity for "care". This is especially the case where the process of talking to the patient is part of the process of diagnosis or the physical exams that support diagnosis. Gathering this sort of information automatically would actually reduce the opportunities for patients to feel that their physician was genuinely concerned for them.¹⁶

Finally, there is a profound tension in the idea that introducing more machines into the medical setting will lead to better relationships between physicians and patients. This is because AI will tend to undermine trust in doctors and because there are connections, both conceptual and empirical, between care and trust. Notoriously, AIs often function as "black boxes", with users – and sometimes even their designers – being unable to understand or explain why the AI produces the output that it does. If doctors start to rely on advice from AI the question will arise whether we should — indeed, how we could — trust our doctors. As Watson and colleagues note, "If doctors do not understand why the algorithm made a diagnosis, then why should patients trust the recommended course of treatment?"¹⁷ If we don't believe that it is our physician who is really making the decisions about our healthcare, it's hard to see how we could feel that they are caring for us. They might care *about* us but that's not the same as caring *for* us.

Indeed, this erosion of trust, with its detrimental impact on care, is likely to happen even if doctors could — if they tried hard enough — explain the outputs of the AI but in practice

don't make the effort to do so. There is a connection here to the question of the likely impact of the introduction of AI on the workload of doctors. If physicians want to retain the trust of their patients and remain the ultimate authority on treatment decisions they will need to supervise and review the operations of AI.¹⁸ At the very least, they will need to be able to assess when the AI is operating properly, which in turn will require being able to access the data on which the AI is relying and check that the conclusions of the AI are plausible in the light of that data. However, the more doctors are expected to do this, the more AI will add to their burden and take their attention away from the patient in front of them.¹⁹ Alternatively, doctors could take the results of the prognostications of AI on faith in the same way they do existing algorithms used in medicines or the conclusions of the peer-reviewed literature. But while patients are used to doctors relying, as we all do, on other people, doctors' reliance on AI is likely to be more disconcerting, especially as AI comes to take over roles, such as diagnosis, that have traditionally thought to be central to the profession of the physician.²⁰ If I come to see my doctor as the handmaiden to an AI, which is actually deciding on my treatment, then it may be difficult for me to understand my doctor as providing care.

None of this is to deny the potential of AI to promote any number of other goods in medicine, including, most importantly more timely and accurate diagnosis of a wide range of conditions. Advances in these areas are to be welcomed. Nevertheless, we should be conscious that they may come at a cost to care, given the current pressures on physicians and healthcare providers.

Topol hopes that AI will be used to expand opportunities for care but wishing for something does not make it so. The factors that have led to the decline in human contact in medicine are economic — which is to say, ultimately political — and it is naïve to think that technological change alone is likely to reverse this. If we want to ensure that AI increases the opportunities for, rather than erodes, care in medicine we will need to think deeper, not about AI but about the business of medicine and the institutional and economic contexts in which it is practised today.

¹ E. Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (New York: Basic Books, 2019).

² Later in the book he suggests that robots featuring AI might even perform some surgery (p. 161-162). See also A. Darzi, H. Quilter-Pinner, and T. Kibasi, *Better Health and Care for All: A 10-Point Plan for the 2020s. The Final Report of the Lord Darzi Review of Health and Care* (London: Institute of Public Policy Research, 2018), available at <https://www.ippr.org/research/publications/better-health-and-care-for-all>; X. Liu, P. A. Keane, and A. K. Denniston, “Time to Regenerate: The Doctor in the Age of Artificial Intelligence,” *Journal of the Royal Society of Medicine* 111, no. 4 (2018): 113-116; A. Verghese, N. H. Shah, and R. A. Harrington, “What This Computer Needs Is A Physician: Humanism and Artificial Intelligence,” *JAMA* 319, no. 1 (2018): 19-20; R. M. Wachter, *The Digital Doctor: Hope, Hype, and Harm at the Dawn of Medicine’s Computer Age* (New York: McGraw-Hill Education, 2015).

³ S. T. Israni and A. Verghese, “Humanizing Artificial Intelligence,” *JAMA* 321, no. 1 (2019): 29-30; B. Meskó, G. Hetényi, and Z. Györfy, “Will Artificial Intelligence Solve the Human Resource Crisis in Healthcare?” *BMC Health Services Research* 18, no. 1 (2018): 1 – 4.

⁴ J. B. McKinley and L. D. Marceau, “The End of the Golden Age of Doctoring,” *International Journal of Health Services* 32, no. 2 (2002): 379 – 416.

⁵ There is more than a whiff of what Broussard calls “technochauvinism” – the belief that smarter, faster, and cleaner technology is the solution to every problem – in Topol’s book, although it also includes passages that acknowledge that the impact of AI on medicine might be less positive than most of the text pretends, as we discuss further below. See M. Broussard, *Artificial Unintelligence: How Computers Misunderstand the World* (Cambridge, MA: MIT Press, 2018).

⁶ T. Diefenbach, “New Public Management in Public Sector Organizations: The Dark Sides of Managerialistic ‘Enlightenment’,” *Public Administration* 87, no. 4 (2009): 892 – 909.

⁷ J. A. Blythe and Farr A. Curlin, “‘Just Do Your Job’”: Technology, Bureaucracy, and the Eclipse of Conscience in Contemporary Medicine,” *Theoretical Medicine & Bioethics* 39, no. 6 (2018): 431 – 52.

⁸ Topol, *Deep Medicine*, 285.

⁹ Topol, *Deep Medicine*; Verghese, Shah, and Harrington, “What This Computer Needs Is A Physician.”

¹⁰ Compare, Organisation for Economic Co-Operation and Development (OECD), *OECD Science, Technology, and Innovation Outlook 2016* (Paris: OECD Publishing, 2016), 85, and E. J. Cassell, *The Nature of Suffering and the Goals of Medicine*, 2nd Edition (Oxford: Oxford University Press, 2004), 76. See also Wachter, *The Digital Doctor*.

¹¹ E. J. Cassell, *Doctoring: The Nature of Primary Care Medicine* (Oxford, UK: Oxford University Press, 1997), 79-80.

¹² R. D. Dias, A. Gupta, and S. J. Yule, “Using Machine Learning to Assess Physician Competence: A Systematic Review,” *Academic Medicine* 94, no. 3 (2018): 427 – 439.

¹³ J.-F. R. Lu and W. C. Hsiao, “Does Universal Health Insurance Make Health Care Unaffordable? Lessons from Taiwan,” *Health Affairs* 22, no. 3 (2003): 77-88.

¹⁴ D. Duijmelinck and W. van de Ven, “What can Europe Learn from the Managed Care Backlash in the United States?” *Health Policy* 120, no. 5 (2016): 509-518; M. W. Friedberg et al., *Factors Affecting Physician Professional Satisfaction and Their Implications for Patient Care, Health Systems, and Health Policy* (Santa Monica, California: The RAND Corporation, 2013); R. G. Hill Jr, L. M. Sears, and S. W. Melanson, “4000 Clicks: A Productivity Analysis of Electronic Medical Records in a Community Hospital ED,” *The American Journal of Emergency Medicine* 31, no. 11 (2013): 1591-1594; A. Verghese, “Culture Shock - Patient as Icon, Icon as Patient,” *The New England Journal of Medicine* 359, no. 26 (2008): 2748 – 2751.

¹⁵ T. M. Maddox, J. S. Rumsfeld, and P. R. O. Payne, “Questions for Artificial Intelligence in Health Care,” *JAMA* 321, no. 1 (2019): 31-32.

¹⁶ R. D. Truog, “Of Slide Rules and Stethoscopes: AI and the Future of Doctoring,” *Hastings Center Report* 49, no. 5 (2019): 3.

¹⁷ D. S. Watson et al., “Clinical Applications of Machine Learning Algorithms: Beyond the Black Box,” *BMJ* 364 (2019): I886.

¹⁸ Verghese, Shah, and Harrington, “What This Computer Needs Is A Physician.”

¹⁹ Maddox, Rumsfeld, and Payne, “Questions for Artificial Intelligence in Health Care.”

²⁰ Wachter, *The Digital Doctor*.