Paper

CLINICAL ETHICS

Clinical Ethics
1–4

DOI: 10.1177/14777509221138750
journals.sagepub.com/home/cet
SAGE

Why algorithmic speed can be more important than algorithmic accuracy

Jakob Mainz[1], Lauritz Munch[1], Jens Christian Bjerring[1] and Sissel Godtfredsen[2]

Abstract

Artificial Intelligence (AI) often outperforms human doctors in terms of decisional speed. For some diseases, the expected benefit of a fast but less accurate decision exceeds the benefit of a slow but more accurate one. In such cases, we argue, it is often justified to rely on a medical AI to maximise decision speed – even if the AI is *less* accurate than human doctors.

Keywords

Bioethics and medical ethics, clinical ethics, enhancement technologies, philosophical aspects

Introduction

Many have claimed that once medical AI systems become more accurate in their decision-making than human doctors, we should rely on them for diagnostic- and treatment purposes.[1] People have disagreed with this. Some have questioned whether medical AI systems are in fact more accurate than human doctors in clinical practice.[2] Others maintain that even if medical AI systems outperform humans in terms of accuracy, it might still be impermissible to rely on them because accuracy is not all that matters for medical decision-making. For instance, it has been argued that medical AI systems are objectionably opaque,[3,4] that they make biased decisions,[5] that replacing humans with AI systems creates 'responsibility gaps',[6] and that medical AI systems are insensitive to patients' values.[7]

Despite these well-known complaints about the use of AI systems in medicine, accuracy is generally considered the most important feature of artificial medical decision-making. After all, receiving the correct diagnosis is often the central criterion of medical success. In this comment, however, we will argue that there is another feature that is sometimes even more important than the accuracy: namely *decisional speed*. When talking about speed in medical decision-making, we have two things in mind. First, we can measure speed in terms of how *early* an AI can detect a disease. Second, we can measure speed in terms of how *fast* an AI can generate a medical verdict from the time at which the relevant input data is available. Although the two notions of algorithmic speed are distinct, we often care about the latter type of speed because we care to achieve the former type. While the importance of the accuracy of medical AI systems is widely appreciated, the importance of their decisional speed has – with few exceptions[8] – largely gone unnoticed in the literature. Here we argue that there are cases in which it is justified to rely on medical AI systems for reasons related to speed, even if the AI in question is less accurate than human doctors.

When speed is key

A key insight from decision theory is that there may be opportunity costs associated with slow decision-making.[9] That is, sometimes making a slow decision can mean that we miss out on options that are available only within a narrow window of opportunity. This insight is especially important for medical decision-making. For not only can faster decision-making free up resources for other medical purposes, thereby increasing cost-effectiveness. Early diagnosis and treatment are also essential for the treatment prospects of many medical conditions.

[1]Aarhus Universitet, Aarhus, Denmark
[2]Aalborg Universitetshospital, Aalborg, Denmark

Corresponding author:
Jakob Mainz, Aarhus Universitet, Aarhus, Denmark.
Email: jtm@cas.au.dk

While fast decision-making is, all else equal, better than slow decision-making, there are medical decisions for which the net benefits of a faster, but slightly less accurate decision are *particularly* substantial. Such cases tend to have two distinguishing features: (1) the opportunity costs of a slow decision-making process are severe and (2) the consequences of unnecessary treatment are acceptable compared to late, and even no treatment.

To give a simple example, think of the current recommendations for first responders to cardiac arrest The American Heart Association's general recommendation is to initiate cardiopulmonary resuscitation (CPR) no matter what.[10] In slogan form: it is better to do *something* rather than *nothing*. Diagnostic accuracy is not of foremost importance here. Although CPR may result in a broken ribcage, doing nothing may often result in the unconscious person dying in case she does have a cardiac arrest.

Another example is sepsis. Sepsis is an extreme, inappropriate reaction to an infection that can result in tissue damage, organ failure, or death.[11] It is one of the most common causes of death worldwide,[12] and its global mortality rate is roughly 30%.[13] Sepsis must be treated quickly. Every hour of delay in treatment increases the risk of death by 4–8%.[14] The opportunity costs of a slow diagnosis of sepsis are therefore severe. Since the primary treatment of sepsis is antibiotics and intravenous fluids, the consequences of unnecessary treatment are normally acceptable for the individual patient. Unfortunately, onset sepsis is notoriously difficult to detect because the symptoms are often vague. As such, many patients risk dying from sepsis because they receive treatment too late.

Accordingly, in cases like sepsis it seems uncontroversial that we should opt for the faster but slightly less accurate decision-making process. When the opportunity costs associated with making a slow but accurate decision exceed the costs associated with making a faster but slightly less accurate decision, we should often pick the latter. For how improper would it not be – in light of the fatal opportunity costs associated with sepsis – to insist on a slower but more accurate diagnosis, when a faster one is available that allows us to initiate a simple treatment involving antibiotics and intravenous fluids? Indeed, even if there is only inconclusive evidence that we are in fact dealing with sepsis, faster diagnosis might still be preferable because of the severe consequences of a slower diagnosis.

Medical AI to the rescue

When fast decision-making is a valuable end, relying on medical AI systems can be a sensible means to that end. After all, there are many medical contexts in which AI systems can arrive at sensible decisions much faster than human doctors. This has been documented for many diseases such as tuberculosis,[15] Alzheimer's disease,[16] and, interestingly, sepsis.[17]

Medical AI systems that can detect sepsis onset have been trained and tested with great success.[18] In fact, some sepsis detection models can detect onset of sepsis several hours before human doctors, and, as we saw above, this advantage in decisional speed can make the whole difference between life and death for the patient concerned. A model called *Sepsis Watch* has attracted a lot of attention.[19] It can predict onset sepsis a median 5 hours before clinical presentation, and it generally performs very well both in terms of triaging and monitoring of patients. Perhaps the most interesting results coming from the team behind *Sepsis Watch* is that the model has been successfully implemented.[20] And in contrast to most cases where attempts have been made to implement AI systems in a clinical setting, the stakeholders involved with the *Sepsis Watch* model have generally welcomed it. Not only does the model outperform humans in terms of decisional speed. The team behind it also claims that they have found ways to mitigate concerns from model opaqueness, fairness, and so on.[21] So in cases of diseases like sepsis where conditions (1) and (2) from above are easily satisfied, we are often justified to implement an AI in clinical practice and to rely on it for purposes of medical decision-making.

What does it mean to 'rely' on an AI system for purposes of medical decision-making? We can distinguish between 'pure reliance' and 'partial reliance'. When we *purely* rely on an AI system, we allow the AI to make independent medical decisions within some predefined domain. By contrast, when we *partially* rely on an AI system, a human doctor has executive decision power but uses the AI as a decision support tool, or as a tool for diagnosing and initiating treatment until the doctor has time to see the patient. Both types of reliance can speed up medical decision-making processes. But pure reliance may often be much faster because disagreements between the AI system and the human doctors – in contrast to cases of partial reliance – will not arise when full decision power is allocated to the AI system. Resolving disagreements between AI systems and human doctors can be time consuming because it can require second opinions from other doctors.[22] In the case of sepsis, the risk of a time-consuming disagreement is high because there is much disagreement among doctors about how sepsis is defined, and what the relevant symptoms are.[23]

Accuracy versus speed

There is an ongoing debate about whether we should hold medical AI systems to the *same* standard as human doctors, or if we should hold them to a *higher* standard.[24,25] In lieu of the above, however, we believe that the speed of medical AI systems provides a reason to hold them to a *lower* standard than human doctors when it comes to accuracy in decision-making. As such it can be justified to increase the speed of an AI system by deliberately reducing its accuracy – at least to an extent.

But how much accuracy should we be willing to trade for a speedier decision-making process? It depends in part on the distribution of errors that comes with a decrease in accuracy. If a decrease in accuracy results mostly in a larger number of *false positives*, then in cases like sepsis, we can set the accuracy threshold relatively low. The costs that individual patients have to endure in cases of unnecessary treatment are low: people rarely suffer from receiving antibiotics and intravenous fluids. But in cases where a decrease in accuracy results mostly in an increase in *false negatives*, we need to proceed more cautiously. After all, an increase in false negatives implies that more people go untreated and therefore risk suffering severe consequences.

Still, even if the accuracy level of the AI system is lower than that of the doctor, and even if the AI produces both more false negatives and more false positives than the doctor, we can nevertheless reduce the total number of deaths caused by sepsis by employing the AI system. The reason is that a significant portion of the human doctors' *true positives* – patients correctly diagnosed with sepsis – end up dying because they are diagnosed and treated too late in the process.[26]

As illustrated by figure 1, early diagnosis by an AI can lead to fewer patients dying of sepsis even though the AI system is less accurate than the human doctor, and even though the AI produces both more false positives and false negatives than the human doctor.

Note that it is generally not a viable option to have human doctors diagnose sepsis at T_1 because the patterns in the input data at T_1 revealing symptoms of sepsis are typically too complex for human minds to grasp – but not too complex for an AI system. Indeed, if typical human doctors tried to diagnose patients at T_1 instead of T_2, their diagnostic accuracy would likely drop way below that of the AI at T_1.

Thus, we have argued in this article, when the benefits of a fast but less accurate medical decision exceed that of a slower but more accurate decision, we are often justified in relying on AI systems for purposes of medical decision-making. This conclusion applies even if the AI systems are less accurate than human doctors. Although one must be cautious of the AI system's distribution of errors – especially false negatives – reliance on a speedy AI system can result in fewer fatal cases than reliance on a human doctor even when the doctor is more accurate and produces fewer false positives and false negatives than the AI system.

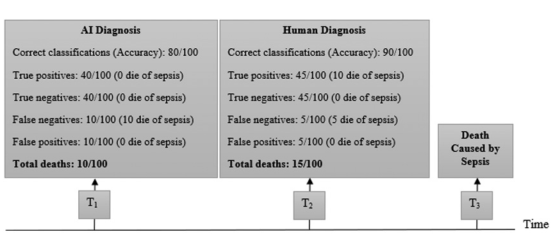

Figure 1. Comparison between an AI diagnosing 100 patients at time T_1, and a human doctor diagnosing 100 patients later at T_2. Even though the AI performs *worse* in terms of both accuracy, false negative rates, and false positive rates, the total number of deaths caused by sepsis are still lower for the AI because some of the human doctor's true positives end up dying of sepsis at T_3 due to delayed diagnosis and treatment.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article. This work was supported by the Carlsbergfondet (grant number CF20-0257).

ORCID iD

Jakob Mainz https://orcid.org/0000-0002-7766-6439

References

1. Norgeot B, Glicksberg B and Butte A. A call for deep-learning healthcare. *Nat Med.* 2019; **25**: 14–15.
2. Grote T. Trustworthy medical AI systems need to know when they don't know. *Jour Med Eth.* 2021; **37**: 337–338.
3. Alvarado R. Should we replace radiologists with deep learning? Pigeons, error and trust in medical AI. *Bioeth* 2021; **36**: 121–133.
4. Bjerring B and Busch J. Artificial intelligence and patient centered decision-making. *Phil Tech* 2021; **34**: 349–371.
5. Owens K and Walker A. Those designing healthcare algorithms must become actively anti-racist. *Nat Med* 2020; **26**: 1327–1328.
6. Santoni de Sio F and Mecacci G. Four responsibility gaps with artificial intelligence: Why they matter and how to address them. *Phil Tech* 2021; **34**: 1057–1084.
7. Birch J, Creel K, Jha A, *et al.* Clinical decisions using AI must consider patient values. *Nat Med* 2022; **28**: 229–232.
8. Sorell T, Rajpoot N and Verrill C. Ethical issues in computational pathology. *Jour Med Eth* 2022; **48**: 278–284.
9. Payne J and Bettman J. When time Is money: Decision behavior under opportunity-cost time pressure. *Org Beh Hum Dec Pro* 1996; **66**: 131–152.
10. https://cpr.heart.org/en/resources/bystander-cpr
11. Gyawali B, Ramakrishna K and Dhamoon A. Sepsis: The evolution in definition, pathophysiology, and management. *Open Med* 2019; **7**: 1–13.
12. Komorowski M, Celi L, Badawi O, *et al.* The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. *Nat Med* 2018; **24**: 1716–1720.
13. Vincent JL, *et al*. Assessment of the worldwide burden of critical illness: the intensive care over nations (ICON) audit. *The Lanc Respir Med*. 2014; **2**: 380–386.

14. Zhao X, Shen W and Wang G. Early prediction of sepsis based on machine learning algorithm. *Comp Int Neu*. 2021; 6522633.
15. Pasa F, Golkov V, Pfeiffer F, *et al.* Efficient deep network architectures for fast chest X-ray tuberculosis screening and visualization. *Sci Rep.* 2019; **9**: 6268.
16. Jo T, Nho K and Saykin A. Deep learning in Alzheimer's disease: Diagnostic classification and prognostic prediction using neuroimaging data. *Front Agi Neu.* 2019; **11**: 220–220.
17. Islam M, *et al.*. Prediction of sepsis patients using machine learning approach: A meta-analysis. *Comp Met Prog Bio.* 2019; **170**: 1–9.
18. Fleuren L, *et al.* Machine learning for the prediction of sepsis: a systematic review and meta-analysis of diagnostic test accuracy. *Int Care Med.* 2020; **46**: 383–400.
19. https://dihi.org/project/sepsiswatch/
20. Sendak, *et al.* Real-world integration of a sepsis deep learning technology into routine clinical care: Implementation study. *JMIR Med Inform*. 2020; **8**: e15182.
21. https://dihi.org/outcome/the-human-body-is-a-black-box-supporting-clinical-decision-making-with-deep-learning-2/
22. Kempt H and Nagel S. Responsibility, second opinions and peerdisagreement: ethical and epistemological challenges of using AI in clinical diagnostic contexts. *Jour Med Eth*. 2022; **48**: 222–229.
23. Lopansri B, *et al.* Physician agreement on the diagnosis of sepsis in the intensive care unit: estimation of concordance and analysis of underlying factors in a multicenter cohort. *Jour Int Care*. 2019; **7**: 1–17.
24. Zerilli J, Knott A, Maclaurin J, *et al.*. Transparency in algorithmic and human decision-making: Is there a double standard?. *Phil Tech.* 2019; **32**: 661–683.
25. Günther M and Kasirzadeh A. Algorithmic and human decision making: for a double standard of transparency. *AI & Soc.* 2021; **37**: 375–381.
26. Prescott H, *et al.* Late mortality after sepsis: propensity matched cohort study. *Br Med J.* 2016; **353**: i2375.