

This is a pre-print version of an article published as:

Hatherley, Joshua. 2020. Limits of trust in medical AI. *Journal of Medical Ethics* 46(7): 478-481. [10.1136/medethics-2019-105935](https://doi.org/10.1136/medethics-2019-105935).

Please cite that version.

LIMITS OF TRUST IN MEDICAL AI

Joshua Hatherley

Abstract: Artificial intelligence (AI) is expected to revolutionise the practice of medicine. Recent advancements in the field of deep learning have demonstrated success in variety of clinical tasks: detecting diabetic retinopathy from images, predicting hospital readmissions, aiding in the discovery of new drugs, etc. AI's progress in medicine, however, has led to concerns regarding the potential effects of this technology upon relationships of trust in clinical practice. In this paper, I will argue that there is merit to these concerns, since AI systems can be relied upon, and are capable of reliability, but cannot be trusted, and are not capable of trustworthiness. Insofar as patients are required to rely upon AI systems for their medical decision-making, there is potential for this to produce a deficit of trust in relationships in clinical practice.

Introduction

Artificial intelligence (AI) is expected to revolutionise the practice of medicine. Recent advancements in the field of deep learning have demonstrated success in variety of clinical tasks; for instance, detecting diabetic retinopathy from images,¹ predicting hospital readmissions,² and aiding in the discovery of new drugs.³ It has been suggested that AI will facilitate a variety of improvements in medical practice, ranging from economic savings to the improvement of empathetic communication between doctors and patients, from increased productivity to greater professional satisfaction, and from improved health outcomes to an amplified rate of discovery in medical research.⁴ AI's progress in medicine, however, has led to concerns regarding the potential effects of this technology upon relationships of trust, particularly between doctors and patients.⁵ In this paper, I will argue that there is merit to these concerns, since AI systems are not the appropriate objects of trust under any familiar philosophical accounts of trust. This is critical, since, as I will argue in section 3, AI systems are likely to displace

the epistemic authority of human clinicians if they come to exceed them in performance. As such, I will argue that insofar as patients are required to rely upon AI systems for their medical decision-making, AI threatens to produce a deficit in trusting clinical relationships between doctors and patients.

Trust in Medicine

Trust has both intrinsic and instrumental significance in medicine.ⁱ Intrinsically, trust is what imbues the doctor-patient relationship with its uniqueness and importance. A patient comes to a physician in a state of sickness and vulnerability, and is thereby forced to place their trust in another person to treat them with competence and, ideally, empathy and care. This vulnerability of the patient is what imbues the relationship with inherent value, since “trust is inseparable from *vulnerability*, in that there is no need for trust in the absence of vulnerability.”⁶ The vulnerability of the patient, and the resulting power of the physician, imbue the physician with a fiduciary obligation to behave in a morally upright and appropriate manner, to use their authority in the service of the patient as opposed to themselves or some other end.

In contrast, trust also has instrumental value in medicine. Firstly, because patients are more likely to accept and behave in accordance with their physician’s judgement if they have a trusting relationship with them. They are more likely to demonstrate “willingness to seek care, reveal sensitive information, submit to treatment, participate in research, adhere to treatment regimens, remain with a physician, and recommend physicians to others”.⁶ Secondly, it is speculated that trusting doctor-patient relationships have a number of therapeutically valuable effects upon patients – improved patient outcomes and placebo effects, for example. Finally, a good physician is one that can demonstrate care for their patients, and patients are more likely to feel that they have been adequately cared for when they trust the person caring for them.

ⁱ Two kinds of trust are discussed in relation to medicine and clinical practice: interpersonal and social.²⁷ Interpersonal trust concerns trust between persons (e.g. between doctors and patients), while social trust is more general and abstract, directed towards groups and institutions as opposed to individuals (e.g. between patients and a particular hospital or the medical institution more generally). In this paper, I leave the issue of social trust to one side in order to focus in on medical AI and interpersonal trust. All references to trust will henceforth refer exclusively to interpersonal trust.

AI in Medicine

AI's effect on relations of trust between doctors and patients is bound up with the precise role that AI may come to occupy in medical practice and the epistemic authority that it comes to hold in clinical decision-making procedures. If AI systems are eventually adopted as merely another tool at the clinician's disposal – akin to a stethoscope, thermometer, or blood pressure monitor – the effect of these systems on trust would likely be minimal. Patients, of course, would rely on the accuracy of these tools, but their trust would be staked in the judgement of the human physician who interprets their outputs and incorporates them into their own clinical judgements. However, recent developments in areas such as deep learning suggest that the epistemic authority of human clinicians in clinical decision-making will be challenged by the use of AI in medicine.

Researchers in AI are working busily to develop AI systems that can surpass the performance of human clinicians in diagnosis, prognosis, and treatment selection⁴ – three of the four fundamental tasks of the clinician, according to Eric Cassell.^{7,ii} Indeed, a recent systematic review and meta-analysis comparing the performance of deep learning AI systems to human clinicians found that deep learning AI systems already match the accuracy of human clinicians in the performance of certain diagnostic tasks.^{8,iii} If AI succeeds in surpassing the performance of human clinicians in such principal medical tasks, how might this effect the epistemic authority of human clinicians in clinical practice?

The prospect gestures at an important problem currently faced in the sciences, which Paul Humphreys has called our 'anthropocentric predicament'. Humphreys argues that advanced technologies have produced a situation in which

“an exclusively anthropocentric epistemology is longer appropriate because there now exist superior, non-human, epistemic authorities. So we are now faced with a problem, which

ⁱⁱ The fourth is the identification of causes. Given that AI systems based on neural networks learn from correlations alone, their capacity to illuminate underlying causes of illness is limited.²⁸

ⁱⁱⁱ Importantly, the study identified a number of troubling methodological limitations in the broader literature comparing the performance of human clinicians to deep learning AI systems, so this finding ought to be taken with a grain of salt. Most alarmingly, of the 31,587 scholarly articles returned on a search for articles comparing the performance of deep learning systems and human clinicians, only 14 compared performance between the two groups on the same test dataset.

we can call the *anthropocentric predicament*, of how we, as humans, can understand and evaluate computationally based scientific methods that transcend our own abilities.”⁹

There have been two principal kinds of response to medicine’s anthropocentric predicament in the wake of medical AI, which I will refer to as substitutionism and extensionism. Substitutionists argue that advanced AI will eventually make doctors obsolete by surpassing them in the performance of key clinical tasks and roles.¹⁰ Extensionists, in contrast, argue that AI will simply extend and improve upon the capabilities and competencies of human clinicians without replacing them outright. In particular, this is because AI systems lack emotional intelligence and empathy, abilities that are essential in the delivery of healthcare, meaning that a human presence will still be essential.¹¹ Yet among both camps, the likely disruptive impact – what Liu and colleagues have labelled a “seismic shift”¹² – that AI will have on medicine is widely undisputed. Although extensionists rally against the substitution of clinicians, the likelihood of their displacement in key clinical roles is often acknowledged. For instance, Eric Topol, a principal physician advocate for the use of AI in medicine, along with his colleague Saurabh Jha, claim that “[j]obs are not lost; rather, roles are redefined; humans are displaced to tasks needing a human element.”¹³

This displacement of the roles of human clinicians in the wake of advanced medical AI reflects a displacement of their epistemic authority. Indeed, if AI surpasses the performance of clinicians in key clinical tasks, doctors will have an epistemic obligation to defer to the judgements of the machine or align their judgements with the AI in their clinical decision-making.¹⁴ As Bjerring and Busch have argued, “if a practitioner knows of an epistemic source that is more knowledgeable, more accurate, and more reliable in decision-making, she should treat it as an expert and align her verdicts with those of the source.”¹⁵ This displacement of the epistemic authority of clinicians would be necessary to realise some of the goals of the introduction of AI in medicine. Aside from the possible reduction of burdensome administrative tasks and the improvement of cost-efficiency in medicine, a primary motivation for research into medical AI is the potential to reduce the alarming prevalence of wastefulness and human error in medical practice.^{4,16} In order to achieve this, it would be necessary in most instances for human clinicians to give more weight to the outputs of a supremely reliable AI system over their own clinical intuitions and judgements.

The displacement of clinicians from a position of epistemic authority in clinical decision-making has important implications for relations of trust between patients and doctors, since it implies a displacement of patient trust from human clinicians to AI systems. In the next section, I will argue that this displacement of trust from humans to machines could lead to shallow relations of trust in clinical practice that are lacking in important respects.

Trust in AI

Trust has been a central topic of concern in the debate over AI and its many applications, with some private corporations and research organisations releasing guidelines for the development of trust and trustworthiness in AI.^{17,18} Concerns over the ‘black box’ nature of some AI systems – particular deep learning – along with the threat of algorithmic bias have pushed the issue of trust to the forefront of debate.¹⁹ But what does it mean to say that one trusts an AI, or that an AI is trustworthy? A key response to this question has been to emphasise the centrality of reliability in trust. Alex John London claims that “[i]f the goal is to secure trust among stakeholders, then the accuracy of a system relative to viable alternatives must be a central concern.”²⁰ Similarly, Zachary Lipton claims that if trust is “simply confidence that a model will perform well [... then] a sufficiently accurate model should be demonstrably trustworthy.”²¹

But is confidence in someone or something’s accuracy or reliability sufficient for trust? According to many accounts of interpersonal trust have been proposed in the philosophical literature, the answer to this question is no. According to these accounts, trusting someone to do *x* is more than merely relying on them to do *x*. Consider the following two scenarios:

- (1) Stan, a thief, is planning a burglary. He has observed a wealthy homeowner, Jane, leaving her home at 9am and returning at 7pm every Monday for the past month. Stan is hoping to go through with his planned burglary next Monday, and is relying upon Jane to continue her pattern in order for his burglary to be successful.
- (2) Brendan has a chronic illness that causes him significant pain and suffering. His illness is managed by his regular GP, Dr. Smith. Dr. Smith has supported Brendan through his illness for 15 years. Brendan has recently been experiencing significantly more pain than usual, which is causing him extreme discomfort. He makes an appointment

with Dr. Smith, confident that she will be able to help him relieve this pain in some way.

In the scenario (1), although the thief relies upon Jane to leave her house at 9am, it seems inappropriate to say that the thief trusts Jane to do so in the same way that Brendan trusts Dr. Smith to successfully treat his illness in scenario (2), despite the fact that Brendan also relies upon Dr. Smith. How do we explain this intuition? What makes trusting someone more than merely relying upon them?

Russell Hardin argues that reliance is insufficient for trust because trusting someone also requires a belief that one's interests are encapsulated in the interests of the trusted person.²² "What matters", claims Hardin, "[..] is not merely my expectation that you will act in certain ways but also my belief that you have the relevant motivations to act in those ways, that you deliberately take my interests into account because they are mine."²² For Hardin, trust requires not only a predictive expectation on the part of the truster, but also a belief that one's interests are encapsulated in the interests of the trusted person and that the trusted person has the right motivations for action. Indeed, Hardin claims that "I would not, in our usual sense, trust a fully programmed automaton, even if it were programmed to discover and attempt to serve my interests – although I might come to rely heavily on it."²²

Having the right kind of motivations for action is an important part of many other influential accounts of trust. Annette Baier, for instance, argues that reliance underdetermines trust because trust "seems to be reliance on [the trusted person's] good will toward one, as distinct from their dependable habits, or only on their dependably exhibited fear, anger, or other motives compatible with ill will toward one, or on motives not directed to one at all."²³ This emphasis upon the good will of the trusted person is also central to Karen Jones' account, wherein she claims that "to trust someone is to have an attitude of optimism about her goodwill and to have the confident expectation that, when the need arises, the one trusted will be directly and favorably moved by the thought that you are counting on her."²⁴ If the right kind of motivations are necessary for the kind of trust that we would usually recognise as interpersonal trust, then AI systems would not appear to be the appropriate objects of this kind of trust. Unlike a human clinician, AI systems have no goodwill towards us, nor any motivation to act in our interests. This may be at least part of the reason that some people may be

uncomfortable with the idea of placing their trust in an AI for important medical decisions or tasks.

Additionally, other philosophical accounts of trust distinguish between trust and reliance on the basis of normative and descriptive expectations: I *rely* on you when I *predict* that you will behave in a certain way, though I *trust* you when I judge that you *ought* to behave in a certain way.²⁵ Trusting someone, that is, generates an obligation on behalf of the trusted person to (at least genuinely attempt to) do what one is trusting them to do. There are some important limitations to this claim, e.g. in circumstances where the trust that one has in another is misguided or unwelcome. Suppose, for instance, that one were to place their trust in a friend who is a dermatologist to remove their wisdom teeth. Trusting the dermatologist for this procedure would appear quite mistaken, given that the dermatologist does not have the expertise or competency to perform this task. Nor, presumably, would the dermatologist welcome this trust in any way.

But outside of this and other somewhat fanciful scenarios, clinicians do in fact have an obligation to perform those tasks that have been entrusted to them, providing of course that this trust has been communicated to them. This is precisely the nature of fiduciary obligations in medicine. If this is true, another limitation of trusting AI would also be demonstrated, since AI systems are not the appropriate objects of moral responsibility. In order for an agent to be morally responsible for an action, they must be blameworthy when they fail to come through on that action. But if an AI system were to incorrectly diagnose a patient, leading to their avoidable death, it would appear misguided or inappropriate to blame the AI for its error. Rather, one would generally look to the designers, the supervising clinician, the hospital, etc. in order to apportion blame. Trusting a clinician generates a moral responsibility on behalf of the clinician, while trusting an AI system generates a moral responsibility on behalf of seemingly anyone but the AI system.

These considerations highlight two important deficits in relations between patients and medical AI systems that each stem from a lack of agency on the part of the AI. Firstly, AI systems lack the right kind of motivation for trust - either in the form of encapsulated interest or a sense of good will - since they lack motivation entirely. Secondly, relations with AI systems cannot be said to be trusting relations, as one might have with a human clinician, since trust generates normative obligations that cannot be borne by an AI. To say that one can trust an

AI is thus akin to saying that one can trust a naturally occurring phenomenon. Although I am supremely confident that tomorrow the sun will rise in the east and set in the west, there is not familiar sense in which I could reasonably said to *trust* the sun to do so. Trusting relations, in other words, are exclusive to beings with agency, meaning that the displacement of human clinicians from a position of epistemic authority and privilege in the clinical encounter threatens to lead to relations of trust that are shallow or deficient in important respects within medical practice.

Conclusion

To say that one can trust an AI system, or that the AI is trustworthy, is merely to say that one can rely on the AI system, or that the system is reliable. Yet as we have seen, reliability is insufficient to generate a relation of trust under any of its familiar philosophical notions, which all require characteristics essential and exclusive to beings with a form of agency. What does this mean for the pursuit of ‘Trustworthy AI’ initiated by the European Union’s High Level Expert Group on Artificial Intelligence (HLEG AI)?¹⁷ Although valuable, the pursuit of trustworthy AI represents a notable conceptual misunderstanding, since AI systems are not the appropriate objects of trust or trustworthiness. Interestingly, this has also been suggested by a key member of the HLEG AI, Thomas Metzinger.²⁶ Rather than trustworthy AI, this pursuit may be better served by being reframed in terms of reliable AI, reserving the label of ‘trust’ for reciprocal relations between beings with agency.

In contrast to AI, therefore, human clinicians can offer their patients the kind of rich interpersonal trust that imbues the doctor-patient relationship with its uniqueness and significance. Insofar as patients come to rely upon AI systems for important medical assessments and decisions as opposed to human clinicians, they may be sacrificing opportunities for trusting relationships in medicine. A more thoughtful engagement is needed with the potential effects of AI on medical practice to further understand the implications of this technology, so that it can be deployed in such a way as to reap its potential benefits whilst retaining those aspects of medicine – such as trust – that are particularly valuable for its functioning.

REFERENCES

1. Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA* -

- J Am Med Assoc.* 2016;316(22):2402-2410. doi:10.1001/jama.2016.17216
2. Caruana R, Lou Y, Microsoft JG, Koch P, Sturm M, Elhadad N. Intelligible Models for Health Care: Predicting Pneumonia Risk and Hospital 30-day Readmission. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* New York, NY: ACM; 2015:1721-1730. doi:10.1145/2783258.2788613
 3. Ramsundar B, Kearnes S, Riley P, Webster DR, Konerding D, Pande V. Massively Multitask Networks for Drug Discovery. *arXiv.org.* 2015;1502.02072.
 4. Topol EJ. *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again.* New York, NY: Basic Books; 2019.
 5. Nundy S, Montgomery T, Wachter RM. Promoting Trust Between Patients and Physicians in the Era of Artificial Intelligence. *JAMA - J Am Med Assoc.* 2019:1-2. doi:10.1001/jama.2018.20563
 6. Hall MA, Dugan E, Zheng B, Mishra AK. Trust in physicians and medical institutions: what is it, can it be measured, and does it matter? *Milbank Q.* 2001;79(4):613-639. <https://anzjsurg.onlinelibrary.wiley.com/doi/full/10.1111/1468-0009.00223>.
 7. Cassell EJ. *The Nature of Suffering: And the Goals of Medicine.* Vol 53. 2nd ed. Oxford, UK: Oxford University Press; 2004. doi:10.1017/CBO9781107415324.004
 8. Liu X, Faes L, Kale AU, et al. A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. *Lancet Digit Heal.* 2019;1(6):e271-e297. doi:10.1016/S2589-7500(19)30123-2
 9. Humphreys P. The philosophical novelty of computer simulation methods. *Synthese.* 2009;169(3):615-626. doi:10.1007/s11229-008-9435-2
 10. Goldhahn J, Rampton V, Spinax GA. Could artificial intelligence make doctors obsolete? *BMJ.* 2018;363(November):1-3. doi:10.1136/bmj.k4563
 11. Krittanawong C. The rise of artificial intelligence and the uncertain future for physicians. *Eur J Intern Med.* 2018;48:e13-e14. doi:10.1016/j.ejim.2017.06.017
 12. Liu X, Keane PA, Denniston AK. Time to regenerate: the doctor in the age of artificial intelligence. *J R Soc Med.* 2018;111(4):113-116. doi:10.1177/0141076818762648
 13. Jha S, Topol EJ. Adapting to Artificial Intelligence: Radiologists and Pathologists as Information Specialists. *JAMA - J Am Med Assoc.* 2016;316(22):2353-2354.

- doi:10.1001/jama.2016.17438
14. Grote T, Berens P. On the ethics of algorithmic decision-making in healthcare. *J Med Ethics*. 2019;1-7. doi:10.1136/medethics-2019-105586
 15. Bjerring JC, Busch J. Artificial Intelligence and Patient-Centered Decision-Making. *Philos Technol*. 2020. doi:10.1007/s13347-019-00391-6
 16. Makary MA, Daniel M. Medical error-the third leading cause of death in the US. *BMJ*. 2016;353(May):1-5. doi:10.1136/bmj.i2139
 17. High-Level Expert Group on Artificial Intelligence. *Ethics Guidelines for Trustworthy AI*. Brussels: European Commission; 2019.
 18. IBM. IBM's Principles for Data Trust and Transparency. <https://www.ibm.com/blogs/policy/trust-principles/>. Published 2018. Accessed January 13, 2020.
 19. Sparrow R, Hatherley J. High Hopes for "Deep Medicine"? AI, Economics, and the Future of Care. *Hastings Cent Rep*. 2020;50(1):14-17. doi:10.1002/hast.1079
 20. London AJ. Artificial Intelligence and Black-Box Medical Decisions: Accuracy versus Explainability. *Hastings Cent Rep*. 2019;49(1):15-21. doi:10.1002/hast.973
 21. Lipton ZC. The Mythos of Model Interpretability. *Commun ACM*. 2016;61(10):36-43. <http://arxiv.org/abs/1606.03490>.
 22. Hardin R. *Trust and Trustworthiness*. New York: Russell Sage Foundation; 2002.
 23. Baier A. Trust and Antitrust. *Ethics*. 1986;96:231-260. doi:10.2307/j.ctt22rbkfw.29
 24. Jones K. Trust as an Affective Attitude. *Ethics*. 1996;107(1):4-25.
 25. Nickel PJ, Vaesen K. Risk and Trust. In: Roeser S, Hillerbrand R, Sandin P, Peterson M, eds. *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*. Springer Netherlands; 2012:857-876. doi:10.1007/978-94-007-1433-5
 26. Metzinger T. EU guidelines: Ethics washing made in Europe. *Der Tagesspiegel*. <https://www.tagesspiegel.de/politik/eu-guidelines-ethics-washing-made-in-europe/24195496.html>. Published 2019.
 27. Mechanic D. Changing Medical Organization and the Erosion of Trust. *Milbank Q*. 1996;74(2):171-189.

28. Marcus G. Deep Learning: A Critical Appraisal. *arXiv.org*. 2018:1-27.
<http://arxiv.org/abs/1801.00631>.