Artificial Intelligence to Inform Clinical Decision Making: A Practical Solution to An Ethical And Legal Challenge

Helen Smith
Centre for Ethics in Medicine, University of Bristol, England, UK
helen.smith@bristol.ac.uk

Abstract

This position paper discusses the challenges of allocating legal and ethical responsibility to stakeholders when artificially intelligent systems (AISs) are used in clinical decision making and offers one possible solution. Clinicians have been identified as at risk of being subject to the tort of negligence if a patient is harmed as a result of their using an AIS in clinical decision making. An ethical model of prospective and retrospective personal moral responsibility is suggested to avoid clinicians being treated as a 'moral crumple zone'. The adoption of risk pooling could support a shared model of responsibility that could promote both prospective and retrospective personal moral responsibility whist avoiding the need for negligence claims.

1 Introduction

My PhD considers the ethical and legal challenges which arise due to using artificial intelligence to inform clinical decision making.

I started by performing a literature review [Smith, 2020] which found that there was no body of legal case law regarding who would be allocated legal responsibility if patient harm resulted due to AIS use in clinical decision making. This review also found that the allocation of ethical responsibility is also somewhat unclear. The remainder of my PhD has been spent addressing these two areas.

2 Legal Analysis

I recently published legal analysis with Kit Fotheringham regarding the use of artificially intelligent systems (AIS) and tort of negligence in England and Wales [Smith and Fotheringham, 2020]. Our investigations found that software developing companies (SDCs) might be protected from claims as clinicians are the users of AIS; that clinical users are vulnerable to negligence claims due to *novus actus interveniens.* We speculated that SDCs could be burdened with duty of care if harm resulted due to a clinician using an AIS but that proving causation might be challenging.

The legal analysis felt unfair to clinicians and made me consider if the legal position could be challenged ethically. I felt that if the outcome of this speculative legal analysis were to eventuate, there is a risk that clinicians would become legal and moral crumple zones [Elish, 2019] for the SDCs; i.e. that clinicians would absorb the legal and moral penalties when an AIS provides a clinician with recommendations for patient care which are harmful rather than helpful.

3 Can the Legal Analysis be Ethically Challenged?

Rather than accepting that clinicians ought to be held conclusively responsible as per the legal analysis, I employed ethical theories to determine where responsibility for AISs could also be allocated.

I noted a key aim: that SDCs want their AIS used and the clinician is the user. What interested me was how that dynamic was controlled by SDCs.

I found a key quote about the famous AIS, IBM's Watson for Oncolgy, from an unnamed Executive Consultant for IBM in Hengstler *et al* [2016, p.115]:

> *"I underline that Watson does not make decisions on what a doctor should do. It makes recommendations based on hypothesis and evidence based [sic]"*

Here, the AIS is being presented by the SDC; the SDC has stipulated the way that the AIS's recommendation is to be considered when used. An SDC may dictate to users how their AIS is to be used, which ensures that they achieve their aim of someone using their AIS whilst implying that the SDC does not carry the burden of responsibility if the clinician's use of that AIS results in patient harm.

I found it thought provoking that I have found no consultations in the literature or the media when researching this thesis which describe clinicians and technologists bearing

responsibility for the AIS use jointly. This seems strange when one considers that the AIS has been presented by the SDC to advise clinical practice. NHSX [2020] recently announced a Multi-Agency Advisory Service which will help SDCs and clinicians to evaluate AISs and navigate regulatory issues in the context of the National Health Service in England and Wales; yet whilst an AIS could be deemed acceptable for clinical use by regulators, I have found no examples of cross-disciplinary negotiations where SDCs and clinicians collectively decide how much value may be placed on an AIS's recommendations when the clinician makes their decisions. Instead, as indicated by the earlier quote captured by Hengstler *el al*, the SDC may release their AIS, pronounce its limits, and dictate how the clinician should act when using their AIS. This prescriptive activity by the SDCs places clinical users in Elish's [2019] moral crumple zones.

4 Prospective and Retrospective Moral Responsibility

To establish how responsibility could be allocated more fairly between these two stakeholders, I explored several ethical theories. This paper is limited to personal moral responsibility, as discussed by Zimmerman [1992].

Personal moral responsibility can be broken down into two types: forwards looking and backwards looking [Zimmerman, 1992, p.1089]:

1) Prospective (forwards looking) personal moral responsibility can be identified as moral or legal or defined by a set of rules. Duty of care (having personal moral responsibility towards an individual) is an example of prospective responsibility.

2) Retrospective (backwards looking) personal moral responsibility is to be personally morally responsible for an outcome. Here approval or disapproval for actions would be expressed morally or legally.

I argue that, alongside clinicians, it is not unreasonable for SDCs to be allocated prospective personal moral responsibility for the effects of their AIS used in clinical decision making. I make this claim as, according to Fuscaldo [2006, p.70], an agent is responsible to the consequence of an action if (if and only if) their actions are voluntary and have foreseeable consequences. No one is forcing SDCs to create AISs and it is foreseeable that the consequences of the use of AIS will affect patients as AISs are designed specifically to influence clinical decision making. Because of this foreseeability, it is reasonable to allocate a duty of care to SDCs who develop and deploy AISs to be used in clinical decision making. Also, because of foreseeability, it is unfair to burden clinicians with sole personal moral responsibility; SDC's ought to carry their own responsibility for the effects of the use of their AISs. However, it is also unfair also to allocate sole responsibility of the use of AISs to SDCs alone. If SDCs were to accept their duty of care to patients, they would join the clinical professions who have historically embraced their duty of care. Whitby [2015] suggests that responsibility could be shared between SDCS and clinicians; I agree as both actors have contributed to the use of AISs in clinical decision making. Thus, I argue that a shared model of responsibility between clinicians and SDCs should be sought.

5 A Shared Model of Responsibility

To identify a shared model of responsibility, I consider Zimmerman's [1992] prospective and retrospective models of personal moral responsibility. Whilst Zimmerman discusses how prospective and retrospective models are considered separately, I argue that, in this scenario, they could be considered together. Thus: whilst an agent can assume a duty of care and takes steps to reduce that harm, they may also simultaneously plan how to address a foreseeable potential/possible harm should it eventuate. These actions taken together before the agent acts demonstrates a mixed approach to prospective and retrospective responsibility. To be at their most effective, i.e. by both looking forward to the consequences of one's actions and to look back to take responsibility for an outcome which they generated, an actor is able to consider both their prospective and retrospective personal moral responsibility obligations together.

Applying this to the use of AIS's in clinical decision making, I suggest that stakeholders could consider how the use of an AIS could cause harm, attempt to reduce the risk of that harm eventuating, and have a plan to address patient harm should it eventuate. This plan needs to be developed and in place before the AIS is used by clinicians. A shared model of responsibility could be designed to contain prospective and retrospective characteristics and thus may provide a practical solution which could be fairer to stakeholders.

6 Risk Pooling: a solution?

Risk pooling was offered as a solution in my publication with Kit Fotheringham [Smith & Fotheringham, 2020]. Risk Pooling is a prospective arrangement between all stakeholders involved in the design, deployment and use of an AIS who have a duty of care to retrospectively address possible harms which could eventuate to a patient. It allows a platform for clinical users and SDCs to discuss the use of an AIS and the risks and benefits which would arise from that use. For reasons of transparency and reducing complexity, a single risk pool could be constructed for each individual

AIS in use, rather than a larger risk pool which would spread the risk of several AISs from different SDCs.

Risk pooling is similar to enterprise liability (offered by Allain, 2013) in that a risk pooling insurance scheme would cover harms when they arise, but differs as the defendants only contribute in accordance with the extent of their liability. Risk pooling can be engineered so that it relieves a large portion of a harmed patient's burden of making a negligence claim. With a well-defined and planned risk pooling arrangement in place, an injured patient can make a claim easily, commence their recovery journey sooner, and avoid the additional distress of pursuing a potentially lengthy and uncertain legal process. After a claim is made against the risk pool, a court may engage in detailed consideration of the legal issues as they apply to the defendants.

Risk pooling is beneficial here as it directly serves the patient's interests in the absence of case law which confirms where legal responsibility lies. As negligence law appears to place SDCs in a more favorable position than clinical users, risk pooling could provide a path to restoring justice through both actors contributing to the risk pool. Once case law has been generated and legal processes are more certain, then the cost and conditions of membership for each actor to the risk pool may be reconsidered and adjusted. However, such contribution adjustments may be calculated using ethical considerations which take into account the actions of all stakeholders, and not just legal precedent.

Whilst risk pooling does not itself make AISs safer, it's requirement would create a financial incentive for an SDC to develop the AIS to be as safe as possible prior to deployment and for a clinician to use it as safely as possible. This could be achieved by strategies such as relieving the 'black box' nature of an AIS by deploying AIS which are interpretable to stakeholders and ensuring that the AIS's failings were visible to the end user. SDCs could additionally proactively and prospectively train users to safely use AISs in clinical practice. NHSX's [2020] new Multi-Agency Advisory Service may wish to consider providing new AIS-specific guidelines for end-user clinicians who need to evaluate whether tools meet 'minimum safety' requirements.

6 Conclusion

Risk pooling could be an answer to ethically sharing responsibility for harms caused due to AIS use. This model provides an example of how stakeholders can work together to utilise AISs at the bedside having confidently made the forward planning needed to account for the possible consequences of that AISs use.

Acknowledgements

This PhD has only been possible due to the wise, kind, and skilled supervision of Dr Jonathan Ives, Dr Giles Birchley and Professor Andrew Charlesworth. The legal analysis which I performed with Kit Fotheringham laid the foundation for my subsequent ethical considerations to be made. I am grateful for the kind and constructive comments made by this workshop's reviewers. My husband, Mr Chris Martin, has been a constant source of support, encouragement, and coffee throughout my PhD.

Thank you all! □

References

[Allain, 2013] Jessica S. Allain. 'From Jeopardy! to Jaundice: The Medical Liability Implications of Dr Watson and Other Artificial Intelligence Systems', *Louisiana Law Review*. Available at: https://digitalcommons.law.lsu.edu/lalrev/vol73/iss4/7/

[Elish, 2019] Madeleine Clare Elish. Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction. *Engaging Science, Technology and Society*. DOI: https://doi.org/10.17351/ests2019.260

[Fuscaldo, 2006] Giuliana Fuscaldo. Genetic ties: are they morally binding? *Bioethics*. DOI: https://doi.org/10.1111/j.1467-8519.2006.00478.x

[Hengstler *et a*l, 2016] Monika Hengstler, Ellen Enkel, and Selina Duelli. Applied artificial intelligence and trust—the case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*. DOI: https://doi.org/10.1016/j.techfore.2015.12.014

[NHSX, 2020] NHSX. Regulating the AI ecosystem. Available at: https://www.nhsx.nhs.uk/ai-lab/ai-lab-programmes/regulating-the-ai-ecosystem/

[Smith, 2020] Helen Smith. Clinical AI: opacity, accountability, responsibility and liability. *AI & Society*. DOI: https://doi.org/10.1007/s00146-020-01019-6

[Smith and Fotheringham, 2020] Helen Smith and Kit Fotheringham. Artificial intelligence in clinical decision-making: Rethinking liability. *Medical Law International*. DOI: https://doi.org/10.1177/0968533220945766

[Whitby, (2015] Blay Whitby. 'Automating medicine the ethical way', in Machine Medical Ethics: Intelligent Systems, Control and Automation: Science and Engineering. Kluwer Academic Publishers, pp. 223-232.

[Zimmerman, 1992] Michael J. Zimmerman. 'Responsibility', in Lawrence C. Becker and Charlotte B. Becker. (eds.) *Encyclopaedia of Ethics Volume II: L-Index*. London: Garland Publishing Inc.