

Artificial Intelligence and Resource Allocation in Health Care: The Process-Outcome Divide in Perspectives on Moral Decision-Making

Sonia Jawaid Shaikh

Annenberg School for Communication, University of Pennsylvania
3620 Walnut St, Philadelphia, PA 19104, United States
sjshaikh@asc.upenn.edu

Abstract

Pandemics or health emergencies create situations where the demand for clinical resources greatly exceeds the supply leading to health providers making morally complex resource allocation decisions. To help with these types of decisions, health care providers are increasingly deploying artificial intelligence (AI)-enabled intelligent decision support systems. This paper presents a synopsis of the current debate on these AI-enabled tools to suggest that the existing commentary is outcome-centric i.e. it presents competing narratives where AI is described as a cause for problematic or solution-oriented abstract and material outcomes. Human decision-making processes such as empathy, intuition, and structural and agentic knowledge that go into making moral decisions in clinical settings are largely ignored in this discussion. It is argued here that this process-outcome divide in our understanding of moral decision-making can prevent us from taking the long view on consequences such as conflicted intuition, moral outsourcing and deskilling, and provider-patient relationships that can emerge from the long-term deployment of technology devoid of human processes. To preempt some of these effects and create improved systems, researchers, providers, designers, and policymakers should bridge the process-outcome divide by moving toward human-centered resource allocation AI systems. Recommendations on bringing the human-centered perspective to the development of AI systems are discussed in this paper.

Introduction

Resource allocation is a type of decision-making which involves the procurement, assignment, and distribution of resources between actors. Typically, decision-making pertaining to resource allocation is considered to be a difficult enterprise because access or ownership of resources is competitive and can not only affect individuals and group health,

but a set of interconnected socio-economic variables. Furthermore, this type of decision-making has a moral dimension as it requires humans (e.g. physicians, hospital administrators, nurses, etc.) to make tradeoffs such as those involving utilitarian and egalitarian parameters which exacerbates its complexity (Robert et al. 2020). In a pandemic such as COVID-19, the demand for resources (e.g. beds, ventilators, ICU units, etc.) is many times greater than their supply which complicates the problem of allocation. This situation forces health care providers to set up a variety of triage protocols to determine if and to what extent someone qualifies for one or more resources.

One of the ways providers are making resource allocation decisions to deal with the COVID-19 pandemic is to use AI-enabled decision support systems. These systems can use patient's electronic health records (EHR) and/or clinical measurements (e.g. blood pressure, fever, health conditions) to make diagnoses and prognoses (Lamanna and Byrne 2018; Medlej 2018) which can be subsequently used to make decisions about the level of care and allocation of resources to patients (Debnath et al. 2020).

Since the technology can affect people's lives in multiple ways, it renders itself as a problem concerning social good, and thus deserves our attention. Furthermore, as it becomes more sophisticated and widely deployed, it needs to be understood now to preempt any negative consequences on human decision-making practices and to generate helpful policies and guidelines. To enhance our understanding on this matter, this paper presents a brief synopsis of the debates surrounding the use of AI to help humans make moral decisions pertaining to resource allocation during health crises. The core argument made here is that the current debate on

the deployment of AI-enabled decision support systems comprises competing narratives that are mostly outcome-centric i.e. they focus on the abstract (e.g. fairness) and material (e.g. saving costs) outcomes AI technology can yield. By emphasizing upon these outcomes, we fundamentally ignore the human decision-making processes that health care providers also use in addition to established clinical guidelines when allocating resources. The disregard of human processes of decision-making in the pursuit of AI's use in resource allocation might have long-term consequences on the design of technology and providers' abilities to make decisions. My hope is that this paper will provide insights to researchers, designers, and policymakers to bridge process-outcome divide to implement more human-centered technology for moral decision-making in clinical settings.

To support these arguments, the following sections begin by presenting a brief synthesis of current perspectives on using AI for resource allocation during COVID-19. We observe that the predominant ideas on making resource allocation decisions reflect a process-outcome divide where the current debate is heavily tilted toward framing AI as an entity that can create outcomes which are either problematic or solution-oriented in nature. This approach disregards human processes such as empathy, intuition, and structural and agentic knowledge that health care providers rely upon to make resource allocation decisions. This is followed by noting the possible effects of long-term deployment of outcome-centric technology on human decision-making such as disruptions in intuitive processes, moral outsourcing and deskilling, and relationships with patients. It is recommended here that the process-outcome divide must be bridged by creating human-centered AI systems which can preempt adverse consequences for providers and patients. Human-centered AI systems in health care can be incorporated by having providers work as co-developers of technology, building in-house AI capabilities, and developing regulations pertaining to the use of AI in health care settings.

AI and Resource Allocation: An Outcome-Centric Approach to Decision-Making

Fairness is a recurrent theme in the discussion on the use of AI to make moral decisions. The concept of fairness when applied to any resource allocation process refers to the instance whereby a decision results in reduction of biases that may prioritize one group over another or where it increases equity between different stakeholders. The goal of any decision-maker, therefore, is to increase fairness as an outcome. However, the debate on the use of AI with reference to fairness is a debate comprising competing narratives. Some argue that AI-based tools in health care resource allocation during COVID-19 pandemic are useful because machines are driven by complex logic and predetermined parameters,

and therefore, can be fair and less biased resource allocators (Shea et al. 2020) . This perspective frames AI-enabled intelligent decision support systems as solutions that can amplify fairness.

However, a competing narrative suggests the opposite. It argues that AI systems can create problems as they may reflect existing systemic biases and thus, are more likely to make unfair appraisals exacerbating inequality between different racial and socioeconomic groups (Röööslı et al. 2020). For instance, a study found that a commercial algorithm factored in health care costs more heavily compared to physiological symptoms of illness which led to sicker Black patients being provided with fewer services (Obermeyer et al. 2019). In the case of a pandemic such as COVID-19, an AI-based tool might use underlying health conditions (e.g. obesity, heart problems) or disability to predict a lower chance of recovery for a patient suffering from virus-induced complications which would affect the likelihood of them receiving a hospital bed. An AI system such as this is likely to perpetuate unfair outcomes by giving people who suffer from ailments due to socioeconomic inequalities a lower chance of accessing a health resource.

In addition to fairness, another theme concerning the use of AI tools in health care resource allocation pertains to the effects of AI's computational prowess on various abstract and material outcomes. It has been argued that unlike humans, machines have extraordinary computing and information processing power which allows them to gather, analyze, and interpret data quickly which ultimately helps with timely diagnosis and provision of care (Shea et al. 2020). This leads to not only determining who gets resources, but helps save health care costs, effort, and time (Adly et al. 2020). A differential take suggests that AI's very computational reach and scope can backfire if a corrupt algorithm incorrectly diagnoses or distributes resources to many people in a very short amount of time.

It is evident that current perspectives discussed above on deploying AI to make resource allocation decisions are concerned with tackling abstract (e.g. fairness) or material outputs (e.g. costs). Such framing is not only outcome-centric, but overly simplifies complex phenomena concerning moral decision-making such as resource allocation. The application of AI-enabled technology to allocate resources occurs in human contexts where decision-makers use distinct and identifiable processes to make moral decisions. An AI-enabled system neither accounts nor substitutes for these processes and therefore, the need for identifying and discussing human processes becomes ever more important to not only fill theoretical voids, but also affect how we design and implement technology.

The Role of Human Decision-Making Processes in Diagnoses and Resource Allocation

Clinical decision-making in the realm of resource allocation is a multifaceted activity which requires providers to use complex decision-making processes in addition to predefined scientific and well-established protocols. The sections below present a brief discussion on the distinct processes which are regularly employed in moral decision-making by health care providers but are yet to be explored within the context of resource allocation especially during a pandemic.

Intuition

Intuition refers to an information processing mode which lacks conscious reasoning but incorporates affective and cognitive elements to make decisions (Sinclair and Ashkanasy 2005). Doctors and caregivers often use their intuitions as a part of their clinical decision-making processes in addition to using the guidelines and medical procedure (Van de Brink et al. 2019; Rew 2000). Evidence for the use of intuition or 'gut feelings' to allocate resources is evidenced across other cultures (Le Reste et al. 2013; Ruzca et al. 2020). Intuitive decision-making process can affect how health care providers make diagnostic recommendations which lead to allocation of services. For instance, findings from a large-scale study on has shown that doctors' sentiments affected the number of tests their patients received in an ICU setting (Ghassemi et al. 2018). This suggests that providers' intuition plays an important role in making diagnostic and subsequently resource allocation decisions.

Structural and Agentic Knowledge

Health providers often have a deep understanding of structural and agentic variables that underlie and affect their day-to-day operations. For instance, they may know how their hospital's location affects resource procurement, patient arrival and admission. Providers are also more are more likely to be aware of differences in hospital personnel personality types, work ethics, interpersonal relationships, cultural values and political beliefs, bureaucratic procedures and administrative conduct, equipment issues, etc. Together, this conscious and subconscious knowledge of structural and agentic factors can guide providers' moral decision-making and thus, how they allocate resources (Lemieux-Charles et al. 1993). The current technology can hardly substitute this knowledge as it exists beyond the purview of an AI-enabled tool.

Empathetic Concern

Many diagnostic and moral decisions are driven by empathetic concern for others (Selph et al. 2008). Hence, it is no surprise that care providers often take an empathetic approach to identify illnesses or allocate specific resources.

Empathetic concern plays a very important role in identifying biases in any system, policy, and practice. For instance, when we reflect on a process empathetically, we are more likely to understand how it affects people which can therefore allow us to intervene to help and make changes for them (Batson 2016). To illustrate this further, let's reconsider the findings on the use of algorithm which led to Black patients being given access to fewer resources despite them being sicker since the tool factored the costs of health care more heavily when allocating services (Obermeyer et al. 2019). If the AI program were developed using an empathetic approach, then it might have accounted for the fact that Blacks on average have lower income than White patients and thus, are less likely to spend on hospital services despite having more physiological symptoms. This shows that the design and use of technology to make moral decisions without any concern for empathetic human processes can reflect in the outcomes technology produces.

While intuition, structural and agentic knowledge, and empathetic concern are important processes that help and guide moral decision-making in clinical (and non-clinical settings), they are largely ignored in the debate on the use of AI to allocate resources since competing narratives are mostly focused on the outcomes the technology produces. This raises the question of what the deployment of AI means for human processes in decision-making. To help researchers, administrators, and policy makers engaged in long-term planning and thinking, I present some reflections on the possible effects of AI-enabled tools devoid of elements concerning human processes.

Deployment of AI-Enabled Tools for Resource Allocation: A Note on Potential Consequences

The focus of this paper is on moral decisions pertaining to resource allocation especially within pandemic-related settings. Since many people compete for limited resources and there is little time to make decisions; it is tempting, and in some cases, advantageous to apply AI-based tools. However, when humans use AI-enabled tools to make moral decisions, their internal decision-making processes are likely to be affected or influenced by technology. This could affect how providers assess, analyze, and treat patients. However, short-term solutions can potentially have long-term unintended and unwanted effects. The following sections note some of the consequences on providers' decision-making processes concerning intuition, knowledge, and empathetic concern which may occur as a function of long-term deployment of AI.

Disrupted and Conflicted Intuition

As AI continues to be incorporated in moral decision-making, providers will have to divide their attention between the

AI's recommendations and their own intuitive judgement especially if they have different or opposing ideas. They will face the added tension of determining tradeoffs between the machine and their own morals (Grote and Berens 2020) especially when applied to moral decision-making such as resource allocation. Such scenarios will require additional human cognitive input and will be more likely to interrupt the intuitive approaches doctors already use to make allocate scarce resources decisions.

It is arguable that the addition of AI may could facilitate the doctors' decision-making processes by sharing the cognitive burden pertaining to diagnostic evaluation. However, it is important to note that the process of moral decision-making comprises more than a more than a mechanical diagnostic endeavor. It also includes how users react to and accept or react suggestions from AI. Prior research has shown people's tendency to both accept and reject advice from algorithms (Dietvorst et al. 2014; Logg et al. 2019) and therefore, it is likely that such judgements pertaining to the recommendations made by AI will also be made by doctors in conjunction with their own intuitive responses.

Disrupted and conflicted intuition is likely to affect the internal moral compass decision-makers use to organize their worlds. It is also going to reflect in how they allocate resources where some may exclusively rely on technology to mitigate their internal tensions and others may develop their own course of action. Although it is possible that providers use a combination where they select when to choose intuitive or machine judgement to make decisions, this will be a hard skill to learn, and thus, difficult to use especially in emergencies where decisions have to be made quickly.

Moral Outsourcing and Deskilling

Assigning and rationing resources between people is an issue that is directly tied to the issues of ethics and morality. Making moral decisions can be a difficult and distressing process because it typically involves trade-offs concerning self-interests and group needs, personal and cultural values, and immediate and future rewards within the context of health care (McCarthy and Deady, 2008; Wright et al. 1997). As such, it requires that a decision-maker gives them careful attention, thought, deliberation, along with engaging in interactions with others. Moral decision-making is thus, a skill that is learned over time and with consistent practice by care providers. The deployment of AI-based tools to help with moral decision-making creates an increased risk for moral outsourcing i.e. the tendency to allow machines to make moral decisions for us (see Danaher 2015). This is especially likely due to the human bias where machines are often considered fairer (Lee 2018). Thus, while the use of machine-based tools to help us make difficult decisions is inevitable, an over reliance on AI to make ethical and moral

decisions is problematic because it may lead to moral deskilling.

Health Care Provider-Patient Relationships

A patient's relationship with their health provider is based on several factors including the providers' abilities to empathize and make decisions that show off their competence, expertise, and clinical prowess (Larson and Yao 2005). With the incorporation of AI in clinical settings, some have argued that the use of AI could augment providers' competency by helping them think about alternative diagnostic options or providing them with feedback on their performance. These factors could amplify the trust patients place in physicians (Nundy et al. 2019). This could be one of the outcomes of AI application in a regular clinical setting. However, pandemics with high-mortality rates where resource shortages affect day-to-day functioning of hospitals and clinics, the use of AI to make critical diagnostic and subsequently allocation decisions could theoretically be viewed differently by patients. Reliance on AI could lead patients and their families to question providers' competency to and treat and care for patients along with their ability to be fair. Patient doubts on providers' competence could amplify if the technology commits errors or is found to be biased (Nundy et al. 2019). Thus, we can imagine that situations such as these could easily erode the trust and belief patients and their families place in health care providers.

Bridging the Process-Outcome Divide: Toward the Development of Human-Centered AI Resource Allocation Systems in Health Care

Now that we have identified that there is a process-outcome divide on how moral decision-making is conceptualized, discussed, and applied within clinical settings, the next question is what we can do about it? The bridging of the process-outcome divide in moral decision-making can occur with the development of human-centered resource allocation AI systems as applied to clinical settings. A human-centered approach to AI development incorporates the perspectives and processes of users that use intelligent systems (see Xu 2019). Thus, AI systems using a human-centered approach are more likely to create synergy between both human decision-making processes and machine outcomes to positively affect and amplify both physician and patient welfare. That being said, the challenge remains as to how we can achieve human-centered design in the deployment and development of AI tools.

To overcome this challenge, we first need to further unpack the process-outcome divide in the context of moral decision-making pertaining to resource allocation within a pandemic (or non-pandemic) setting. The "process" in the

process-outcome divide refers to the *human* decision-making processes such as empathy, intuition, and structural and agentic knowledge which health care providers use (in addition to pre-determined clinical protocols and guidelines) to make diagnostic judgements and allocation decisions. While the “outcome” refers to the machine-driven or related consequences or functionality such as maximizing fairness or computational capacity. Hence, process-outcome divide by nature can be said to also imply a *human-machine divide*. Note that here the term human-machine divide is not meant in the same way as its prior use in the context of technical features of the machine and how they are informed by human biology (e.g. neurons) (see Warwick 2015). The focus here is on how the process-outcome divide in perspectives on moral decision-making pertaining to resource allocation reflect the split between human processes of decision-making and machine-generated outcomes pertaining to resource allocation decisions. I argue that this divide could potentially be addressed by moving toward human-centered AI systems which will require recognizing and iteratively integrating human *processes* in the development, deployment, management, and regulation of AI-enabled systems. To this end, the following sections present some recommendations on how human processes can be weaved into the development of AI systems.

Health Care Providers as Co-Developers of AI Technology

An admittedly simplified way of understanding how AI-enabled technology is scaled is to reflect on two stages: development and deployment. More often than not, these tools are developed either independently (i.e. by manufacturer/company or within academic settings) or in some consultation with health care providers. Once developed, they may be pitched to various health care providers where the technology is customized to their needs. Sometimes, the technology is rolled out in phases where it is tested on a smaller level and subsequently expanded to include more patients or units (Gago et al. 2005). Thus, by and large, development of AI technology is followed by its deployment with lagged or punctuated feedback from the user to the developer. This practice indicates a bifurcation between the developers and users (here: health care providers).

This approach to the scaling of an AI-enabled decisions support system may seem natural and functional. However, I argue that for including human processes of decision-making in how technology is used; it is best to see development and deployment linked together in an interactive and iterative process where they inform each other. This is particularly important in health care settings where the availability of and access to resources vary and the environment (e.g. infection and mortality rates, deaths, policy, information) change rapidly and often unpredictably.

To create an iterative loop between development and deployment, the lines between users of technology have to be blurred. While developers of technology can understand its various technical aspects and have the requisite knowledge and skills to build it; the users can often envisage its effects and uses more deeply due to their day-to-day experience, exposure to patients’ needs, and structural and personnel-based issues in clinical settings. To understand this further, let us imagine that an AI program is designed to help decide if patient gets a bed during a pandemic. The program conducts a risk assessment of the severity of patient’s condition by assigning scores on a pre-determined set of factors. One of the factors relates to prior health condition where the AI assigns a score in case a patient has any (e.g. heart problem). However, health care providers may know via their day-to-day experiences that it is likely for a patient without a prior health record in the hospital’s system and yet having an underlying condition to arrive in the emergency department. The patient might be unaccompanied and unable to report their medical history due to physical ailment or language barrier. They could also be unaware of their underlying medical condition. In such a scenario, the use of an AI program that determines if a resource (e.g. bed.) can be allocated to a patient based on the above-mentioned criterion may not be the most appropriate option. If providers and developers work in an interactive and iterative fashion, then these observations could be passed on to the developers who may be able to account for these issues i.e. a lack of prior medical record, or language barrier, or being unaccompanied along with obvious severity in symptoms when assessing risk. An AI program could then use a different scoring system which accounts for these variables to allocate beds. Thus, continual integration of human processes via relevant updates and modifications is more advantageous than one-time testing or multi-phase testing with a pre-determined end.

Prior research has shown that users and developers can co-create technology by contributing their differing expertise in a process called cooperative prototyping (Bødker and Grønbæk 1991). However, the rapidly changing health environments require us to leap from the cooperative prototyping approach to iterative cooperative development and management of technology. When health care providers will act as co-developers of technology, it will allow them to fuse human processes (e.g. empathy, intuition) involved in moral decision-making to build and update AI systems.

That being said, it must be mentioned that from normative and prescriptive perspectives, the extension of providers as co-developers of technology may sound like an appealing and useful idea. However, from a practical point of view, this may prove to be a difficult enterprise because it would require health care facilities to dedicate personnel and their time toward the development of such systems. Presumably, this thought may be a deterrent for some to adopt such

measures and protocols. However, I argue that in the long run, this will be a small cost to bear. A team of health professionals who are dedicated to testing AI-enabled intelligent decision support systems can only better the technology which in turn will produce superior outcomes and decrease the cost of day-to-day operations as well as reduce risks associated with poorly designed systems. Consider this with reference to the latest developments in space science. IBM developed a robot called CIMON which was tested for its efficacy by an astronaut aboard the International Space Station (ISS) (CIMON brings AI to the International Space Station n.d.). The feedback and testing allowed for a new and upgraded robot CIMON-2 to be sent to the ISS (IBM 2020). The developers and user (i.e. astronaut) of the space robot played important roles in both the development and deployment of the technology before it could be used in a high-stakes environment such as a space mission. Health care settings should be treated no less than a space mission as they are high-stakes and expense-laden environments which affect socioeconomic and mortal outcomes for billions of people around the world. It logically follows then AI-enabled decision support systems within health care should not only be tested regularly but be informed by the very users who employ it making critical resource allocation decisions.

Developing AI-Focused In-House Capacities

As decision-makers, people are managed by others such as human resource departments, administrative procedures or protocols, upper managements, etc. These institutional actors and protocols manage human activities, solve issues, and recommend further actions. Management also extends to medical equipment as hospitals and clinics often have technical staff or support teams on site or procured via third-party contracts. However, such staff or administrative units are often missing when it comes to overseeing AI-enabled technology. Many health care settings deploy AI with little to no oversight of these systems since their management requires particular skill sets. Increasing sophistication of AI-enabled systems and their authority to pass judgement (assign risks to patients/calculate scores) makes them not only tools but also decision-making actors to some extent. As actors and tools, they too, need supervision. Therefore, health care administrators will need to develop in-house expertise and create departments that are specifically dedicated to the monitoring, modification, and management of AI-enabled systems or embodied intelligent assistants such as robots which have increasingly become a part of health care settings.

Such an endeavor would have several benefits as it would: a. allow the integration of providers' perspectives within the AI system, b. identify any issues quickly, and c. make modifications to the system potentially within the clinical

settings or outsource them to the third-party or original developers within a short period of time.

Developing AI-Specific Regulations, Protocols, and Ethical Guidelines and Educating Providers

It was argued above that one of the long-term consequences of using AI to make moral decisions could manifest as humans putting in more cognitive effort and challenging their intuition especially if personal judgement were at odds with the advice given by an AI program. It was also suggested that an over-reliance on intelligent systems could lead to moral outsourcing and deskilling. To preempt such scenarios, health care providers will need to create specific protocols, regulatory, and ethical guidelines to regulate the use of AI within their premises. These guidelines and directives will need to specify whose judgment—i.e. human or AI—will be considered as the final say when making a diagnostic or allocation decision. These guidelines will also have to delineate parameters of assigning culpability and responsibility if choices made by a human in conjunction with or against the advice of AI results in adverse outcomes. These regulations will help providers understand their roles in moral decision-making and allow them to continue sharpening their skills when it comes to moral decision-making in the presence of AI.

Additionally, medical schools and educational programs will also need to train providers and students on how to interact with AI, evaluate its judgement, and effects on human decision-making. Together, these practices will allow providers to better understand AI, its management, and relationship with humans within clinical settings.

Conclusion

The core argument presented in this paper is that the discussion on AI decision support tools used in moral decision-making such as resource allocation within clinical settings provides competing narratives which delineate the pros and cons of AI in terms of the material and abstract outcomes the technology produces. Such perspectives distract us from focusing on the role human decision-making processes such as empathy, intuition, and structural and agentic knowledge play in resource allocation decisions. This scenario reflects process-outcome divide in the current perspectives on moral decision-making within health care settings. If these human processes are disregarded while AI is used to make moral decisions, it may result in long-term consequences such as conflicted intuition, moral outsourcing and deskilling, and poor patient-provider relationships. To preempt some of these consequences and create better health outcomes; researchers, developers and policymakers seriously consider the importance of human processes along with machine-driven outcomes. One of the ways we can bridge the

process-outcome divide is to create human-centered AI systems specific to health care. To this end, some recommendations are proposed: a. health care providers work with developers of technology as co-developers in an iterative and interactive fashion, b. health care facilities should develop in-house AI expertise and create a specific department to manage, regulate, and modify the technology, and c. regulatory protocols and guidelines specific to the use of AI in making moral decisions should be developed. These guidelines should be able to specify how and when humans should override AI decisions. It should also delineate rules on culpability should a decision made in conjunction or against AI advice produce adverse effects. Furthermore, providers and students should be trained on understanding the effects of AI on their decision-making. Together, these endeavors could help with taking and implementing a broader and more human-centered perspective on the use of AI in health care to advance social good.

References

- Adly, A. S., Adly, A. S., and Adly, M. S. 2020. Approaches Based on Artificial Intelligence and the Internet of Intelligent Things to Prevent the Spread of COVID-19: Scoping Review. *Journal of Medical Internet Research*. 22(8): e19104. doi.org/10.2196/19104
- Batson, C. D. 2016. Empathy and Altruism. In *The Oxford Handbook of Hypo-egoic Phenomena*, edited by K. W. Brown, and M. R. Leary, 161-174. New York: Oxford University Press
- Bødker, Susanne, and Grønbæk, K. 1991. Cooperative Prototyping: Users and Designers in Mutual Activity. *International Journal of Man-Machine Studies*. 34(3): 453–78. doi.org/10.1016/0020-7373(91)90030-B
- CIMON Brings AI to the International Space Station. Accessed October 15, 2020. Retrieved from https://www.ibm.com/thought-leadership/innovation_explanations/article/cimon-ai-in-space.html.
- CIMON-2 Masters Its Debut on the International Space Station, IBM. April 15, 2020. Retrieved from <https://newsroom.ibm.com/2020-04-15-CIMON-2-Masters-Its-Debut-on-the-International-Space-Station>.
- Danaher, J. 2018. Toward an Ethics of AI assistants: An Initial Framework. *Philosophy & Technology*. 31(4): 629-653. doi.org/10.1007/s13347-018-0317-3
- Debnath, S., Barnaby, D. P., Coppa, K., Makhnevich, A., Kim, E. J., Chatterjee, S., Tóth, V., Levy, T. J., Paradis, M. D., Cohen, S. L., Hirsch, J. S., Zanos, T. P., & The Northwell COVID-19 Research Consortium. 2020. Machine Learning to Assist Clinical Decision-Making During the COVID-19 Pandemic. *Bioelectronic Medicine*. 6(1): 1–8. doi.org/10.11186.s4234-020-00050-8
- Dietvorst, B. J., Simmons, J. P., and Massey, C. 2015. Algorithm Aversion: People Erroneously Avoid Algorithms after Seeing Them Err. *Journal of Experimental Psychology: General* 144 (1): 114–26. doi.org/10.1037/xge0000033
- Gago, P., Santos, M. F., Silva, A., Cortez, P., Neves, J., and Gomes, L. 2005. INTCare: A Knowledge Discovery Based Intelligent Decision Support System for Intensive Care Medicine. *Journal of Decision Systems*. 14 (3): 241–59. doi.org/10.3166/jds.14.241-259
- Ghassemi, M. M., Al-Hanai, T., Raffa, J. D., Mark, R. G., Nemati, S., and Chokshi, F. H. 2018. How is the Doctor Feeling? ICU Provider Sentiment Is Associated with Diagnostic Imaging Utilization. *2018 40th Annual International Conference of the IEEE Engineering In Medicine And Biology Society (EMBC)*, Honolulu. 4058-4064. IEEE. doi.org/10.1109/EMBC.2018.8513325.
- Grote, T., & Berens, P. 2020. On the Ethics of Algorithmic Decision-Making in Healthcare. *Journal of Medical Ethics*. 46(3): 205–211. doi.org/10.1136/medethics-2019-105586
- Lamanna, C., and Byrne, L. (2018). Should Artificial Intelligence Augment Medical Decision Making? The Case for an Autonomy Algorithm. *AMA Journal of Ethics*, 20(9): E902–E910. doi.org/10.1001/amajethics.2018.902
- Larson, E. B., & Yao, X. 2005. Clinical Empathy as Emotional Labor in the Patient-Physician Relationship. *The Journal of the American Medical Association*. 293(9): 1100–1106. doi.org/10.1001/jama.293.9.1100
- Lemieux-Charles, L., Meslin, E. M., Aird, C., Baker, R., and Leatt, P. 1993. Ethical Issues Faced by Clinician/managers in Resource-Allocation Decisions. *Hospital & Health Services Administration* 38(2): 267–85. Doi.org
- Le Reste, J.-Y., Coppens, M., Barais, M., Nabbe, P., Le Floch, B., Chiron, B., Dinant, G. J., Berkhout, C., Stolper, E., & Barraine, P. 2013. The Transculturality of “Gut Feelings”. Results from a French Delphi Consensus Survey. *The European Journal of General Practice*. 19(4): 237–243. doi.org/10.3109/13814788.2013.779662
- Logg, J. M., Minson, J. A., and Moore, D. A. 2019. Algorithm Appreciation: People Prefer Algorithmic to Human Judgment. *Organizational Behavior and Human Decision Processes*. 151: 90–103. doi.org/10.1016/j.obhdp.2018.12.005
- McCarthy, J., & Deady, R. 2008. Moral Distress Reconsidered. *Nursing Ethics*. 15(2): 254–262. doi.org/10.1177/0969733007086023
- Medlej, K. (2018). Calculated Decisions: Sequential Organ Failure Assessment (SOFA) Score. *Emergency Medicine Practice*. 20(10): CD1–CD2.
- Nundy, Shantanu, Tara Montgomery, and Robert M. Wachter. 2019. Promoting Trust Between Patients and Physicians in the Era of Artificial Intelligence. *The Journal of the American Medical Association*. doi.org/10.1001/jama.2018.20563
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. 2019. Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*. 366(6464): 447–453. doi.org/10.1126/science.aax2342
- Rew, L. (2000). Acknowledging Intuition in Clinical Decision Making. *Journal of Holistic Nursing*. 18(2): 94–113. doi.org/10.1177/089801010001800202
- Robert, R., Kentish-Barnes, N., Boyer, A., Laurent, A., Azoulay, E., & Reignier, J. 2020. Ethical Dilemmas Due to the COVID-19 Pandemic. *Annals of Intensive Care*. 10(1): 84. doi.org/10.1186/s13613-020-00702-7
- Röösli, E., & Rice, B. (2020). Bias at Warp Speed: How AI may Contribute to the Disparities Gap in the Time of COVID-19. *Journal of the American Medical Informatics Association*. 1–3. doi.org/10.1093/jamia/ocaa210
- Selph, R. B., Shiang, J., Engelberg, R., Curtis, J. R., & White, D. B. (2008). Empathy and Life Support Decisions in Intensive Care Units. *Journal of General Internal Medicine*, 23(9): 1311–1317. doi.org/10.1007/s11606-008-0643-8

- Shea, G. P., Laudansky, K. K., & Solomon, C. A. 2020, March 27). Triage in a Pandemic: Can AI Help Ration Access to Care? Retrieved from: <https://knowledge.wharton.upenn.edu/article/triage-in-a-pandemic-can-ai-help-ration-access-to-care/>
- Sinclair, M., & Ashkanasy, N. M. (2005). Intuition: Myth or a Decision-Making Tool? *Management Learning*. 36(3): 353–370. doi.org/10.1177/1350507605055351
- Van de Brink, N., Holbrechts, B., Brand, P. L. P. P., Stolper, E. C. F., and Van Royen, P. 2019. Role of Intuitive Knowledge in the Diagnostic Reasoning of Hospital Specialists: A Focus Group Study. *BMJ Open*. 9 (1): e022724. doi.org/10.1136/bmjopen-2018-022724
- Warwick, K. 2015. The Disappearing Human-Machine Divide. In *Beyond Artificial Intelligence*, edited by J. Romport, and E. Zackova, J. Kelemen, 1–10. Switzerland: Springer
- Wright, F., Cohen, S., & Caroselli, C. 1997. Diverse Decisions. How Culture Affects Ethical Decision Making. *Critical Care Nursing Clinics of North America*. (9)1: 63–74. doi.org/10.1016/S0899-5885(18)30292-2
- Xu, W. 2019. Toward Human-Centered AI: A Perspective from Human-Computer Interaction. *interactions*. 26(4): 42–46. doi.org/10.1145/3328485