Wrongful birth: AI-tools for moral decisions in clinical care in the absence of disability ethics

Maya Sabatello

Columbia University

Meier et al. (2022) describe a pilot study that developed METHAD, an AI-based Medical Ethics Advisor tool that draws on the principlism approach and was tested using text-book cases and clinical ethics experts. There are several reasons why such tools should raise concerns, including the distrust they are likely to sow in patient-clinician relationships, and the upshot of delegating *moral* decision-making—presumably what parts humans from other animals—to a machine. But with METHAD already in place, and likelihood that other similar AI-tools will follow, it is critical to ensure that these technologies integrate, and prioritize, equitable implementation. I argue that developing AI-tools for ethical medical decision-making in the absence of disability ethics fails to uphold this fundamental responsibility.

People with disabilities comprise 15% of the population worldwide (pre-pandemic) and commonly experience health disparities and low social determinants of health. Yet, except for an anecdotal use of the word "disabilities", the authors provide no information about how people with disabilities were considered in the algorithm development and what measures were taken to mitigate the disability bias that each step and the ultimate outcome of METHAD are otherwise likely to yield. (Note: although the article considers individuals who lack or have lost decision-making capacity, these circumstances *do not* immediately apply to most people with disabilities). Below I highlight three key areas where the use of METHAD is likely to cause undue burden and harm to patients with disabilities and suggest measures to promote disability/intersectional justice.

Contextual and relational autonomy

In developing METHAD's algorithm, the authors highlight the value of respect for autonomy and its exercise by patients with decisional capacity. They endorsed a capacity scale, starting with the clinical binary construct (capable/incapable) and expanding it to include marginally and moderately capable individuals. They then trained the algorithm to ascribe more weight to patients' preferences if they score higher on the capacity scale. When the user (i.e., a clinician) rates the patient's decisional capacity as impaired, METHAD was trained to follow written sources for establishing patient preferences and shift to surrogate decision-making and best-interest approaches.

These steps are widely accepted in mainstream bioethics and follow the emphasis on individuals as independent decision-makers. However, they do not include parameters to assess how structural and environmental factors affect capabilities and disregard how ableist social power determines the scope and application of autonomy in clinical care.

For example, ineffective communication modalities and inaccessible medical information affect the ability of blind patients to independently make healthcare decision for themselves (Agaronnik et al. 2019). Commonly used psychological measures to assess capacity are infrequently designed or adapted for deaf patients, which result in low evaluation of cognitive abilities (Morere et al. 2019). In addition, research indicates that clinicians express discomfort communicating with patients with disability and have limited knowledge of disability accessibility and accommodation (Agaronnik et al. 2019), which can affect capacity determinations. More generally, people with disabilities are also often presumed to lack capacity—without individualized assessment and regardless of abilities—and to have a guardian who makes medical decisions on their behalf (NCD 2019), although such arrangements (which require a judicial determination) are exceedingly few. Without remedial steps, there is a risk that METHAD's algorithm will consistently produce lower capability scores for patients with disability.

Disability ethics builds on such examples and could inform METHAD (and similar AI-tools). It challenges the prevalent moral construct of autonomy that reflects elitist ideals of independency but overlooks the lived experiences of many people with disabilities for whom accessibility is intertwined with the exercise of autonomy. Instead, disability ethics highlights relational autonomy, which recognizes that individuals' matrix of relationships and responsibilities provide the conditions for self-determination (Scully 2008). Translating these insights into algorithm development would require creating disability-responsive scoring systems for capacity determinations. It should incorporate assessment of clinicians' knowledge about the needs of patients with disabilities alongside actual provision of tailored accommodations to patients; ascribe equal moral value to supported decision-making, which recognizes both abilities based on what one can do alone *and with supports*; and train the algorithms to make relevant distinctions between caregiving, guardianship (and its types) and relational supports.

Disability cultural humility

To operationalize the principle of beneficence and non-maleficence, METHAD was trained to consider the patient's life expectancy (presumed to be a quantifiable value) and quality of life (a more subjective parameter) given the proposed intervention. It calculates potential and likely improvement or harm and allows for a continuous slide switch to incorporate patient preferences in instances of competing aims. Per the authors, this "formula" was developed in an effort to address common discrepancies between subjective and objective perceptions of quality of life "in the context of disabilities" (without a subject!) and a third-person clinical evaluator. Yet, there is no indication that the algorithm was trained to identify and rectify disability bias, which is likely to affect each of these assessments.

The article's discussion of examples used for METHAD's evaluation does not elaborate on the most illustrative case category for patients with disabilities: clinicians' decisions to withhold or withdraw treatment that they deem nonbeneficial. But as some triage policies during the COVID-19 pandemic have shown (Sabatello et al. 2020), the threshold for futile care decisions is often far lower for people with disabilities. And, in practice, *both* prognosis for survival *and* quality of life include subjective components that are prone to

disability bias. Disability is often conflated with a disease and adverse health outcomes, regardless of medical justifications. For example, many transplant programs in the US still use group-level categorization of patients with neurodevelopmental genetic conditions and intellectual disability as contraindication to transplant listing, despite evidence for medical benefit and improved quality of life (Wall et al. 2020). Another study found that 48.5% of death certificates of adults with development disability in the US had a developmental disability coded as the underlying cause of death rather than the direct cause of death such as a respiratory disease or pneumonia (Landes, Stevens, and Turk 2019). Research also found that many clinicians perceive worse quality of life for people with disabilities—contrary to the views of most people with disabilities themselves (NCD 2019)—and report low confidence in their ability to provide the same quality of care to patients with disabilities (Iezzoni et al. 2021), which may subconsciously impact prognosis decisions.

Bioethicists have long highlighted the importance of cultural humility in clinical care. Cultural humility enables clinicians and patients to better understand each another, build trust, and improve health outcomes; it is a process for clinicians to self-reflect on their biases and develop empathetic openness to others' views and preferences. *Disability* cultural humility has received little attention but is critical for developing a moral algorithm in clinical decision-making. Initial steps can include engaging with people with disabilities *as experts*, requiring disability competency training in relevant fields (e.g., medicine, data scientists, bioethics), and diversifying the respective workforces that are markedly underrepresented by people with disabilities.

Disability/intersectional (in)justice

METHAD's incorporation of the principle of justice (defined as a fair distribution of healthcare resources) is particularly limited and only passes judgement at the individual level. Yet, by neglecting to operationalize this principle, it is all but guaranteed that existing *injustices* will remain. From a disability ethics lens, two issues must be highlighted.

The first is representational harms. For AI-tools to assure fairness in resource allocation, algorithms must be trained on datasets that are bias-free and representative of diverse populations. For METHAD, text-book cases and medical ethics committees' decisions were selected as most promising for enabling the algorithm to replicate typical recommendations. Yet, this dataset is likely to compromise disability fairness in AI. Although ethics committees are instrumental in mitigating difficult medical cases, their decisions are bound by the demographically non-diverse populations that comprise them (both clinicians and bioethicists) and policies that reflect institutional interests and an ableist value-system. People with disabilities are rarely included in medical ethics committees or consulted in the development of such policies (NCD 2019). A disability-responsive training dataset could be developed by revisiting text-book cases through a dialogue with disability ethicists; developing a clear, transparent, and unbiased process for medical ethics committees' decisions; and assembling interdisciplinary and demographically diverse medical ethics committees that include clinicians with disabilities and disability experts (Sabatello et al. 2020).

The second issue centers on disability and intersectional justice. Although disability is a universal phenomenon, its prevalence is disproportionately high among marginalized racial, ethnic, gender, and low-income communities. The extensive experiences of disability prejudice and system-level discrimination that preclude social inclusion, equal opportunity, and self-determination—i.e., structural ableism—are intertwined with and compounded by structural racism, sexism, and classicism, all of which are well-established and ingrained in our healthcare system (Chin 2021). A disability ethics approach thus highlights the need for moral AI-tools to expand on the traditional focus on distributive (including procedural) justice to promote social justice. Doing so requires sociopolitical will to reflect on and redress existing social injustices and algorithms that are designed to actively implement social equity parameters in ethical medical decision-making.

AI-tools are powerful and can be harnessed to promote the common good. Training machines to do so in ethical medical decision-making must be preceded by hard conversations about structural ableism and assurance that disability/intersectional justice are upheld.

Acknowledgement

References

Agaronnik N, Cambell EG, Ressalam L, et al. (2019). Communicating with patients with disability: perspectives of practicing physicians. Journal of General Internal Medicine 34: 1139–1145. [PubMed: 30887435]

Chin NM 2021. Centering disability justice. Syracuse Law Review 71: 684–749.

Iezzoni LL, Rao SR, Ressalam J, et al. 2021. Physicians' perceptions of people with disability and their health care. Health Affairs 40(2): 297–306. [PubMed: 33523739]

Landes SD, Stevens JD, & Turk MA (2019). Obscuring effect of coding developmental disability as the underlying cause of death on mortality trends for adults with developmental disability: A cross-sectional study using us mortality data from 2012 to 2016. BMJ Open, 9:e026614. doi:10.1136/bmjopen-2018-026614

Morere DA, Dean PM, & Mompremoer L (2019). Mental Health Assessment of Deaf Clients: Issues with Interpreters use and Assessment of Person with Diminished Capacity and Psychiatric Populations. JADARA, 42(4): 241–258.Retrieved from https://repository.wcsu.edu/jadara/vol42/iss4/9

National Council on Disability (NAD). Medical futility and disability bias: Part of the bioethics and disability series (Washington DC; 2019).

Sabatello M, Blankmayer-Burke T, McDonald K, Appelbaum PS (2020). Disability, Ethics and Health Care in the Covid-19 Pandemic. Am J Public Health, 110: 1523–527. [PubMed: 32816541]

Scully JL Disability Bioethics: Moral Bodies, Moral Difference (New York: Rowman & Littlefield, 2008).

Wall A, Lee GH, Maldonado J, et al. 2020. Genetic diseases and intellectual disability as contraindications to transplant listing in the United States: A survey of heart, kidney, liver, and lung transplant programs. Pediatric Transplantation 24(7):e13837 [PubMed: 32997378]