



What Counts as “Clinical Data” in Machine Learning Healthcare Applications?

Joshua August Skorburg

To cite this article: Joshua August Skorburg (2020) What Counts as “Clinical Data” in Machine Learning Healthcare Applications?, The American Journal of Bioethics, 20:11, 27-30, DOI: [10.1080/15265161.2020.1820107](https://doi.org/10.1080/15265161.2020.1820107)

To link to this article: <https://doi.org/10.1080/15265161.2020.1820107>



Published online: 26 Oct 2020.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)

unreasonable for practitioners to feel bewildered. In the next decade, with the insights developed from additional experimentation, integration, and implementation of deep learning models within health care delivery, editorials will be published on how wrong most of the current assumptions and assessments of artificial intelligence were. That is the nature of the evolution of knowledge. What we ought not to be wrong about are the societal values that are driving these discussions. In all our deliberations, we must aspire to working out solutions that promote patient benefit and minimize harm, while facilitating an environment where technology can improve patient care. This will be vital—technological innovations are ephemeral, human values are enduring.

DISCLOSURE STATEMENT

Dr. Junaid Nabi is a consultant for early-stage technical groups involved in developing machine learning based digital applications.

REFERENCES

- Agniel, D., I. S. Kohane, and G. M. Weber. 2018. Biases in electronic health record data due to processes within the healthcare system: Retrospective observational study. *British Medical Journal* 361:k1479.
- Brummel, E. S. 2017. Confronting natural conflicts of interest and artificial intelligence. *Journal of Law and the Biosciences* 4 (2):435–444.
- Char, D. S., M. D. Abràmoff, and C. Feudtner. 2020. Identifying ethical considerations for machine learning healthcare applications. *The American Journal of Bioethics* 20 (11):7–17. doi: [10.1080/15265161.2020.1819469](https://doi.org/10.1080/15265161.2020.1819469).
- COVID-19 Screening Tool. 2020. Apple Inc. <https://covid19.apple.com/screening>.
- Ghahramani, Z. 2015. Probabilistic machine learning and artificial intelligence. *Nature* 521 (7553):452–459.
- Hinton, G. 2018. Deep learning—A technology with the potential to transform health care. *Journal of American Medical Association* 320 (11):1101–1102.
- LeCun, Y., Y. Bengio, and G. Hinton. 2015. Deep learning. *Nature* 521 (7553):436–444.
- Maxmen, J. S. 1976. *The post-physician era: Medicine in the twenty-first century*. New York, NY: Wiley and Sons.
- McCarthy, A. M., M. Bristol, S. M. Domchek, et al. 2016. Health care segregation, physician recommendation, and racial disparities in BRCA1/2 testing among women with breast cancer. *Journal of Clinical Oncology* 34 (22):2610–2618.
- Nabi, J. 2018. How bioethics can shape artificial intelligence and machine learning. *The Hastings Center Report* 48 (5):10–13.
- Obermeyer, Z., B. Powers, C. Vogeli, and S. Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366 (6464):447–453.
- Parikh, R. B., S. Teeple, and A. S. Navathe. 2019. Addressing bias in artificial intelligence in health care. *Journal of American Medical Association* 322 (24):2377–2378.
- Wittbold, K. A., C. Carroll, M. Iansiti, H. M. Zhang, and A. B. Landman. 2020. How hospitals are using AI to battle Covid-19. *Harvard Business Review*, April 3, 2020. <https://hbr.org/2020/04/how-hospitals-are-using-ai-to-battle-covid-19>.

THE AMERICAN JOURNAL OF BIOETHICS
2020, VOL. 20, NO. 11, 27–30
<https://doi.org/10.1080/15265161.2020.1820107>



OPEN PEER COMMENTARIES



What Counts as “Clinical Data” in Machine Learning Healthcare Applications?

Joshua August Skorburg^{a,b}

^aUniversity of Guelph; ^bCentre for Advancing Responsible and Ethical Artificial Intelligence (CARE-AI)

In developing a systematic framework for identifying ethical issues in Machine Learning Health Care Applications (ML-HCAs), Char et al. (2020) have helpfully synthesized heretofore distinct stands of ethical inquiry. By linking the ML developmental pipeline—from conception to implementation—with a parallel

ethical pipeline of evaluation and oversight, the authors also bring to the fore a crucial insight that ethical issues are thoroughly embedded within the developmental process, and not merely afterthoughts or add-ons.

Throughout, Char et al. (2020) push back against the increasingly common trope of artificial intelligence

CONTACT Joshua August Skorburg ✉ g.skorburg@gmail.com 📍 University of Guelph, Guelph, Ontario N1G 2W1, Canada; Centre for Advancing Responsible and Ethical Artificial Intelligence (CARE-AI), Guelph, Ontario, Canada.

© 2020 Taylor & Francis Group, LLC

exceptionalism. For example: “No sound reason as yet exists to believe that the health applications powered by ML are, in and of themselves, exceptional,” and “the technology itself is also built from essentially standard clinical information,” such that “standard ethical considerations about these data also likely apply to ML-HCAs.” As a result, we “do not need to be focused on exceptions, even as it should leave space for exceptional considerations to be identified” (9).

To be sure, there are many good reasons to avoid ahistorical exceptionalism about emerging technologies. But in this commentary, I will raise a few questions and examine a few cases which reveal potential shortcomings and oversights in the proposed framework.

The first question is whether standard bioethical frameworks are indeed able to capture the variety of relevant ethical considerations in ML-HCAs as the authors suggest. Metcalf and Crawford (2016), in the context of Big Data approaches more generally, have convincingly argued that they are not. For example, traditional biomedical research ethics frameworks have tended to assume that there is a relevant distinction between public and private data, such that the former is often exempt or excluded from IRB review. But a litany of cases has shown how the combination of publicly available datasets with machine learning methods can reveal sensitive information and cause various forms of physical, emotional, and economic harms.

To take just one example, using a publicly available dataset of New York City taxi rides, researchers were able to infer taxi drivers’ incomes, addresses, and religiosity. And when combined with other publicly available information from blogs, they were able to track celebrities (Metcalf and Crawford 2016, 9). It is not difficult to imagine how sensitive health information, such as addictions, mental health visits, or pregnancy status could be revealed from combinations of seemingly disparate public data. Indeed, the most touted feature of ML-HCAs is precisely this ability to uncover such hidden and unintuitive relationships.

Similar arguments have been developed for the inadequacy of other traditional bioethical concepts such as interventions, research, consent, harm, and indeed, the very idea of human subjects, in light of recent advances in ML/AI (Metcalf and Crawford 2016). Additionally, given the many unique ethical considerations for AI healthcare applications raised by Morley et al. (2020, see esp. Table 2), there are sound reasons to believe that many healthcare applications

powered by ML are exceptional, such that standard ethical considerations are inadequate. This is not to say that the traditional bioethical frameworks are hopelessly outdated, but rather, that these unique issues raised by ML-HCAs might be closer to rules than exceptions.

A second set of considerations in this vein concerns the distinction between “standard” (roughly, data generated and collected through interactions with healthcare professionals) and “non-standard” (roughly, data generated and collected outside of interactions with healthcare professionals) clinical data. While there are many ML-HCAs that employ “essentially standard clinical information” as Char et al. (2020, 9) suggest, there are a rapidly increasing number of ML-HCAs that do not fit this mold. Many scholars, the author included (Skorburg and Friesen *forthcoming*), have highlighted how, in relying on nonstandard clinical data, many digital health applications powered by AI/ML raise novel and pressing ethical issues. Casting these as mere exceptions risks overlooking an emerging field which enjoyed an unprecedented \$5.4B of venture funding in the first six months of this year (RockHealth 2020).

Reviewing the literature on digital phenotyping approaches to mental health, for example, reveals hundreds, if not thousands of published papers envisioning a health ecosystem where social media posts are analyzed in real-time to detect changes in mood; photoplethysmograms from smartwatches reveal tiny fluctuations in heart rate variability associated with stress; tapping and scrolling patterns from smartphone screens indicate anxiety symptoms; fitness trackers identify movement patterns associated with depression; automated speech analysis identifies symptoms of cognitive impairment; chatbots powered by artificial intelligence provide emotional support for hard-to-reach populations; smart toilets automatically analyze stool to detect the presence of gut microbes associated with autism; smart mattresses assess sleep quality for patients with borderline personality disorder; smart pill boxes monitor adherence with anti-psychotic medications.

Combined with machine learning techniques, these data promise to “transform the diagnosis and treatment of mental illness globally by enabling passive, continuous, quantitative, and ecological measurement-based care” (Martinez-Martin et al. 2018, 4).

On the one hand, most of these data are “clinical data” in the sense that they provide clinically relevant information. But on the other hand, most of these data are not obtained through standard interactions

with healthcare professionals. And while this may be (slowly) changing, many of the data sources described above are not subject to the same regulations and protections as more standard forms of clinical information such as patient demographics, laboratory values, or diagnostic images (Char et al. 2020, 9).

Why is this a problem? After all, many of these nonstandard forms of clinical data have the potential to improve both the personalization of medical interventions and also the scale at which such interventions can be delivered. Take the example of social media posts. In a recent review, Chancellor and De Choudhury (2020) describe 75 studies which use various kinds of ML-HCAs to detect subtle mental health-relevant signals in both the content of posts and associated meta-data to predict depression, anxiety, stress, suicidality, eating disorders, PTSD, and many other conditions.

However, the ethical risks with collecting and analyzing mental health data in this way were made clear when reporting from Australia revealed that advertisers on social media can determine, in real time, when teenagers feel “insecure,” “worthless” and “need a confidence boost,” “stressed,” “defeated,” “overwhelmed,” “anxious,” “nervous,” “stupid,” “silly,” “useless” or a “failure,” and target their advertisements accordingly (Levin 2017). Many scholars have rightly worried that the usual way of thinking about informed consent, the research/practice distinction, expectations of privacy, etc. are inadequate in these contexts, given the scale and sophistication of emerging digital mental health tools powered by AI/ML.

Are these examples just cherry-picked exceptions? A recent meta-review by LeComte et al. (2020) suggests not. They found 24 meta-analyses comprising thousands of studies related to the use of digital mental health applications. There are hundreds of thousands of health and wellness apps available to consumers, and their use has been skyrocketing since the beginning of the COVID-19 pandemic. The flood of data generated in this space (along with the new ML/AI techniques developed in response) make it increasingly difficult to distinguish between clinical and non-clinical data. Indeed, a recent qualitative study with digital health experts “uniformly indicated that all data can be health data, particularly when aggregated across sources and time” (Grande et al. 2020).

In the end, the framework proposed by Char et al. (2020) is well-suited to identify the kinds of ethical issues that arise in the development and implementation of AI/ML applications involving training data

generated in traditional healthcare contexts (e.g., diagnostic imaging, electronic health records). The examples here are meant to raise the question of how well this framework will fare in nontraditional contexts involving data generated from wearables, social media, and the Internet of Things. While there is certainly some acknowledgement of these contexts (Char et al. 2020, 12), I have suggested that they are likely to become more central than peripheral.

Presumably, the issues I have been discussing would fall under the “Conception” and “Development” cells in Figure 1 (Char et al. 2020, 9). My primary concern is that the subsequent focus on overly broad questions like “what is the overall goal?” or overly narrow questions like “could the training datasets perpetuate bias?” are not adequate to capture the diversity of ethical issues likely to arise from the pervasive use of nonstandard forms of clinical data in ML-HCAs described above.

The increasing use of nonstandard forms of clinical data in ML-HCAs has been a concern in the literature on the ethics of mHealth and digital phenotyping for many years. The COVID-19 pandemic has made it overwhelmingly likely that these considerations will apply across healthcare systems more broadly. As more medicine becomes tele-medicine and as more AI/ML companies penetrate the healthcare sector, more health-relevant data will be generated, analyzed, and monetized. We need to make sure that a systematic framework for identifying ethical issues in ML-HCAs is ready for this.

REFERENCES

- Chancellor, S., and M. De Choudhury. 2020. Methods in predictive techniques for mental health status on social media: A critical review. *NPJ Digital Medicine* 3 (1):1–11. doi: 10.1038/s41746-020-0233-7.
- Char, D. S., M. D. Abràmoff, and C. Feudtner. 2020. Identifying ethical considerations for machine learning healthcare applications. *The American Journal of Bioethics* 20 (11):7–17. doi: 10.1080/15265161.2020.1819469.
- Grande, D., X. L. Marti, R. Feuerstein-Simon, et al. 2020. Health policy and privacy challenges associated with digital technology. *JAMA Network Open* 3 (7):e208285. doi: 10.1001/jamanetworkopen.2020.8285.
- Lecomte, T., S. Potvin, M. Corbière, et al. 2020. Mobile apps for mental health issues: Meta-review of meta-analyses. *JMIR mHealth and uHealth* 8 (5):e17458. doi: 10.2196/17458.
- Levin, S. 2017. Facebook told advertisers it can identify teens feeling ‘insecure’ and ‘worthless’. *The Guardian*. <https://www.theguardian.com/technology/2017/may/01/facebook-advertising-data-insecure-teens>
- Martinez-Martin, N., T. R. Insel, P. Dagum, H. T. Greely, and M. K. Cho. 2018. Data mining for health:

Staking out the ethical territory of digital phenotyping. *NPJ Digital Medicine* 1 (1):1–5. doi: [10.1038/s41746-018-0075-8](https://doi.org/10.1038/s41746-018-0075-8).

Metcalfe, J., and K. Crawford. 2016. Where are human subjects in big data research? The emerging ethics divide. *Big Data & Society* 3 (1):2053951716650211. doi: [10.1177/2053951716650211](https://doi.org/10.1177/2053951716650211).

Morley, J., C. C. Machado, C. Burr, et al. 2020. The ethics of AI in health care: A mapping review. *Social Science & Medicine* 260:113172. doi: [10.1016/j.socscimed.2020.113172](https://doi.org/10.1016/j.socscimed.2020.113172).

RockHealth. 2020. Midyear digital health market update: Unprecedented funding in an unprecedented time. RockHealth Reports. <https://rockhealth.com/reports/2020-midyear-digital-health-market-update-unprecedented-funding-in-an-unprecedented-time/>

Skorburg, J. A., and P. Friesen. forthcoming. Ethical issues in text mining for mental health. In *The atlas of language analysis in psychology*, eds. M. Deghani and R. Boyd. New York: Guilford Press. Preprint available at: <https://philarchive.org/archive/AUGEII>.

THE AMERICAN JOURNAL OF BIOETHICS
2020, VOL. 20, NO. 11, 30–32
<https://doi.org/10.1080/15265161.2020.1820109>



Taylor & Francis
Taylor & Francis Group

OPEN PEER COMMENTARIES



An Ethical Framework to Nowhere



Eric S. Swirsky^a , Carol Gu^a, and Andrew D. Boyd^a 

University of Illinois at Chicago

In their article, Char et al. (2020) have created a model intended to tidy up the messy landscape of ethical concerns arising from machine-learning health care applications (ML-HCAs). The novel conceptual framework depicts the pipeline through which these applications are developed, implemented, and evaluated as part of a larger system, and it is offered along with the claim that the domain is “bereft of any conceptual map” that charts the entirety of the process (Char et al. 2020, 7). We challenge this claim. There are frameworks already used to synthesize ethical issues related to machine learning and the use of autonomous agents, and they are associated with the sociotechnical perspective of technological development. These frameworks consider aspects of design, development, adaptation, governance, and incorporate a diverse array of philosophical and theoretical dimensions of overlapping disciplines in framing ethical approaches; these include, bioethics, economics ethics, business and organizational ethics, social ethics, environmental ethics, computer and technology ethics, and professional ethics for the varied disciplines involved in the evolution of the technology (Fiore 2020). Value sensitive design methodologies, for example, have been in development for over 20 years and provide for an integrative, theoretically grounded approach that

considers human values in technology design and development (Friedman et al. 2002). These approaches do not merely look backwards at what has previously been reported, but they also consider the unanticipated consequences that arise in sociotechnical systems where humans and computers interact. As such, value sensitive design models are recognized by many as the most comprehensive in accounting for human values throughout the process of design (Winkler and Spiekermann 2018).

As for the framework promoted in the target article, we question whether it is adequate for its intended purpose. Char et al.’s pipeline framework for ML-HCAs has tied ethical concerns to various steps of the development-implementation stage in a linear fashion that will most likely be outpaced by technological innovation. We recommend the authors adopt a more agile approach to ethics that incorporates greater adaptability and implements rapid feedback loops. Agile ethics iteratively addresses concerns in ethical, legal, and social issues (ELSI) with a flexible approach to adapt to the needs of stakeholders and end users as the technology develops. Crisis management information technologies have been confronted with similar concerns to ML-HCAs in that rapidly designed and deployed systems do not give enough

CONTACT Eric S. Swirsky  eswir@uic.edu  Department of Biomedical and Health Information Sciences, College of Applied Health Sciences, University of Illinois at Chicago, 1919 W. Taylor Street, Mail Code 530, Chicago, 60612, USA.

© 2020 Taylor & Francis Group, LLC