

Healthsheet: Development of a Transparency Artifact for Health Datasets

Negar Rostamzadeh
Google Research
Montreal, Canada
nrostamzadeh@google.com

Diana Mincu
Google Research
London, UK

Subhrajit Roy
Google Research
London, UK

Andrew Smart
Google Research
San Francisco, US

Lauren Wilcox
Google Research
Mountain View, US

Mahima Pushkarna
Google Research
Toronto, Canada

Jessica Schrouff
Google Research
London, UK

Razvan Amironesei
Google Research
San Francisco, US

Nyalleng Moorosi
Google Research
Ghana

Katherine Heller
Google Research
San Francisco, US

ABSTRACT

Machine learning (ML) approaches have demonstrated promising results in a wide range of healthcare applications. Data plays a crucial role in developing ML-based healthcare systems that directly affect people's lives. Many of the ethical issues surrounding the use of ML in healthcare stem from structural inequalities underlying the way we collect, use, and handle data. Developing guidelines to improve documentation practices regarding the creation, use, and maintenance of ML healthcare datasets is therefore of critical importance. In this work, we introduce *Healthsheet*, a contextualized adaptation of the original datasheet questionnaire [22] for health-specific applications. Through a series of semi-structured interviews, we adapt the datasheets for healthcare data documentation. As part of the Healthsheet development process and to understand the obstacles researchers face in creating datasheets, we worked with three publicly-available healthcare datasets as our case studies, each with different types of structured data: Electronic health Records (EHR), clinical trial study data, and smartphone-based performance outcome measures. Our findings from the interview study and case studies show 1) that datasheets should be contextualized for healthcare, 2) that despite incentives to adopt accountability practices such as datasheets, there is a lack of consistency in the

broader use of these practices 3) how the ML for health community views datasheets and particularly *Healthsheets* as diagnostic tool to surface the limitations and strength of datasets and 4) the relative importance of different fields in the datasheet to healthcare concerns.

CCS CONCEPTS

• **Human-centered computing** → *Heuristic evaluations*.

ACM Reference Format:

Negar Rostamzadeh, Diana Mincu, Subhrajit Roy, Andrew Smart, Lauren Wilcox, Mahima Pushkarna, Jessica Schrouff, Razvan Amironesei, Nyalleng Moorosi, and Katherine Heller. 2022. Healthsheet: Development of a Transparency Artifact for Health Datasets. *J. ACM* 37, 4, Article 111 (August 2022), 25 pages. <https://doi.org/10.1145/3531146.3533239>

1 INTRODUCTION

The use of machine learning (ML) is rapidly expanding in healthcare as the amount of new data generated improves our capacity to effectively manage complex clinical and diagnostic information [26, 40, 50, 64, 70]. During the last decade, ML has played a central role in high-stakes healthcare problems such as precision medicine [51, 65], survival analysis [36, 43], disease diagnosis [20], and continuous innovations in treatment plans. While ML approaches present opportunities to assist healthcare professionals, streamline the healthcare system, and potentially improve patient outcomes, they bring with them ethical concerns [8, 24, 41] that range from racial and gender disparities, accessibility of clinical studies, to subjectivity in healthcare practices and biases [14].

Many ethical concerns in ML applications can be traced back to the development processes of the underlying datasets [53, 60]. As Parasidi et al [47] argue, societal fairness dictates that health data be used to advance the public good in ways that can avoid causing or exacerbating inequities. These societal goals are reflected in the many declarations of ethical principles from AI firms and research institutions [42]. Limited ethical and social structural deliberation

Authors' addresses: Negar Rostamzadeh, Google Research, Montreal, Canada, nrostamzadeh@google.com; Diana Mincu, Google Research, London, UK; Subhrajit Roy, Google Research, London, UK; Andrew Smart, Google Research, San Francisco, US; Lauren Wilcox, Google Research, Mountain View, US; Mahima Pushkarna, Google Research, Toronto, Canada; Jessica Schrouff, Google Research, London, UK; Razvan Amironesei, Google Research, San Francisco, US; Nyalleng Moorosi, Google Research, Ghana; Katherine Heller, Google Research, San Francisco, US.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

0004-5411/2022/8-ART111

<https://doi.org/10.1145/3531146.3533239>

during data collection can also exacerbate unfair racial biases in healthcare predictions [45]. Therefore, caution and improved accountability in data collection, use and maintenance are critical in the adoption of ML in the health domain.

As a step towards more accountable ML practices[28], several guidelines and frameworks have been proposed around data transparency, collection and use in ML [9, 10, 22, 23, 49, 56]. For example, “Datasheets for datasets” [22] provided a critical new approach for the rigorous and reflective documentation of ML datasets. The fundamental concepts that underlie datasheets have been adopted in various forms across industry and academia, such as Artsheet [62], Data Cards [54], and FactSheets [3]. However, the value of dataset documentation extends beyond the documentation artifact to the process of meeting documentation requirements; thereby increasing opportunities for greater integrity and accountability in data collection practices.

To begin addressing gaps in current practices, frameworks, and standards for the ethical collection of health data for ML, we introduce a contextualized type of datasheet that attends to the needs of healthcare data: *Healthsheet*. The purpose of *Healthsheet* is to contribute to the meaningful ethical review of healthcare data, in addition to existing data governance practices or legal requirements of healthcare data. Further, it aligns with recent initiatives in clinical trials data collection [38, 57], and data-driven digital health technologies [29].

According to [47], existing regulatory frameworks that could cover health data are limited in their applicability to the general use of health datasets for ML. For example, The Health Insurance Portability and Accountability Act (HIPAA) does not mandate ethics review for data collection and downstream use. HIPAA also places no limits on the use of de-identified data, regardless of who controls the information [47]. As [47] argue, the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act focus on notification, consent, and deletion rights, but this doesn’t necessarily address issues about the ethical collection, documentation, and use of data [47].

Besides these limitations in regulatory frameworks, we pose three main questions: (1) *could the current state of health data documentation practices introduce challenges and blockers in equitable health research and practice?*, (2) *could dataset consumers use the Healthsheet to assess the quality of datasets?* and finally (3) *what would be the incentives for health dataset curators to create the Healthsheet documentation?*

To answer these questions and contextualize the datasheet, we first conducted a participatory study by interviewing 21 experts (dataset creators and consumers), with a wide range of expertise and diverse backgrounds, in relation to healthcare data. We then tested our new proposed framework on 3 publicly available datasets, each composed of a different data modality to examine the practicality and blockers in creation of *Healthsheets*. In section 2, we describe the *Healthsheet* questionnaire development processes. In section 3, we define transparency as a step towards accountability. Section 4 details the original datasheet questionnaire [22]. Section 5 discusses our expert interview methodology and findings. In section 6, we discuss the necessity and limitations of the datasheet when employed in healthcare applications of ML. We then present *Healthsheet* case studies in 8 and corresponding findings. Finally, we

conclude by discussing our study takeaways, the broader impacts and limitations of this work.

2 HEALTHSHEET DEVELOPMENT METHODOLOGY

Integrating feedback from a variety of stakeholders meets the needs of more communities [39], particularly in high-stake scenarios such as healthcare [61]. For this work, we created a team with multi-disciplinary expertise ranging from Healthcare, ML for Health, ML Fairness, Applied Ethics and Human-Centered Design.

This study is framed as follows:

- (1) We ground our research by co-defining transparency artifacts, and discuss how to use them as a step towards a more accountable dataset development process.
- (2) Following a systemic review of ML for health literature, we created an initial adaptation of the datasheet [22], the *Primary Healthsheet*, which was introduced as stimuli to expert participants.
- (3) Through a set of expert interviews spanning backgrounds related to ML, health and socio-technical research, we (i) identified documentation shortcomings in health datasets and the impact of these shortcomings on research advancement and equitable healthcare practices, (ii) improved upon and used the *Healthsheet* as a diagnostic tool for health datasets assessment, (iii) discussed the necessity and limitations of the *Primary Healthsheet* questionnaire, and (iv) discussed the potential incentives needed by dataset curators and extractors to create a *Healthsheet* for their datasets.
- (4) We then validated our proposed *Healthsheet* using 3 publicly available datasets, MIMIC-III [30], MSOAC [35] and Floodlight [6], taking note of the process, limitations, challenges, as well as gaps in existing documentation.
- (5) Finally, We made sure the corpus represents our interdisciplinary perspective through deep qualitative engagement across many iterations.

3 CO-DEFINING TRANSPARENCY IN HEALTHCARE DATASET CREATION

Algorithmic decision-making and the role of data has been central to the discourse on algorithmic transparency. There are many definitions and characteristics of transparency that seek to serve the common goals of accountability and trust. [15] define transparency as explaining decisions made by algorithmic systems, which enables the identification of harms introduced by algorithmic decision making, holds the entities in the decision making chain accountable for such practices, and detects errors in input data that cause adverse decisions. In providing explanations for an adverse decision, specific guidance and factors can be provided to reverse such a decision. Annany and Crawford suggest using the limitations of transparency as conceptual tools that pave the way towards accountability [2]. Kizilcec’s work observed a bell-shaped curve between transparency and trust in an experiment in which transparency through explanations was measured as a factor of perceptions of understanding, fairness, accuracy of process, and trust in actors [31].

For the purposes of this work, we describe five salient characteristics of transparency that the authors arrived at through generative participatory methods [55] *prior* to engaging with study participants: (1) **Culture and Aspiration**: Transparency is a property of the systemic and organizational ecosystem around datasets. (2) **Knowledge Sharing and Management**: Transparency enables the fluid movement and absorption of knowledge across upstream, downstream, past, present, and future stakeholders and across projects themselves. (3) **Access Management**: Transparency cannot be conflated with access. For transparency, aspects of access include disclosure about rationales and processes, and ensuring some degree of control to prevent misuse or uncalibrated use. (4) **Beyond Data Creation**: Transparency is not a tail-end consideration occurring towards the end of the *project's* life cycle, but rather across the *dataset's* life cycle, from the stage where requirements are gathered to when it is actively used in experimentation, research, and production. (5) **Motivating Factors**: Transparency should capture and foster respect for both the provenance of data *and* the motivations for collecting and applying the data to solve the problems it was intended to address. Our proposed Healthsheet template seeks to operationalize these characteristics of transparency.

4 HEALTHSHEETS - DATASHEETS FOR HEALTHCARE

Datasheets for datasets [22], was proposed as a means of standard communication between dataset creators and dataset consumers, specifically in the context of ML practitioners. Although generalization is the primary goal of ML, we depart from general-purpose datasheets to address the unique context and requirements of healthcare. For example, in data from clinical studies, inclusion criteria can impact on 'who is represented in datasets', subsequently affecting the model being trained and tested on. *Healthsheets* is meant to extend many of the benefits of datasheets, specifically to the health setting: dataset consumers can make more informed decisions about using the dataset for a specific task, while dataset creators also have a tool for self-reflection on the dataset they create in order to establish the intended use of their dataset, discover hidden characteristics of their data that can impact the outcome, and reduce the potential harms that use of the dataset can create. All of these benefits can eventually lead to the creation and promotion of safer, more reliable and equitable ML datasets in healthcare. We should first investigate the unique challenges that dataset consumers have for choosing and working with Healthcare datasets. How Healthsheets could be used as a diagnostic and auditing tool for a better understanding of datasets, their strength and limitations. Finally, while the datasheets for datasets paper [22] has been widely cited,¹ all but a few actually create a datasheet, instead acknowledging the potential usefulness of datasheets. This shows that although the benefits of using datasheets are evident to the community, there are insufficient incentives for many dataset curators or creators to provide datasheets for their datasets.

In cases where datasheets have been produced, creators have found the results to be extremely beneficial and analyses reported even more surprising. For example, Bandy et al [7] found several

thousand duplicated books and a significant skew in genre representation when they created a datasheet for the BookCorpus dataset, which has been used to train large language models in industry. Several findings like this indicate that there is yet an unmet need to document and examine the myriad datasets which form the basis of ML, which are particularly critical in healthcare contexts. For example, a dataset with thousands of duplicated records could lead to inaccurate results that might impact patient safety.

We find that current research practice in ML for healthcare often lead with model development, following which researchers search for datasets and use cases: "*If I have an ML claim that this method is good for something. I may use an off-the-shelf dataset... such that actually making progress on whatever classes are associated with these datasets, really translates to progress on some real-world problem.*"(P3). Although this might be commonplace in current ML model testing workflows, it is important that nuances of datasets be taken into consideration during model development for positive real-world impact in healthcare.

5 INTERVIEW PROCEDURE AND FINDINGS

We conducted 21 semi-structured interviews with experts with a wide range of applied, industrial and academic expertise. Participants were recruited using snowball sampling [25], in which an initial round of interview participants were recruited through emails targeting experts from clinical, legal, policy, privacy, bioethics, healthcare-related ML research, healthcare-related sociotechnical research, and applied healthcare ML engineering (see table 1 in Appendix B for detailed information on the expertise). Clinical experience included specializations in ophthalmology, medical imaging, nephrology, surgery and cancer; Engineering experience spanned research, production, program management, and product management in healthcare contexts. Additional experts were recruited through referrals from previous interviewees. Consent was obtained and stimuli material was provided to experts at least three days prior to interviews. More detailed discussions on the interviewee is provided in Appendix B.

5.1 Challenges faced by dataset consumers

In general, we found that the lack of centralized and comprehensive documentation exacerbates the challenges that experts face when selecting and using datasets. Of many reported challenges, a majority of participants considered the following to be the most concerning when it comes to impact on outcomes: **availability and clarity of meta-data, labeling and subjectivity** in labeling, and **inclusion/ exclusion criteria**.

Meta-data clarity and availability: The most common issues cited during the interviews was limited or lacking centralized meta-data information. This could lead to additional difficulties and challenges in using datasets:

"How much missingness in data exists... if datasets are already made, ... clarity around whether it's on GitHub, whether you have to contact this particular board, just understanding that saves a lot of time and a lot of research resource... There's a mismatch between a clinical academic understanding in what and how clinical question can be operationalized into processes in terms of where's the data from devices captured, labeling subjectivity,..."- P19, *"it's difficult to understand*

¹as of the day of writing, Datasheets for datasets has 574 citations

what can be done and what cannot be done with the data ... that is under a lot of variables that we need to take into account to answer that question ... and sometimes what we have found is that the knowledge is not centralized in one place"-P5. Beyond identifying the appropriate use of datasets, interview participants raised questions about the inherent credibility and utility: *"A challenge we start with is, which entities should we get in contact with, what kind of content was acquired ... it just requires a lot of back and forth with the data provider to understand this. When they say sex do they mean sex at birth or self-identified sex like gender, what is it?..."* - P4. *"A good chunk of our datasets is more (of a) by-product ... but the metadata was never applied with the intent to research on, but it was more applied with the intent to have a proper understanding of the patient, and the usual ICD codes"* - P8.

Subjectivity and labeling was stated as one of the most critical concerns that could contribute to disparity of outcomes. In many cases, there are disagreements even on definitions of gold standards: *"understanding how the labels are generated is often the most difficult aspect ... It's a big issue with chest x-rays actually"* - P3. *"So, for example, how is a heart attack defined? Is it by blood test? Is it by symptoms? Because those huge amounts of subjectivity and inclusion criteria in that, excludes women and ethnic minorities."* - P2.

Related themes that were frequently raised during the interviews by ML researchers and clinicians pertained to dataset versioning, specific changes that were made, the impact of changes, intentions for the change, labeling, and availability of labeling guidelines: *"The other aspect, that you touched on in the questionnaire on labeling is, it's well understood that agreement between doctors on a given task is way below 100% in which level of granularity would a human specialist go"* - P19.

Inclusion criteria and accessibility was the third most cited theme, often described as a leading cause of disparity of outcomes in clinical studies – a sentiment primarily expressed by clinicians. Geographic location of collected dataset, referring hospitals, criteria to participate in trials, and guidelines are all of immense importance in determining participant eligibility and selection in clinical studies: *"Some hospitals require referrals, some are teaching hospitals, some are general hospitals ... If someone said to me, these are all the x-rays from X's neonatal unit. I know what that means, but most people don't"* - P2.

For example, the EDSS score [33], which is the most commonly used metric for assessing disabilities in Multiple Sclerosis (MS), was defined for men hospitalized for MS in the United States Army during World War II [34]. One participant, in particular pointed to the unique impact of inclusion and exclusion criteria on datasets: *"what are the inclusion criteria? ... very different from a dataset like imagenet which is to some extent curated". "... in healthcare, there might be some particular exclusion criteria like all the clinical data for over 18s ... pretty blunt and pretty obvious ..."* - P2

As stated above, some of the challenges were specifically brought up and discussed by certain roles. For example, details around data composition, and versioning was primarily raised by ML Researchers, whereas issues related to calibration of the devices, inclusion/exclusion criteria, sites of data collections, were mainly discussed by clinicians. Subjectivity in labeling and demographic information were discussed across multiple roles.

5.2 Incentives of dataset curators to fill the Healthsheet questionnaire

Before discussing the benefits of healthsheets for dataset consumer, an important first step is to investigate potential incentives that can encourage dataset curators and extractors to create healthsheets for their datasets. The subjective nature of incentives was a repeated theme mentioned by several, largely dominated by a discussion of the trade-offs between the perceived effort and time required in creating a healthsheet and its impact. *"... if I've done a lot of hard work in releasing datasets, it's fair to ask me to create a datasheet. That's kind of meant to be a well-curated thing. That's ... an artifact ... with real Impact. ... it makes sense in that context to have something relatively detailed. ... this shouldn't be the most time-consuming and energy consuming piece of that project."* -P3 . P5 mentions *"People spend a ton of time when they're launching an open source tool investing in the documentation, right? I spend a ton of time not because I care about documentation itself, but because I know that's best practice for launching a tool."* Notably, one participant described long-term gains over time from reduced overhead, as a result of a one-time investment in the creation of healthsheets: *"that is going to be a peace of mind for them... So, if we don't have this data handy, like, in a single place, ... these teams are going to be on the hook to answer the same questions over and over and over and it's more overhead for them."*

Many participants also acknowledged *"we don't have good incentives in place right now to give people credit for creating datasheets."*-P5. This acknowledgement is key for our purposes because it allows us to move past an analysis of datasheets (and by extension, the modes of creation of datasets), as solely reliant on the subjective intentions of stakeholders producing these artifacts. P4 says *"Imagine I created a datasets that is very weird or very particular and I kind of want to hide that a bit or bury it ... You really don't want to document these things. You'd hope this doesn't happen but I can imagine it happening ... so I suppose the motivation has to be persuading people to use the datasets. I could imagine somewhere down the line, they're being sort of standardized reporting and Regulation and it being a sort of thing you have to do."* Along this line P16 says, *"More processes are harder for people to do given that their focus is on the research itself... So, either a carrot or stick to get someone to fill this out. Right? So it's either a standard, and if you don't use the standard across the board, they are unwilling to ingest your data ... or something that is going to make them really excited to fill it out"*.

P19 touches on valuing the contributions differently to change the incentives of the community *"Publications are valued more than data sharing and maybe that should change. Like, for example, when you get evaluated for promotion, how many datasets have you shared? Because that impacts so much more than just having one publication on a very limited topic. So I think revamping the way we evaluate people, might help with incentivizing."*-P19. In sum, incentives should not be solely based on personal intentions or motivation alone. This should be grounded by both the way the community and institutions rewarding the creation of these artifacts, as well as standardized guidelines.

5.3 Healthsheet as a diagnostic tool

After discussing the challenges faced by dataset consumers, we posed questions about the utility of *Healthsheets* to tackle the specific aforementioned issues, as well as general issues concerning the use of health datasets. All participants stated that they would use *Healthsheets* if datasets were accompanied by them. The most common utilities of *Healthsheets* reported were **understanding nuances and limitations of datasets** prior to their use; and an expected increase in the **efficiency and ease of use of datasets** heralded by accessible, clear and transparent meta-data. Below are quotes from the interviewees around their motivations for using *Healthsheets*.

Understanding nuances and limitations of datasets: *“I would see it as a guide for surfacing the domain specific questions or things that we should be mindful of”*-P1.

“Struggles with data could be technical or non-technical. The technical stuff is easy, the non-technical considerations is not... knowing the important caveats of the datasets ... You need actually documented somewhere.”- P4. *“Understanding the nuances of the datasets before you spend a lot of time working with it..., for example, if there’s a dataset that is incredibly clean, so all the photos in the dataset are from specific angles, always with the same camera always in the same position, that’s fine if that’s the task you’re training the model for. If you’re training the model for consumers it’s terrible”*- P4. P4 continues on the importance of transparency in data documentation by stating *“It feels morally easy to say that all datasets should cover everyone in the world, in a fair and Equitable way, and there’s some truth to that but there’s also practical real-world issues to deal with, and it’s okay if datasets are not fully representative as long as you know that from the beginning.”*- P4.

Efficient and easier use of datasets: *“some datasets may be created 30 years ago, ... science was quite different and the data is very different from how people nowadays think about things ... there’s a paradigm shift in the healthcare data, just like the healthcare practice. Having better documentation ... is definitely super helpful ... We care a lot about long-term follow-up, 10 years, 15 years, 20 years ... and it is non-trivial to find documentation ...it becomes extremely challenging. So if people have this level of details documented, like 30 years ago, I think that would have made our life a lot easier”*- P3.

“We need to understand the data, where it’s been used, or what are the possible usages of the data in order to advise teams, and also for ourselves to be able to store it securely ... we want to take as much of that information out of a written doc and make it enforceable through infrastructure and so teams have to think less about data and can focus on their work. So the more information we have about datasets, definitely can help.”- P11. *“For working on a product launch it is hugely important where the data came from, how large it is, and how to work on it”*- P7.

6 DATASHEET FOR HEALTHCARE DEEP DIVE

In this section, we dive deep into different groups of questions in the original datasheet [22], analyze the effectiveness and adequacy of the questionnaire sections for healthcare applications relying on the case studies, health literature and interviews. The original datasheet consists of questions grouped by stages of dataset creation, from

motivation, to data collection and maintenance. We also address the gaps that led to the development of *Healthsheet*².

6.1 Motivation

The motivation section centres around incentives for the creation of datasets.

“What was the purpose of creating the dataset? Were there specific gaps to be filled?”[†].

Some ML datasets are created to understand an ML research question [32] or examine the power of a learning approach. Some are collected to introduce a new task, or application that can be addressed by ML. It is important to explicitly mention the purpose for the creation of datasets, which helps both dataset creators and consumers to make informed decisions. In ML for healthcare, data is often collected for purposes such as effectiveness and side effects of medications, diagnostic tools, and disease prognosis.

Datasheets also address the implicit motivations and reasons coming from *“funding resources or research institutes”* involved in ML healthcare research. Questions on funding could give a better understanding of underlying incentives for choosing the given research problem and dataset. One of the historical examples of this issue is the tobacco companies funding lung cancer research [12]. In addition, transparency of funding could shed a light on which healthcare problems are prioritized and why.

Studies [14] show, there is a disparity of funding rate on problems that impact lower income and disadvantaged groups compared to the general population that could stem from social injustice. For example, Sickle Cell disease impacts mainly the Black population while cystic fibrosis impacts mainly the white population. Cystic fibrosis receives significantly more funding and research resources, despite both being genetic disorders of similar severity [14, 19, 48]. Multiple studies also show that health research on topics related to women’s health does not get sufficient funding and attention [13, 14, 16, 52].

Studies suggest that disparities in research teams demographics and backgrounds could impact funding priorities and exacerbate existing socioeconomic, racial, and gender injustices [14, 52, 66]. To address this issue, in the Primary Healthsheet questionnaire motivation section, we added a question on the demographic disparity of their research teams behind dataset creation that initiated generative discussions and comments were divided. Most participants were really eager to keep this question. 4 participants suggested to instead highlight and prioritize questions, which were on (i) demographic information in the dataset such as inclusion criteria in clinical studies and data collection, (ii) labeler guidelines and subjectivity in labeling, and (iii) expertise of researchers involved in data collection. From these 4 participants, 2 were in favour of removing or modifying the question.

“I think the demographics of researchers can start to feel quite uncomfortable quite quickly. The classic example is Sickle Cell, which is massively underfunded. I think, though, I would push back on that because I would say, Let’s say there’s no sickle-cell dataset and there should be a better signal so we can explain that for all kinds of socio-cultural ways. But that doesn’t actually change. The fact that

² all questions in this section with dagger sign, represents questions from original datasheet paper [22] or with small modification

what you've got now is a cystic fibrosis dataset, right? So if the entire research team was black, would it be still a cystic fibrosis research datasets or it wouldn't be."-P4.

"If you condition on the dataset itself, the demographic becomes an explanation as to why the data was created, but it adds very little to the actual interpretation of the datasets. And I think one thing that I would, I would be much more interested in the labels and subjectivity. So, for example, how is a heart attack defined? Is it by blood test? Is it by symptoms? Because those huge amounts of subjectivity and inclusion criteria in that, that excludes women and ethnic minorities."-P4. One of the participants also mentioned the subjectivity in the definition of demographic disparities in teams.

"It's tricky because it makes the question quite subjective ... a research team of 10 white men from the same University and similar ages may say they working with two people who are not the same as them. They might think: oh well compared to the rest of the field, we have this wonderfully diverse team. I wouldn't like to blame them for that. But that's also missing important information"-P2.

Multiple participants also mentioned the importance of highlighting the expertise of curators of datasets.

"if you're collecting a dermatological dataset, we'd want to make sure that the composition of the team includes clinicians, and, you know, it's not just ml researchers..."-P6.

On the other hand, most participants were in favor of keeping the question and spoke out about the implications of this. One participant who was also a dataset creator mentioned the correlation that they observed between the gender diversity of the research team and availability of datasets. They also observed penitential relationships to other demographic attributes of researchers and dataset availability to public.

"We have a lot of projects, that I was hoping to convince people to think about this question. So one of them which is partially related is, we're looking at the researchers and publicly available datasets. So we systematically searched ML projects in Healthcare in 2019, ... we show nicely that there is diversity of researchers, if a dataset is publicly available ... So we've also improved the representation of minority researchers. If a dataset is publicly available versus if this is only a accessible to the internal investigators, tend to preserve the order of old white male, professors publishing off those datasets. So I think this is the evidence that we've been wanting like we everyone's talking about open science, but everything is a bit more theoretical when it comes to benefit."-P21.

After analyzing the data on this question, we add the following variation of the questions, that both touches on the demographic disparity and expertise of researchers:

What is the demographic population of dataset curators/ generators? and What is the distribution of backgrounds and experience/expertise of the dataset curators/generators?

In addition, we cover Health-related transparency questions related to the inclusion criteria, subjectivity of labeling and demographic information of people in datasets.

Many Health datasets are released without predefined tasks and labels, or will be later labeled by researchers for specific tasks and problems. In addition, datasets in healthcare usually gets expanded over time. This dynamic updates of datasets for various reasons were brought up with all interviewees who worked directly with

Health related datasets. It is important to report the rationals, motivations, and funding sources for initial data collection process as well as new versioning, tasks, and research problems. This aspect is thoroughly discussed in 7.1 and additional questions are suggested.

6.2 Composition

The composition section investigates the information that helps dataset consumers to make informed decision prior to using the dataset. Questions on the type and characteristics of the data in the dataset, target labeling, data splits, and all compositional aspects of the data are addressed in this section. Having a centralized documentation resource on composition of data was brought up mainly by ML researchers and engineers working with healthcare datasets or working on health-data related infrastructure. This was specifically highlighted as a time-efficient practice for both dataset creators and consumers. *"there may be some information in data composition, for example, ... I received, like, ten emails a day asking me where did you store this information or how I can extract this tables from your data? even from a very superficial level of just like data acquisition, I will save time for myself...(by centralizing documentations)."*-P4

"What do the instances that comprise the dataset represent?†", and *"What data does each instance consist of?†"*.

In healthcare scenarios, data is typically associated with people who are patients or subjects of the experiments. Each subject/patient is associated with a hierarchy of information, which can span multiple time periods. For instance, a patient can have multiple admissions to the hospital, with each admission lasting multiple days. Similarly, a patient might be linked to multiple medical imaging cases, each including multiple images or different types of images (e.g., structural MRI and CT). This hierarchy of the information is an important aspect of the healthcare application, and is currently not captured in datasheets.

"Is there a label or target associated with each instance?, Are there recommended data splits?†"

Target labels of datasets can be used to define a task based on the given data. In many healthcare scenarios, data is collected as a form of database, with data being associated with patients characteristics but not a specific task and consequently no data splits associated with it. Datasets in healthcare are often collected over the span of multiple years and as a result, multiple tasks could be associated with the data and problem definitions that were not of interest when the data collection process was started, may now be of importance. However due to the sensitivity and importance of healthcare applications, it is crucial to update the healthsheet for new tasks and applications that become associated with the dataset.

"Does the dataset contain data that might be considered confidential?†"

For many Public Health datasets, it is important to ensure the confidentiality of data and data should be de-identified. In addition, compliance with data protection rules such as EU's General Data Protection Regulation (GDPR), or comparable regulations in other jurisdictions such as Health Insurance Portability and Accountability Act (HIPAA) depending on the country of data collection and use, is of immense importance in healthcare applications. Data

de-identification in healthcare should follow country-based rules and regulations.

“Is any information missing from individual instances?” †

In healthcare datasets, presence of censoring where the outcome is only partially observable happens frequently. This could be due to multiple reasons, including loss, no follow-up, patients choosing to skip survey questions, and accessibility issues. It is specifically important to identify loss of information due to accessibility issues, and question these aspects during the data collection process.

“Does the dataset identify any subpopulations (e.g., by age, gender)?” †

This question is very important on healthcare applications but it should also be expanded on various aspects, incentives for collection and use of demographic information, the way this information is collected and employed, and compliance with countries regulation should be considered.

6.3 Collection, uses, distributions, and maintenance

Developmental stages of a dataset from collection to use, distributions and suitability of datasets are crucial in health-specific applications.

“Who was involved in the data collection process?” †. In healthcare applications, clinicians, specialists, and patients themselves could be involved in providing information as part of the data collection process. For example, providing MRI reports, answering surveys in app datasets, assessing disability scales (e.g. EDSS in MS), functionality tests and pain scales. In healthcare, in addition to acknowledging the expertise of people involved in the data collection process, it is crucial to note that collected data could be very subjective. Studies show that EDSS scores could be very subjective in estimating disability in MS. It is reported that in many cases pain assessments is biased towards undermining the experiences of Black patients [27].

“Over what time-frame was the data collected?” †. *“Has the dataset been used for any tasks already?” †*, *“What (other) tasks could the dataset be used for?” †*, *“Are there tasks for which the dataset should not be used?” †*

Healthcare datasets are often collected over span of multiple years. This contributes to shift in the tasks, intended use, and new research questions derived from the datasets. For example, MIMIC-III was created over the span of 11 years. Multiple versions of datasets are being produced in the process. Floodlight dataset is constantly being updated by new data entries. Although that it is crucial to enable the use of datasets for new research questions and develop updated versions of them, it is important to have a comprehensive track of development processes, targeted research questions, intended uses and rationale behind them.

7 HEALTHSHEET: AUGMENTING THE DATASHEETS FOR HEALTHCARE

In this section, relying on literature, interviews, and our case studies, we discuss the aforementioned aspects that should be expanded or discussed in health context.

7.1 Dataset versioning

The amount of available Healthcare data has seen significant growth in recent years [1]. Multiple healthcare datasets, specifically datasets created using mobile apps such as Floodlight, are dynamic, meaning a new version of the dataset could be available in a short span of time. For some datasets, multiple versions of the same dataset may be released.

So, it is important to keep track of the properties that evolve with time. In the same way coding documentation is a living artifact, we propose that *Healthsheets* be treated in a similar manner.

Some datasets in healthcare are released without labels and pre-defined tasks, or will be later labeled by researchers for specific tasks and problems, to form sub-versions of the dataset. The following set of questions clarifies the information about the current version of the dataset. It is important to report the rationale on labeling the data in any of the versions and sub-versions that this datasheet addresses, funding resources and motivations behind each released version of the dataset. Versioning aspects and the importance of having access to detailed documentation of records for variant versioning was mentioned with multiple interviewees. *“In many cases, dataset curators don’t really intend to create from the outset a dynamically updating dataset, like they kind of release the datasets and then some time passes and they’re like, oh, we have some more data now. So now we’re going to release again. I think there’s very few institutions that go from the outset with the intention of dynamically updating.”*-P17.

Questions and instructions related to dataset versioning are provided in Healthsheet questionnaire, Appendix A, Dataset Versioning section.

As we move into the world of mobile health [63], dataset creation will move from snapshots to a more fluid form where data points are continuously added. It is important to capture this distinction since it will inform the way a dataset is utilized.

In the case of static datasets, each sampling is done for a purpose. That motivation will surface in the dataset’s properties, how many instances of each data point exist, where they were collected from, during what period. These properties change for each released version of a static dataset and should evolve with the versions.

7.2 Accessibility in data collection

One of the most cited aspects that was addressed during the interview process, was inclusion criteria. It is important to know who is involved in this dataset, and how. As P4 said: *“What were the inclusion criteria? ... If it was only done in, let’s say US hospitals, it’s a limitation. If you could certainly apply to other US, hospitals, could you apply it to a Canadian or UK hospital? Possibly! Could you apply it to a South African Hospital? Maybe! Could you apply it to an Eritrean Hospital? Probably not. So, just understanding that kind of really headline stuff about recruitment where the data was from. It’s going to depend on your deployment target, and I think it’s very easy to say that.”*

In addition, many healthcare data collection systems exclude some patients/subjects due to the inaccessibility of data collection frameworks. Two of the datasets that we worked with are on Multiple Sclerosis patients, i.e. MSOAC [35] and Floodlight. Some patients with MS have movement issues such as tremor and this may make

it harder for them interact with mobile apps. Some also have visual impairments. In order to have better patient representation, and consequently useful benchmarks, it is important to assess accessibility during the data collection process. In addition, most of the data collection processes only cover English language, which is not the default language for many patients. It is very important to be transparent about the limitation of the datasets with regarding to the accessibility aspects. (1) Is there any language-based communications with patients? If yes, describe the choices of language(s) for communication. (e.g. if there is an app for communication, what are the language options in the app? (e.g: English, Spanish)) (2) What are the accessibility measurements and what aspects were considered when the study was designed and implemented? (e.g. in a clinical study if it is accessible to someone with motor function disability? or if there is an app for patients with MS, are all aspects of inaccessibility associated with their potential disabilities considered?) (3) What are the inclusion criteria in the study? What is the recruitment strategy? (e.g., men over 18 years old, patients who are referred from the hospital X)

7.3 Collection and use of demographic information

Racism and social conditions have had a fundamental impact on causes and prognosis of many diseases throughout the history [37, 67]. This has been increasingly evident during the COVID-19 pandemic. Data collected shows that BIPOC populations, specifically, Indigenous, Pacific islanders and Black populations have experienced the highest death toll from COVID-19³ [17]. On the one hand, given that socio-cultural factors can significantly affect the healthcare conditions of BIPOC populations, collection and access to demographic information is critical to conducting fairness analyses of ML healthcare models.

On the other hand, modern American medicine has historical roots in scientific racism and eugenics movements, and racialized conceptions of susceptibility to disease persist to this day [4]. At the structural level, actions by parties ranging from medical schools to providers, insurers, health systems, legislators, and employers have ensured that racially segregated Black communities have limited and substandard care [5]. Within medicine, there is no uniform practice regarding the use of race as a study variable and little to no expectation that authors examine racism as a cause of residual health inequities among racial groups [11]. Likewise for ML, there is no uniform practice regarding the use of race as a variable in predictive models or datasets, nor is there an expectation that ML researchers examine structural racism as a cause of health inequities among racialized groups.

Due to these historical and structural reasons, great care must be taken when collecting and using demographic information and specifically race in healthcare datasets. As Roberts [58] points out, "Using biological terms to define social inequities makes them seem natural—the result of inherent racial differences that can't be changed instead of unjust societal structures that must be dismantled."

[21] recently released updated guidance on reporting of race and ethnicity in medical journals. These suggestions encourage

³<https://www.apmrresearchlab.org/covid/deaths-by-race>

fairness, equity, consistency and clarity in the use and reporting of race and ethnicity in medical and science journals. The guidance explicitly recognizes race and ethnicity as social constructs, and we agree that this understanding is critical. In *healthsheet*, we would like to identify whether the demographic information or variables for a dataset are collected and if yes/no, was there any rational or motivation behind it⁴.

Does the dataset identify any demographic sub-populations (e.g., by age, gender, sex, ethnicity)? If yes, (1) The reasons that these categories were assessed also should be described in the datasheet [21]. (2) Please describe who identified these categories and the source of the classifications used (e.g: self-report or selection, investigator observed, database, electronic health record, survey instrument) [21]. (3) Are patients aware / consent to the collection and use of their demographic information? (4) Is there any task that any of these demographic information (e.g: gender, age, ethnicity) has any biologically proven impact on the outcome? If yes, please provide a link to the study, or publication. (5) Is there any mechanism for updating some of this information through the data collection process or afterwards? For example, if someone wants to change their gender information. Provide a description of the respective distributions of each subgroup populations within the dataset. Is there any disparity of results between different demographics? If no, (1) Is there any regulation that prevents some of the gender / sex data collection in your study (for example, the country that the data is collected in)? (2) Are you employing methods to reduce the disparity of error rate between different demographic subgroups when demographic labels are unavailable? Please describe.

7.4 Devices and Contextual attributes in Data Collection

The modality of a dataset refers to the different types of data it contains. Depending on the domain, healthcare datasets are created using a variety of devices and equipment. These devices come in various forms, ranging from large machines located in the hospitals (Magnetic resonance imaging (MRI), Computer Tomography Scans, X-rays, etc.) to small-scale wearable devices (eg. smart health trackers, wearable blood pressure monitors, etc.). Modalities not only govern the fidelity but also the statistical properties of the dataset. Hence, it is important to capture the details about the modalities of the dataset. Questions related to the data that requires a device or equipment for collection or the context of the experiment, are provided in Appendix A, Healthsheet questionnaire, under the Data Composition section, as "Devices and contextual attributes in data collection"

7.5 Labeling and subjectivity in labeling

In medical domains, researchers usually take a dataset and appropriate it for a defined task. Researchers may have their own guidance. It is important to know what the incentive of the original creators was, if there is a guideline for the current version or sub-versions of the dataset? It is also important to note the impact of subjectivity of labeling and labelers, as well as the human level performance

⁴Please note that the dept. of Health and Human Services (HHS) and FDA usually call out for demographic information to be collected. However, there is not a requirement of informed consent forms required by HHS. IRBs can require that information if they deem that information to be important to an individual's choice as to whether they participate in research or not

on the applications that the dataset is supposed to address. Related questions are provided in Appendix A, Healthsheet questionnaire, under the "Labeling and subjectivity of labeling".

7.6 Challenge tests and confounding factors

When developing models for healthcare, it is crucial to understand what limitations they will have. Such limitations stem partly from the data distribution covered: the narrower the data distribution, the more limited the use cases. This problem does not pertain to healthcare and generalization is a field of research on its own (see [18, 46, 68, 69] for reviews). In this section, we suggest to explicitly report limitations in the data distribution that might prevent generalization: (1) Which factors in the data might limit the generalization of potentially derived models? Is this information available as auxiliary labels for challenge tests? For instance: (a) Number and diversity of devices included in the dataset, (b) Data recording specificities, e.g., the view for a chest x-ray image, (c) Number and diversity of recording sites included in the dataset. (d) Distribution shifts over time (e.g., [44]) (2) What confounding factors might be present in the data? (a) Interactions between demographic or historically marginalized groups and data recordings, e.g., were women patients recorded in one site, and men in another? (b) Interactions between the labels and data recordings, e.g., were healthy patients recorded on one device and diseased patients on another?

8 CASE STUDIES AND FINDINGS

For the purpose of refining our framework, we used 3 publicly available datasets: MIMIC-III [30], MSOAC [35] and Floodlight [6].

Reasons for choices of the datasets: MIMIC-III is a clinical dataset. The reason that we chose MIMIC-III [30], was due to the availability of the dataset, wide range of tasks, and applications on the dataset, availability and depth of documentation and details available from the dataset. In addition, multiple publications, employed the dataset to suggest new sub-version, tasks and labels on the dataset [59]. We also studied MSOAC [35] and Floodlight [6], both on Multiple Sclerosis disease and are respectively on Electronic Health Records (EHR), and smartphone-based performance outcome measures. Floodlight is one of the examples of a dataset that has dynamic updates based on user data, and we chose it to reflect on our versioning and accessibility questions.

We started by answering the original datasheet questionnaire, and identify the missing aspects, to be integrated in our *Primary Healthsheet*. Once we had the finalized questionnaire, we went back to answer all questions from the start and identified some areas where necessary information is not publicly available or easily found. While this is due to the fact that we are not owners or creators of the dataset, it does outline the need for dataset creators to make use of such frameworks as the one we propose.

Funding information: while for MIMIC-III and MSOAC this information was easy to find either on their official website or through online presentations, for Floodlight we could only make an assumption that it was sponsored by Roche based on the dataset website.

Dataset statistics: MIMIC-III was the only dataset in the study that voluntarily presented such information. We recognize this is a time consuming process, but we would like to urge anyone creating a new dataset to consider releasing such information. It can be used as a test for identifying gaps in the data.

Data acquisition: This was one of the hardest questions to find information on. Healthcare data consist of numerous fields, each of which could be collected through various mechanisms and very often these mechanisms are not specified anywhere. Even in the case of MIMIC-III, one of the more detailed datasets we used, for some fields the mechanism of acquisition is unknown.

Demographics of dataset creators: Many researchers focus on diversity in the dataset itself, but when it comes to curating the data we often find that we do not know who was involved or included.

Versioning: Coming up with a common nomenclature for versions was a difficult task. We needed a term that would cover both conventional datasets that are collected and released once and continuously updated datasets. Some datasets such as MIMIC-III become so widely used that some research problems define their own version of a dataset which gains independent popularity. The original Healthsheet creators cannot be expected to maintain all information derived from their original work. For this reason we expect that once the base Healthsheet is completed, the community will step up and add Healthsheet information for any new widely used version.

Another observation from the case studies was how difficult it was to find the information needed for the various sections. This meant that one source of information was usually not enough, we had to dig through official webpages, GitHub repositories, arXiv papers, news articles and more to pull the information and be sure it is accurate. This further outlines the need of a unified place where all this information lives. Further to completing a Healthsheet, we also propose that any entered information is officially verified and confirmed by the dataset owners, creators and maintainers.

Finally the true motivations behind the creation of each dataset can only be filled in by the groups responsible for them. As users of the datasets we can only assume the intention or infer it from publicly available resources, but us answering these questions will not influence the use of the data or the means of acquisition.

9 DISCUSSION

The field is well aware that algorithm design is not enough—we need to find “good” datasets with which to evaluate algorithms. Yet, we have not had a community-wide standard to define what a “high-quality” or “good” health dataset is. Healthsheet aims to bridge this gap and ensure that:

- (1) Data used in ML for health experiments is representative of real applications, with authentic and representative patient data.
 - (2) Biases in datasets are more easily detected and understood, due to better means of describing and treating data types, and mitigating faulty assumptions about data types that arise due to lack of documentation.
 - (3) Benchmark datasets are selected based on completeness and alignment with documentation standards, to make reliable, ethical comparisons of ML performance results possible.
 - (4) Healthsheet could be a tool for retrospective and historical dataset analysis.
 - (5) Future auditing could be conducted more easily.
 - (6) Healthsheet serves as a community-based artifact: with opportunities for review and evolution to meet community needs.
- Limitations:** We only worked with publicly available datasets. Majority of datasets in healthcare are not publicly available and this on its own could add additional concerns around transparency.

Our methodology for defining new questions was grounded on the expertise of the team of authors, interview participants, and datasets we studied.

We hope that our work will start a conversation in the community and will continue to tailor *Healthsheet* to broader healthcare scenarios over time. In addition, we believe that involving more stakeholders, such as patients or patients-representative communities, could immensely improve the *Healthsheet* questionnaire, as a transparency artifact. Finally, we would like to emphasize here that transparency and this questionnaire should not be thought of as ends by themselves but defined as means toward responsible and equitable ends.

ACKNOWLEDGMENTS

We would like to thank Dr Madeleine Elish and Dr Berk Ustun and Noah Broestl, for their contributions to the earlier version of the paper presented at ML4H, Dr Timnit Gebru, Dr Xiao Liu, Dr Viknesh Sounderajah, Dr Leo Anthony Celi, Dr Mark Diaz, Dr Danielle Belgrave and Gerardo Flores for their valuable feedback, and last but not the least, all the interviewees who participated in this study and contributed in the formation of this paper.

REFERENCES

- [1] 2020. Big hopes for big data. *Nature Medicine* 26, 1 (2020), 1. <https://doi.org/10.1038/s41591-019-0740-8>
- [2] Mike Ananny and Kate Crawford. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *new media & society* 20, 3 (2018), 973–989.
- [3] Matthew Arnold, Rachel KE Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, Aleksandra Mojsilović, Ravi Nair, K Natesan Ramamurthy, Alexandra Olteanu, David Piorkowski, et al. 2019. FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (2019), 6–1.
- [4] Zinzi D. Bailey, Justin M. Feldman, and Mary T. Bassett. 2021. How Structural Racism Works — Racist Policies as a Root Cause of U.S. Racial Health Inequities. *New England Journal of Medicine* 384, 8 (2021), 768–773. <https://doi.org/10.1056/NEJMms2025396> arXiv:<https://doi.org/10.1056/NEJMms2025396>
- [5] Zinzi D Bailey, Nancy Krieger, Madina Agénor, Jasmine Graves, Natalia Linos, and Mary T Bassett. 2017. Structural racism and health inequities in the USA: evidence and interventions. *The Lancet* 389, 10077 (2017), 1453–1463.
- [6] Mike Baker, Johan van Beek, and Christian Gossens. 2020. Digital health: Smartphone-based monitoring of multiple sclerosis using Floodlight. *Scientific American*. Accessed June 2 (2020), 2020.
- [7] Jack Bandy and Nicholas Vincent. 2021. Addressing" Documentation Debt" in Machine Learning Research: A Retrospective Datasheet for BookCorpus. *arXiv preprint arXiv:2105.05241* (2021).
- [8] Imon Banerjee, Ananth Reddy Bhimireddy, John L Burns, Leo Anthony Celi, Li-Ching Chen, Ramon Correa, Natalie Dullerud, Marzyeh Ghassemi, Shih-Cheng Huang, Po-Chih Kuo, et al. 2021. Reading Race: AI Recognises Patient's Racial Identity In Medical Images. *arXiv preprint arXiv:2107.10356* (2021).
- [9] Emily M Bender and Batya Friedman. 2018. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics* 6 (2018), 587–604.
- [10] Misha Benjamin, Paul Gagnon, Negar Rostamzadeh, Chris Pal, Joshua Bengio, and Alex Shee. 2019. Towards standardization of data licenses: The montreal data license. *arXiv preprint arXiv:1903.12262* (2019).
- [11] Rhea W Boyd, Edwin G Lindo, Lachelle D Weeks, and Monica R McLemore. 2020. On racism: a new standard for publishing on racial health inequities. *Health Affairs Blog* 10, 10.1377 (2020).
- [12] Allan M Brandt. 2012. Inventing conflicts of interest: a history of tobacco industry tactics. *American journal of public health* 102, 1 (2012), 63–71.
- [13] Shraddha Chakradhar. 2018. Discovery cycle. *Nature medicine* 24, 8 (2018), 1082–1086.
- [14] Irene Y Chen, Emma Pierson, Sherri Rose, Shalmali Joshi, Kadija Ferryman, and Marzyeh Ghassemi. 2020. Ethical Machine Learning in Healthcare. *Annual Review of Biomedical Data Science* 4 (2020).
- [15] Anupam Datta, Shayak Sen, and Yair Zick. 2016. Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. In *2016 IEEE symposium on security and privacy (SP)*. IEEE, 598–617.
- [16] VH Eisenberg, C Weil, G Chodick, and V Shalev. 2018. Epidemiology of endometriosis: a large population-based database study from a healthcare provider with 2 million members. *BJOG: An International Journal of Obstetrics & Gynaecology* 125, 1 (2018), 55–62.
- [17] Gabriel J Escobar, Alyce S Adams, Vincent X Liu, Lauren Soltesz, Yi-Fen Irene Chen, Stephen M Parodi, G Thomas Ray, Laura C Myers, Charulata M Ramaprasad, Richard Dlott, et al. 2021. Racial disparities in COVID-19 testing and outcomes: retrospective cohort study in an integrated health system. *Annals of internal medicine* (2021).
- [18] Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R Arabnia. 2020. A Brief Review of Domain Adaptation. (Oct. 2020). arXiv:[2010.03978](https://arxiv.org/abs/2010.03978) [cs.LG]
- [19] Faheem Farooq and John J Strouse. 2018. Disparities in foundation and Federal Support and development of new therapeutics for sickle cell disease and cystic fibrosis. *Blood* 132 (2018), 4687.
- [20] Oliver Faust, Yuki Hagiwara, Tan Jen Hong, Oh Shu Lih, and U Rajendra Acharya. 2018. Deep learning for healthcare applications based on physiological signals: A review. *Computer methods and programs in biomedicine* 161 (2018), 1–13.
- [21] Annette Flanagan, Tracy Frey, Stacy L Christiansen, AMA Manual of Style Committee, et al. 2021. Updated Guidance on the Reporting of Race and Ethnicity in Medical and Science Journals. *JAMA* 326, 7 (2021), 621–627.
- [22] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).
- [23] R Stuart Geiger, Kevin Yu, Yanlai Yang, Mindy Dai, Jie Qiu, Rebekah Tang, and Jenny Huang. 2020. Garbage in, garbage out? Do machine learning application papers in social computing report where human-labeled training data comes from?. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 325–336.
- [24] Judy Wawira Gichoya, Liam G McCoy, Leo Anthony Celi, and Marzyeh Ghassemi. 2021. Equity in essence: a call for operationalising fairness in machine learning for healthcare. *BMJ Health & Care Informatics* 28, 1 (2021), e100289.
- [25] Leo A Goodman. 1961. Snowball sampling. *The annals of mathematical statistics* (1961), 148–170.
- [26] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cudros, et al. 2016. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama* 316, 22 (2016), 2402–2410.
- [27] Kelly M Hoffman, Sophie Trawalter, Jordan R Axt, and M Norman Oliver. 2016. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proceedings of the National Academy of Sciences* 113, 16 (2016), 4296–4301.
- [28] Ben Hutchinson, Andrew Smart, Alex Hanna, Emily Denton, Christina Greer, Oddur Kjartansson, Parker Barnes, and Margaret Mitchell. 2021. Towards accountability for machine learning datasets: Practices from software engineering and infrastructure. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 560–575.
- [29] Hussein Ibrahim, Xiaoxuan Liu, Nevine Zariffa, Andrew D Morris, and Alastair K Denniston. 2021. Health data poverty: an assailable barrier to equitable digital health care. *The Lancet Digital Health* (2021).
- [30] Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li Wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific Data* 3 (2016), 1–9. <https://doi.org/10.1038/sdata.2016.35>
- [31] René F Kizilcec. 2016. How much information? Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 2390–2395.
- [32] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [33] John F Kurtzke. 1983. Rating neurologic impairment in multiple sclerosis: an expanded disability status scale (EDSS). *Neurology* 33, 11 (1983), 1444–1444.
- [34] John F Kurtzke. 2015. On the origin of EDSS. *Multiple sclerosis and related disorders* 4, 2 (2015), 95–103.
- [35] Nicholas G LaRocca, Lynn D Hudson, Richard Rudick, Dagmar Amtmann, Laura Balcer, Ralph Benedict, Robert Bermel, Ih Chang, Nancy D Chiaravalloti, Peter Chin, et al. 2018. The MSOAC approach to developing performance outcomes to measure and monitor multiple sclerosis disability. *Multiple Sclerosis Journal* 24, 11 (2018), 1469–1484.
- [36] Changhee Lee, Jinsung Yoon, and Mihaela Van Der Schaar. 2019. Dynamic-deepfit: A deep learning approach for dynamic survival analysis with competing risks based on longitudinal data. *IEEE Transactions on Biomedical Engineering* 67, 1 (2019), 122–133.
- [37] Bruce G Link and Jo Phelan. 1995. Social conditions as fundamental causes of disease. *Journal of health and social behavior* (1995), 80–94.
- [38] Xiaoxuan Liu, Samantha Cruz Rivera, David Moher, Melanie J Calvert, and Alastair K Denniston. 2020. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. *bmj* 370 (2020).

- [39] Donald Martin Jr, Vinodkumar Prabhakaran, Jill Kuhlberg, Andrew Smart, and William S Isaac. 2020. Participatory problem formulation for fairer machine learning through community based system dynamics. *arXiv preprint arXiv:2005.07572* (2020).
- [40] Michael E Matheny, Danielle Whicher, and Sonoo Thadaney Israni. 2020. Artificial intelligence in health care: a report from the National Academy of Medicine. *Jama* 323, 6 (2020), 509–510.
- [41] Maciej A Mazurowski. 2020. Artificial intelligence in radiology: some ethical considerations for radiologists and algorithm developers. *Academic radiology* 27, 1 (2020), 127–129.
- [42] Brent Mittelstadt. 2019. Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence* 1, 11 (2019), 501–507.
- [43] Chirag Nagpal, Steve Yadlowsky, Negar Rostamzadeh, and Katherine Heller. 2021. Deep Cox mixtures for survival regression. *arXiv preprint arXiv:2101.06536* (2021).
- [44] Bret Nestor, Matthew B. A. McDermott, Willie Boag, Gabriela Berner, Tristan Naumann, Michael C. Hughes, Anna Goldenberg, and Marzyeh Ghassemi. 2019. Feature Robustness in Non-stationary Health Records: Caveats to Deployable Model Performance in Common Clinical Machine Learning Tasks. In *Proceedings of the 4th Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 106)*, Finale Doshi-Velez, Jim Fackler, Ken Jung, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens (Eds.). PMLR, Ann Arbor, Michigan, 381–405. <http://proceedings.mlr.press/v106/nestor19a.html>
- [45] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [46] Sinno Jialin Pan and Qiang Yang. 2010. A Survey on Transfer Learning. *IEEE Trans. Knowl. Data Eng.* 22, 10 (Oct. 2010), 1345–1359.
- [47] Efthimos Parasisidis, Elizabeth Pike, and Deven McGraw. 2019. A Belmont report for health data. *The New England journal of medicine* 380, 16 (2019), 1493–1495.
- [48] M Park. 2010. NCAA genetic screening rule sparks discrimination concerns. *CNN.com* (2010).
- [49] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M Bender, Emily Denton, and Alex Hanna. 2020. Data and its (dis) contents: A survey of dataset development and use in machine learning research. *arXiv preprint arXiv:2012.05345* (2020).
- [50] Marco V Perez, Kenneth W Mahaffey, Haley Hedlin, John S Rumsfeld, Ariadna Garcia, Todd Ferris, Vidhya Balasubramanian, Andrea M Russo, Amol Rajmane, Lauren Cheung, et al. 2019. Large-scale assessment of a smartwatch to identify atrial fibrillation. *New England Journal of Medicine* 381, 20 (2019), 1909–1917.
- [51] Francesco Piccialli, Francesco Calabrò, Danilo Crisci, Salvatore Cuomo, Edoardo Prezioso, Roberta Mandile, Riccardo Troncone, Luigi Greco, and Renata Auricchio. 2021. Precision medicine and machine learning towards the prediction of the outcome of potential celiac disease. *Scientific Reports* 11, 1 (2021), 1–10.
- [52] Emma Pierson, Tim Althoff, Daniel Thomas, Paula Hillard, and Jure Leskovec. 2019. The menstrual cycle is a primary contributor to cyclic variation in women's mood, behavior, and vital signs. *bioRxiv* (2019), 583153.
- [53] Vinay Uday Prabhu and Abeba Birhane. 2020. Large image datasets: A pyrrhic win for computer vision? *arXiv preprint arXiv:2006.16923* (2020).
- [54] Mahima Pushkarna and Andrew Zaldivar. 2021. Data Cards: Purposeful and Transparent Documentation for Responsible AI. (2021). https://datacentralai.org/papers/112_CameraReady_Data_Cards.pdf
- [55] Mahima Pushkarna, Andrew Zaldivar, and Daniel Nanas. [n.d.]. Data Cards Playbook: Participatory Activities for Dataset Documentation. https://facctconference.org/2021/acceptedcraftsessions.html#data_cards
- [56] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timmit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 33–44.
- [57] Samantha Cruz Rivera, Xiaoxuan Liu, An-Wen Chan, Alastair K Denniston, and Melanie J Calvert. 2020. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the SPIRIT-AI extension. *bmj* 370 (2020).
- [58] Dorothy E Roberts. 2018. The most shocking and inhuman inequality: Thinking structurally about poverty, racism, and health inequities. *U. Mem. L. Rev.* 49 (2018), 167.
- [59] Subhrajit Roy, Diana Mincu, Eric Loreaux, Anne Mottram, Ivan Protsyuk, Natalie Harris, Emily Xue, Jessica Schrouff, Hugh Montgomery, Ali Connell, et al. 2021. Multi-task prediction of organ dysfunction in the ICU using sequential sub-network routing. (2021).
- [60] Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M Aroyo. 2021. “Everyone wants to do the model work, not the data work”: Data Cascades in High-Stakes AI. In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [61] Mark Schurgin, Mark Schlager, Laura Vardoulakis, Laura R Pina, and Lauren Wilcox. 2021. Isolation in Coordination: Challenges of Caregivers in the USA. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [62] Ramya Srinivasan, Emily Denton, Jordan Famularo, Negar Rostamzadeh, Fernando Diaz, and Beth Coleman. 2021. Artsheets for Art Datasets. (2021).
- [63] S. Steinhubl, E. D. Muse, and E. Topol. 2015. The emerging field of mobile health. *Science Translational Medicine* 7 (2015), 283rv3 – 283rv3.
- [64] Nenad Tomašev, Xavier Glorot, Jack W Rae, Michal Zielinski, Harry Askham, Andre Saraiva, Anne Mottram, Clemens Meyer, Suman Ravuri, Ivan Protsyuk, et al. 2019. A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature* 572, 7767 (2019), 116–119.
- [65] Mohammed Uddin, Yujiang Wang, and Marc Woodbury-Smith. 2019. Artificial intelligence for precision medicine in neurodevelopmental disorders. *NPJ digital medicine* 2, 1 (2019), 1–10.
- [66] D Vidyasagar. 2006. Global notes: the 10/90 gap disparities in global health research. *Journal of Perinatology* 26, 1 (2006), 55–56.
- [67] Derick Wade. 2007. Ethics of collecting and using healthcare data.
- [68] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Wenjun Zeng, and Tao Qin. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. (March 2021). [arXiv:2103.03097](https://arxiv.org/abs/2103.03097) [cs.LG]
- [69] Karl Weiss, Taghi M Khoshgofaar, and Dingding Wang. 2016. A survey of transfer learning. *Journal of Big Data* 3, 1 (May 2016), 1–40.
- [70] Jionglin Wu, Jason Roy, and Walter F Stewart. 2010. Prediction modeling using EHR data: challenges, strategies, and a comparison of machine learning approaches. *Medical care* (2010), S106–S113.

APPENDIX A: HEALTHSHEET

The provided answers here are for the MIMIC-III dataset, that was one of our case studies. Several questions in the paper are from original datasheets [22] or an adaptation of datasheets questionnaire. The Healthsheets questionnaire is being updated by ongoing interviews and feedback.

General Information

If the answer to any of the questions in the questionnaire is N/A, please describe why the answer is N/A (e.g: data not being available)

Provide a 2 sentence summary of this dataset.

MIMIC (Medical Information Mart for Intensive Care) is a large, freely-available database comprising deidentified health-related data from patients who were admitted to the critical care units of the Beth Israel Deaconess Medical Center.

Has the dataset been audited before? If yes, by whom and what are the results?

N/A. Information could not be easily found.

Dataset Versioning

Version: A dataset will be considered to have a new version if there are major differences from a previous release. Some examples are a change in the number of patients/participants, or an increase in the data modalities covered.

Sub-version: A sub-version tends to apply smaller scale changes to a given version. Some datasets in healthcare are released without labels and predefined tasks, or will be later labeled by researchers for specific tasks and problems, to form sub-versions of the dataset.

The following set of questions clarifies the information about the current (latest) version of the dataset. It is important to report the rationale for labeling the data in any of the versions and sub-versions that this datasheet addresses, funding resources, and motivations behind each released version of the dataset.

Does the dataset get released as static versions or is it dynamically updated?

- If static, how many versions of the dataset exist?
- If dynamic, how frequently is the dataset updated?

Static versions:

- MIMIC-IV contains data from 2008-2019. The data was collected from Metavision bedside monitors.
- MIMIC-III contains data from 2001-2012. The data was collected from Metavision and CareVue bedside monitors.
- MIMIC-II contains data from 2001-2008. The data was collected from CareVue bedside monitors. MIMIC-II is no longer publicly available but the data can still be obtained from MIMIC-III by only including data from the CareVue monitors.

Is this datasheet created for the original version of

the dataset? If not, which version of the dataset is this datasheet for?

This datasheet is created for version 3 (MIMIC-III). There is no healthsheet for the original dataset version.

Are there any datasheets created for any versions of this dataset?

No, this is the first datasheet being created for this dataset.

Does the current version/sub-version of the dataset come with predefined task(s), labels, and recommended data splits (e.g., for training, development/validation, testing)? If yes, please provide a high-level description of the introduced tasks, data splits, and labeling, and explain the rationale behind them. Please provide the related links and references. If not, is there any resource (website, portal, etc.) to keep track of all defined tasks and/or label definitions?

The original dataset doesn't come with predefined tasks, labels or recommended data splits. There is a community with ongoing contributions (pull requests) and most of the dataset labeling / task definitions fall under those categories. See <https://github.com/MIT-LCP/mimic-code/tree/main/mimic-iii/concepts>.

If the dataset has multiple versions, and this datasheet represents one of them, answer the following questions:

a. What are the characteristics that have been changed between different versions of the dataset?
b. Explain the motivation/rationale for creating the current version of the

dataset.

c. Does this version have more subjects/patients represented in the data, or fewer?

d. Does this version of the dataset have extended data or new data from the same patients as the older versions? Were any patients, data fields, or data points removed? If so, why?

e. Do we expect more versions of the dataset to be released?

f. Is this datasheet for a sub-version of the dataset? If yes, does this sub-version of the dataset introduce a new task, labeling, and/or recommended data splits? If the answer to any of these questions is yes, explain the rationale behind it.

g. Are you aware of any widespread sub-version(s) of the dataset? If yes, what is the addressed task, or application that is addressed?

a. New data is being added (collected between 2008-2012). In addition, many data elements have been regenerated from the raw data in a more robust manner to improve the quality of the underlying data. For a full list of changes, please refer to <https://mimic.mit.edu/iii/about/releasenotes/>

b. MIMIC-III is an extension of MIMIC-II: it incorporates the data contained in MIMIC-II (collected between 2001 - 2008) and augments it based on the changes described above (and in more detail in the release notes).

c. Yes, new patients have been added.

d. Yes, corrections to fields and new data has been included. See release notes for a full list of changes/additions.

e. MIMIC-IV has been released, which is an update to MIMIC-III. While we cannot say with certainty, we expect more versions to be released in the future.

f. It is not, while there are multiple sub-versions of MIMIC-III, we are including information for the overall version since this was a major release.

g. MIMIC extract is widely used by the community. More details can be found at: https://github.com/MLforHealth/MIMIC_Extract

Motivation

Reasons and motivations behind creating the dataset, including but not limited to funding interests.

For any of the following questions, if a healthsheet has already been created for this dataset, then refer to those answers when filling in the below information.

For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

The dataset was created for research and development in electronic health records.

What are the applications that the dataset is meant to address? (e.g., administrative applications, software applications, research)

Broadly healthcare research. Specific tasks were not set up as part of the dataset release.

Are there any types of usage or applications that are discouraged from using this dataset?

No commercialization. From the agreement: "The LICENSEE will use the data for the sole purpose of lawful use in scientific research and no other."

Who created this dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?

MIMIC was created through the work of researchers at the MIT Laboratory for Computational Physiology and their collaborators. See: <https://mimic.mit.edu/iii/about/acknowledgments/>.

Who funded the creation of the dataset? If there is an associated grant, please provide the name of the grantor and the grant name and number.

Dataset creation was supported by grants from the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health (NIH) under award numbers R01-EB001659 (2003-2013) and R01-EB017205 (2014-2018).

What is the distribution of backgrounds and experience/expertise of the dataset curators/generators?

N/A. This information could not be easily found.

Any other comments?

Data Composition

Instances: Refers to the unit of interest. The unit might be different in the healthsheet compared to the downstream use case: an instance might relate to a patient in the database, but will be used to provide predictions for specific events for that patient, treating each event as separate.

What do the instances that comprise the dataset represent (e.g., documents, images, people, countries)? Are there multiple types of instances? Please provide a description.

MIMIC is a relational database containing tables of data relating to patients who stayed within the intensive care units at Beth Israel Deaconess Medical Center. A table is a data storage structure which is similar to a spreadsheet: each column contains consistent information (e.g., patient identifiers), and each row contains an instantiation of that information (e.g., a row could contain the integer 340 in the patient identifier column which would imply that the row's patient identifier is 340). A list of all tables can be found here: <https://mimic.mit.edu/iii/mimictables/>

How many instances are there in total (of each type, if appropriate)? (breakdown based on schema, provide data stats)?

See <https://mit-lcp.github.io/mimic-schema-spy/>.

How many patients / subjects does this dataset represent? Answer this for both the preliminary dataset and the current version of the dataset.

There are 46520 patients in total in the MIMIC-III dataset.

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set? If the dataset is a sample, then what is the larger set? Is the sample representative of the larger

set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable). Answer this question for the preliminary version and the current version of the dataset in question.

MIMIC is a relational database containing tables of data relating to patients who stayed within the intensive care units (ICU) at Beth Israel Deaconess Medical Center between 2001 and 2012. The dataset is designed to be representative of electronic health records in ICUs but may not be representative of general electronic health records data such as in the non-ICU hospital wards.

What data modality does each patient data consist of? If the data is hierarchical, provide the modality details for all levels (e.g: text, image, physiological signal). Break down in all levels and specify the modalities and devices.

The dataset contains different types of clinical data such as:

- Time-stamped nurse-verified physiological measurements (for example, hourly documentation of heart rate, arterial blood pressure, or respiratory rate).
- Documented progress notes by care providers.
- Continuous intravenous drip medications and fluid balances.

What data does each instance consist of? "Raw" data (e.g., unprocessed

text or images) or features? In either case, please provide a description.

No images are linked in MIMIC-III. All data is in the form of structured data, where each field comes from an electronic health record, and therefore has been processed.

Is any information missing from individual instances? If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

The de-identification process for structured data required the removal of all eighteen of the identifying data elements listed in HIPAA, including fields such as patient name, telephone number, address and exact dates. Protected health information was removed from free text fields, such as diagnostic reports and physician notes.

There are other sources of missingness coming from the sparseness of the data, which is the nature of EHR. Not all lab values will be present at all times for a given patient for example.

Are relationships between individual instances made explicit (e.g., They are all part of the same clinical trial, or a patient has multiple hospital visits and each visit is one instance)? If so, please describe how these relationships are made explicit.

All of the subjects are patients who stayed within the intensive care units at Beth Israel Deaconess Medical Center between 2001 and 2012.

Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description. (e.g., losing data due to battery failure, or in survey data subjects skip the question, radiological sources of noise)

There are redundancies when it comes to laboratory values which are repeated in the CHARTEVENTS table. This occurs because it is desirable to display the laboratory values on the patient's electronic chart, and so the values are copied from the database storing laboratory values to the database storing the CHARTEVENTS.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources: a. Are there guarantees that they will exist, and remain constant, over time

b. Are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created)

c. Are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

The dataset is self-contained, and the only information gathered from external resources is the date of death. This is acquired from the social security death registry.

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals non-public communications)? If so, please provide a description.

N/A. We believe it does not, but could not guarantee.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.

We do not believe so.

If the dataset has been de-identified, were any measures taken to avoid the re-identification of individuals? Examples of such measures: removing patients with rare pathologies or shifting time stamps.

The dataset has been de-identified and dates/times have been shifted.

Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)? If so, please provide a description.

It contains health data, but in a non-identifiable way. For each patient, the following fields are specified:

- insurance
- language

- religion
- marital status
- ethnicity
- gender (genotypical sex)

Devices and contextual attributes in data collection

For data that requires a device or equipment for collection or the context of the experiment matters on the outcome, provide details, answer the following additional questions or provide relevant information based on the device or context that is used (for example)

- If there was an MRI machine used, what is the MRI machine and model used?
- If heart rate was measured what is the device for heart rate variation that is used?
- If cortisol measurement is reported at multi site, provide details.
- If smartphones were used to collect the data, provide the names of models.
- Anything else?

N/A. We could not find information related to this.

Challenge in tests and confounding factors

Which factors in the data might limit the generalization of potentially derived models? Is this information available as auxiliary labels for challenge tests? For instance:

- Number and diversity of devices included in the dataset.
- Data recording specificities, e.g., the view for a chest x-ray image.
- Number and diversity of recording sites included in the dataset.
- Distribution shifts over time.

Beth Israel Deaconess Medical Center was the only site from which data was collected. It is based in Boston, MA, USA.

For distribution shifts, in 2008 the EHR system changed from CareVue to MetaVision. Data which could not be merged is given a suffix to denote the data source. There are also smaller shifts in non-transition years as the patient distribution is non-stationary.

For all other questions we could not find the information.

What confounding factors might be present in the data?

a. Interactions between demographic or historically marginalized groups and data recordings, e.g., were women patients recorded in one site, and men in another?

b. Interactions between the labels and data recordings, e.g., were healthy patients recorded on one device and diseased patients on another?

a. As there is a single recording site, all groups have had data recorded from the same place. However, we are aware of work that looks into site-wise demographic differences.

b. N/A, as labels aren't associated with this version.

Any other comments?

categories and the source of the classifications used (e.g: self-report or selection, investigator observed, database, electronic health record, survey instrument).

c. If patients' demographic data is included, are patients aware / did they consent to the collection and use of their demographic information?

d. In some cases, there have been biologically proven associations between demographics and the outcome. Are you aware of similar associations in the tasks covered by this dataset? Should users be wary of specific proxies or associations when using the dataset? If yes, please provide a link to the study, or publication.

e. Is there any mechanism for updating some of this demographic information after its initial collection? For example, if someone wants to change their gender information, what are the mechanisms to do so?

f. Provide a description of the respective distributions of each subgroup population within the dataset.

Most information could not be easily found in official sources.

For f. Subgroup information:

- Female: 44.1%, Male 55.9%
- Asian 2.3%, Black 9.4%, Hispanic 3.27%, White 70.89%, Other 14.04%

Does the dataset identify any demographic sub-populations (e.g., by age, gender, sex, ethnicity)? If yes:

a. The reasons that these categories were assessed also should be described in the datasheet.

b. How was this information acquired? Please describe who identified these

If no, is there any regulation that prevents demographic data collection in your study (for example, the country that the data is collected in)?

N/A, as the dataset contains demographic information.

Collection and use of demographic information

Pre-processing / de-identification

Was there any pre-processing for the de-identification of the patients? Provide the answer for the preliminary and the current version of the dataset.

The data was de-identified in accordance with Health Insurance Portability and Accountability Act (HIPAA) standards using structured data cleansing and date shifting.

Was there any pre-processing for cleaning the data? Provide the answer for the preliminary and the current version of the dataset

Various fields have been cleaned through harmonization or de-duplication. The following release notes provide ample information as to what was changed and how: <https://mimic.mit.edu/docs/iii/about/releasenotes/>. Dates have also been date shifted as part of the deidentification process.

Was the “raw” data (post de-identification) saved in addition to the pre-processed/cleaned data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the “raw” data.

N/A. This information could not be found.

Were instances excluded from the dataset at the time of preprocessing? If so, why? For example, instances related to patients under 18 might be discarded.

In the MIMIC-IV cohort, patients who underwent extubation during ICU stays were included. The exclusion criteria were as follows: (i) age <18

years, (ii) unplanned extubation, (iii) not the first extubation during the hospital stay, or (iv) no MV records before extubation.

Reference: <https://www.frontiersin.org/articles/10.3389/fmed.2021.676343/full>

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)? Answer this question for both the preliminary dataset and the current version of the dataset.

The data is sliced based on time and collected from medical records information systems. No specific sampling has been used.

Any other comments?

Labeling and subjectivity of labeling

Is there an explicit label or target associated with each data instance?

Please respond for both the preliminary dataset and the current version.

a. If yes:

i) What are the labels provided?

ii) Who performed the labeling? For example, was the labeling done by a clinician, ML researcher, university or hospital?

b. What labeling strategy was used?

i) Gold standard label available in the data (e.g., cancers validated by biopsies)

ii) Proxy label computed from available data:

1. Which label definition was used? (e.g., Acute Kidney Injury has multiple definitions)

2. Which tables and features were considered to

compute the label?
 iii) Which proportion of the data has gold standard labels?
 c. Human-labeled data
 i) How many labellers were considered?
 ii) What is the demographic of the labellers? (countries of residence, of origin, number of years of experience, age, gender, race, ethnicity, ...)
 iii) What guidelines did they follow?
 iv) How many labellers provide a label per instance? If multiple labellers per instance:
 1. What is the rater agreement? How was disagreement handled?
 2. Are all labels provided, or summaries (e.g., maximum vote)?
 v) Is there any subjective source of information that may lead to inconsistencies in the responses? (e.g: multiple people answering a survey having different interpretation of scales, multiple clinicians using scores, or notes)
 vi) On average, how much time was required to annotate each instance?
 vii) Were the raters compensated for their time? If so, by whom and what amount? What was the compensation strategy (e.g., fixed number of cases, compensated per hour, per cases per hour)?

What are the human level performances in the applications that the dataset is supposed to address?

Is the software used to preprocess/clean/label the instances available?

If so, please provide a link or other access point.

Is there any guideline that the future researchers are recommended to follow when creating new labels / defining new tasks?

Are there recommended data splits (e.g., training, development/validation, testing)? Are there units of data to consider, whatever the task? If so, please provide a description of these splits, explaining the rationale behind them. Please provide the answer for both the preliminary dataset and the current version or any sub-version that is widely used.

Any other comments?

No questions were answered in this section, as this version does not come with labels.

Collection Process

Were any REB/IRB approval (e.g., by an institutional review board or research ethics board) received? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

The project was approved by the Institutional Review Boards of Beth Israel Deaconess Medical Center (Boston, MA) and the Massachusetts Institute of Technology (Cambridge, MA).

How was the data associated with each instance acquired? Was the data directly observable (e.g., medical images,

labs or vitals), reported by subjects (e.g., survey responses, pain levels, itching/burning sensations), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

Data was collected in-hospital by clinical staff, collected in the critical care unit or from the hospital record system. For example, for labs, a member of the clinical staff acquires a fluid from a site in the patient's body, labels it and sends it for processing.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)? How were these mechanisms or procedures validated? Provide the answer for all modalities and collected data. Has this information been changed through the process? If so, explain why.

Two different critical care information systems were used for collecting the dataset: Philips CareVue Clinical Information System (models M2331A and M1215A; Philips Health-care, Andover, MA) and iMDsoft MetaVision ICU (iMDsoft, Needham, MA).

Who was involved in the data collection process (e.g., patients, clinicians, doctors, ML researchers, hospital staff, vendors, etc) and how were they compensated (e.g., how much were contributors paid)?

We could only find high-level information regarding the groups involved.

MIMIC is made available largely through the work of researchers at the MIT Laboratory for Computational Physiology and the following research groups:

- Beth Israel Deaconess Medical Center
- MIT Clinical Decision Making
- MIT Computational Physiology and Clinical Inference Group
- MIT Lab for Computational Physiology
- Philips Health Care

Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.

The data was collected between 2001 and 2012, however the dates in the dataset have been time shifted in order to help with de-identification.

Does the dataset relate to people? If not, you may skip the remaining questions in this section.

Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., hospitals, app company)?

Third parties or other sources. MIMIC-III consists of data collected from two different clinical information systems, CareVue and MetaVision.

Were the individuals in question notified about the data collection? If

so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.

N/A. This information could not be found.

Did the individuals in question consent to the collection and use of their data? If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.

Requirement for individual patient consent was waived because the project did not impact clinical care and all protected health information was de-identified.

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses?

If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate).

N/A, as consent was not obtained.

In which countries was the data collected?

The dataset was collected in Boston, Massachusetts, United States of America.

Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted? If

so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

N/A. This information could not be found.

Inclusion Criteria- Accessibility in Data Collection

Is there any language-based communication with patients? If yes, describe the choices of language(s) for communication. (for example, if there is an app used for communication, what are the language options?)

To the best of our knowledge, verbal communication was used with the patients. We could not find information on the various language accommodations that were made.

What are the accessibility measurements and what aspects were considered when the study was designed and implemented?

N/A. We could not find this information in official sources.

If data is part of a clinical study, what are the inclusion criteria?

N/A. We could not find this information in official sources.

Any other comments?

Uses

Has the dataset been used for any tasks already? If so, please provide a description.

Yes. A common example includes “Length of stay in the ICU”.

Does using the dataset require the citation of the paper or any other forms of acknowledgement? If yes, is it easily accessible

through google scholar or other repositories.

es, citations are required and the citation format depends on the version used. The following page provides the citation format: <https://mimic.mit.edu/docs/about/acknowledgments/> depending on the version.

The paper itself is also linked in the documentation website and easily accessible through Google Scholar or online sources.

Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.

There is no official repository available, but one could find papers through looking up the citation (e.g., <https://read.qxmd.com/keyword/229497>)

Is there anything about the composition of the dataset or the way it was collected and pre-processed/cleaned/labeled that might impact future uses? For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

We do not have enough information to accurately answer this question. However the methods through which various fields were collected could lead to a bias in results, especially if there is a difference in care for an individual or a group.

Are there tasks for which the dataset should not be used? If so, please provide a description. (for example, dataset creators could recommend against using the dataset for considering immigration cases, as part of insurance policies)

There is no official banning of any specific task, however research that could lead to the enforcement of any rule or law should be validated through other means.

Any other comments?

Dataset Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.

Yes, anyone can access the dataset provided they meet the following criteria:

- Become a credentialed user on PhysioNet. This involves completion of a training course in human subjects research.
- Sign the data use agreement (DUA). Adherence to the terms of the DUA is paramount.
- Follow the tutorials for direct cloud access (recommended), or download the data locally.

How will the dataset be distributed (e.g., tarball on website, API, GitHub) Does the dataset have a digital object identifier (DOI)?

The data can be accessed on Cloud through either BigQuery or AWS; Alternatively it can be downloaded locally from PhysioNet.

When will the dataset be distributed?

The dataset was released on the 2nd of September 2016. Anyone can download it after they are approved for access.

Assuming the dataset is available, will it be/is the dataset distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

Yes, there is a user agreement that a researcher must agree to before getting access. The terms and conditions can be found in the PhysioNet account required during setup, and need to be accepted for each version of the dataset that you intend to use.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

Any results obtained from the work on the dataset should be open sourced to benefit the community. There are no restrictions concerning who owns the IP.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions,

and provide a link or other access point to, or otherwise reproduce, any supporting documentation.

N/A. The information required to answer this question could not be easily found.

Any other comments?

Maintenance

Who will be supporting/hosting/maintaining the dataset?

MIT Laboratory for Computational Physiology and their collaborators (<https://mimic.mit.edu/docs/about/acknowledgments/>).

How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

It is generally recommended to raise any issues with MIMIC Code Repository GitHub. If there are issues related to sensitive information then one can contact phi-report@physionet.org (for PHI) and mimic-support@physionet.org (for general private issues)

Is there an erratum? If so, please provide a link or other access point.

No erratum as of now, but release notes with corrections are published periodically (<https://mimic.mit.edu/docs/iii/about/releasenotes/>).

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

Yes. The maintainers of the dataset provide regular updates. These are communicated through the release notes page.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.

N/A. We could not find information about this topic.

Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

As of the time of writing older versions are still supported.

We could not find information related to what happens when a version is turned down, or when this would take place.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

Yes, community contributions are welcomed and are typically performed through GitHub. More information can be found on <https://mimic.mit.edu/docs/community/contributing/>.

Any other comments?

Participants Expertise	Number of participants
Legal and Regulatory	1
Clinical	7
ML for Health Research	6
Health Data Infrastructure	2
Product	4
Bioethics	1
Equity	2
Privacy	2
Sociotechnical Research	1
Total number of participants	21

Table 1: Expertise of the interview participants. There were in total 21 participants. Some participants expertise are beyond one category and is on the intersection of 2 or more expertise.

APPENDIX B: SEMI-STRUCTURED EXPERT INTERVIEW

We recruited experts with wide range of expertise using health data in relation to their job. Expertise are listed in Table 1. During the interview, study objectives and research questions were presented. Interviews time varied between 30 and 180 minutes over one to two sessions, depending on the experts' availability and engagement. In certain cases, interviewees chose to offer additional comments after the interviews.

We first asked questions pertaining to the participants expertise in health data and professional background. We then conducted a semi-structured interview using the protocol described below. Finally, participants and interviewers together studied the *Primary Healthsheet*, discussing⁵ specific questions and potential opportunities for improvement.

The following questions were asked from the expert interview participants, and then followed by a deep dive in questionnaire:

- Do you use healthcare data in relation to your job? If yes, which kinds of healthcare applications have you worked with?
- What are the challenges you have faced when using or choosing a healthcare dataset?
- If you have access to the healthsheets of all datasets, would you use it to decide which dataset to work with and if yes, how do you use Healthsheets to help you decide if you want to use a dataset or not?
- Among the given categories/ questions in the Primary Healthsheet, are there any categories or questions that you would prioritize having access to?
- Which questions are the most important questions for you and your assessment of a dataset?
- If I am a curator of a dataset, what should be my incentive for creating something like this? Is it worth it to spend time and effort on filling this long questionnaire?
- It's a very large questioner, how can we make healthsheet more personalized for your role?

⁵The interview transcripts and notes are over 300 pages, providing comprehensive insights into a variety of aspects to consider when evolving future versions of the Healthsheet.