

Datasheets for Healthcare AI: A Framework for Transparency and Bias Mitigation

Marjia Siddik, Harshvardhan J. Pandit

ADAPT Centre, School of Computing, Dublin City University, Dublin, Ireland

Abstract

The use of AI in healthcare has the potential to improve patient care, optimize clinical workflows, and enhance decision-making. However, bias, data incompleteness, and inaccuracies in training datasets can lead to unfair outcomes and amplify existing disparities. This research investigates the current state of dataset documentation practices, focusing on their ability to address these challenges and support ethical AI development. We identify shortcomings in existing documentation methods, which limit the recognition and mitigation of bias, incompleteness, and other issues in datasets. We propose the 'Healthcare AI Datasheet' to address these gaps, a dataset documentation framework that promotes transparency and ensures alignment with regulatory requirements. Additionally, we demonstrate how it can be expressed in a machine-readable format, facilitating its integration with datasets and enabling automated risk assessments. The findings emphasise the importance of dataset documentation in fostering responsible AI development.

Keywords

AI in Healthcare, Dataset Documentation, Data Transparency, Bias Mitigation, Ethical AI, GDPR, EU AI Act, Risk Assessment

1. Introduction

AI has the potential to enhance diagnostics, treatment planning, patient monitoring, and overall care delivery [1]. By utilising medical records and other data collected over time, AI can make healthcare more efficient, personalized, and accessible [2]. However, its deployment also introduces ethical and legal challenges, particularly due to the use of sensitive data and the risk of harms. Prominent amongst known issues are biases [3] which exacerbate existing health disparities and create unequal treatment outcomes. Such biases can arise from training datasets, algorithmic practices, and incorrect deployments, and act to compromise the effectiveness, thereby threatening patient safety and undermining trust in healthcare systems [4]. As a result, the use of AI poses challenges to the fairness and integrity of healthcare delivery, particularly when it relies on unrepresentative datasets [4], flawed data collection [5], and reflects existing societal prejudices [6]. These issues can lead to AI models that perform poorly for certain demographic groups, amplifying health disparities and compromising patient care [7].

Beyond bias, other concerns include data incompleteness [8], inaccuracies [9], and outdated information [10], all of which can undermine the effectiveness and reliability of AI systems. To

AICS'24: 32nd Irish Conference on Artificial Intelligence and Cognitive Science, December 09–10, 2024, Dublin, Ireland

✉ marjia.siddik3@mail.dcu.ie (M. Siddik); me@harshp.com (H. J. Pandit)

🆔 0009-0008-6295-6979 (M. Siddik); 0000-0002-5068-3714 (H. J. Pandit)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

assess whether such issues exist in the development and use of AI, it is essential to have comprehensive documentation regarding the origins, composition, limitations, and other contexts for how the AI system was developed - in particular the data used to train it [11]. Without this information, AI systems cannot be reliably assessed for suitability of use, and can inadvertently cause harm or fail to deliver effective care [12].

With data and AI enriched healthcare poised to significantly progress in the next decade based on legal advancements such as the EU Health Data Space regulation, it is important for uses of data and AI in healthcare to adhere to existing regulatory frameworks such as General Data Protection Regulation (GDPR) [13] and the Artificial Intelligence Act (AI Act) [14]. Which means that dataset documentation practices should also incorporate information to support compliance with GDPR and AI Act so that issues such as bias and potential harms can be evaluated and enforced through legal mechanisms. At the same time, healthcare is highly dependant on local contexts, where different regions and countries have differing frameworks and legislations for how the data gets generated and utilised in health research. Ireland recently published its Health Information Bill [15] in 2023 which enables the reuse of data for healthcare research. In such cases, it is also vital to evaluate whether dataset documentation practices are sufficient to support such initiatives.

This study therefore investigates the question: *"How can dataset documentation support mitigation of bias and promote ethical AI in healthcare systems?"* and explores the answer through the following objectives:

- RO1** Identify categories of bias which should be documented (Section 2.1).
- RO2** Identify legal considerations which should be documented for datasets (Section 2.2).
- RO3** Evaluate existing dataset documentation practices regarding representation of identified bias categories and legal considerations (Section 2.3).
- RO4** Develop a machine-readable dataset documentation method that incorporates identified requirements and fills in gaps in current practices (Section 3).
- RO5** Discuss how the solution will work within the Irish healthcare context (Section 4).

2. Literature Review

2.1. Categorisation of Bias

Bias in AI refers to the tendency of AI systems to produce results that reflect societal inequities, which is especially concerning in healthcare, where biased AI tools can have serious consequences. AI models trained on non-diverse datasets often fail to generalize across populations, potentially worsening healthcare disparities. For example, Celi et al. [16] observed that many AI datasets come from high-income countries like the US and China, reducing their relevance in low- and middle-income countries with distinct healthcare challenges. Despite the importance of mitigating these biases, existing dataset documentation practices often overlook these disparities, revealing major shortcomings in current frameworks [17].

Several types of bias impact healthcare AI, leading to inequitable outcomes. *Sample bias* occurs when training data does not adequately represent the target population, resulting in skewed predictions. *Annotator bias* emerges when individuals labeling the data introduce their prejudices, further distorting AI outputs. *Temporal bias* arises from changes in data patterns over

time, impacting AI model relevance [18]. For example, *gender bias* in diagnostic AI, such as chest X-ray interpretation, has shown to skew results based on biological differences that arise over time [19]. Unfortunately, current documentation rarely addresses these biases comprehensively, pointing to the need for improved strategies.

Biases such as *data-driven* and *algorithmic bias* also contribute to healthcare inequality. Data-driven bias occurs when certain demographics are over-represented, while algorithmic bias reinforces patterns favoring majority groups [19]. *Human bias* introduced by researchers or clinicians adds further complexity to fairness [19]. If dataset documentation practices do not distinguish between these different kinds of bias, or only report on specific ones (e.g. gender bias) - then it risks creating a false sense of security by assuming biases have been identified and addressed. Further, by not incorporating the information required to identify and address such biases, the dataset documentation also limits its usefulness.

Based on these, we identify a gap in current practices that needs to be address by having dataset documentation *distinguish between the different categories of bias and should record that aids in identifying and addressing them*. This addresses **RO1**.

2.2. Legal Frameworks Governing Bias

The development and deployment of data-based AI systems in healthcare requires strict adherence to laws to ensure fairness, accountability, and transparency. Laws such as GDPR and AI Act require conducting impact assessments to protect human rights and prevent harms based on a risk-based approach where certain data and technologies are considered as *high-risk* based on their sensitivity and potential risks. If dataset documentation approaches do not incorporate such requirements, or do not provide sufficient information to support implementing them, it leads to legal uncertainties, risks, and makes assessing *liability* difficult - which is a vital incentive to ensure safety and security in technology.

The GDPR, which regulates processing of personal data, establishes specific categories, which includes health, of data as being *special* (Article 9) - meaning they are more sensitive and merit a higher degree of consideration in risk management (Article 32) and impact assessments (Article 35). GDPR also establishes accountability based on the role of 'Controller' where an entity determines the 'means and purposes' of processing data, where processing covers any collection, storage, use, sharing, and erasure of personal data. Further, the GDPR also establishes rights (Articles 12-23) associated with data - such as the requirement to provide notices, ability to opt-out of automated decision making, right to be forgotten, rectification, and erasure. When datasets constitute personal data (as defined by GDPR Article 4), their collection and use, as well as potentially the AI systems developed using them are likely to be subject to the GDPR. Without sufficient information, users of data miss out on the *safety net* provided by the GDPR in terms of safety and security obligations when reusing data, and end up creating complications and potential liabilities for themselves as they do not have documented evidence of the dataset's quality and GDPR compliance.

The AI Act, a recent development, establishes risk levels for use of AI, and has obligations regarding transparency, human oversight, and data management. Similar to the GDPR, documented evidence is vital for obligations under the AI Act to assess and demonstrate that data, or AI developed using data, is compliant with safety and reliability standards. More specific

to Ireland, the Health Information Bill (2023) establishes the creation and sharing of digital health records and provides a framework for the reuse data for scientific research and public benefit. In this, it requires specific documented assessments of data and uses of AI similar to the obligations of GDPR to ensure appropriate practices regarding security and safety.

From this, we establish that dataset documentation practices should incorporate *information to support legal obligations regarding transparency and accountability - most specifically the provenance and legal categorisation of data, risk and impact assessments, and involvement of entities in specific legal roles* This addresses **RO2**.

2.3. Current Dataset Documentation Practices

Comprehensive documentation practices in the machine learning community often receive limited attention beyond immediate technical information, and no standardised processes exist to ensure transparency, accountability, reproducibility, interoperability, and quality - especially in contexts such as healthcare where non-reporting of issues such as bias can cause harms [11]. To address this, several proposals have been published, of which we focus on notable ones that are widely known or are in the scope of our work.

'Datasheets for Datasets' [11] is a seminal work that defines information requirements to record motivation, composition, and usage aspects to enhance transparency and reproducibility. While it acknowledges regulations such as GDPR and issues such as bias, it does not provide for recording specific risks such as different categories of biases and does not align its information with regulatory requirements. 'Dataset Nutrition Label' [20], developed by the Data Nutrition Project, aims to "enhance context, contents, and legibility" by "providing at-a-glance information". It contains information required for bias identification but does not address measures for mitigation or specifying further risks or regulatory requirements.

'Open Datasheets' [21, 22] provide a machine-readable format designed to improve dataset discoverability and usability, but does not expand upon risk assessment and regulatory information. 'Data Statements for NLP' [23] enables recording information with the goal of supporting bias mitigations, but does not account for different risks or regulations. Tools such as DataDoc Analyzer [24] and MetaReader [25] support ensuring completeness and bias identification based on existing approaches, but do not explore expanding their information requirements.

These existing approaches¹ show a necessity to document information regarding datasets so as to inform and support the 'data value chain' in addressing risks - such as biases - and avoiding harms. However, they have limitations in terms of acknowledging different risks beyond a few bias categories (e.g. gender or sex), do not support systematic risk assessments, and more critically are not aligned with regulatory requirements - which makes their enforcement and use in accountability difficult. In the context of healthcare, we could not find any specific approach which adapts or explores the specific collection and (re-)use of data in a clinical or medical research context. Further, healthcare settings typically have additional policies and guidelines established within the institution, consortium, or as sectorial regulations - which can only be supported in dataset documentation practices if they are extensible. Additionally, the information required to be documented can come from different entities - including inter-

¹Due to spacial limitations, an overview of existing approaches is provided later as part of our proposed approach.

organisational units - which necessitates standardisation and interoperability to ensure its effectiveness. We could not find any approaches which tackled these aspects.

From this analysis, we determined that the use of data for AI in healthcare settings requires a solution that addresses existing gaps regarding expanded bias categorisations, risk assessments, and compliance with regulations. This addresses **RO3**.

3. Developing an Improved Machine-Readable Datasheet

Based on the analysis of the state of the art, we identified the need to create an improved dataset documentation approach that supports information requirements regarding bias categorisation and risk assessment, and is aligned with the GDPR (**RO4**). We also identified the need for structured machine-readable representations to support maintaining and providing datasheets alongside the data. For this, we selected the ‘Datasheets for Datasets’ [11] approach as a baseline given its existing prevalence and impact, and extended it to incorporate our additional requirements. Our proposed datasheet is **available online**² with an example schema.

3.1. Methodology

We first identified and analysed the information that could be recorded from using existing approaches for datasheet documentation and found three gaps (bias categories, risk assessment, regulations), for which we then developed specific requirements to document information. Through an iterative process, we developed a structured datasheet by starting with 18 identified information fields from the ‘Datasheets for Datasets’ [11] approach, and extended it to over 50 fields in the final iterations. The additional fields were developed based on requirements to document information associated with identified categories of bias - such as temporal characteristics and demographics, risk assessment information - such as provenance of data, completeness, existing or potential measures, and regulatory information - such as applicable laws and impact assessments. In addition to this, we also made explicit the information fields associated with purposes for which the dataset was created, and its intended uses, and usage restrictions - which can aid the process of determining suitable data reuses and avoid misuses. For existing fields, we focused on specificity where vague fields were refined for clarity (e.g. usage restrictions and data characteristics).

In order to evaluate the effectiveness of developed information requirements, we sought to identify existing datasets on popular platforms such as Kaggle and Hugging Face whose documentation contained this information. We could not identify a suitable dataset as most datasets did not contain even the preliminary information required by existing dataset documentation practices, and going through their associated publications and reports would have required exorbitant amounts of time³. Therefore, we undertook manual exercises to create documentation for hypothetical datasets based on different scenarios such that all information fields would be populated. Through this process, we identified several refinements based on

²<https://github.com/marjiasdk/Healthcare-AI-Datasheet>

³The lack of findability mechanisms based on machine-readable data also contributed to these difficulties.

ambiguity in information (e.g. date format), necessity to provide a controlled vocabulary (e.g. to express likelihood), and additional fields (e.g. usage prohibitions derived from risk assessments).

We then developed a JSON based structure to represent the datasheet in a machine-readable format. We chose JSON as it is a popular data format that is natively supported in all major programming languages, is easily communicated on the web, and enables a structured schemas that can be validated for completeness, correctness, and compliance. We also chose JSON as it is easy to learn and iterate prototypes for a developing schema. For future interoperability and standardisation, we recommend using existing standards such as DCAT⁴ with ODRL⁵ for expressing usage policies and DPV [26] to represent regulatory information.

3.2. Description of Information Fields

There are 55 information fields broadly categorised in 10 sections as follows.

Metadata These fields contain information describing the dataset in terms of its title, version, publisher, and license.

Purpose These fields describe the purposes for which the dataset was created (e.g. clinical research), what was the intended benefit (e.g. improve diagnostic accuracy), and the intended beneficiaries (e.g. healthcare providers, patients).

Source Information These fields describe the source and origin of data, and whether the collection process had (ethical) approval (e.g. from organisation) and its funding sources.

Temporal Information These fields describe temporal aspects of the data in the dataset, such as which period it covers (e.g. 2019-2023) and last updated (e.g. December 2023).

Demographic Information These fields describe demographic information for individuals whose data is present, such as age and age ranges (e.g. 18-65), gender, and ethnicity, as well as fields to indicate the *likelihood* of specific kinds of bias due to the demographic distributions.

Data Characteristics These fields describe the media type for data (e.g. images), and also indicate whether the data is *incomplete* along with the missing elements and reasons.

Bias Mitigation Methods These fields describe bias mitigation methods that have already been applied as well as suggested measures to adopters.

Personal Data These fields describe whether the data constitutes as *personal* (e.g. non-anonymized patient records), specific categories of personal data (e.g. name, age), its sensitivity (e.g. low), and for (partially-)anonymised data - which anonymisation techniques were used and its risk of reidentification.

Risk and Compliance These fields provide a way to indicate the risk levels (separately for generic and legal), jurisdiction and applicable laws (e.g. EU and GDPR), existence of impact assessments (e.g. a GDPR DPIA), and suggested mitigation measures (e.g. auditing security risks prior to data reuse).

Usage Restriction These fields define limitations and constraints on the (re-)use of the dataset, such as through access restrictions (e.g. only use within organisation), specific permissions (e.g. only used for cancer research) or prohibitions (e.g. no third party sharing), and obligations (e.g. reciprocity to share results back with data provider).

⁴<https://www.w3.org/TR/vocab-dcat/>

⁵<https://www.w3.org/TR/odrl-model/>

3.3. Comparison with Existing Approaches

Table 1 compares our developed Datasheet approach with existing established approaches from Section 2.3 - namely the Datasheets for Datasets [11], Dataset Nutrition Labels [20], and Data Statements for NLP [23]. Due to spatial limitations of this article, we only provide a summary overview of this comparison, with the full analysis available online⁶. In the table, the rows represent the information in sections (as described in Section 3.2, and the values represent whether the information is present (●), has some fields missing (○), or is not present (×). The last two rows consider whether the approach requires structured information (e.g. a consistent vocabulary) and is interoperable by being machine-readable (e.g. using JSON).

Table 1

Comparative analysis of our proposed Datasheet with existing documentation approaches

Category	This Approach	Datasheets for Datasets [11]	Dataset Nutrition Label [20]	Data Statements for NLP [23]
Metadata	●	●	●	●
Purpose	●	●	●	●
Source Information	●	●	●	○
Temporal Information	●	○	○	×
Demographics	●	×	○	●
Data Characteristics	●	○	○	○
Bias Mitigations	●	×	×	×
Personal Data	●	○	○	○
Risk and Compliance	●	○	○	×
Usage Restriction	●	○	●	×
Machine-readable	●	×	●	×
Interoperability	○	×	○	×

From this comparison, we see how the proposed advances the state of the art by highlighting important gaps in current approaches and providing our solution as the path forward. Most prominently, we address the crucial issue of missing information in dataset documentation practices regarding data characteristics and temporal information which is necessary to identify and mitigate commonly found biases. We also addressed risk and (legal) compliance more thoroughly which enables existing legal mechanisms and obligations to be used to support and enforce accountability and prevent harms that may arise from data collection and (re-)use. Finally, we also address the lack of providing documentation as structured information and ensuring it is interoperable and machine-readable. In this, our approach does not provide the best possible solution as it does not propose a standardised or standards-based representation of information - though it does show why this is required and how it can be achieved (e.g. using semantic web standards of DCAT, ODRL, and DPV as mentioned in Section 3.2) which are promising areas for future work.

Though not visible from the overview table, the legal and ethical risks documented in the extended datasheet set it apart from existing approaches in a crucial and important manner. For example, our approach explicitly considers potential legal risks for re-identification risks and data sensitivity which are essential considerations when dealing with sensitive healthcare

⁶<https://github.com/marjiasdk/Healthcare-AI-Datasheet/blob/main/comparison-sota.csv>

data, especially under regulations like GDPR. These risks are often only briefly touched upon in existing frameworks but are critical to be documented as having been considered for ensuring legal and policy compliance. In healthcare contexts, the sensitivity of patient data combined with the severe consequences of non-compliance can result in harm to patients, privacy violations, and legal repercussions. Therefore, our datasheets specifically support the higher level of scrutiny required in health and biomedical research to ensure patient safety and regulatory compliance by providing information fields that must be documented alongside the dataset.

Another important distinguishing factor is the demographic information section which goes beyond simple demographic documentation by requiring both information of distributions reflected in the dataset and also how such distributions may introduce bias into AI models. We feel this is a more comprehensive and in-depth approach that is also pragmatic for the healthcare context as compared to frameworks like Datasheets for Datasets which do not assess potential demographic bias in this way. In our datasheets, as the fields for distributions are explicit, the onus of identifying potential risks starts from the creation of the dataset, and any missing value (e.g. gender distribution unknown) becomes a risk in itself. To address this, we also provide a mitigations field which can enable a data provider to recommend that before using the dataset, an assessment of the bias should be carried out to eliminate or reduce potential issues.

4. Application in Irish Healthcare Context

The Irish healthcare system has been slow to adopt universal healthcare and has struggled with inconsistent data management practices across hospitals and clinics, influenced by historical resistance from the Catholic Church and private medical practitioners [27]. This has led to fragmented reforms and slow legislative and infrastructural progress, creating challenges in healthcare data infrastructure, and impacting the implementation of effective AI systems. Gaps in leadership and strategy in health information management have also slowed the development of a connected system, hindering the fairness and effectiveness of AI systems [28].

Despite initiatives like IHIs and ePrescribing, investment in health information systems and ICT remains low in Ireland [29]. Efforts such as the DQI framework seek to improve data quality, ensuring it is complete and reliable for creating fair AI systems [30]. While improving infrastructure is pivotal for reliable healthcare data, unaddressed biases could lead to legal and ethical issues, such as discrimination or violations of regulations like the GDPR. The two-tier healthcare model further complicates universal data-sharing standards as each organisation relies on its own methods and data management practices.

While measures like the Health Identifiers Act 2014 and the establishment of Health Information and Quality Authority (HIQA) aimed to improve health data governance, the absence of well-defined and regulated dataset documentation practices continues to be a challenge. Areas such as data governance, legal uncertainties, and the development of central platforms to support health data research remain underexplored, making it difficult to create a unified system for equitable healthcare access [29]. The lack of centralised policies that mandate uniform data practices hinders the development of unified datasets, which are critical for the implementation of fair and effective healthcare AI systems.

The Health Information Bill [15] (HIB) aims to resolve some of these challenges by creating a

legal obligation for the state to set up specific healthcare research infrastructures and facilitate the reuse of data for research. It is expected to tie in to the European Health Data Spaces regulation which is currently under the legislative process and is expected to be finalised in 2025. More prominently, the HIB addresses the current healthcare system's lack of coordination and disconnect between data management processes among organisations. While not explicitly addressing AI, the HIB does provide a legal framework for the reuse of healthcare data for research purposes and in conjunction with the recently published AI Act will guide the use of AI in healthcare for the near future.

Our developed datasheet provides a necessary and timely approach to address both the technical and legal challenges by establishing a structured and uniform approach to dataset documentation that is aligned with legal requirements from GDPR and which will support the requirements of Irish and European regulations regarding risk and impact assessments. Through this work, we have shown that the current prevalent dataset documentation practices are not sufficient for the current legal landscape in Ireland and also do not support the practicalities of intra-organisational interactions which require documented information to avoid uncertainties and liabilities.

By promoting transparent, accountable, and bias-conscious dataset documentation, our datasheet can serve as an important tool in helping Ireland's healthcare sector prepare for these regulations and cultivate a practice of risk assessment across the data and AI lifecycles. Additionally, it supports the existing documentation requirements in current and upcoming regulations as well as policy frameworks such as HIQA's guidelines. The possibility of representing the information in machine-readable form also opens up opportunities to automate the processes associated with creating the datasheets up to date with changes in AI systems, and to automate risk assessments by using datasheets as inputs that are passed along to stakeholders downstream within the AI value chain. Therefore, we recommend further developing and requiring the use of datasheet documentation practices along with supporting infrastructure and policies based on our work to enable and promote the ethical and legal reuse of data using AI across the Irish healthcare system. This completes our **(RO5)**.

5. Conclusion & Future Work

Our research highlights the importance of comprehensive dataset documentation in mitigating biases and promoting ethical AI in healthcare. Existing frameworks, such as Datasheets for Datasets and Dataset Nutrition Labels, often lack a specific focus on bias mitigation, particularly in the healthcare context. To address these gaps, we developed the Healthcare AI Datasheet, which incorporates detailed demographic information, data collection methods, and explicit bias mitigation strategies. This approach enhances transparency, accountability, and fairness in AI development, ensuring that systems reflect diverse populations and contribute to reducing healthcare disparities. Additionally, the machine-readable version of the datasheet facilitates integration into AI workflows, promoting responsible and ethical practices. By thoroughly documenting dataset characteristics and potential biases, healthcare providers can make more informed decisions when deploying AI systems, ensuring more equitable care. This is especially vital in preventing biased datasets from leading to misdiagnoses or unequal treatment outcomes.

While the Healthcare AI Datasheet represents notable progress, it has limitations. Its primary focus on healthcare datasets may restrict its broader applicability to other domains. Furthermore, its effectiveness relies on the accuracy and completeness of the information provided by dataset creators, which can introduce variability. Future research should prioritize real-world testing across healthcare settings to validate its effectiveness and refine its components. Implementing the datasheet in ongoing AI projects will help assess its ability to mitigate bias and improve compliance with regulations such as GDPR and the EU AI Act. Expanding its evaluation to AI systems deployed in other sectors could reveal its broader applicability and effectiveness. Additionally, refining the machine-readable version for better integration with diverse AI systems and exploring its use beyond healthcare could further enhance its impact, providing a scalable solution for ethical AI development across industries.

For the Irish healthcare context, there is a push for implementing a national framework that enables the wider sharing and reuse of healthcare data - especially for secondary purposes - and which is aligned with the future implementation of European Health Data Spaces (EHDS). We believe our approach for creating datasheets based on GDPR will facilitate this approach, and therefore would benefit from its application and refinement through use in real-world use-cases such as in hospitals and other clinical research settings.

Acknowledgments

This work was funded by the Health Research Board (HRB) through the Summer Student Scholarships (SS) scheme awarded to Marjia Siddik. The ADAPT SFI Centre for Digital Media Technology is funded by Science Foundation Ireland through the SFI Research Centres Programme and is co-funded under the European Regional Development Fund (ERDF) through Grant#13/RC/2106_P2.

References

- [1] E. J. Topol, High-performance medicine: the convergence of human and artificial intelligence, *Nature medicine* 25 (2019) 44–56.
- [2] P. Shah, F. Kendall, S. Khozin, R. Goosen, J. Hu, J. Laramie, M. Ringel, N. Schork, Artificial intelligence and machine learning in clinical development: a translational perspective, *NPJ digital medicine* 2 (2019) 69.
- [3] Z. Obermeyer, E. J. Emanuel, Predicting the future—big data, machine learning, and clinical medicine, *New England Journal of Medicine* 375 (2016) 1216–1219.
- [4] Z. Obermeyer, B. Powers, C. Vogeli, S. Mullainathan, Dissecting racial bias in an algorithm used to manage the health of populations, *Science* 366 (2019) 447–453.
- [5] D. S. Char, N. H. Shah, D. Magnus, Implementing machine learning in health care—addressing ethical challenges, *New England Journal of Medicine* 378 (2018) 981–983.
- [6] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, A. Galstyan, A survey on bias and fairness in machine learning, *ACM computing surveys (CSUR)* 54 (2021) 1–35.
- [7] R. Benjamin, Assessing risk, automating racism, *Science* 366 (2019) 421–422.

- [8] E. Vayena, A. Blasimme, I. G. Cohen, Machine learning in medicine: addressing ethical challenges, *PLoS medicine* 15 (2018) e1002689.
- [9] A. Rajkomar, M. Hardt, M. D. Howell, G. Corrado, M. H. Chin, Ensuring fairness in machine learning to advance health equity, *Annals of internal medicine* 169 (2018) 866–872.
- [10] F. Jiang, Y. Jiang, H. Zhi, Y. Dong, H. Li, S. Ma, Y. Wang, Q. Dong, H. Shen, Y. Wang, Artificial intelligence in healthcare: past, present and future, *Stroke and vascular neurology* 2 (2017).
- [11] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. Iii, K. Crawford, Datasheets for datasets, *Communications of the ACM* 64 (2021) 86–92.
- [12] A. Paullada, I. D. Raji, E. M. Bender, E. Denton, A. Hanna, Data and its (dis) contents: A survey of dataset development and use in machine learning research, *Patterns* 2 (2021).
- [13] Regulation2016, Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation), *Official Journal of the European Union* L119 (2016). URL: <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:L:2016:119:TOC>.
- [14] Regulation2024, Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence and amending regulations (ec) no 300/2008, (eu) no 167/2013, (eu) no 168/2013, (eu) 2018/858, (eu) 2018/1139 and (eu) 2019/2144 and directives 2014/90/eu, (eu) 2016/797 and (eu) 2020/1828 (artificial intelligence act), *Official Journal of the European Union* L (2024). URL: <http://data.europa.eu/eli/reg/2024/1689/oj>.
- [15] G. of Ireland, Health information bill, *Oireachtas (No. 61 of 2024)* (2024). URL: <https://www.oireachtas.ie/en/bills/bill/2024/61/>.
- [16] L. A. Celi, J. Cellini, M.-L. Charpignon, E. C. Dee, F. Dernoncourt, R. Eber, W. G. Mitchell, L. Moukheiber, J. Schirmer, J. Situ, et al., Sources of bias in artificial intelligence that perpetuate healthcare disparities—a global review, *PLOS Digital Health* 1 (2022) e0000022.
- [17] M. Russo, M.-E. Vidal, Leveraging ontologies to document bias in data, *arXiv preprint arXiv:2407.00509* (2024).
- [18] B. Gaonkar, K. Cook, L. Macyszyn, Ethical issues arising due to bias in training ai algorithms in healthcare and data sharing as a potential solution, *The AI Ethics Journal* 1 (2020).
- [19] M. Ganz, S. H. Holm, A. Feragen, Assessing bias in medical ai, in: *Workshop on Interpretable ML in Healthcare at International Conference on Machine Learning (ICML)*, 2021.
- [20] K. S. Chmielinski, S. Newman, M. Taylor, J. Joseph, K. Thomas, J. Yurkofsky, Y. C. Qiu, The dataset nutrition label (2nd gen): Leveraging context to mitigate harms in artificial intelligence, *arXiv preprint arXiv:2201.03954* (2022).
- [21] A. C. Roman, J. W. Vaughan, V. See, S. Ballard, N. Schifano, J. Torres, C. Robinson, J. M. L. Ferres, Open datasheets: Machine-readable documentation for open datasets and responsible ai assessments, *arXiv preprint arXiv:2312.06153* (2023).
- [22] A. K. Heger, L. B. Marquis, M. Vorvoreanu, H. Wallach, J. Wortman Vaughan, Understanding machine learning practitioners’ data documentation perceptions, needs, challenges, and desiderata, *Proceedings of the ACM on Human-Computer Interaction* 6 (2022) 1–29.
- [23] E. M. Bender, B. Friedman, 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.

- [24] J. Giner-Miguel, A. Gómez, J. Cabot, Datadoc analyzer: A tool for analyzing the documentation of scientific datasets, in: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, 2023, pp. 5046–5050.
- [25] H. M. Jannah, Metareader: A dataset meta-exploration and documentation tool, 2014.
- [26] H. J. Pandit, B. Esteves, G. P. Krog, P. Ryan, D. Golpayegani, J. Flake, Data privacy vocabulary (dpv) – version 2, The 23rd International Semantic Web Conference (ISWC) (in-press) (2024). URL: <https://arxiv.org/abs/2404.13426>.
- [27] M.-A. Wren, S. Connolly, A european late starter: lessons from the history of reform in irish health care, Health Economics, Policy and Law 14 (2019) 355–373.
- [28] S. Craig, N. Kodate, Understanding the state of health information in ireland: A qualitative study using a socio-technical approach, International journal of medical informatics 114 (2018) 1–5.
- [29] B. Walsh, C. Mac Domhnaill, G. Mohan, Developments in healthcare information systems in ireland and internationally, Dublin: The Economic and Social Research Institute (2021).
- [30] D. Hickey, R. O’Connor, P. McCormack, P. Kearney, R. Rosti, R. Brennan, The data quality index: improving data quality in irish healthcare records, ICEIS, 2021.