



HAL
open science

Ethical implications of AI-generated synthetic health data

Juho Vaiste

► **To cite this version:**

| Juho Vaiste. Ethical implications of AI-generated synthetic health data. 2023. hal-04216538

HAL Id: hal-04216538

<https://hal.science/hal-04216538v1>

Preprint submitted on 25 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Ethical implications of AI-generated synthetic health data

Juho Vaiste

September 2023

1 Purpose

This article explores ethical implications surrounding the AI-supported generation of synthetic health data. Ethical implications are assessed through the European Union's high-level expert group's guidelines "Ethics guidelines for trustworthy AI" (AI HLEG) and the set of 7 key requirements presented. Generation of synthetic data has been highlighted as a potential solution for data availability and privacy issues in healthcare sector, but a broad discussion on ethical consequences is necessary.

2 Design/methodology/approach

The research was conducted as a case study within a project investigating the generation and use of AI-generated health data. Ethical considerations were formulated based on earlier literature, insights from other organizations involved in the project, discussions with the technical research team, and the project's technical results.

This paper is divided into three parts. The first part offers a brief background on synthetic health data and its current methods of synthesis. The second part delves into ethical considerations, categorized under seven key requirements as outlined by AI HLEG guidelines, and recaps previously related literature. The final part suggests pertinent regulatory and policy implications.

3 Findings

Synthetic health data can provide significant benefits, especially when used as training data for health technology or for educational initiatives. Ethical challenges arising from generating and using synthetic health data can be multifaceted, as many ethical decisions intersect with technical ones. Synthetic health data is an intriguing topic of operationalizing technology ethics into practice. Nonetheless, synthetic data doesn't introduce entirely new ethical dilemmas;

many of these ethical themes are already recognized in other domains of information technology.

4 Originality/value

Ethical discussion of synthetic health data is limited to a few articles in major medical journals. Although the medical community has a major say in health-related issues, the perspective of technology ethics can be useful in filling the gaps.

5 Keywords

Synthetic data, health data, machine learning Paper type: conceptual

6 Introduction

Health data plays a crucial role in research, diagnostics, and treatment, and the demand for health data continues to grow. However, data availability is challenging due in part to privacy restrictions. Data availability is an essential issue for producing innovations and technology solutions in the healthcare sector (Chen et al., 2021). Data availability challenges have been identified, for example, in the analysis of imaging services (Islam, Wijewickrema O’Leary, 2021), health data analytics (Tayefi et al., 2021), and the production of infectious disease surveillance solutions (Kamel Boulos, 2022). Private healthcare and information technology companies are also interested in developing their products and services to meet the needs of the self-care and home monitoring sectors, among others. The issue of data availability is also linked to the development of new artificial intelligence solutions (Mahmood et al., 2019; Wan Jones, 2020).

One proposed solution to the data availability is synthetic data. While not an invention, synthetic data generation is the process of creating new data samples that mimic the statistical properties of an existing data set. The goal of generating synthetic data is to increase the size of the data set or to augment the data set with new samples that are representative of the original data distribution. Different contemporary methods to create synthetic data are for example embedding, generative adversarial networks, sequential synthesis and tree or copula based methods (James et al., 2021). By processing synthetic data, it is not only possible to expand or fulfill data sets, but synthetic data offers properties helping data sharing and strengthening privacy protection.

Synthetic health data holds the potential to offer a viable solution to maintaining patient privacy while facilitating research and development in healthcare. Generally, the medical community seems receptive to these innovative approaches, anticipating considerable potential benefits (Wan Jones, 2020; Chen

et al., 2021; Braddon et al., 2023). However, from a technical perspective, several questions remain unresolved, and the technology is still in the early stages of adoption. Incorporating ethical analysis from the initial phases can assist in ensuring that synthetic data is trustworthy, precise, and utilized responsibly.

Ethical considerations might influence how synthetic medical data is utilized, favoring some use cases over others. A pivotal question for synthetic data is its intended purpose. The 'silver bullet'—sparking the excitement around synthetic health data—is the potential for synthetic data to replace real data, either wholly or partially, in health and medical research. This advanced application of synthetic health data may still be a distant goal or turn out even entirely unrealistic. However, more attainable use scenarios can be found in the development of health technology or for educational purposes.

6.1 Benefits and use cases of synthetic medical data

Medical data is an essential resource to the advancement of research, diagnostics, and treatment in healthcare. With the ongoing shift from traditional, human-centric healthcare methodologies to advanced, technology-driven practices, the importance and value of data are accentuating, a trend highlighted by Murtaza et al. (2023). Ethical considerations, such as patient privacy and data security, are invariably linked to the use of medical data. Strict privacy regulations, which are crucial, often limit the ability to share data, impacting its availability. Other reasons for limited data availability include concerns over consent, fear of potential pitfalls, organizational barriers, and resource constraints (Schwendicke Krois, 2022).

The fundamental advantage of synthetic medical data is its capacity to mitigate issues related to data availability, offering a novel solution for privacy protection. It maintains statistical properties comparable to actual data, ensuring the robust protection of patient privacy. Synthetic health data is optimally utilized when perceived as a supplement to real-world health data, aptly filling or enhancing numerous existing data availability gaps (James et al., 2021).

The application of synthetic data is multifaceted. In their review article, Gonzales, A., Guruswamy, G., Smith, S. R. (2023) delineated seven distinct categories of use for synthetic medical data, including: 1) Simulation and Prediction Research 2) Hypothesis, Methods, and Algorithm Testing 3) Epidemiological Study/Public Health Research 4) Health IT Development and Testing 5) Education and Training 6) Public Release of Datasets 7) Linking data.

Beyond generating new data internally, the process of data synthesis facilitates data sharing among medical entities due to its inherent privacy-preserving characteristics. While it may not address every concern related to proprietary information, synthesis enhances data exchange between willing and authorized entities. Consequently, synthetic data supports improved model development and the adoption of innovative techniques by third parties, enables the consolidation of data without disclosing sensitive information, and promotes research reproducibility in fields like bioinformatics and healthcare.” (Nikolenko, 2019).

6.2 Grounds for the ethical assessment

Despite the potential benefits of synthetic medical data, there are also ethical concerns associated with its use. This article will examine the ethical implications of synthetic medical data, dividing into seven subsections: human agency and oversight, technical robustness and safety, privacy and data governance, transparency, diversity, non-discrimination and fairness, societal and environmental well-being, and accountability.

The subsections of this ethical review are grounded in the guidelines for AI use proposed by the European Union’s High-Level Expert Group (AI HLEG). The choice of this specific framework may prompt inquiries; however, the AI HLEG guidelines have gained substantial recognition. This acknowledgment is evidenced by the existence of principles or requirements in AI HLEG that are comparable to those found in other AI guidelines published recently, as highlighted by Jobin, A., Ienca, M., Vayena, E. (2019). A distinct discussion might involve whether guidelines explicitly devised for data processing or use would be more appropriate. However, it is pertinent to mention within this discourse that methods of generation categorized under AI technologies are presumed to constitute a substantial portion of the entity under assessment.

Ethical concerns regarding synthetic data generation are intertwined. Generating synthetic data demands careful selection of generation methods and statistical parameters, each influencing trade-offs between desired attributes in distinct ways. Central to synthetic data are the trade-offs between utility-privacy, utility-fairness, and privacy-fairness. Decisions on these trade-offs directly impact the need for informed consent. Furthermore, the intricate nature of these trade-offs might pose challenges for transparency.

However, this ethical investigation ought to be viewed as exploratory research. Existing literature on synthetic health data often gravitates towards either high levels of abstraction or concentrates predominantly on technical aspects. For a more nuanced ethical evaluation, we need concrete use cases of synthetic health data and in-depth accounts of the synthesization processes. Ultimately, generating and utilizing synthetic health data will necessitate assessments that are specific to each case and organization.

7 Ethical considerations on AI-generated synthetic health data

7.1 Human agency and oversight

Informed consent is a cornerstone in medical research and clinical care, acting as a safeguard to ensure individuals are fully aware of, and agree to, the collection, use, and sharing of their personal information. Although the construct of informed consent has its imperfections, it is considered the gold standard in data collection and utilization policies (Froomkin, 2019). Obtaining informed consent is crucial when using real patient data, serving to uphold patients’ autonomy

and privacy rights.

The proposition for the utilization of synthetic health data often encounters reservations due to potential reductions or omissions in the required informed consent for data use. It is debatable if the prime motivation behind synthesizing data seems to alleviate or negate consent protocols. Some argue that the predominant consent-based model for data processing is argued to be one central reason for the data availability issues and propose reconsideration of informed consent based on, for example, ‘easy rescue’, ‘no harm’ and ‘consent bias’ (Ploug, 2020).

The discussion surrounding informed consent and synthetic data usually oscillates between two polarizing viewpoints:

- 1) The creation of synthetic medical data involves the transformation or simulation of real patient data to generate a new dataset that retains the original data’s statistical properties without revealing individual patients’ information. This preservation of privacy prompts some to assert that informed consent might not be requisite, as stated by Shi et al. (2022).
- 2) Conversely, the connection between synthetic and original patient data can differ based on the generation method employed. In certain scenarios, synthetic data might still possess identifiable information or could be traced back to the original data, posing subsequent privacy risks. Consequently, the relevance of informed consent persists in the utilization of synthetic medical data.

The need for informed consent in creating synthetic health data, and more broadly in big data practices, remains a subject of debate, succinctly articulated by Ploug (2020). However, modern methods of synthesizing AI-generated data do not inherently resolve these differing views. Nonetheless, attaining advanced levels of data protection could shift some of the discourse, suggesting that consent may not always be universally required. Still, before further progress in this debate, it can be considered valid that the indirect involvement of human participants in data generation mandates sustaining transparency and adherence to individual rights, maintaining the indispensability of informed consent.

Individual cases and organizational policies significantly influence this discourse, with regulations and governmental policies often exhibiting diverse perspectives on this matter. Regardless of regulatory determinations, ethical conduct and provisions present adaptable, secure, and responsible solutions to informed consent requirements. For example, exploration of varied consent models, including broad or dynamic consent, can amplify individual control over data use in synthetic data generation processes. Also, the integration of privacy protection mechanisms within synthetic data might presumably encourage more individuals to allow the secondary use of their data.

7.2 Technical Robustness and safety

Data accuracy is pivotal for ensuring technical robustness and safety in AI-generated synthetic health data. As depicted in 'Nature Outlook: Robotics and Artificial Intelligence,' "machine-generated datasets have the potential to enhance privacy and representation in artificial intelligence, if a balanced alignment between accuracy and artificiality is achieved" (Savage, 2023).

Evaluating the precision of synthetic medical data is feasible through a myriad of statistical tools (Yale et al., 2020), which include comparing the statistical attributes of synthetic and real-world medical data (Chen et al., 2019). Moreover, analyzing the predictive proficiency of models trained on synthetic data can shed light on the data's reliability (El Emam et al., 2022). Given current technological advancements, Azizi et al. (2021) affirm that synthetic data can adequately represent real medical data, although there are detractors who...

Reliability is intertwined with accuracy and examines the longitudinal stability of the data, assessing whether it can consistently yield dependable outcomes across varied applications and contexts. While accuracy is paramount in validating synthesization techniques against specific datasets, the reliability of the same technique when applied to different datasets is crucial. The pursuit of reliability is incessant, yet diverse synthetic data generation methods have demonstrated encouraging outcomes, with minimal compromises in precision (Rankin et al., 2020).

Beyond technical considerations, scrutinizing accuracy and reliability from an ethical lens is imperative. It raises ethical queries about the acceptable thresholds of accuracy and reliability before endorsing specific methodologies, as elucidated in Article 4 of UNESCO's Universal Declaration on Bioethics and Human Rights. Establishing clinical quality indicators and evaluative metrics for synthetic data poses inherent challenges, particularly when dealing with novel or rare medical conditions (Chen et al., 2021). Given the nascent stage of synthetic medical data utilization, meticulous assessments on a case-by-case and organizational level appear to be the pragmatic approach, necessitating profound expertise from both medical and technical domains.

7.3 Privacy and data governance

While synthetic medical data is tailored to uphold patient privacy, it remains susceptible to unauthorized access and breaches. The pursuit of advanced data protection techniques has driven the healthcare sector to explore synthetic data since the early 1990s (Nikolenko, 2019). The prevailing discourse on data protection scrutinizes the frailty of conventional anonymization and statistical disclosure control methods and argues that they are susceptible to various adversarial attacks, such as linkage or re-identification attacks (Nieminen, 2022).

Creating synthetic data bolsters privacy protection by generating data that mirrors the original dataset without exposing sensitive individual information. However, the degree of privacy safeguarded is contingent on the efficacy of the synthetic data generation methods and their capability to thwart re-identification

or data leakage. Synthesizing data alone does not fully fortify data against re-identification (James, 2021), and different synthesis methods exhibit varied efficacy. For instance, Yoon et al. (2022) proposed a data synthesis method named EHR-Safe, claiming it offers “near-optimal privacy preservation” in producing synthetic medical data.

Implementing privacy-preserving techniques is essential to shield sensitive patient data during synthetic medical data generation. Several strategies, including Differential Privacy, k-Anonymity, l-Diversity, t-Closeness, Secure Multi-Party Computation (SMPC), and Federated Learning, can be applied to achieve this objective. However, the myriad of methods available for data protection in the realm of synthetic data has led to a fragmented field. From an ethical and non-technical standpoint, this makes evaluating the levels of privacy protection among different options quite challenging. Although there are some more clarifying models that elucidate data protection methods and parameter selection, such as models for epsilon in differential privacy (Hsu et al., 2014), clear gold standards seem still be lacking.

If synthetic data retains linkages to individual patients or the foundational dataset, it may lead to privacy infringements. To avoid such breaches, it is crucial for researchers and organizations to integrate stringent security protocols, such as stringent access controls, encryption, intrusion detection systems, and frequent security assessments, to safeguard synthetic medical data. While these processes may seem obvious, the introduction of synthetic data can create new data leakage scenarios for which preparation is necessary.

In conclusion, data protection and mitigation potential misuse of synthetic medical data necessitates a multifaceted strategy. Through meticulous evaluation and enhancement of data generation techniques, coupled with the promotion of ethical consciousness and the enforcement of robust security protocols, researchers and healthcare practitioners can collaborate to minimize risks related to data misuse and uphold responsible synthetic medical data utilization. Employing technical privacy-preserving methods is pivotal when the situation demands, but further research and golden standards in the technical community is needed for deeper ethical evaluation.

7.4 Transparency

Trust, transparency, and explainability are critical ethical considerations in employing synthetic medical data for research, diagnostics, and treatment purposes. The review by Tucci, V., Saary, J., Doyle, T. E. (2022) provides an in-depth analysis of the factors influencing trust in healthcare artificial intelligence. Trust is a multifaceted entity, nurtured by several elements, notably including transparency and explainability, which are pivotal for its establishment. Tucci, Saary, Doyle dissect 35 factors that contribute to forming trust. To instill trust in synthetic data, a minimum requirement is the transparency in data generation and utilization processes, coupled with clear and coherent communication and collaboration among all stakeholders—patients, researchers, healthcare professionals, and policymakers.

Explainability, like trust and transparency, is a trending topic in discussions surrounding ethics of AI, especially in healthcare. While the need for explainability is universally acknowledged, it often remains a nebulous ethical requirement subject to varying interpretations. The consensus is that explainability pertains to the ability to elucidate the workings of a technological application, but the extent and target audience for such explanations are subjects of ongoing debate. Amann et al. (2020) identify five different perspectives on explainability: technological, legal, medical, patient, and ethical. Ghassemi, Oakden-Rayner, Beam (2021), for instance, caution against a limited interpretation of explainable AI, where systems merely generate visual artifacts to illustrate results.

Transparency serves as another foundational element, necessary to achieve optimal levels of trust and explainability. Kiseleva, A., Kotzinos, D., De Hert, P. (2022) propose a robust framework for understanding transparency: "In the healthcare context, we advocate that transparency should be regarded as a system of accountabilities involving various actors (AI developers, healthcare professionals, and patients) positioned at different layers (insider, internal, and external layers, respectively)."

In conclusion, the significance of trust, explainability, and transparency cannot be understated in the context of emerging technologies. However, a pertinent question remains—do we necessitate extraordinary efforts towards these principles specifically for synthetic data? This proposition stems from the presumption that if synthetic data are to be leveraged, they will be integrated into existing data processes and procedures. Certainly, information regarding the data's origin will be disclosed (e.g., medical sample y vs. synthesized medical sample y), but it is crucial to deliberate whether additional, specialized measures are indispensable in the context of synthetic data. Balancing these ethical considerations is imperative to harness the potential of synthetic data responsibly and effectively.

7.5 Diversity, non-discrimination and fairness

Bias in data generation manifests as systematic errors or inaccuracies that can skew results and lead to unfair or discriminatory outcomes. In the realm of synthetic medical data, it's imperative that the methods employed neither perpetuate existing biases nor introduce new ones. Such biases can stem from numerous sources, including but not limited to, biased original data, flawed algorithms, or underrepresentation of certain demographic groups. Disproportionate representation of classes, especially concerning rare diseases, can hinder the AI models' diagnostic and prognostic accuracy (Chen et al., 2021).

As highlighted by Bhanot et al. (2021), quantifying bias is a pivotal first step towards its mitigation. They argue that conventional synthetic data generation methods often inaccurately reflect proportions in real data, inadvertently introducing biases, given their lack of consideration for fairness. They advocate for generative models and covariate fairness metrics, which contemplate fairness, particularly for non-temporal and temporal healthcare records. They

further recommend the integration of fairness metrics into generative methodologies and the utilization of conditional GANs to amplify under-represented subgroups.

Addressing biases mandates meticulous data preparation, utilization of diverse and representative datasets, and scrutiny of data generation algorithms. Implementing fairness metrics, enhancing model interpretability, and conducting algorithmic audits are crucial to curtail biases inherent in synthetic data generation methods. The inclusion of varied stakeholders and transparent disclosure of methods, algorithms, and data sources play a critical role in cultivating trust and fostering collaborative ventures within the medical and research domains. This ensures comprehensive perspective consideration and helps in identifying potential biases.

7.6 Societal and environmental well-being

The criterion of 'social well-being,' as delineated in the AI HLEG guidelines, is inherently expansive and a bit abstract. Nevertheless, this can be construed as an endorsement of benevolent objectives in the realm of synthetic data. The paramount aim in the exploration and advancement of synthetic data utilization is to catalyze enhanced results in medical and health technology fields. Contrary to dystopian projections of rampant exploitation of health data cloaked in synthetic data technologies, it's more plausible that synthetic data will evolve as a meticulously regulated and wisely utilized resource in medical and healthcare industries.

As for environmental repercussions, the impact of synthetic data is potentially inconclusive. The encompassing discourse on environmental and climatic implications predominantly intersects with deliberations surrounding information technology and artificial intelligence. For a more in-depth exploration of this subject, refer to studies like those by Higón, D. A., Gholami, R., Shirazi, F. (2017).

7.7 Accountability

From our viewpoint, the generation and utilization of synthetic health data do not inherently carry significant implications for accountability. The responsibility for creating, employing, and disseminating synthetic data predominantly aligns with those actors and organizational entities already held accountable for the organization's existing data practices. Nonetheless, in the context of synthetic health data, potential collaborations with third parties need recognition. Several commercial entities supply synthetic data in varied formats, which often tend to be of subpar quality. Engagements with such entities might invoke novel challenges concerning accountability and liability.

On the other hand, when specific types of health data, for instance, health registry data, become accessible for wider collaborative endeavors, the accountability of the original data proprietor becomes more prominent. The proprietor remains liable for any potential misuse or breaches of the data. This trajectory

might be favorable, with the onus of accountability ultimately residing with large public health institutions.

8 Regulatory and Policy Implications

The regulatory landscape for synthetic medical data varies across countries and regions, with some jurisdictions having more comprehensive regulations than others. In general, the regulatory frameworks for synthetic medical data are guided by broader data protection and privacy laws, such as the General Data Protection Regulation (GDPR) in the European Union or the Health Insurance Portability and Accountability Act (HIPAA) in the United States. These regulations typically address issues such as data anonymization, pseudonymization, and data processing, which are relevant to synthetic data generation and usage.

However, specific regulations that directly address synthetic medical data are lacking or still in development. In the short term, this may lead to uncertainty regarding the legal and ethical requirements for generating, using, and sharing synthetic medical data. At the same time it is possible that more technical research and knowledge on synthetic health data is needed before further regulative steps can be taken.

To address ethical concerns and ensure the responsible use of synthetic medical data, a variety of potential policy changes could be considered. First, regulatory bodies should develop specific regulations and guidelines that address the generation, use, and sharing of this data. These guidelines would need to cover issues such as recommended methods for creating synthetic medical data, metrics to assess data accuracy and privacy protection, and guidelines for required consent. Second, national and regional regulatory frameworks should aim to harmonize data protection and privacy regulations. This would help to reduce inconsistencies and confusion regarding synthetic medical data usage, facilitating cross-border data sharing and collaboration while still addressing ethical concerns.

Given the global nature of medical research and the increasing use of synthetic medical data, international collaboration and harmonization of standards are essential. Efforts to achieve harmonization in synthetic medical data involve collaboration with international organizations, such as the World Health Organization (WHO) or the International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH), to develop global guidelines and best practices.

9 Conclusion

Synthetic medical data brings remarkable promises to data availability issues in the field of data-intensive healthcare and medical research. Synthetic medical data could enable various developments in healthcare still without compromising privacy (Coutinho-Almeida, Rodrigues Cruz-Correia, 2021). Synthetic medical

data is a worth of well-resourced research and RD activities at medical organizations, universities and companies. Utilization of synthetic medical data is in early-phase and there are no universal methods or metrics to create or evaluate it (Hernandez et al., 2022). Additional research is needed to reach common best practices, and in the meanwhile, all existing generative models are not yet prepared for straightforward creation of synthetic data, and if used without caution in healthcare systems, they might introduce weaknesses that could result in patient re-identification.

Responsible use of synthetic medical data requires a multifaceted approach that includes the development of guidelines, the implementation of privacy-preserving techniques, and the promotion of transparency in data generation and sharing processes. Ethical considerations concerning to synthetic medical data share a lot of common grounds with existing data ethics and AI ethics literature, but consists some novel elements. Utilization of synthetic medical data might add responsibilities for work roles which are responsible of the generation process, such as data scientists or data managers.

Establishing commonly agreed-upon methods and measures for evaluating synthetic medical data is a crucial subsequent step. While a considerable amount of review and validation literature exists, it appears that shared best practices are still lacking. Regulatory updates could potentially aid in this process as well. Developing measures to assess data accuracy and reliability, privacy protection, and bias is a fundamental step in progressing and broadening the scope of synthetic medical data.

However, there are ethical concerns as well. Synthetic health data should not be employed as a mechanism to bypass informed consent or to compromise any ethical principles in pursuit of abstract notions such as 'public good' or 'technological progress.' There may be instances where more liberal approaches toward the secondary use of health data are warranted, but such approaches must also be justified from ethical standpoints. Concrete evidence of the performance and outcomes of utilizing synthetic health data is requisite before pursuing such approaches.

In conclusion, synthetic medical data presents both challenges and opportunities. By addressing ethical concerns, promoting transparency, and fostering collaboration among stakeholders, we can ensure the responsible use of synthetic medical data and unlock its potential to contribute to a better understanding of diseases, unbiased and accurate diagnostics, and more effective treatments for the benefit of patients worldwide.

10 References

Amann, J., Blasimme, A., Vayena, E., Frey, D., Madai, V. I., Precise4Q Consortium. (2020). Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC medical informatics and decision making*, 20, 1-9.

Azizi, Z., Zheng, C., Mosquera, L., Pilote, L., El Emam, K. (2021). Can

synthetic data be a proxy for real clinical trial data? A validation study. *BMJ open*, 11(4), e043497.

Barbierato, E., Vedova, M. L. D., Tessera, D., Toti, D., Vanoli, N. (2022). A Methodology for Controlling Bias and Fairness in Synthetic Data Generation. *Applied Sciences*, 12(9), 4619.

Bhanot, K., Qi, M., Erickson, J. S., Guyon, I., Bennett, K. P. (2021). The problem of fairness in synthetic healthcare data. *Entropy*, 23(9), 1165.

Braddon, A. E., Robinson, S., Alati, R., Betts, K. S. (2023). Exploring the utility of synthetic data to extract more value from sensitive health data assets: a focused example in perinatal epidemiology. *Paediatric and Perinatal Epidemiology*, 37(4), 292-300.

Chen, J., Chun, D., Patel, M., Chiang, E., James, J. (2019). The validity of synthetic clinical data: a validation study of a leading synthetic data generator (Synthea) using clinical quality measures. *BMC medical informatics and decision making*, 19(1), 1-9.

Chen, R. J., Lu, M. Y., Chen, T. Y., Williamson, D. F., Mahmood, F. (2021). Synthetic data in machine learning for medicine and healthcare. *Nature Biomedical Engineering*, 5(6), 493-497.

Coutinho-Almeida, J., Rodrigues, P. P., Cruz-Correia, R. J. (2021). GANs for Tabular Healthcare Data Generation: A Review on Utility and Privacy. In *Discovery Science: 24th International Conference, DS 2021, Halifax, NS, Canada, October 11–13, 2021, Proceedings 24* (pp. 282-291). Springer International Publishing.

El Emam, K., Mosquera, L., Fang, X., El-Hussuna, A. (2022). Utility metrics for evaluating synthetic health data generation methods: validation study. *JMIR medical informatics*, 10(4), e35734.

Froomkin, A. M. (2019). Big data: Destroyer of informed consent. *Yale JL Tech.*, 21, 27.

Hernandez, M., Epelde, G., Alberdi, A., Cilla, R., Rankin, D. (2022). Synthetic data generation for tabular health records: A systematic review. *Neuro-computing*.

Higón, D. A., Gholami, R., Shirazi, F. (2017). ICT and environmental sustainability: A global perspective. *Telematics and Informatics*, 34(4), 85-95.

Ghassemi, M., Oakden-Rayner, L., Beam, A. L. (2021). The false hope of current approaches to explainable artificial intelligence in health care. *The Lancet Digital Health*, 3(11), e745-e750.

Islam, K. T., Wijewickrema, S., O'Leary, S. (2021). A deep learning based framework for the registration of three dimensional multi-modal medical images of the head. *Scientific Reports*, 11(1), 1-13.

James, S., Harbron, C., Branson, J., Sundler, M. (2021). Synthetic data use: exploring use cases to optimise data utility. *Discover Artificial Intelligence*, 1(1), 15.

Jobin, A., Ienca, M., Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature machine intelligence*, 1(9), 389-399.

Jordon, J. et al. (2022) Synthetic Data - what, why and how? Commissioned by the Royal Society. <https://doi.org/10.48550/arXiv.2205.03257>.

Jordon, J., Yoon, J., Van Der Schaar, M. (2018, September). PATE-GAN: Generating synthetic data with differential privacy guarantees. In International conference on learning representations.

Jälkö, J., Lagerspetz, E., Haukka, J., Tarkoma, S., Honkela, A., Kaski, S. (2021). Privacy-preserving data sharing via probabilistic modeling. *Patterns*, 2(7), 100271.

Kamel Boulos, M. N., Kwan, M. P., El Emam, K., Chung, A. L. L., Gao, S., Richardson, D. B. (2022). Reconciling public health common good and individual privacy: new methods and issues in geoprivacy. *International Journal of Health Geographics*, 21(1), 1-9.

Kiseleva, A., Kotzinos, D., De Hert, P. (2022). Transparency of AI in health-care as a multilayered system of accountabilities: Between legal requirements and technical limitations. *Frontiers in Artificial Intelligence*, 5, 82.

Mahmood, F., Borders, D., Chen, R. J., McKay, G. N., Salimian, K. J., Baras, A., Durr, N. J. (2019). Deep adversarial training for multi-organ nuclei segmentation in histopathology images. *IEEE transactions on medical imaging*, 39(11), 3257-3267.

Murtaza, H., Ahmed, M., Khan, N. F., Murtaza, G., Zafar, S., Bano, A. (2023). Synthetic data generation: State of the art in health care domain. *Computer Science Review*, 48, 100546.

Nieminen, V. (2022). Differentially private synthetic tabular data generation with a generative adversarial network and privacy amplification by subsampling. Master's Thesis University of Turku, Department of Computing, Data Analytics. <https://urn.fi/URN:NBN:fi-fe2022082956643>.

Nikolenko, S. I. (2019). Synthetic data for deep learning. arXiv preprint arXiv:1909.11512.

Ploug, T. (2020). In Defence of informed consent for health record research—why arguments from ‘easy rescue’, ‘no harm’ and ‘consent bias’ fail. *BMC medical ethics*, 21, 1-13.

Rajotte, J. F., Mukherjee, S., Robinson, C., Ortiz, A., West, C., Ferres, J. M. L., Ng, R. T. (2021). Reducing bias and increasing utility by federated generative modeling of medical images using a centralized adversary. In *Proceedings of the Conference on Information Technology for Social Good* (pp. 79-84).

Rankin, D., Black, M., Bond, R., Wallace, J., Mulvenna, M., Epelde, G. (2020). Reliability of supervised machine learning using synthetic data in health care: Model to preserve privacy for data sharing. *JMIR medical informatics*, 8(7), e18910.

Savage, N. (2023). Synthetic data could be better than real data. *Nature Outlook: Robotics and artificial intelligence*. <https://doi.org/10.1038/d41586-023-01445-8>.

Schwendicke, F., Krois, J. (2022). Data dentistry: how data are changing clinical care and research. *Journal of dental research*, 101(1), 21-29.

Shi, J., Wang, D., Tesei, G., Norgeot, B. (2022). Generating high-fidelity privacy-conscious synthetic patient data for causal effect estimation with multiple treatments. *Frontiers in Artificial Intelligence*, 195.

Tayefi, M., Ngo, P., Chomutare, T., Dalianis, H., Salvi, E., Budrionis, A., Godtlielsen, F. (2021). Challenges and opportunities beyond structured data in analysis of electronic health records. *Wiley Interdisciplinary Reviews: Computational Statistics*, 13(6), e1549.

Tucci, V., Saary, J., Doyle, T. E. (2022). Factors influencing trust in medical artificial intelligence for healthcare professionals: A narrative review. *J. Med. Artif. Intell*, 5(4).

Vu, D., Slavkovic, A. (2009, December). Differential privacy for clinical trial data: Preliminary evaluations. In *2009 IEEE International Conference on Data Mining Workshops* (pp. 138-143). IEEE.

Wan, C., Jones, D. T. (2020). Protein function prediction is improved by creating synthetic feature samples with generative adversarial networks. *Nature Machine Intelligence*, 2(9), 540-550.

Yale, A., Dash, S., Dutta, R., Guyon, I., Pavao, A., Bennett, K. P. (2020). Generation and evaluation of privacy preserving synthetic health data. *Neurocomputing*, 416, 244-255.

Jinsung Yoon, Michel Mizrahi, Nahid Ghalaty et al. EHR-Safe: Generating High-Fidelity and Privacy-Preserving Synthetic Electronic Health Records, 07 December 2022, PREPRINT (Version 1) available at Research Square. <https://doi.org/10.21203/rs.3.rs-2347130/v1>.

Ziller, A., Usynin, D., Braren, R., Makowski, M., Rueckert, D., Kaissis, G. (2021). Medical imaging deep learning with differential privacy. *Scientific Reports*, 11(1), 1-8.