
Synthetic datasets can provide the health service with better AI models

PERSPECTIVES

VIBEKE BINZ VALLEVIK

vibeke.binz@gmail.com

Vibeke Binz Vallevik, PhD candidate at the University of Oslo and senior researcher at Group Research and Development Healthcare, DNV. The author has completed the ICMJE form and declares no conflicts of interest.

ANNE KJERSTI C. BEFRING

Anne Kjersti C. Befring, jurist and professor at the University of Oslo. She is deputy chair of Nordic Permed Law and the National Committee of Medical and Health Research Ethics. The author has completed the ICMJE form and declares no conflicts of interest.

HELGA MARIA BERGEM BRØGGER

Helga Maria Bergem Brøgger, specialist in radiology and senior researcher at Group Research and Development Healthcare, DNV. She is a member of the committee for research, quality and innovation of the Norwegian Medical Association and the expert group in Nordic Innovation (under the Nordic Council of Ministers). She is a board member of the Norwegian Society of Radiology and member of the Norwegian Council for Digital Ethics. The author has completed the ICMJE form and declares no conflicts of interest.

COURTNEY NADEAU

Courtney Nadeau, senior researcher at Group Research and Development Healthcare, DNV. The author has completed the ICMJE form and declares no conflicts of interest.

Producing synthetic patient data reduces the risk of privacy violations when AI tools are introduced. However, synthetic data can involve other types of risks.

EU's new Artificial Intelligence Act [\(1\)](#) defines artificial intelligence (AI) as a machine-based system that operates with varying levels of autonomy and infers, from the input it receives, predictions, content and recommendations. Article 2 [\(7\)](#) of the AI Act refers to the GDPR for its provisions on processing of personal data [\(2\)](#).

To ensure that AI tools for use in patient treatment are accurate and safe, providers need large amounts of data to develop and train their products. In addition, the health services need data to validate and test the AI tools, or to adapt them to local populations [\(3\)](#).

«Synthetic data can be produced on the basis of real patient data with the aid of generative AI models and can consist of test results, radiological images and patient record notes that look real, even though they are not»

Synthetic datasets are now being produced for these objectives, and can consist of pictures, sounds, tables, time series [\(2\)](#). Synthetic data can be produced on the basis of real patient data with the aid of generative AI models and can consist of test results, radiological images and patient record notes that look real, even though they are not. There is widespread optimism in this area when it comes to deep generative models that most of us now know from ChatGPT.

In theory, synthetic data cannot be linked directly to individuals, allowing fewer restrictions on processing [\(4\)](#). The EU General Data Protection Regulation (GDPR) [\(5\)](#), which has been incorporated into Norwegian legislation, defines all information that can be linked to individuals as personal data. All personal data that are used for the development of AI, including those used to produce synthetic data, are therefore regulated by the GDPR. Other relevant regulations include the duty to ensure the right to privacy and confidentiality, set out in Section 102 of the Norwegian Constitution, and Section 3 - 6 of the Patient and User Rights Act [\(6\)](#).

Moreover, AI models for the health services must be trustworthy and ensure patient safety, for example by making the correct diagnosis or proposing an effective treatment. These are arguments in favour of permitting access to personal data and health information.

However, the access to personal and health data is restricted by legal provisions as well as scarcity of data. The lack of health data tends to be most conspicuous when it comes to rare conditions and illness in children. For these, synthetic data may represent the best as well as the fastest opportunity to obtain sufficient data for the development of new AI models [\(7\)](#).

The duty of confidentiality

The purpose of the duty of patient confidentiality in the health services is to permit individuals to seek out health assistance with confidence that their personal information will not be divulged or made available to unauthorised parties [\(8\)](#). Public health registries could be a basis for generating synthetic data that represent far less of a challenge to the duty of confidentiality and the trust in the health services. Considerable amounts of health data are currently stored in the national quality registries for medicine and other health registries. Registry data with patient information can be transferred for secondary purposes as defined by the Health Registries Act, and these data are also subject to the duty of confidentiality [\(9\)](#).

Data that are encompassed by the duty of confidentiality can be exempted from this duty and put to use if it is unlikely that they can be linked to the person in question, and if the benefits and risks of this are proportionate. Synthetic data that are not related to individuals are not personal data and hence not subject to the duty of confidentiality.

To enable assessment of the identification risk and risk-mitigating measures, the AI model must be transparent and open about the data that have been used [\(10, 11\)](#).

Risks and challenges

Synthetic datasets must be representative of the population on which they are intended to be used. Quality control is required throughout the entire synthetisation process, from control of the original dataset, measurement of statistical similarities between the training dataset and the synthetic dataset, to testing of the performance of a model which has been trained on the dataset [\(11\)](#). Otherwise, the AI models can introduce new risks for patients when used in diagnostics and treatment.

The use of generative AI models to produce synthetic data is resource-intensive and energy-consuming, and this has consequences in terms of the environment and sustainability [\(12\)](#).

Biases in the training dataset can inadvertently be reinforced through the synthetisation process. Generative methods that fail to capture underrepresented minority groups in the original data may lead to discrimination of persons and patient groups, because the algorithms in the AI models are insufficiently accurate and reliable for specific sub-groups [\(13\)](#). On the other hand, synthetisation may be used to counteract discrimination, for example by correcting for original biases in the generative process.

Protecting personal data

The risk of identification of individuals can be minimised through the use of synthetic data, and the opportunity for freer use and sharing can be extremely beneficial for the health services. When generative AI models are optimised to produce synthetic data with the greatest possible similarity to the original dataset, there may still be a risk of identification. Some models can produce exact copies of parts of the original datasets or data points that are materially identical, even when there is no one-to-one relationship between the real personal data in the training dataset and the synthetic data points. This is referred to as residual risk of identification in the synthetic dataset (11), where the residual risk increases with the degree of knowledge of the generative methods (14).

A pivotal question in the generation of synthetic data is identifying the line between when the data are to be considered personal data and when they are not, allowing processing outside the material scope of the GDPR or the duty of confidentiality. After all, a key objective of synthetic data generation is to reduce the data privacy risk in order to facilitate processing. The legal limit to identification risk is based on Section 4 (1) and recital 26 of the GDPR and various other legal sources (15–17). The likelihood of identification of persons is a key issue.

«The possibility of identification when the data are generated into synthetic data will vary in terms of the complexity of the dataset – the number of variables and patients, statistical outliers, or the generative method used as well as the access to other relevant information»

Unfortunately, it may be possible to infer personal data from synthetic datasets (18). The possibility of identification from synthetic datasets will vary depending on the complexity of the dataset – the number of variables and patients, statistical outliers and the generative method used, as well as the access to other relevant information. The risk of identification also depends on how resource-intensive that process is, and on whether safety precautions have been applied during the generative process.

Safety precautions that are intended to prevent identification from synthetic data may increase the risk that the dataset is no longer representative, which in turn may reduce the adequacy and utility of the data and the AI tool (19).

«In light of consequentialist ethics, the social utility of introducing AI tools in the health services could justify a higher risk of identifying individuals»

Ethical concerns

Ethical concerns and principles form the basis for parts of the legislative acts and can be included in legal deliberations, for example in assessments of reliability. Consequentialist ethics assumes that a decision is ethical if it optimises the level of positive consequences as a whole (20, 21). This implies that decisions must be assessed holistically for an opinion to be formed about their ethical tenability. The totality of positive effects must be balanced against the total damage that an action has caused (21). From the perspective of consequentialist ethics, the social utility of introducing AI tools in the health services could justify a higher risk of identifying individuals. A further consideration relates to the utility that the data may have for future generations – the consequentialists ascribe considerably higher importance to future generations (22). This contrasts with the world view of classical economics, where future utility has lower value than current utility, so-called discounting of future benefits.

From a duty ethics perspective, the duty toward future generations may also have a bearing. In a legal context, such perspectives can be encompassed by social considerations. Theories of duty ethics differ from consequentialist ethics by attaching more importance to the value, autonomy and rights of the individual than to social utility as a whole (21). A framework of duty ethics stipulates a positive duty to help others and a negative duty not to inflict damage. In the case of AI models, utility must be viewed more broadly than concerns for the individual, since this may benefit many people, including in the future. Social attitudes and the risk that people are willing to accept to achieve a collective good may also be relevant considerations. The acceptance of risk is expressed, for example, in surveys on the willingness to make personal data available for research purposes in order to benefit others (23, 24).

For synthetic data to help address the challenge involved in access to health data for development of AI for the benefit of our patients, it is presumed that these data can be processed either as anonymous data or under a legal exemption. If the data are treated as personal information, their use will be restricted and their utility greatly reduced. The fact that the data are synthetic could in itself be regarded as a risk-mitigating measure. The access to large amounts of data through the use of synthetic datasets could also make the health sector less dependent on global technology companies that possess vast amounts of locked proprietary data for use in the development of AI.

Special regulations for synthetic data in the health sector, as well as clarity in the government requirements for approval of equipment trained on such data, could provide more predictability for developers and producers of AI models, and also for users and patients in the health sector.

Patients already need to accept a certain level of risk when benefitting from health services. A topic that so far has been insufficiently focused concerns the kind of data protection risk we as a society believe individuals should accept as

a consequence of these data being used as a basis for improving access to future modern medical treatment, for the benefit of society in general.

REFERENCES

1. The Artificial Intelligence Act - Regulation (EU) 2024/1689.
<https://www.artificial-intelligence-act.com/> Accessed 12.9.2024.
2. Oussidi A, Elhassouny A. Deep generative models: Survey. 2018 International conference on intelligent systems and computer vision (ISCV).
<https://ieeexplore.ieee.org/document/8354080> Accessed 12.9.2024.
3. Sarker IH. Deep learning: a comprehensive overview on techniques, taxonomy, applications and research directions. *SN Comput Sci* 2021; 2: 420. [PubMed][CrossRef]
4. Gonzales A, Guruswamy G, Smith SR. Synthetic data in health care: A narrative review. *PLOS Digit Health* 2023; 2. doi: 10.1371/journal.pdig.0000082. [PubMed][CrossRef]
5. EU. 2016/679 The General Data Protection Regulation (GDPR).
<https://eur-lex.europa.eu/eli/reg/2016/679/oj> Accessed 12.9.2024.
6. Helse- og omsorgsdepartementet. LOV-1999-07-02-63. Lov om pasient- og brukerrettigheter (pasient- og brukerrettighetsloven).
<https://lovdata.no/dokument/NL/lov/1999-07-02-63?q=Pasient%20og%20brukerrettighetsloven> Accessed 12.9.2024.
7. Giuffrè M, Shung DL. Harnessing the power of synthetic data in healthcare: innovation, application, and privacy. *NPJ Digit Med* 2023; 6: 186. [PubMed][CrossRef]
8. Helse- og omsorgsdepartementet. LOV-1999-07-02-64. Lov om helsepersonell m.v. (helsepersonelloven).
<https://lovdata.no/dokument/NL/lov/1999-07-02-64?q=helsepersonelloven> Accessed 12.9.2024.
9. Helse- og omsorgsdepartementet. LOV-2014-06-20-43. Lov om helseregistre og behandling av helseopplysninger (helseregisterloven).
<https://lovdata.no/dokument/NL/lov/2014-06-20-43> Accessed 12.9.2024.
10. Rajotte JF, Bergen R, Buckeridge DL et al. Synthetic data as an enabler for machine learning applications in medicine. *iScience* 2022; 25. doi: 10.1016/j.isci.2022.105331. [PubMed][CrossRef]
11. Vallevik VB, Babic A, Marshall SE et al. Can I trust my fake data - A comprehensive quality assessment framework for synthetic tabular data in healthcare. *Int J Med Inform* 2024; 185. doi: 10.1016/j.ijmedinf.2024.105413. [PubMed][CrossRef]
12. Lavi H. Measuring greenhouse gas emissions in data centres: the environmental impact of cloud computing. *Climatiq*.

<https://www.climatiq.io/blog/measure-greenhouse-gas-emissions-carbon-data-centres-cloud-computing> Accessed 12.9.2024.

13. Bhanot K, Qi M, Erickson JS et al. The problem of fairness in synthetic healthcare data. *Entropy (Basel)* 2021; 23: 1165. [PubMed][CrossRef]
14. Wang Z, Myles P, Tucker A. Generating and evaluating cross-sectional synthetic electronic healthcare data: preserving data utility and patient privacy. *Comput Intell* 2021; 37: 819–51. [CrossRef]
15. Befring AK, Sand IJ. *Big data og kunstig intelligens i helsesektoren*. Oslo: Gyldendal, 2020.
16. EUR-Lex. Patrick Breyer v. Bundesrepublik Deutschland. Case C-582/14. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62014CJ0582> Accessed 12.9.2024.
17. EUR-Lex. Single Resolution Board v. EDPS. Case T-557/20. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A62020TJ0557> Accessed 12.9.2024.
18. Sun C, van Soest J, Dumontier M. Generating synthetic personal health data using conditional generative adversarial networks combining with differential privacy. *J Biomed Inform* 2023; 143. doi: 10.1016/j.jbi.2023.104404. [PubMed][CrossRef]
19. Ghatak D, Sakurai K. A Survey on Privacy Preserving Synthetic Data Generation and a Discussion on a Privacy-Utility Trade-off Problem. I: *Science of Cyber Security - SciSec 2022 Workshops*. New York, NY: Springer, 2022.
20. Beauchamp TL, Childress JF. *Principles of Biomedical Ethics*. Oxford: Oxford University Press, 2001.
21. Peach L. An introduction to ethical theory. I: Penslar RL. *Research Ethics: Cases and Materials*. Bloomington, IN: Indiana University Press, 1995: s. 13–26.
22. Parfit D. *Reasons and persons*. Oxford: Oxford University Press, 1987.
23. Kamm FM. Harming some to save others. *Philos Stud* 1989; 57: 227–60. [CrossRef]
24. Tosoni S, Voruganti I, Lajkosz K et al. Patient consent preferences on sharing personal health information during the COVID-19 pandemic: "the more informed we are, the more likely we are to help". *BMC Med Ethics* 2022; 23: 53. [PubMed][CrossRef]

Publisert: 25 November 2024. Tidsskr Nor Legeforen. DOI: 10.4045/tidsskr.24.0328
Received 10.6.2024, first revision submitted 14.8.2024, accepted 12.9.2024.
Copyright: © Tidsskriftet 2026 Downloaded from tidsskriftet.no 21 June 2026.