



Non-imaging Medical Data Synthesis for Trustworthy AI: A Comprehensive Survey

XIAODAN XING, HUANJUN WU, and LICHAO WANG, Imperial College London, UK

IAIN STENSON and MAY YONG, Alan Turing Institute, UK

JAVIER DEL SER, TECNALIA, Basque Research & Technology Alliance (BRTA), Spain

SIMON WALSH and GUANG YANG, Imperial College London, UK

Data quality is a key factor in the development of trustworthy AI in healthcare. A large volume of curated datasets with controlled confounding factors can improve the accuracy, robustness, and privacy of downstream AI algorithms. However, access to high-quality datasets is limited by the technical difficulties of data acquisition, and large-scale sharing of healthcare data is hindered by strict ethical restrictions. Data synthesis algorithms, which generate data with distributions similar to real clinical data, can serve as a potential solution to address the scarcity of good quality data during the development of trustworthy AI. However, state-of-the-art data synthesis algorithms, especially deep learning algorithms, focus more on imaging data while neglecting the synthesis of non-imaging healthcare data, including clinical measurements, medical signals and waveforms, and electronic healthcare records (EHRs). Therefore, in this article, we will review synthesis algorithms, particularly for non-imaging medical data, with the aim of providing trustworthy AI in this domain. This tutorial-style review article will provide comprehensive descriptions of non-imaging medical data synthesis, covering aspects such as algorithms, evaluations, limitations, and future research directions.

CCS Concepts: • **Computing methodologies** → **Modeling methodologies**;

Additional Key Words and Phrases: Medical data synthesis, electronic healthcare records

ACM Reference format:

Xiaodan Xing, Huanjun Wu, Lichao Wang, Iain Stenson, May Yong, Javier Del Ser, Simon Walsh, and Guang Yang. 2024. Non-imaging Medical Data Synthesis for Trustworthy AI: A Comprehensive Survey. *ACM Comput. Surv.* 56, 7, Article 165 (April 2024), 35 pages.

<https://doi.org/10.1145/3614425>

X. Xing, H. Wu, and L. Wang contributed equally to this research.

S. Walsh and G. Yang co-last senior authors.

This study was supported in part by the ERC IMI (101005122), the H2020 (952172), the MRC (MC/PC/21013), the Royal Society (IEC/NSFC/211235), the NVIDIA Academic Hardware Grant Program, the SABER project supported by Boehringer Ingelheim Ltd, and the UKRI Future Leaders Fellowship (MR/V023799/1). J. Del Ser also acknowledges support from the Department of Education of the Basque Government via the Consolidated Research Group MATHMODE (IT1456-22).

Authors' addresses: X. Xing (corresponding author), H. Wu, L. Wang, S. Walsh, and G. Yang (corresponding author), Imperial College London, London, UK; emails: x.xing@imperial.ac.uk, hw1019@imperial.ac.uk, lw422@imperial.ac.uk, s.walsh@imperial.ac.uk, g.yang@imperial.ac.uk; I. Stenson and M. Yong, Alan Turing Institute, London, UK; emails: istenson@turing.ac.uk, myong@turing.ac.uk; J. Del Ser, TECNALIA, Basque Research & Technology Alliance (BRTA), Derio, Spain; email: javier.delser@tecnalia.com.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

0360-0300/2024/04-ART165 \$15.00

<https://doi.org/10.1145/3614425>

1 INTRODUCTION

The use of **Artificial Intelligence (AI)** on health data is creating promising tools to assist clinicians in fields such as automatic evaluation of diseases and prognosis management [61]. However, AI algorithms can be biased, unfair, or unethical, with a high risk of privacy breaches [151]. These AI algorithms, failing to win human trust [79], hinder the development and large-scale applications of AI in healthcare scenarios. Over the past few decades, researchers have been working on developing *trustworthy AI* [71, 151] by improving robustness, variety, and transparency throughout the AI lifecycle, where the training data of AI algorithms [71] is identified as a key factor in achieving trustworthy AI.

A trustworthy training dataset should have a set of (overlapping) properties: (1) a large quantity, (2) an unbiased variety, (3) strict ethical regulations during collection [102, 145], and (4) a low risk of privacy breaches. Considering the complex procedure and strict protocols of medical data acquisition, practitioners face enormous difficulty acquiring a large quantity of high-quality medical data in the real world. In addition to the challenges of data acquisition, healthcare data are sensitive, and sharing or working with it can easily lead to a violation of patients' privacy.

To address these problems, researchers have developed data synthesis algorithms. By synthesizing medical data instead of acquiring it from the real world, they can improve the size and variety of training datasets, impute missing values, and protect patients' privacy. These synthetic data can serve as a qualified training set for trustworthy AI algorithms [83, 99]. Conventional data synthesis algorithms are based on sampling, where data are generated by sampling from the data distribution after modeling it. However, this statistical sampling requires prior knowledge of the distributional functions. Recently, deep learning models have been developed for data synthesis, as they do not require an explicit selection of distributional functions, and their performance has been widely validated, especially on medical images [49, 98]. These algorithms can generate various types of synthetic medical data, which can be used in various AI algorithms [161].

Synthetic medical data have two forms: imaging data and non-imaging data. While medical imaging data, such as X-ray, CT, and MR images, are important for lesion observation and quantification, they cannot be relied on solely for diagnosis and prognosis. Therefore, analyses of non-imaging medical data are crucial in large-scale *in silico* clinical trials. However, despite the importance of both imaging and non-imaging data, we have observed that non-imaging data have received little attention in data synthesis algorithms and computer-aided medical applications in recent decades. Furthermore, we have found that the development of synthesis algorithms has been dominated by deep learning models, and there has been a lack of detailed explanations on non-deep learning methods and their relationships in the literature.

To bridge the research gaps and provide guidance on trustworthy AI-based non-imaging medical data synthesis, this article presents a comprehensive survey of algorithms, evaluation metrics, and related datasets. During our literature review, we identified seven papers focusing specifically on non-imaging medical data synthesis algorithms (Table 1). While some empirical studies [45, 149, 159] conducted comparative analyses of EHR synthesis using different algorithms, their scope was limited and did not provide a comprehensive comparison of all algorithms and evaluation methods. Review papers such as [47, 54] only focused on specific aspects of data synthesis algorithms and did not cover the entire workflow of non-imaging medical data synthesis. Although comprehensive reviews like References [55, 95] provided literature reviews of non-imaging data synthesis algorithms, they, too, lacked a comprehensive dataset review and did not consider the trustworthy aspects of data synthesis algorithms. Furthermore, these comprehensive reviews mainly listed the algorithms without providing proper explanations or discussing their relationships with one another.

Table 1. Comparison of Existing Non-imaging Data Review Studies

Paper title	Number of studies	Period	Dataset review	Evaluation metrics	Contents
[149]	2	~2021	no	Fidelity and privacy	Empirical study
[45]	8	~2020	no	Fidelity* and privacy	Empirical study
[55]	34	~2022	no	Fidelity, utility and privacy	Comprehensive review paper
[47]	72	~2022	yes**	NaN	Review paper focusing on applications and use cases
[95]	70	~2022	no	Fidelity, utility and privacy	Comprehensive review paper
[54]	NA	~2022	no	Fidelity, utility and privacy	Review paper focusing on evaluation metrics
[159]	NA	~2022	no	Fidelity, utility and privacy	Empirical study
Ours	82	~2023	yes	Fidelity, utility, privacy and fairness	Comprehensive review paper

*In Reference [45], the term “data utility” was used to refer to data fidelity metrics, which could be confusing. **The paper only included a review on fully synthetic datasets.

In contrast, our review aims to provide a more comprehensive and up-to-date survey of non-imaging medical data synthesis algorithms, covering a broader range of synthesis techniques, evaluation metrics and open sourced datasets that contribute to the development of trustworthy AI. We also include a detailed description of the mathematical and statistical foundation of non-imaging data synthesis and provide a tutorial for non-imaging synthesis algorithms. Moreover, we address the limitations and research issues related to non-imaging medical data synthesis and provide guidance for data synthesis algorithms that contribute to trustworthy AI. This survey offers three contributions:

- Providing a comprehensive and up-to-date survey on non-imaging medical data synthesis algorithms;
- Clearly defining and describing non-imaging medical data, including open source datasets and pre-processing methods;
- Explaining the mathematical and statistical principles behind non-imaging data synthesis and providing a step-by-step guide to non-imaging synthesis algorithms.

The remainder of this review article is structured as follows: Section 2 outlines the criteria used to collect and filter the literature reviewed in this survey and presents the taxonomy used to systematically and coherently organize it. In Sections 3–5, we will introduce three major types of data generation algorithms. In Section 6, we will discuss three critical aspects of trustworthy synthetic data quality and their corresponding measurements. In Section 7, we will present several datasets categorized by their data types and provide commonly practiced pre-processing procedures for these data types. Finally, in Section 8, we will analyze the limitations of non-imaging medical data synthesis and propose potential research directions for non-imaging data synthesis that consider trustworthy AI.

2 LITERATURE COLLECTION AND TAXONOMY

All papers included in this review were obtained by a three-stage searching strategy.

First, we selected the papers regarding non-imaging healthcare data synthesis from January 1, 2000, to July 1, 2022, with the keywords “data synthesis,” “synthetic data,” “data generation,” “data

augmentation,” and “oversampling.” They were concatenated in an “or” logic relation. We confined our search to computer sciences area and deleted the papers that are not related according to their abstracts. We focus on two types of non-imaging medical data during our searching process: tabular data and sequential data. Other non-imaging data medical types, although provide crucial information in healthcare analysis, are not discussed. It is because (1) the synthesis and applications of these data types are scarce, i.e., social networks for infectious [8] or family inherited diseases, or (2) the synthetic data does not provide new information for downstream tasks, e.g., medical reports [62]. At this stage, we used Scopus (<https://www.scopus.com>) as our search engine, because it is the largest database for peer-reviewed literature. After the first screening, 988 papers were selected.

During the initial stage of our reference search, we did not specifically focus on healthcare data synthesis algorithms. For instance, we included **Synthetic Minority Over-sampling Technique (SMOTE)** [20], which proposes an algorithm from an application-agnostic perspective, without identifying any specific use cases. Therefore, after the preliminary screening based on abstracts, we read the papers more thoroughly and eliminated those that were not healthcare related or did not propose innovative algorithms. This left us with a total of 67 papers.

In the final stage of our literature collection process, we thoroughly read the selected papers and summarized their methods and applications. Additionally, we checked the reference lists of these papers to ensure we had not missed any relevant literature. We did not include papers from arXiv (<https://arxiv.org/>) in the first two stages, as they are not peer reviewed. However, highly cited arXiv papers (with a citation count >20) were included in the reference stage. After the final stage, we selected and analyzed 82 papers, as shown in Figure 1.

3 SIMULATION-BASED ALGORITHMS

In this section, we will introduce simulation-based algorithms for data synthesis. Simulation-based methods aim to generate synthetic data by simulating underlying real-world mechanisms. We have identified two categories of simulation-based algorithms based on their target non-imaging data types.

3.1 Medical Signal Simulation

Simulation-based algorithms have been widely used, particularly for generating medical signals, such as MEG, EEG, and fMRI [150]. For medical signal simulation, simulation-based methods use the summation of three basic components of the signals: a baseline signal, a signal of interest (in fMRI simulation, this signal of interest is the BOLD signal; in EEG simulation, the signal of interest is the electrical signals produced by the brain.) and noises, and the final simulated signals are the summation of these three basic components.

The baseline signal in simulated signals represents their basic numerical level, and the baseline values can be tissue specific [39]. However, some algorithms simply set the baseline value to zero [3]. For the synthesis of signals of interest, most medical signal methods consider the correlations among signals from different brain regions. Multivariate autoregressive modeling [9, 128] provides spatial-related correlations for medical signal simulation. The noises can be modeled by Gaussian distributions or mixture Gaussian distributions [84]. Motion noise can also be added to simulate patient motion during scanning.

3.2 EHR Simulation

One of the most well-known simulation-based synthesis algorithms for EHR is the **Synthetic Electronic Medical Records Generator (EMERGE)** project [14, 92]. This project synthesizes time series of events, which are referred to as “patient care models,” “care flow,” or “caremaps,” for both

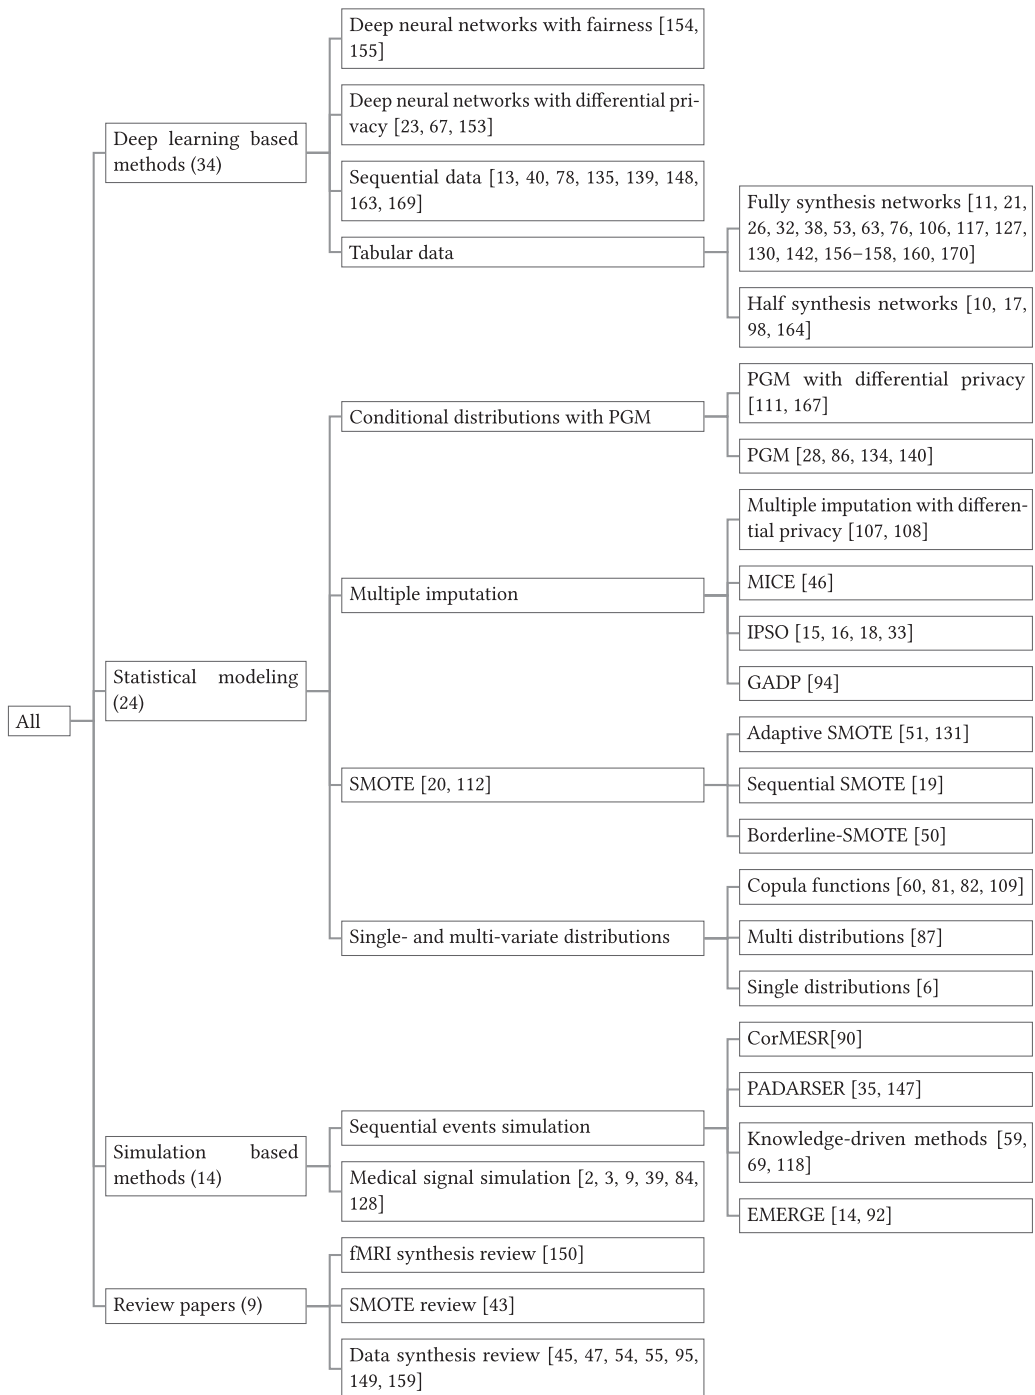


Fig. 1. All papers included in this review. These papers are categorized by their synthesis algorithms.

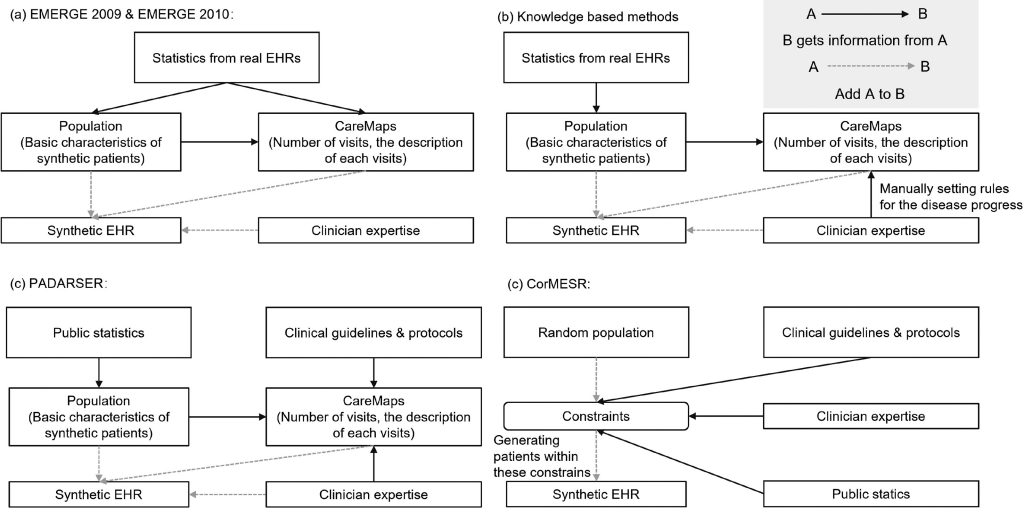


Fig. 2. A detailed development of the care model synthesis.

general populations and populations with specific diseases. Patient care models consist of a series of care-related tasks involved in managing a patient trajectory and provide a workflow guidance for patients with specific diseases [48].

EMERGE began by synthesizing a group of basic demographic information and symptoms reported during the first visits. Then a series of timestamps for each synthetic patient were synthesized. For each timestamp of each synthetic patient, EMERGE selected the closest healthcare record from the real dataset based on weighted Euclidean and Jaccard distances. Finally, a human expert was invited to modify the care models for each synthetic patient.

Here, we provide further details regarding the timestamp synthesis utilized in the EMERGE project. It is worth noting that the goal was to generate a population-frequency distribution of visits rather than individual visit timestamps. To illustrate this strategy, let us consider an example where the EMERGE project has a real dataset of patients who visited the hospital for Viral Enteritis (ICD 008) before January 1, and 100 such patients were generated. The EMERGE project first calculates the percentage of patients who returned to the hospital on January 2 in the real dataset, let us say 10%. Based on this percentage, 10 visiting records will be generated for the synthetic dataset, and these records will be randomly assigned to the synthetic patient population.

The data-driven approach of EMERGE has been criticized for its potential risk of patient re-identification [5]. To address this issue, knowledge-based algorithms have been proposed [59, 69, 118]. The PADARSER [35] and CorMESR [90] methods combine expert knowledge and data-driven approaches by utilizing information from public statistics, Clinical Practice Guidelines, and Health Incidence Statistics. A toolkit called Synthea [147] provides a well-engineered implementation of PADARSER. The detailed development of the care model synthesis can be found in Figure 2.

Unlike the methods discussed in Sections 4 and 5, simulation-based algorithms do not necessarily require a reference dataset from the real world. This reduces the risk of potential privacy breaches. However, these algorithms rely heavily on expert knowledge during the simulation of data generation mechanisms, leading to a significant increase in human workload. In the next section, we will introduce statistical modeling-based algorithms that do not require extensive manual guidance.

Table 2. Papers Using Single- and Multi-variate Distribution Sampling

Paper reference	Year	Distributions	Medical data applications
[87]	2008	Multinomial sampling with a dirichlet prior	Demographics (Census data)
DPCopula [81]	2014	Copula functions with differential privacy	
DPSynthesizer [82]	2014	Copula functions with differential privacy	Demographics (Census data)
COCOA [6]	2016	11 common data distributions	NaN*
[60]	2016	Copula functions	Hospital emergency population
SyntheticDataVault [109]	2016	Copula functions	NaN*

*Although these papers did not report synthesis performance on medical data, they are open sourced and easily implemented.

4 STATISTICAL MODELING

This section will introduce data synthesis algorithms, as listed in Table 2, that utilize statistical modeling and sampling strategies. A common characteristic of these algorithms is their ability to approximate attribute distributions and synthesize data by sampling. Attribute distributions can be independent (as discussed in Section 4.1), jointly attributed (using Copula functions, as discussed in Section 4.1, or SMOTE, as discussed in Section 4.2), conditionally distributed based on selected attributes (as discussed in Section 4.3), or conditioned on attribute relations (as discussed in Section 4.4). It should be noted that EHR simulation methods, as mentioned in Section 3.2, may use statistical modeling in the synthesis pipeline; however, the performance of these methods relies more on prior knowledge, such as population information and treatment workflows of diseases, rather than sophisticated modeling of attribute distributions.

We will denote the motivation as synthesized data $Y = \{y_1^T, y_2^T, \dots, y_N^T | y_i^T \in R^M, i \in [1, 2, \dots, N]\}$ from real data $X = \{x_1^T, x_2^T, \dots, x_N^T | x_i \in R^M, i \in [1, 2, \dots, N]\}$. Here x_i^T and y_i^T are attributes for both datasets. For tabular data, the indexing of all N attributes is permutable. For a series of events, the indexing follows a chronological order, where i indicates the i th events.

4.1 Sampling from Single- and Multi-variate Distributions

The simplest method for data generation is to generate each variable independently from the corresponding pre-defined distributions. These distributions can be denoted as $Pr(X_i; \theta)$, where θ is the parameter estimated from the occurrence of values in the real-world datasets and X_i is the i th attribute.

This independent attribute distribution modeling has been widely used in many data-driven clinical applications, such as the EMERGE project [14, 92] we mentioned before. Gaussian distributions are used for continuous variables, and for discrete variables, binomial distributions are used. However, the selections of distributions can be more varied. COCOA [6] is a framework for generating relational tables, and it has 11 common data distributions, including normal, beta, chi, chi-square, exponential (exp), gamma, geometric, logarithmic (log), Poisson, t -student (Tstu), and uniform (uni) [146].

The major limitation of this independent variable synthesis is that the intrinsic pattern between variables is discarded during training, leading to unmatched variables for synthetic populations. Thus, multivariate distributions of all attributes were proposed [87]. However, the multivariate distributions rely heavily on the type of distributions chosen, and for high-dimensional datasets, the computation efficiency is low, and the multivariate distributions might be sparse.

Copula functions provide a way to model the correlations between features, as well as avoid intensive parameter searching. For a two-dimensional dataset $X = \{x_1^T, x_2^T\}$, and their marginal **cumulative distribution function (CDF)** of each attribute $F_1(i) = Pr(x_1^T \leq i)$ and $F_2(i) = Pr(x_2^T \leq i)$. There exists a two-dimension Copula function C such that $F(i_1, i_2) = C(F_1(i_1), F_2(i_2))$ [126].

Thus, the multivariate distribution modeling can then be simplified into Copula function estimating and marginal computation. The DPCopula [81] and its extension DPSynthesizer [82] used Gaussian copula functions, which further disentangled the multivariate Gaussian distributions into the product of the Gaussian dependence and margins.

The Copula functions have been proved to be efficient in population synthesis [60], i.e., generating basic demographics for the target population. An open sourced Copula modeling implementation can be found in Reference [109]. However, since the Copula is defined on the CDF, the Copula function-based modeling can only be applied to continuous variables.

4.2 A Special Multi-variate Distribution Modeling: SMOTE

The statistical methods described in Section 4.1 explicitly compute the joint distributions of all variables, which can require a tricky parameter selection process. To avoid this, the SMOTE [20] was proposed to generate samples through interpolation. In SMOTE, each piece of data from an individual is treated as a point in the data space, and the distribution of data is not explicitly modeled. SMOTE approximates the distribution by assuming that the data space can be spanned by all existing data points and samples from the distribution by interpolating existing data points. SMOTE is particularly useful for addressing data imbalance problems by creating data points that belong to the minority class.

A detailed SMOTE family review can be found in Reference [43]. Being well engineered and implemented in many toolkits [80], SMOTE and its variants has proved their efficiency in the medical data analysis domain, especially for disease classification where the number of patients is much less than the number of normal controls. Applications include Parkinson Disease classification [112], Alzheimer's Disease Classification [131], and so on.

The SMOTE strategy considers all attributes together, while it fails to model the relations between attributes. Mathematically, considering x_1^T as the *label* vector for the combination of all other attributes, the synthetic samples in SMOTE only consider the conditional distribution $Pr(x_{-1}^T | x_1^T)$, rigidly inheriting the marginal of non-label attributes from real datasets by interpolation. Moreover, the interpolation nature of these methods is cursed with the mode collapse problem, where the synthetic data generated lacks diversity within the minority class.

4.3 Sampling from Conditional Distributions: Multiple Imputation

A further improvement for the attribute relation modeling is multiple imputation, shown in Table 3. Initially proposed for missing data imputation, the concept of multiple imputation proposed by Rubin [119] has also been widely used in privacy protection data releases. The main concept of multiple imputation is to produce partially synthetic datasets, where the missing attributes are predicted by other non-missing attributes. In the scenarios of privacy protection, the sensitive attributes are treated as “missing attributes”: Sensitive attributes are replaced by the values conditionally synthesized from non-sensitive attributes.

For example, let us consider dataset $X = \{x_1^T, x_2^T, \dots, x_N^T\}$ with N attributes, and the subset $X_C = \{x_1^T, x_2^T, \dots, x_C^T | C < N\}$ contains missing values (or sensitive values in privacy protection scenarios). The basic idea of multiple imputation is to sample each attribute x_i^T from conditional distributions $\{x_i^T \sim P(x_i^T | x_1^T, x_2^T, \dots, x_N^T), i \in [1, C]\}$, respectively. We will address all non-confidential (or non-missing) attributes as $X_U = \{x_{C+1}, x_{C+2}, \dots, x_N\}$ for a clear statement.

The conditional distributions can be modeled explicitly with a specific mean and an variance. For example, **General Additive Data Perturbation (GADP)** [94] was proposed to define means and variances for the conditional distributions. Later in 2003, **Information Preserving Statistical Obfuscation (IPSO)** [15] was proposed. In the IPSO, the synthetic confidential attributes were obtained by multiple regression of X_C on non-confidential attributes X_U ; and this multiple

Table 3. Papers Using Multiple Imputation

Paper reference	Year	Methods	Medical data applications
GADP [94]	1999	Defining mean and variances for the distributions of X_C conditioned on X_U	NaN
IPSO [15]	2003	General linear models for X_C from X_U	NaN
CART [16]	2010	Random forests for X_C on X_U (only applicable to discrete sensitive attributes)	Demographics (Census data)
[18]	2009	Fuzzy c-means for X_C on X_U	Demographics (Census data)
[33]	2010	Support vector machines for X_C on X_U	Health insurances data
[46]	2020	MICE	Cancer registry data from the Surveillance Epidemiology and End Results program
PeGS [107]	2013	General linear models with differential privacy for X_C from X_U (only applicable to discrete sensitive attributes)	Public Patient Discharge Data from California Office of Statewide Health Planning and Development
PeGS applications [108]	2013	General linear models with differential privacy for X_C from X_U (only applicable to discrete sensitive attributes)	Public-use data files from Centers for Medicare and Medicaid Services

regression model, also known as general linear model, has then been improved by regression trees [16] and fuzzy c-regression [18]. Thus, after IPSO, the multiple imputation algorithm using SVM [33] allowed the synthesis of discrete variables (but it can only be applied to discrete variables).

Multiple Imputation by Chained Equations (MICE) [152] also used regression models to synthesize sensitive attributes from nonsensitive attributes, but it is also featured by an iterative synthesis strategy, and the final synthesis results were pooled by all results synthesized during iterations. Applications can be found in breast cancer data synthesis [46].

Another extension of multiple imputation methods is to improve the privacy-protection during data release. The multiple imputation does not protect data privacy by nature. PeGS [107] then introduced the differential privacy concept in multiple imputation algorithms, and an application of their algorithm on healthcare data can be found in Reference [108].

Multiple imputation algorithms modelled the condition of missing attributes (or sensitive attributes) on existing attributes (or all attributes). However, for fully synthetic data, one needs to traverse all variables to investigate the cross-conditional distributions, which is tedious and time-consuming.

4.4 Sampling from Conditional Distributions with Attribute Relationships: Probabilistic Graphical Model

To model the relations among attributes, a **probabilistic graphical model (PGM)** can be used. The edges for the PGM are relationships between attributes, while each node represents a

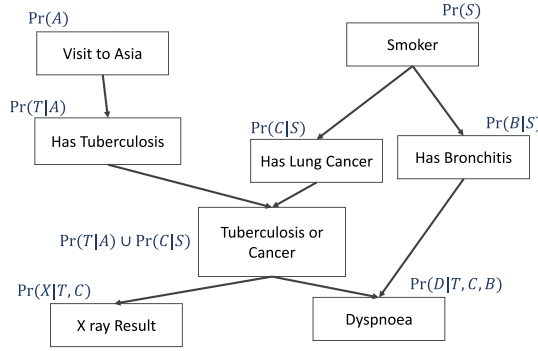


Fig. 3. An example for Bayesian networks.

conditional distribution of one attribute. A Bayesian network, also known as a Bayesian belief network, is a graphical model that represents the joint probability distribution of variables of interest in a directed acyclic graph. The directions of the edges in the graph indicate causal relations between variables. Figure 3 presents a popular healthcare Bayesian network, the Asia Network, which was defined in 1988 [77]. The Asia Network assumes a Bernoulli distribution for all attributes. To synthesize a patient's data from the Asia Network, one typically starts by sampling from $Pr(S)$ and $Pr(A)$ as the root nodes and then obtains the values of other attributes from their conditional distributions, as shown in Figure 3. Bayesian networks can also be used to infer the conditional distributions of variables given other variables. In summary, Bayesian networks provide a graphical representation of causal relations among variables and enable efficient calculation of the conditional distributions of each variable [77].

Three steps are required to use Bayesian network to generate synthetic samples:

- First is structure learning, i.e., identifying the causal relations between attributes.
- Second is parameter learning, i.e., learning the conditional distribution of each variable.
- The third step is the inference: Values of each attribute are first sampled from sets of initial attributes and propagated to values in other attributes according to conditional distributions.

It should be noted that the structure and parameter learning of Bayesian networks can always be knowledge driven, i.e., human experts can pre-define the relations between each attribute by their experiences. For data-driven algorithms, the structure and parameters are learned simultaneously. Data-driven methods for constructing a Bayesian network include the following:

- Constraint-based algorithms. In these algorithms, conditional independence tests were used to evaluate the dependency between each pair of attributes, and a BN was constructed using related attribute pairs.
- Score-based algorithms. The score-based algorithms first searched all possible structures and then used a score function [75, 123] to evaluate these graphs.
- Hybrid algorithms. These algorithms used constraint-based algorithms to generate a subspace of all possible structures and then used score-based algorithms to evaluate and select these graphs.

We summarized all Bayesian network-based algorithms in Table 4. Bayesian networks have been widely used in medical data synthesis, because the inherent suitability of Bayesian networks for knowledge-driven structure design and parameter learning enable improved trustworthiness of medical experts in this kind of models.

Table 4. Bayesian Network-based Algorithms for Data Synthesis

Paper reference	Year	Structural and parameter learning	Inference	Medical data applications
[134]	2015	Score-based (tabu search by Python Package bnlearn [137])	Global sampling	Demographics
PrivBayes [167]	2017	Constraint-based (Mutual Information and differential privacy)	Global sampling	NaN*
DataSynthesizer [111]	2017	PrivBayes	Global sampling	NaN*
[28]	2020	Score-based (AIC by Python Package pomegranate [122])	Global sampling	Demographics
[140]	2020	Constraint-based (FCI with EM for missing data)	Global sampling	CPRD Aurum data synthesis
[70]	2021	Score-based (by Python Package bnlearn [137])	Heart Disease (UCI), Diabetes datasets (UCI), MIMIC-III	
[86]	2021	Constraint-based (G^2 -test)	Global sampling from the label attribute	Breast cancer (UCI), Diabetes (UCI)
PrivSyn [168]	2021	Constraint-based (Independent Difference (InDif for short))	Gradually Update Method	NaN*

*Although these papers did not report synthesis performance on medical data, they are open sourced and easily implemented.

5 DEEP LEARNING

Two types of deep neural networks (also referred to as Deep Learning) have been widely used in data synthesis: **Auto-Encoder (AE)** [72] and **generative adversarial networks (GAN)** [49]. They are all composed of stacked linear or non-linear functions, and the major difference between these two types of methods is their target functions. An AE usually has two basic components, an encoder \mathcal{E} that maps vectors in the data space into a latent space and a decoder \mathcal{D} that maps the latent space features into the data space. Mathematically, the objective function for an AE with an input real data x is defined as

$$L(X, \mathcal{E}, \mathcal{D}) = \|\mathcal{D}(\mathcal{E}(x)) - x\|_p, \quad (1)$$

where $\|\cdot\|_p$ is the p-norm. Synthetic data can be generated from AE by first sampling vectors from the latent space and then mapping the sampled vectors into the data space.

The GAN method, however, uses an additional network, discriminator, to optimize the performance of the data synthesizer. A GAN is also composed of two components: a generator \mathcal{G} and a discriminator \mathcal{D} . The inputs for generators are usually noises z , which are transformed into meaningful data vectors by deep learning models and can improve the variety of synthetic data. The inputs for the discriminator are both synthetic data from generators $\mathcal{G}(z)$ and real data x for reference. The objective function of a GAN is

$$L(X, \mathcal{G}, \mathcal{D}) = E_x[\log(\mathcal{D}(x))] + E_z[\log(1 - \mathcal{D}(\mathcal{G}(z)))], \quad (2)$$

where E_x and E_z are expectations over all data instances. The generator can have many variants [4, 132], and some researches even use AE as an generator [135, 139], and the loss functions are then a combined value of both Equations (1) and (2).

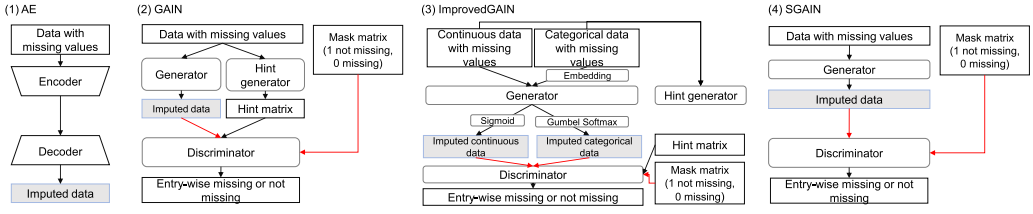


Fig. 4. Architectures for tabular data imputation.

In this section, we categorize deep learning-based algorithms based on their target data types, considering the shared characteristics of AEs and GANs. While **differential privacy (DP)** and fairness techniques are will be discussed separately to emphasize their unique capabilities. Additionally, we have included diagrams of these models to aid in understanding their structures, with red arrows highlighting inputs to the discriminator and gray boxes with blue boundaries representing the target synthesis object.

5.1 Deep Neural Networks for Tabular Data

According to the output of deep neural networks, we further divided the tabular data deep learning models into two categories: half synthesis networks whose target is to impute values and fully synthesis networks.

5.1.1 Half Synthesis Networks. We plotted the structures of half synthesis networks in Figure 4. A use case of AE in clinical data imputation can be found in Reference [10]. **Denoising autoencoders (DAE)** [144] was proposed to extract robust feature representations, but its structure has also been used in missing data imputation. During training, the DAE first set several elements of inputs to zero randomly and then was trained to reconstruct the values of these elements. Once well trained, the DAE could be used for missing data imputation. The objective function for DAE is the **mean squared error (MSE)** between the corrupted input and the corresponding ground truth. Thus, the training of DAE is fully supervised and requires a large scale of complete ground truths. To improve the fully supervised training of DAE, **Generative Adversarial Imputation Nets (GAIN)** [164] used a discriminator to optimize the synthesis performance during training. In data imputation tasks, the discriminators identify the real and fake elements in the imputed data. This imputation discriminator is similar to the patch discriminator [30] used in image synthesis. Instead of outputting a single binary value indicating whether the entire synthesized vector is real or fake, the discriminator produces a vector indicating the confidence of each value in the synthesized data. This “patchwise” discriminator has been widely used in imputation GAN models.

The GAIN models were further improved by GAN training tricks such as Gradient Penalty and Wasserstein Loss in SGAIN [98]. To further specify the imputation of categorical and numerical missing variables, improved GAIN [17] split these two kinds into variables and imputes them separately.

5.1.2 Fully Synthesis Networks. For the fully synthesis networks, we will discuss these networks according to their architectures.

Encoder–decoder or not: A debate on tabular data synthesis. AE models, which are featured by the encoder–decoder structure, can be applied to tabular data synthesis. Synthetic data can be generated by manipulating the hidden feature vectors derived from real data. Examples can be found in AE-ELM [53] and OVAE [142], which used VAE to model the distributions of latent feature vectors.

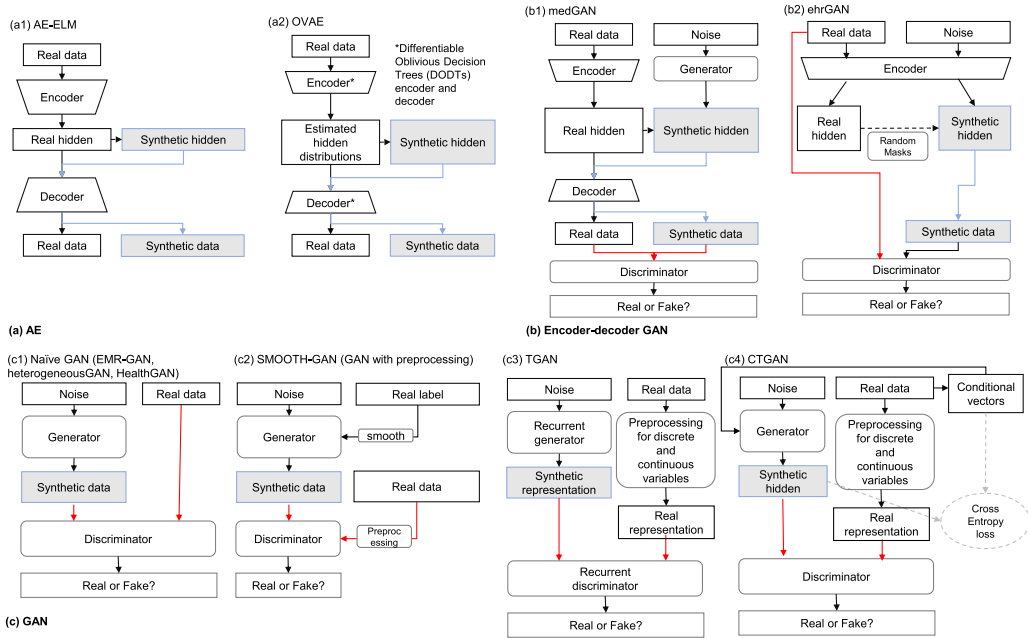


Fig. 5. Architectures for tabular data synthesis, grouped by the usage of an encoder-decoder structure.

As for the GAN models, we discovered a long-lasting debate over the object to be generated in the tabular GAN generators: Some algorithms [21, 26, 127] used an encoder-decoder structure to map the original data into a latent space, while others generated the vectors of data directly [170]. We plotted their architectures in Figure 5.

Algorithms that used encoder-decoders to map feature vectors into a latent space tried to avoid synthesizing tabular data directly [21, 26]. For example, medGAN [26] used an encoder network that mapped the original values into a hidden space, and the generator generated values in the hidden space instead of in the original space. The encoder-decoder in medGAN was also pre-trained on real datasets, and the target for the pre-training was to reconstruct the input real datasets. EhrGAN [21] used the encoder-decoder structure to map a transition distribution of the form $P(\tilde{x}|x)$, where \tilde{x} and x are synthetic and real data.

However, EMR-GAN [170] has a conflicting conclusion claiming that the direct synthesis of values works better than the synthesis of hidden feature maps. They argue that because “these GANs (medGAN [26]) rely on an autoencoder, they may be led to a biased model, because noise is introduced into the learning process.” SPRINT-GAN [11], heterogeneousGAN [160], healthGAN [158], and SMOOTH-GAN [117] also synthesized data values directly.

TGAN [157] also discovered that “Simply normalizing numerical feature to $[-1, 1]$ and using tanh activation to generate these (numerical) features does not work well.” However, TGAN did not turn to latent space synthesis and did not use encoder-decoder structures. TGAN introduced a statistical representation synthesis strategy known as mode-specific normalization. Instead of synthesizing numerical values directly, the authors of TGAN used a Gaussian Mixture Model with m Gaussian distributions to model each feature, and they synthesized the parameters for GMMs. In their improved CTGAN [156], they replaced the **Long Short-Term Memory (LSTM)** generator in TGAN and used a conditional synthesis strategy to model the categorical features. Other pre-processing algorithms can be found at smoothGAN [117], where continuous data were pre-processed by deleting outliers and scaling, and discrete data were mapped into continuous scores.

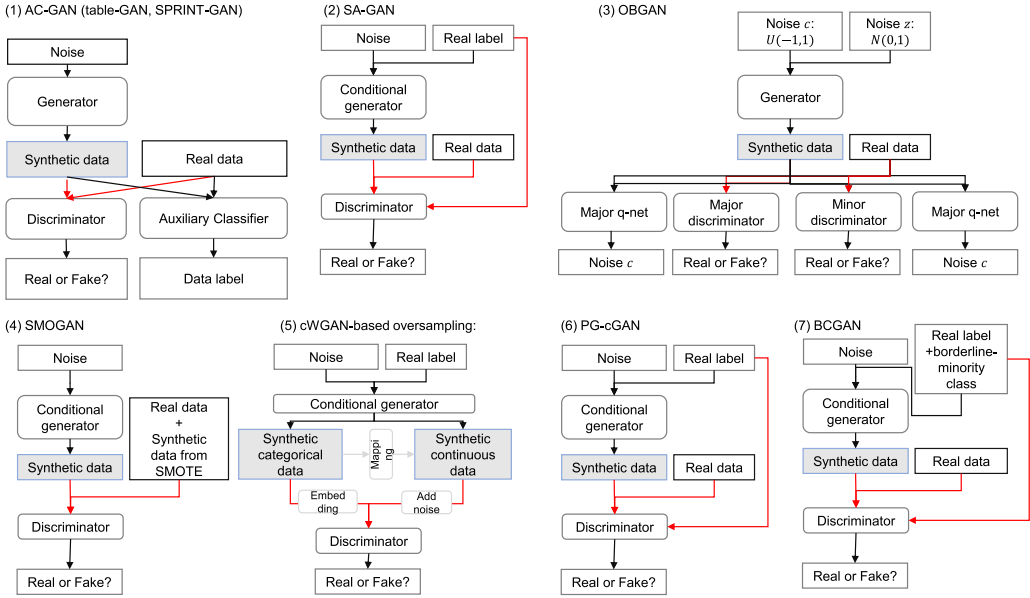


Fig. 6. Architectures of tabular data synthesis with labels.

Synthesis with labels. For downstream prediction tasks, the labels of data should be synthesized alongside the data. ACGAN [101] used an auxiliary classifier to synthesize data labels, and has successful implementations as SPRINT-GAN [11] and table-GAN [106].

Some medical data synthesis applications aim to synthesize data within a target group, e.g., synthesizing data with the minority class. For minority class data synthesis, the conditional generator has been widely used. BCGAN [130] and OBGAN [63] synthesized data points that were close to the decision boundary, and the former algorithm achieved borderline synthesis by introducing an additional borderline minority class; the latter used the Q-Net concept from InfoGAN [24] to allow output editing in GAN models. SMOGAN [76] used a SMOTE algorithm before GAN to augment the data points for GAN training. SAGAN [32] brought the relation between single attributes and data labels to GAN models. To better analyze discrete variables, cWGAN-based oversampling [38] adjusted the GAN model and embedded the discrete attributes.

In addition to network architectures as shown in Figure 6, different loss functions have also been introduced in deep generative models. ADS-GAN [162] introduced the indetifiability loss, which maximizes the Euclidean distance between real and synthetic data. HealthGAN adopted the Wasserstein loss [4] during the training of GAN models.

5.2 Deep Neural Networks for Sequential Data

For sequential data, recurrent structures have been utilized. The recurrent structures have many variants, such as **Recurrent Neural Networks (RNN)** [120], Gated Recurrent Unit [25], and LSTM [57], but they share a same basic structure shown in Figure 7. Here, this recurrent mapping maps input $\{x_1, x_2, \dots, x_T\}$ into an output $\{y_1, y_2, \dots, y_T\}$. Each cell in the recurrent mapping receives two inputs: one from the present x_T and one from the latest past h_{T-1} . For the first cell, the h_0 is addressed as the initial state, which, for many RNN implementations [115], are assigned by a vector of zeros. Each cell in the recurrent mapping outputs two outputs, but in some algorithms, such as RNN and GRU, $h_T = o_T$.

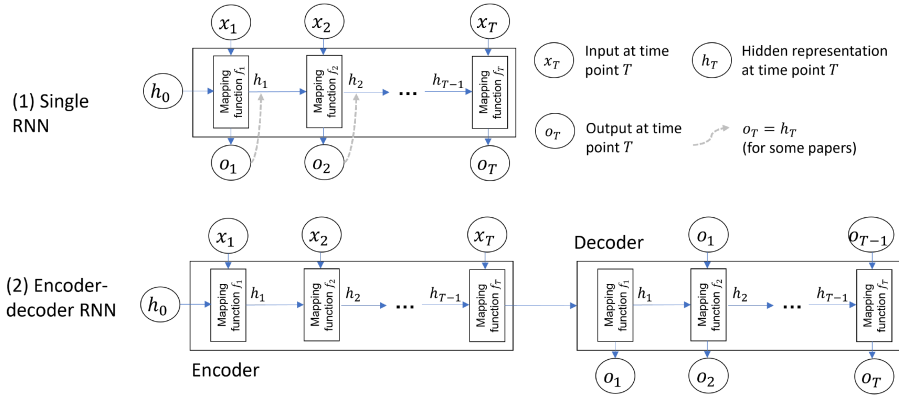


Fig. 7. Architecture of a recurrent network (1) and an encoder-decoder recurrent network (2).

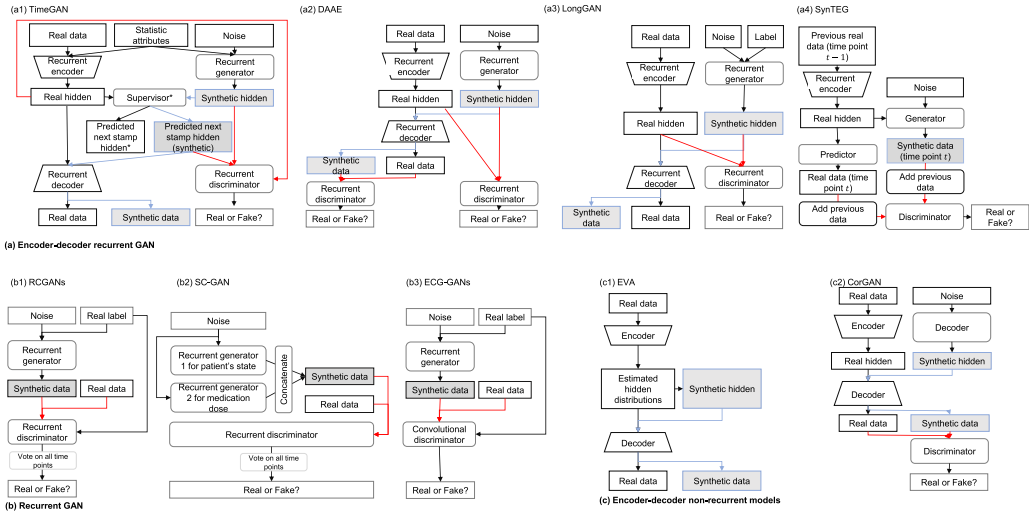


Fig. 8. Architectures of deep neural networks for sequential data, grouped by the use of encoder-decoder structure and recurrent structure. *An additional MSE loss is added between real $[0, t-1]$ and predicted next stamp hidden $[1, t]$.

There is an implied convention of using encoder-decoder structures in sequential synthesis models [97]. As shown in Figure 7, the encoder-decoder structure in sequential synthesis allows the network to read all time points before producing outputs. The decoupling of data reading and data generating has been commonly practiced in sequential data synthesis. Thus, we can easily find successful applications of encoder-decoder recurrent structures for sequential EHR generation, including TimeGAN [163], DAAE [78], LongGAN [135], and SynTEG [169], as is shown in Figure 8.

However, some algorithms break this convention. We noticed that, in sequential EHR synthesis, RCGAN [40] and SC-GAN [148] generate data without the encoder-decoder structures. In addition, for ECG synthesis, we found another recurrent synthesis application without encoder-decoder structures [29]. These algorithms generate sequences directly from noises. However, these papers did not claim that their direct synthesis is better than latent synthesis.

In addition to recurrent networks, convolutional operations have also been used to investigate and preserve the inner correlations among timestamps, and papers including EVA [13] and CorGAN [139] used one-dimensional convolutions on temporal datasets. The experiments in ECG-GAN [29] also demonstrated the best synthesis performance using a recurrent generator and a non-recurrent discriminator. Despite the fact that they do not use recurrent structures, the EVA and CorGAN also decouple data reading and data generating with encoder–decoder structures.

Here, we elaborate on loss functions improvements for sequential data synthesis. In addition to the network architectures, TimeGAN introduced a supervised loss in the training of recurrent GAN models, where the temporal relationships between timestamps were used for supervised criteria. The recurrent supervisor in TimeGAN was proposed to “explicitly encouraging the model to capture the stepwise conditional distributions in the data.” Practically, TimeGAN used an additional recurrent network named supervisor. The recurrent supervisor received the latent feature extracted from the real data $\{h_1, h_2, \dots, h_t\}$ (real hidden in Figure 8) and outputted the latent features in the next timestamps $\{h'_2, h'_3, \dots, h'_{t+1}\}$. The loss function of the supervisor is thus

$$L_{\text{MSE}} = \sum_{i=0}^t ||h_i - h'_i||^2. \quad (3)$$

This inner sequential supervision proved its efficacy in medical data synthesis using a large private lung cancer pathways dataset.

5.3 Deep Neural Networks with Additional Targets

In this section, we will discuss deep neural networks with additional targets, such as differential privacy and fairness. The concept of differential privacy provides a well-defined solution to data privacy protection, and detailed definitions of differential privacy will be elaborated on in Section 6.3.3. To incorporate the concept of DP into deep learning algorithms, **Differentially Private Stochastic Gradient Descent (DP-SGD)** [1] was proposed in 2016, which adds noises in the gradient during training stages. In 2018, DPGAN [153] was introduced as the first GAN model incorporating differential privacy. It also adopts the noisy gradient strategy as in DP-SGD. DP-AuGM and DP-VAEGM [23] also use the DP-gradient descent strategy during training and propose to use AE- and VAE-based strategies for data synthesis.

PATE-GAN [67] incorporates the concept of DP into GAN models by using a **private aggregation of teacher ensembles (PATE)** mechanism during the training of discriminators to synthesize data according to the level of DP. PATE-GAN uses n teacher discriminators and splits the real data into n subsets, with each teacher discriminator only discriminating between synthetic data and its corresponding subset of real data. Additionally, PATE-GAN implements a student discriminator that does not rely on any public data and only receives synthetic data as inputs, with the labels of these data assigned by the teacher discriminators. Successful applications of PATE-GAN include implementations on the Kaggle cervical cancer dataset [42], UCI ISOLET dataset [34], and UCI Epileptic Seizure Recognition dataset [34].

FairGAN [154] and FairGAN+ [155] aim to improve data fairness by using conditional GANs. In FairGAN, an additional discriminator is used to minimize attribute disclosure, i.e., to minimize the predictive ability of nonsensitive attributes on sensitive attributes. FairGAN+ then introduces classification fairness. We include the architecture mentioned in this subsection in Figure 9. Although fairness is an important quality required in a trustworthy training dataset, these fairness generative models have not yet been applied to healthcare data.

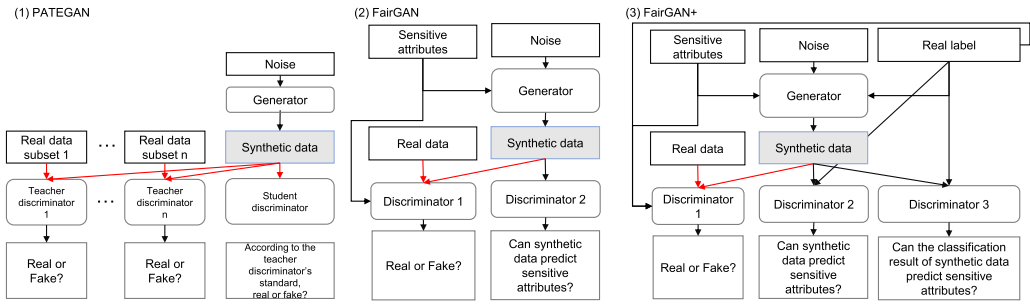


Fig. 9. Architectures of deep neural networks for differential privacy and fairness.

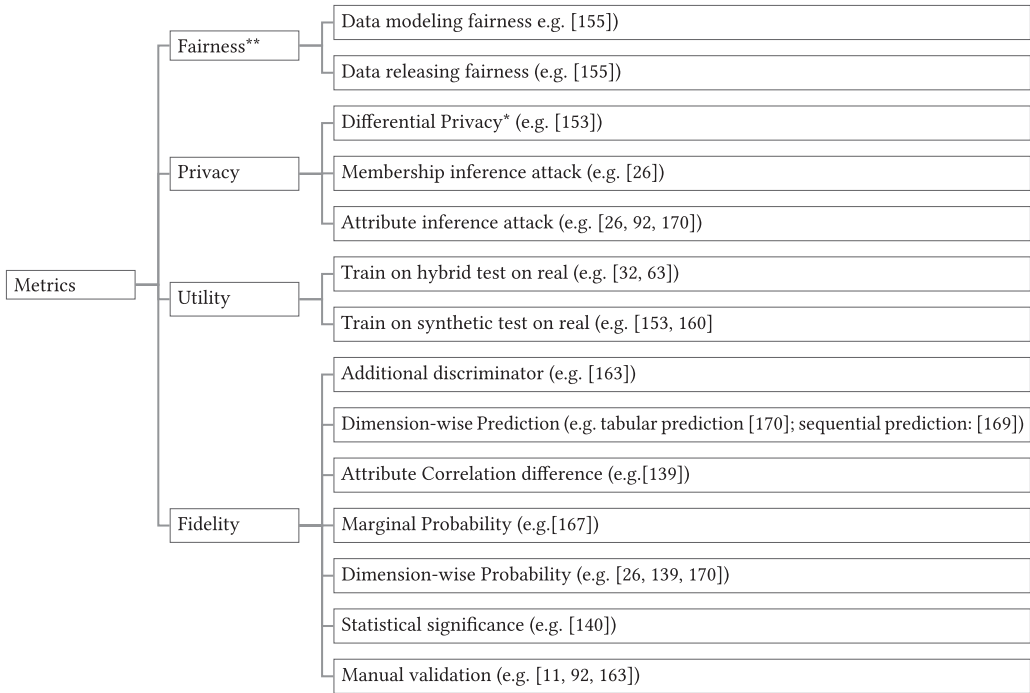


Fig. 10. Tree diagram of evaluation metrics. *It should be noted that differential privacy is not strictly a metric by definition. **Fairness is not commonly evaluated in medical data synthesis, while we still introduce this metric due to its high correlation with trustworthiness.

6 METRICS

We have categorized the evaluation metrics of the selected papers into four classes, as shown in Figure 10. Importantly, the evaluation methods outlined in this are applicable to all data synthesis algorithms. These evaluation criteria are independent of the specific algorithm used and only require synthetic data and real reference data as input for assessment. All of these qualities contribute to the trustworthiness of the data. High-quality synthetic data are a key component of achieving trustworthy AI.

6.1 Fidelity

The fidelity of synthetic data can be evaluated by a panel of experts, such as clinicians who work closely with real data, who can inspect the synthetic data. For instance, SPRINT-GAN [11] employed ACGAN to generate a group of synthetic sequential health data, including blood pressure and medication counts, and then mixed real and synthetic patients. They invited three experienced physicians to determine whether the sequential data were real. EMERGE [92] invited a medical expert to review the synthetic healthcare records to identify any records that had content problems or inconsistencies, such as the disease category not matching the symptoms of synthetic patients. To help clinicians investigate data fidelity further, dimensionality reduction algorithms were used to visualize the data distribution. TimeGAN [163] used the T-SNE [56] algorithm to visualize the distributions of real and synthetic data points.

In addition to manual validation, the statistical closeness between real and synthetic data can also be quantitatively measured to assess synthesis performance. Statistical closeness can be measured in several ways:

Statistical significance. Two of the most commonly used textbook measurement of the similarity between two datasets in medical data synthesis are the chi-squared test and the two-sample **Kolmogorov–Smirnov (KS)** test. The chi-squared test is commonly used for discrete variables. In this test, a p value is obtained indicating the significance of the test, and, commonly, a $p > 0.05$ indicates that there is no significant difference between the real and synthetic variables [136]. However, in the KS test, the significance is measured under a specific confidence level, and each confidence level has a unique threshold for the statistic. One can claim that the synthetic data have similar distributions to the real data [88] if the statistic is lower than the threshold.

Dimension-wise Probability. The **Kullback–Leibler (KL)** divergence [73] can be used measure the distribution difference between real attributes and synthetic attributes. For discrete variables, the KL divergence from the synthetic distribution Q to the real distribution P is defined as

$$D_{KL}(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right), \quad (4)$$

where the \mathcal{X} denotes all possible values in both datasets. For continuous variables, the KL divergence from Q to P is defined by an integral:

$$D_{KL}(P||Q) = \int_{-\infty}^{\infty} p(x) \log \left(\frac{p(x)}{q(x)} \right) dx, \quad (5)$$

where p and q are the probability densities of P and Q .

Especially for the ICD9 embedded EHR data, Bernoulli success possibility of each dimension can be used for a dimension-wise probability metric (an example can be found in CorGAN [139]). For each ICD9 code, the “success” in the Bernoulli trail is defined by the frequency of this code in all time points from all patients. The frequency for each ICD9 code should be similar between the real and the synthetic dataset.

Marginal Probability. Dimension-wise probability measures the distribution difference on individual attributes, while the marginal probability measures the joint probabilities among attributes. A k -way marginal is associated with a subset of k attributes. For example, $Pr(x_1, x_2)$ is a two-way marginal. The KL divergence can also be used to measure the difference in k -way marginal distributions. In PrivBayes [167], the total variation distance between the synthetic and real marginals is used to evaluate the marginal distribution. Mathematically, the total variation distance between the synthetic distribution Q and the real distribution P is calculated as

$$\mu(P, Q) = \sup_{\mathcal{X}} |P - Q|, \quad (6)$$

where the supremum runs over all possible values or finite grids (for continuous variables) [113]. As with KL divergence, a lower distance indicates a better synthesis performance.

Attribute Correlation difference. The correlation between attributes is also an important feature of datasets. Preserving the correlation among attributes ensures the clinical utility of the synthetic dataset. To evaluate the preservation of feature correlation, TGAN [157] computed the mutual information between each pair of feature columns for single tables. Pearson’s correlation [139] is another function that measures the intrinsic pattern among features. However, it does not provide a qualitative comparison of the correlation differences.

Dimension-wise Prediction. This evaluation also measures how well synthesis models capture the inner correlations in real datasets. First, a random attribute $x_{(k,real)}$ is selected from the real dataset, and the remaining attributes $x_{(\neg k,real)}$ in the real dataset and corresponding synthetic data $x_{(\neg k,syn)}$ are used to train two classifiers for predicting the value of the selected attribute. If the synthesis model captures the correlation between attributes well, then the performance of these two classifiers should be similar. In the case of sequential EHR data, the value for the last time point is often chosen as the target attribute [78].

Additional discriminator. Some articles [78, 163] introduced an additional discriminator to classify synthetic data from real data. Different from the discriminator in GAN models, this additional discriminator is not optimized alongside the generator. A lower accuracy of the additional discriminator indicates a better synthesis performance.

6.2 Utility

The utility of synthetic datasets is measured by the performance of synthetic datasets on downstream tasks. Usually, a label is assigned to each record in the datasets. For privacy preserved data synthesis, the Train on Synthetic Test on Real strategy has been widely used. A good synthetic dataset should achieve similar performance as the real dataset, and this classification similarity can be measured by KS tests. For data imputation and data augmentation, the utility is measured by classification improvement. The imputed data and augmented data should be able to improve the classification accuracy compared to the original dataset.

6.3 Privacy

Rubin proposed the use of fully synthetic data for privacy-preserved data sharing in 1993 [119]. However, it should be noted that synthetic data are not inherently private [66]. Therefore, if a “privacy-preserved” synthesis algorithm is proposed, then it is recommended to conduct either an empirical analysis of privacy breach risks or an analytic proof of privacy. Of all the papers reviewed, three types of metrics have been used to evaluate the ability of privacy preservation, namely (1) attribute inference attack, which involves obtaining specific values and statistical properties of the dataset; (2) membership inference attack, which aims to identify the presence of a specific record; and (3) differential privacy, which provides analytic proofs of privacy.

6.3.1 Attribute Inference Attack. The attribute inference attack assumes that the attacker has a compromised dataset (values of some non-sensitive attributes) and uses these attributes, along with the synthetic dataset, to speculate the sensitive attributes [89]. Most algorithms use a similarity attack strategy [26, 92, 170]. First, for each individual in the real dataset, n attributes are randomly chosen and provided to the attacker. Next, for each record in the compromised dataset, the attacker computes the similarity of its publicized n attributes with all records in the synthetic dataset. The similarity measurements include mean absolute value difference and the Euclidean distance [15]. Finally, the attacker obtains the sensitive attributes of the most similar synthetic record as the values of sensitive attributes of this real individual.

Specifically, as most EHR synthesis relies on ICD9 codes, the data typically consist of discrete values. To preserve privacy, several studies [26, 92, 170] employed k-nearest neighbor classifications on the synthetic dataset for each sensitive attribute. By applying the classifiers on the compromised dataset, they could obtain the discrete values for the sensitive attributes. The classification accuracy metrics for each sensitive attribute were then reported, with lower accuracies indicating a higher level of privacy preservation.

6.3.2 Membership Inference Attack. The membership inference attack assumes that the attacker only has access to the data and not to the generative models. The attacker obtains a set of complete records where all attributes are publicized. By observing the synthetic dataset S , the attacker will determine whether a given data record in P was part of the synthetic model's training dataset [58]. Most membership inference attacks can be viewed as a classification task where the goal is to classify whether each record in P is 0 (not present in the training dataset) or 1 (present in the training dataset).

Most data generative models use a metric-based membership inference attack and perform it in an unsupervised manner. These attacks presume that synthetic records must bear similarities with the records that were used to generate them [22]. Using different distance metrics, the similarity between the target record p_i and all records from S can be measured. If the mean similarity between this record and all synthetic records is below a specific threshold, then this record is considered as 1 (presented in the training dataset). The metrics include Hamming distance (for discrete variables) [26] and the Euclidean distance (for continuous variables).

In 2017, a shadow model-based membership inference attack was proposed by Shokri et al. [124], which was later adapted for generative models by Stadler et al. [133]. For each record p_i , the shadow model attack trains two generative models: one with p_i and one without p_i . The synthetic records generated from the model trained with p_i are assigned the label 1, while the synthetic data from the model without p_i are assigned the label 0. Then, a classifier is trained on the synthetic records and their corresponding labels, and by inferring this classifier with the publicly available synthetic records S , the attacker can identify the presence of p_i in the training dataset.

6.3.3 Differential Privacy. Differential privacy is an important solution in the context of membership inference. The term "differential" here refers to preserving patient privacy by tracking the difference between a released dataset and a modified version of it.

Suppose we have a dataset D with sensitive information about patients, including their age, gender, and medical conditions. The dataset contains four males. Now suppose we add a new patient to the dataset, resulting in a new version called D' . If an attacker gains access to both versions of the dataset, then they could easily figure out information about the new patient by comparing the two datasets. For example, they might notice that there is now one more male in the dataset, which could help them guess the new patient's medical condition. Differential privacy aims to protect individual privacy in such situations by blurring the true values of the released data. In this case, DP would blur the gender information by adding random noise instead of revealing the exact number of males in the dataset.

Mathematically, assuming a function K where attackers can only obtain the information by applying K to D , the goal of DP is to minimize the difference between the distributions of $K(D)$ and $K(D')$.

As mentioned previously, KL-Divergence is used to measure distribution differences. Since the concept of DP is to measure the maximum privacy breach, the maximum divergence is used in [37]. Thus, the target of DP algorithms can be mathematically defined as follows: An algorithm K provides ϵ -differential privacy if the maximum divergence between the distributions of $K(D)$ and

$K(D')$ is bounded by ϵ , i.e.,

$$\text{Div}_\infty(K(D)||K(D')) = \max_{d \in D} \left[\ln \frac{\Pr[K(D) \in S]}{\Pr[K(D') \in S]} \right] \leq \epsilon. \quad (7)$$

More commonly, this definition is written as

$$\Pr[K(D) \in S] \leq e^\epsilon \Pr[K(D') \in S] \quad (8)$$

if K gives ϵ -differential privacy. The DP can be achieved by adding specific patterns of noises when releasing the dataset, i.e.,

$$K(D) = D + \text{Noise}. \quad (9)$$

The distributions of noises include Laplacian mechanism [36] and the exponential mechanism [91]. For numerical data, the Laplacian mechanism outputs synthetic data with noises from a Laplacian distribution. For categorical data, the exponential mechanism introduces a scoring function and produces the possibility of each value.

Unlike the two types of inference attacks mentioned before, differential privacy does not provide a typical “evaluation metric,” while it provides analytic proof of privacy in models, which means that DP can provide a guarantee for privacy protection. Thus, DP is, at most of the time, achieved by introducing the DP mechanisms into the synthesis algorithm. For example, DP-SGD [1] introduced the concept of DP in the SGD optimizer in deep learning, and Reference [167] introduced DP in Bayesian networks.

6.4 Fairness

Fairness is a crucial aspect of trustworthy AI. A fair dataset should prevent exacerbating the differential impact among different groups, particularly among “protected groups”—a category of individuals protected by law, policy, or similar authority [125]. In the medical domain, fairness in datasets ensures equitable healthcare access and outcomes for people of different races and genders.

Considering the patients’ identities as sensitive attributes, a fair dataset [155] should accomplish two goals: (1) not reveal the sensitive attributes through insensitive attributes (data releasing fairness) and (2) avoid biased downstream predictions with respect to the sensitive attributes (data modeling fairness).

Considering a dataset composed of three components $D = X_U, X_C, Y$, where X_U, X_C, Y are insensitive attributes, sensitive attributes (race or sex), and data labeling, respectively. For data releasing fairness, two measurements can be used as follows:

Risk difference (RD) [155]. The RD for data releasing is defined by

$$RD_r = \Pr(Y = 1|X_C = 1) - \Pr(Y = 1|X_C = 0). \quad (10)$$

Balanced error rate (BER) [41]. A trust model $f : X_U \rightarrow X_C$ is built to predict sensitive variables X_C from insensitive variables X_U [41]. The BER of f is defined as

$$\text{BER}(f(X_U), X_C) = \Pr[f(X_U) = 0|X_C = 1] + \Pr[f(X_U) = 1|X_C = 0]. \quad (11)$$

According to BER, a synthetic dataset (X_U, X_C) is ϵ -fair if for any trust models,

$$\text{BER}(f(X_U), X_C) > \epsilon. \quad (12)$$

For the measurement of data modeling fairness, a classification model $\eta : X_U \rightarrow Y$ is built to predict data labels Y from insensitive variables X_U . The data modeling fairness requires that the prediction of η is unbiased with respect to X_C . Mathematically, three metrics are defined to measure the data modeling fairness

(a) A single table


Patient Name	Sex	Age/Y	Temperature/°C	Smoking History
Ann	Female	27	37.8	Y
Mary	Female	20	NaN	N
Joe	Male	21	40.1	N
...

(b) Relational tables

Patient Name	Sex	Age/Y
Ann	Female	27
Mary	Female	20
Joe	Male	21
...

Patient Name	Cancer Grade	Surgery	White cell count/ 10^9ml^{-1}
Ann	III	True	4.1
Joe	V	False	5.1
...

(c) Medical Signals



(d) Time series of events

Patient Name	Sex	Age/Y	Smoking History	White cell count/ 10^9ml^{-1}	Force Vital Capacity/ml	X-ray grading
Ann	Female	27	Y	10.1	1700	Severe

Username	Device	Exercise	Time/min	Heart rate/ min^{-1}
Ann	Watch	Running	21	80
Ann	Watch	HIIT	10	NaN
Ann	Phone	HIIT	4	112
...

Medication	Dose
NSAID	2

Time point 0: Patient admission Time point 1: 1st laboratory test Time point 2: 1st Medication Time point 3: Patient Discharge

Fig. 11. Three major non-imaging medical data identified by this review article: tabular data ((a) and (b)), medical signals (c), and time series of events (d).

RD [155]. The RD for modeling, which is also known as demographic parity, is defined by

$$RD_m = Pr(\eta(X_U) = 1 | X_C = 1) - Pr((\eta(X_U) = 1 | X_C = 0)). \quad (13)$$

Odds difference (OD) [155]. The equality of odds requires the classifier to have equal true positive rates and equal false positive rates between two subgroups $X_C = 1$ and $X_C = 0$. Mathematically, the odds difference is defined by

$$OD_m = \sum_{y \in \{0,1\}} Pr(\eta(X_U) = 1 | Y = y, X_C = 1) - Pr((\eta(X_U) = 1 | Y = y, X_C = 0)). \quad (14)$$

7 COMMON PRE-PROCESSING PRACTICES AND DATASETS

In this section, we will provide common pre-processing practices for different data types as in (Figure 11), and we will also provide open sourced datasets in Table 5 available for the synthesis algorithms development. The datasets we included in this survey are open access for research purposes. However, due to the high sensitiveness of healthcare data, accesses to these dataset may require formal inquiries. In addition, we will also share three released synthetic datasets whose synthesis procedures produce hands-on experiences for data synthesis practitioners.

7.1 Pre-processing Methods for Tabular Data

For tabular data, each row of the table represents a patient, and columns of the table are features describing the patient. Tabular data are straightforward for data analysis, and the statistical properties, such as mean values and standard deviations, of the population can be derived. A sub-type of tabular data is multiple-tables, or relational tables, where information from different sites and under different levels is linked with one column from the tables. It should be noted that these relational tables can be merged into a meta table by a unique combination of linkage variables.

Table 5. Open Sourced Datasets Used for Non-imaging Medical Data Synthesis

Dataset name	Patient number	Data type	Data information	Disease Category
MIMIC-I (or MIMIC) [93]	100	Medical signals and sequential EHR	Patient monitor data, patient-descriptive data (gender, age, record duration), symptoms, fluid balance, diagnoses, progress notes, medications, and laboratory results	Potential hemodynamically unstable
MIMIC-II [121]	33,000	Medical signals and sequential EHR	Patient monitor data, patient-descriptive data (demographics, admissions, transfers, discharge times, dates of death), diagnoses, notes, reports, procedure data, medications, fluid balances, and laboratory test data	diseases of the circulatory system; trauma; diseases of the digestive system; pulmonary diseases; infectious diseases; and neoplasms
MIMIC-III [65]	46,520	Medical signals and sequential EHR	Patient monitor data, patient-descriptive data, diagnoses, reports, notes, interventions, medications, and laboratory tests data.	Diseases of the circulatory system, pulmonary diseases, infectious and parasitic diseases, diseases of the digestive system, diseases of the genitourinary system, neoplasms, diseases of the genitourinary system, and trauma
MIMIC-IV [64]	383,220	Medical signals and sequential EHR	Hosp module contains patient-descriptive data, basic health data (blood pressure, height, weight...), medication, procedure data, and diagnoses. Icu module contains timing information data, patient monitor data, fluid balance, and procedure data.	Diseases of the circulatory system, pulmonary diseases, infectious and parasitic diseases, diseases of the digestive system, diseases of the genitourinary system, neoplasms, diseases of the genitourinary system, and trauma
eICU-CRD [114]	139,367	Sequential EHR	Vital signs, laboratory measurements, medications, APACHE components, care plan information, admission diagnosis, patient history, and time-stamped diagnoses.	pulmonary sepsis, acute myocardial infarction, cerebrovascular accident, congestive heart failure, renal sepsis, diabetic ketoacidosis, coronary artery bypass graft, atrial rhythm disturbance, cardiac arrest, and emphysema

(Continued)

Table 5. Continued

Dataset name	Patient number	Data type	Data information	Disease Category
Amsterdam UMCdb [138]	20,109	Medical signals and sequential EHR	Patient monitor and life support device data, laboratory measurements, clinical observation and scores, medical procedures and tasks, medication, fluid balance, diagnosis groups and clinical patient outcomes	Not specified
UT Physicians clinical database [141]	5,501,776	Sequential EHR	Demographic data, vital signs, immunization data (body site, dose), laboratory data, transaction data (evaluation and management, radiology, medicine, surgery, anesthesia), appointment data, medications, and invoices	diabetes mellitus, hyperlipidemia, hypertension, and unspecified chest pain
Breast Cancer Wisconsin dataset (UCI) [34]	569	Tabular data	Diagnoses, radiuses, texture data, perimeters, areas, smoothness data, compactness data, concavity data, concave points data, symmetry data, and fractal dimensions.	Breast cancer
Heart Disease dataset (UCI) [34]	303	Tabular data	Demographic data, smoking status data, disease history data, exercise protocols, chart data (blood pressure, heart rate, ECG), pain status data, and diagnoses	Heart disease
Diabete dataset (UCI) [34]	70	Sequential data	Insulin dose, blood glucose measurement, hypoglycemic symptoms, meal ingestion, exercise activity	Diabete

Tabular data are composed of two sets of variables, continuous variables and categorical variables. Continuous variables include age, blood pressure, and temperature. In many statistical modeling-based algorithm, continuous variables are often pre-processed into discrete variables. It is because the finding a suitable prior distribution for continuous variables can be complex, and the joint or conditional distributions among multiple continuous variables are difficult to derive from data-driven methods. Thus, for most statistical modeling synthesis algorithms, continuous variables are classified into categorical variables according to their value. For example, in PrivBayes [167], continuous variables were first discretized into a fixed number l of equi-width bins and then binarized into $logl$ classes.

7.2 Pre-processing Methods for Sequential Data

The sequential data majorly has two forms, the medical signals and the EHRs. Many datasets provide both data, while during synthesis and analysis, these two data forms have different pre-processing steps.

Medical signals include neurological signals such as EEG and fMRI, and physiological signals such as continuous blood pressure waveforms or continuous heart rates. For each time point, these signals only contain one value. These medical signals are often periodic, so pre-processing methods of these signals can focus on time domain [105, 129] and frequency domain [96]; and quantities such as amplitude [110] and frequency [100] for these signals are considered for the synthesis and analysis. To medical signals, pre-processing [12, 85] include de-noising, artifact removing, and normalizations. Specific procedures depend on the modality of medical signals.

Another sequential data form in the medical context is sequential events that collected in EHR. Although sequential events data could be technically merged into one meta table, with date stamps as an attribute, the tabular structure of the meta table fails to investigate the chronological order of events. Some sequential event synthesis algorithms, particularly the simulation-based algorithms [14, 35, 90, 92] in Section 3, preserve the complex multi-table structure of EHR and the data structure is addressed as “patient care maps” or “careflow” in their algorithms. They would first synthesize several time points for each synthetic patient and then would add random tables for each time point. Each time point would have different numbers and types of tables, representing different events happened at this time point.

Other algorithms, however, normalize the events at each time point into a fix-size vector, because the fix-size input for each time point is required for algorithms, particularly the deep learning-based algorithms [26, 139, 169]. In these algorithms, the number of time points may vary, but the length of vectors at each time point must be fixed. Since all events contained in the datasets can span an event space, the fix-size vector is then a one-hot vector, and each entry of this vector represents the presence of the corresponding event at this specific time point. For example, considering an event space contains white cell abnormality, FVC abnormality, X-ray abnormality, and medication usage, the fixed size vector [1, 1, 1, 0], indicates a presence of white cell count abnormality, FVC abnormality, and X-ray abnormality, while no presence of medication usage.

To standardize the event space and provide a standardized dictionary for different symptoms, many algorithms will first use **International Classification of Diseases (ICD)** format to describe the events and then use the space of ICD codes to generate fix-size vectors for each time point. The ICD format [44] is a set of diagnostic encoding rules for clinical signs, symptoms, abnormal findings, complaints, social circumstances, and external causes of injury or diseases. Under this format, two versions, ICD9 [44] and ICD10 [104], are commonly used in synthesis algorithms. By encoding each symptom into the ICD format, an attribute space containing all ICD codes in the dataset is constructed, and for each timestamp, a fix-size, one-hot encoded vector representing the coordinators in the ICD space is derived for each timestamp of EHR.

7.3 Synthetic Datasets for EHR

We also discovered three released synthetic datasets, including Vanderbilt **Synthetic Derivative (SD)** dataset, **Data Entrepreneurs’ Synthetic Public Use File (DE-SynPUF)** dataset, and EMR Bots. These datasets, using different synthesis algorithms, provide hands-on practices for data synthesis.

7.3.1 Vanderbilt SD Dataset. The Vanderbilt SD dataset [27] is a synthetic dataset derived from a real database containing over 2.2 million patients. In SD dataset, demographic information, ICD-9 codes (diagnoses), CPT procedure codes, medications, vital signs, registry, patient histories, and

lab values is included. The dataset was de-identified by altering the records with their closest neighbors.

7.3.2 CMS 2008-2010 DE-SynPUF. The DE-SynPUF dataset [103] is also a synthetic EHR dataset containing data from over 2 millions synthetic patients in five domains: Beneficiary Summary, which includes the demographic information and hospital enrollment reasons; Inpatient Claims, such as the presence of a surgery and other clinical measurements for patients admitted in hospitals; Outpatient Claims, which is the procedure of examinations happened outside of hospitals; Carrier Claims, which are derived of the bill information of all medical services and include the name and date of the billed services, as well as the reimbursement amount related to this bill; and Prescription Drug Events, which contains the medication information for each patient.

To derive this large synthetic dataset, a combination of simulations and multiple imputation algorithms were used. Five steps of data synthesis were used: (1) variable reduction, where only clinical useful attributes were selected to be released; (2) suppression, where the rare data that had disclosure risk were removed; (3) substitution, where, similarly to the SD dataset [27], the attribute values for each patient were replaced by its nearest neighbors; (4) imputation, where the collected values of single variables were replaced by values synthesized from conditional distributions on key variables; (5) perturbation, where timelines of patient records were altered by changing dates; and (6) coarsening, where the continuous variables were coarsened into discrete variables.

7.3.3 EMRBots. The EMRBots [69] dataset is also an artificially generated EHR dataset contains three sub-datasets with 100, 10,000, and 100,000 synthetic patients. Unlike the Vanderbilt SD and De-SynPUF datasets, the data from EMRBots are generated by a set of pre-defined criteria, which was set by an experienced clinician.

8 TAKEAWAY MESSAGES

8.1 A Dilemma between Data-driven and Knowledge-driven Approaches

Deep neural networks have been providing efficient tools in medical image synthesis. In the field of non-imaging data synthesis, deep generative models are also on the top of our recommendation list. Deep neural networks, such as CTGAN [156], medGAN [26], and temporal gans such as TimeGAN [163], are easy to implement, and they do not require any prior knowledge during the training. However, the data-driven nature of deep generative models makes these methods suffer from overfitting (or mode collapse in data synthesis field). Methods such as Wasserstein loss [4] ease the pain of overfitting. Moreover, deep generative models are notorious for their difficulty in interpretation. If there are any unreal output values, then it is nearly impossible for ones to adjust accordingly in deep generative models. Attention mechanisms [7, 143] have been widely used to improve the network explainability, while they have been rarely discussed in the non-imaging data synthesis field.

The EMERGE family [35, 92] and Bayesian networks [167] allow ones to bring experts' prior knowledge in data synthesis. However, both of these algorithms are not good at handling high-dimensional data (tabular data with >1,000 attributes) [167]. For EMERGE-based methods, building patients' caremap models from scratch is laborious and time-consuming. Synthea [147] provides over 35 modules for different patient caremap modeling regarding different diseases, while customized caremap modeling is still difficult. A new type of data modeling named theory-driven modeling [68] is proposed to find a balance between knowledge-driven and data-driven methods. In addition, incorporating prior knowledge into deep generative networks [116] shall be a potential solution towards the synthesis efficiency and network explainability.

8.2 A Dilemma between Data Utility and Data Privacy

In 2021, a comparative study [133] was performed on different data synthesis algorithms, and the authors evaluated the risk of privacy violation and the utility of synthetic data from these algorithms. They discover that synthetic data do not protect the patients' privacy naturally. The synthetic data with high utility are vulnerable to the membership inference attack and, thus, have a high risk of individual re-identification. Meanwhile, they also perform a utility assessment on DP-based generative models and the results further support the dilemma between data utility and data privacy.

This dilemma is also reported in many articles in this review, especially those with DP mechanism [23, 153, 167]. With more noise added (lower ϵ in DP), the statistical closeness and prediction accuracy for downstream tasks get lower. Thus, it is important to re-think the rationale of *only* using data synthesis to protect privacy during data release. In 2004, a paper [2] was proposed to combine three approaches for confidential data release: (1) synthetic data releasing, (2) real analytical data releasing, and (3) restricted access to data. This complex method combining techniques and policy might be a solution to this dilemma, yet research investigating this dilemma is still ongoing.

8.3 Beyond the Debate over Encoder-decoder: Uncertainty and Diversity

For tabular data, where rows represent individuals and columns represent attributes, various deep learning architectures have been proposed. In Section 5.1.2, a debate was presented regarding the use of encoder-decoder structures in tabular data synthesis. Some algorithms [11, 158, 170] synthesize data directly, while others [21, 26] synthesize and sampling in a latent space and use a decoder to generate target data from the latent space. While we cannot determine a clear winner in terms of synthesis performance in this review, we can provide two perspectives that will help readers evaluate these two synthesis strategies.

First is the uncertainty of the synthetic data. The encoder-decoder approach aims to reconstruct the input as accurately as possible, but even with a well-trained model, the MSE loss between the input and the reconstruction can never reach zero. This means that there will always be some additional noise in the synthetic data, which can increase uncertainty and potentially introduce a shift between the synthetic and real data domains.

Second is the diversity, where latent space synthesis has its own advantages. For instance, energy-based models [74] introduce a new perspective that allows us to understand the diversities of data in GAN models. Additionally, these models provide theoretical proof for the diversities in the latent space synthesis.

8.4 Other Potential Research Directions

In addition to investigating the solutions to two aforementioned dilemmas, we will provide other potential research directions for non-imaging data synthesis in this subsection.

New models and new metrics. Diffusion models have shown superior synthesis performance compared to GANs in the domain of image synthesis [31]. Although some studies, such as References [52, 165], have explored the potential of diffusion models in non-imaging data synthesis, further research in this area is still relatively scarce.

As for the metrics, we noticed that fairness is not commonly evaluated in medical data synthesis even though this metric has a high correlation with trustworthiness. Additional metrics related to privacy, including k -anonymity, l -diversity, t -closeness, and p -indistinguishability [89], are also not commonly practiced in the data synthesis field, and implementations of these metrics have not been investigated yet.

Multi-modality synthesis. The concept of multi-modality [166] is widely used in medical image synthesis, where different image modalities such as CT, MRI and X-ray images are synthesized

together for a comprehensive representation of patients. Since non-imaging medical data are also an important clinical modality, the hybrid synthesis of imaging and non-imaging data should be considered. However, many algorithms only investigate the independent synthesis of imaging data and non-imaging data and coupled synthesis is still scarce in the literature.

Data harmonization using data synthesis. Data collection in healthcare research involves gathering data from various sources, each with its own unique formats and properties. Non-imaging data synthesis algorithms offer a solution by generating synthetic data that follow a standardized format while retaining the statistical characteristics of the original data sources. Despite this potential, the application of non-imaging data synthesis algorithms in data harmonization has not received adequate attention in the development of non-imaging data. Therefore, this section aims to highlight the potential of non-imaging data synthesis algorithms in achieving data harmonization, thereby contributing to the overall improvement in the reliability of AI applications.

9 CONCLUSION

The trustworthiness represents a set of essential qualities required for an AI algorithm: privacy, robustness, explainability, and fairness. To develop the trustworthy AI on non-imaging medical data, data synthesis algorithms have been proposed. By improving the number, variety and privacy of training samples, data synthesis algorithms are able to help AI models with a better accuracy at a lower cost. Trustworthy AI algorithms should cover all tasks in the AI field, including prediction and generation. However, most works so far concentrate on trustworthy “predictive modeling,” whereas the AI generation model raises many concerns as those exposed in this manuscript. Thus, this survey aims to be a referential point of discussion and a motivating catalyst of research around trustworthy synthetic data generation.

In this article, we identified three major types of non-imaging medical data synthesis algorithms and provided a comprehensive literature review about them and their evaluations. We also identified two challenges faced by all data synthesis algorithms: finding the balance between the utilization of data and knowledge and finding the balance between data utility and data privacy. We revealed some limitations existing in non-imaging data synthesis and called for new architectures, new evaluation metrics and multi-modality strategies to drive future efforts in this exciting research area.

REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, 308–318.
- [2] John M. Abowd and Julia Lane. 2004. New approaches to confidentiality protection: Synthetic data, remote access and research data centers. In *Privacy in Statistical Databases*, Josep Domingo-Ferrer and Vicenç Torra (Eds.). Springer, Berlin, 282–289.
- [3] Babak Afshin-Pour, Hamid Soltanian-Zadeh, Gholam-Ali Hossein-Zadeh, Cheryl L. Grady, and Stephen C. Strother. 2011. A mutual information-based metric for evaluation of fMRI data-processing approaches. *Hum. Brain Map.* 32, 5 (2011), 699–715.
- [4] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*, Doina Precup and Yee Whye Teh (Eds.). PMLR, 214–223.
- [5] Ravi V. Atreya, Joshua C. Smith, Allison B. McCoy, Bradley Malin, and Randolph A. Miller. 2013. Reducing patient re-identification risk for laboratory results within research datasets. *J. Am. Med. Inf. Assoc.* 20, 1 (2013), 95–101.
- [6] Vanessa Ayala-Rivera, A. Omar Portillo-Dominguez, Liam Murphy, and Christina Thorpe. 2016. COCOA: A synthetic data generator for testing anonymization techniques. In *Privacy in Statistical Databases*, Josep Domingo-Ferrer and Mirjana Pejić-Bach (Eds.). Springer International Publishing, Cham, Switzerland, 163–177.
- [7] Dmirty Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. <https://doi.org/10.48550/ARXIV.1409.0473>

- [8] Christopher L. Barrett, Richard J. Beckman, Maleq Khan, V. S. Anil Kumar, Madhav V. Marathe, Paula E. Stretz, Tridib Dutta, and Bryan Lewis. 2009. Generation and analysis of large synthetic social contact networks. In *Proceedings of the Winter Simulation Conference (WSC'09)*. IEEE, 1003–1014. <https://doi.org/10.1109/WSC.2009.5429425>
- [9] Elham Barzegaran, Sebastian Bosse, Peter J. Kohler, and Anthony M. Norcia. 2019. EEGSourceSim: A framework for realistic simulation of EEG scalp data using MRI-based forward models and biologically plausible signals and noise. *J. Neurosci. Methods* 328 (2019), 108377.
- [10] Brett K. Beaulieu-Jones, Jason H. Moore, and Pooled Resource Open-Access ALS Clinical Trials Consortium. 2017. Missing data imputation in the electronic health record using deeply learned autoencoders. In *Proceedings of the Pacific Symposium on Biocomputing*. World Scientific, 207–218.
- [11] Brett K. Beaulieu-Jones, Zhiwei Steven Wu, Chris Williams, Ran Lee, Sanjeev P. Bhavnani, James Brian Byrd, and Casey S. Greene. 2019. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation* 12, 7 (2019), e005122.
- [12] Nima Bigdely-Shamlo, Tim Mullen, Christian Kothe, Kyung-Min Su, and Kay A. Robbins. 2015. The PREP pipeline: Standardized preprocessing for large-scale EEG analysis. *Front. Neuroinf.* 9 (2015), 16.
- [13] Siddharth Biswal, Soumya Ghosh, Jon Duke, Bradley Malin, Walter Stewart, Cao Xiao, and Jimeng Sun. 2021. EVA: Generating longitudinal electronic health records using conditional variational autoencoders. In *Proceedings of the 6th Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 149)*, Ken Jung, Serena Yeung, Mark Sendak, Michael Sjoding, and Rajesh Ranganath (Eds.). PMLR, 260–282.
- [14] Anna L. Buczak, Steven Babin, and Linda Moniz. 2010. Data-driven approach for creating synthetic electronic medical records. *BMC Med. Inf. Decis. Mak.* 10, 1 (2010), 1–28.
- [15] Jim Burridge. 2003. Information preserving statistical obfuscation. *Stat. Comput.* 13, 4 (2003), 321–327.
- [16] Gregory Caiola and Jerome P. Reiter. 2010. Random forests for generating partially synthetic, categorical data. *Trans. Data Priv.* 3, 1 (2010), 27–42.
- [17] Ramiro D. Camino, Christian A. Hammerschmidt, and Radu State. 2019. Improving missing data imputation with deep generative models. [arXiv:1902.10666](https://arxiv.org/abs/1902.10666). Retrieved from <https://arxiv.org/abs/1902.10666>
- [18] Isaac Cano and Vicenç Torra. 2009. Generation of synthetic data by means of fuzzy c-Regression. In *Proceedings of the IEEE International Conference on Fuzzy Systems*. IEEE, 1145–1150.
- [19] Hong Cao, Xiao-Li Li, David Yew-Kwong Woon, and See-Kiong Ng. 2013. Integrated oversampling for imbalanced time series classification. *IEEE Trans. Knowl. Data Eng.* 25, 12 (2013), 2809–2822.
- [20] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. *J. Artif. Intell. Res.* 16 (2002), 321–357.
- [21] Zhengping Che, Yu Cheng, Shuangfei Zhai, Zhaonan Sun, and Yan Liu. 2017. Boosting deep learning risk prediction with generative adversarial networks for electronic health records. In *Proceedings of the IEEE International Conference on Data Mining (ICDM'17)*. IEEE, 787–792. <https://doi.org/10.1109/ICDM.2017.93>
- [22] Dingfan Chen, Ning Yu, Yang Zhang, and Mario Fritz. 2020. GAN-Leaks: A taxonomy of membership inference attacks against generative models. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS'20)*. Association for Computing Machinery, New York, NY, 343–362. <https://doi.org/10.1145/3372297.3417238>
- [23] Qingrong Chen, Chong Xiang, Minhui Xue, Bo Li, Nikita Borisov, Dali Kaarfar, and Haojin Zhu. 2018. Differentially private data generative models. [arXiv:1812.02274](https://arxiv.org/abs/1812.02274). Retrieved from <https://arxiv.org/abs/1812.02274>
- [24] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Eds.), Vol. 29. Curran Associates, Inc., Barcelona, Spain.
- [25] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. [arXiv:1409.1259](https://arxiv.org/abs/1409.1259). Retrieved from <https://arxiv.org/abs/1409.1259>
- [26] Edward Choi, Siddharth Biswal, Bradley Malin, Jon Duke, Walter F. Stewart, and Jimeng Sun. 2017. Generating multi-label discrete patient records using generative adversarial networks. In *Proceedings of the 2nd Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 68)*, Finale Doshi-Velez, Jim Fackler, David Kale, Rajesh Ranganath, Byron Wallace, and Jenna Wiens (Eds.). PMLR, 286–305.
- [27] Ioana Danciu, James D. Cowan, Melissa Basford, Xiaoming Wang, Alexander Saip, Susan Osgood, Jana Shirey-Rice, Jacqueline Kirby, and Paul A. Harris. 2014. Secondary use of clinical data: The Vanderbilt approach. *J. Biomed. Inf.* 52 (2014), 28–35.
- [28] Irina Deeva, Petr D. Andriushchenko, Anna V. Kalyuzhnaya, and Alexander V. Boukhanovsky. 2020. Bayesian networks-based personal data synthesis. In *Proceedings of the 6th EAI International Conference on Smart Objects and Technologies for Social Good (GoodTechs'20)*. Association for Computing Machinery, New York, NY, 6–11. <https://doi.org/10.1145/3411170.3411243>

- [29] Anne Marie Delaney, Eoin Brophy, and Tomas E. Ward. 2019. Synthesis of realistic ECG using generative adversarial networks. [arXiv:1909.09150](https://arxiv.org/abs/1909.09150). Retrieved from <https://arxiv.org/abs/1909.09150>
- [30] Ugur Demir and Gozde Unal. 2018. Patch-based image inpainting with generative adversarial networks. [arXiv:1803.07422](https://arxiv.org/abs/1803.07422). Retrieved from <https://arxiv.org/abs/1803.07422>
- [31] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Adv. Neural Inf. Process. Syst.* 34 (2021), 8780–8794.
- [32] Yongfeng Dong, Huaxin Xiao, and Yao Dong. 2022. SA-CGAN: An oversampling method based on single attribute guided conditional GAN for multi-class imbalanced learning. *Neurocomputing* 472 (2022), 326–337.
- [33] Jörg Drechsler. 2010. Using support vector machines for generating synthetic datasets. In *Privacy in Statistical Databases*, Josep Domingo-Ferrer and Emmanouil Magkos (Eds.). Springer, Berlin, 148–161.
- [34] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. Retrieved from <http://archive.ics.uci.edu/ml>
- [35] Kudakwashe Dube and Thomas Gallagher. 2014. Approach and method for generating realistic synthetic electronic healthcare records for secondary use. In *Foundations of Health Information Engineering and Systems*, Jeremy Gibbons and Wendy MacCaull (Eds.). Springer, Berlin, 69–86.
- [36] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer, Berlin and Heidelberg, 265–284.
- [37] Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (2014), 211–407.
- [38] Justin Engelmann and Stefan Lessmann. 2021. Conditional wasserstein GAN-based oversampling of tabular data for imbalanced learning. *Expert Syst. Appl.* 174 (2021), 114582.
- [39] Erik B. Erhardt, Elena A. Allen, Yonghua Wei, Tom Eichele, and Vince D. Calhoun. 2012. SimTB, a simulation toolbox for fMRI data under a model of spatiotemporal separability. *Neuroimage* 59, 4 (2012), 4160–4167.
- [40] Cristóbal Esteban, Stephanie L. Hyland, and Gunnar Rätsch. 2017. Real-valued (medical) time series generation with recurrent conditional GANs. [arXiv:1706.02633](https://arxiv.org/abs/1706.02633). Retrieved from <https://arxiv.org/abs/1706.02633>
- [41] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'15)*. Association for Computing Machinery, New York, NY, 259–268. <https://doi.org/10.1145/2783258.2783311>
- [42] Kelwin Fernandes, Jaime S. Cardoso, and Jessica Fernandes. 2017. Transfer learning with partial observability applied to cervical cancer screening. In *Pattern Recognition and Image Analysis*, Luís A. Alexandre, José Salvador Sánchez, and João M. F. Rodrigues (Eds.). Springer International Publishing, Cham, 243–250.
- [43] Alberto Fernández, Salvador García, Francisco Herrera, and Nitesh V. Chawla. 2018. SMOTE for learning from imbalanced data: Progress and challenges, marking the 15-year anniversary. *J. Artif. Intell. Res.* 61 (2018), 863–905.
- [44] National Center for Health Statistics (US) and Council on Clinical Classifications. 1980. *The International Classification of Diseases, 9th Revision, Clinical Modification: ICD-9-CM*. Vol. 2. US Department of Health and Human Services, Public Health Service, Health Care Financing Administration.
- [45] Andre Goncalves, Priyadip Ray, Braden Soper, Jennifer Stevens, Linda Coyle, and Ana Paula Sales. 2020. Generation and evaluation of synthetic patient data. *BMC Med. Res. Methodol.* 20, 1 (2020), 1–40.
- [46] Andre Goncalves, Priyadip Ray, Braden Soper, Jennifer Stevens, Linda Coyle, and Ana Paula Sales. 2020. Generation and evaluation of synthetic patient data. *BMC Med. Res. Methodol.* 20, 1 (2020), 1–40.
- [47] Aldren Gonzales, Guruprabha Guruswamy, and Scott R. Smith. 2023. Synthetic data in health care: A narrative review. *PLOS Digit. Health* 2, 1 (2023), e0000082.
- [48] Phil Gooch and Abdul Roudsari. 2011. Computerization of workflows, guidelines, and care pathways: A review of implementation challenges for process-oriented health information systems. *J. Am. Med. Inf. Assoc.* 18, 6 (2011), 738–748.
- [49] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc.
- [50] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. 2005. Borderline-SMOTE: A new over-sampling method in imbalanced data sets learning. In *Advances in Intelligent Computing*, De-Shuang Huang, Xiao-Ping Zhang, and Guang-Bin Huang (Eds.). Springer, Berlin, 878–887.
- [51] Haibo He, Yang Bai, Eduardo A. Garcia, and Shutao Li. 2008. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *Proceedings of the IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*. IEEE, 1322–1328.
- [52] Huan He, Shifan Zhao, Yuanzhe Xi, and Joyce C. Ho. 2023. MedDiff: Generating electronic health records using accelerated denoising diffusion model. [arXiv:2302.04355](https://arxiv.org/abs/2302.04355) [cs.LG]. Retrieved from <https://arxiv.org/abs/2302.04355>

- [53] Yu-Lin He, Sheng-Sheng Xu, and Joshua Zhexue Huang. 2022. Creating synthetic minority class samples based on autoencoder extreme learning machine. *Pattern Recogn.* 121 (2022), 108191.
- [54] Mikel Hernadez, Gorka Epelde, Ane Alberdi, Rodrigo Cilla, and Debbie Rankin. 2023. Synthetic tabular data evaluation in the health domain covering resemblance, utility, and privacy dimensions. *Methods Inf. Med.* S01 (2023), e19–e38.
- [55] Mikel Hernandez, Gorka Epelde, Ane Alberdi, Rodrigo Cilla, and Debbie Rankin. 2022. Synthetic data generation for tabular health records: A systematic review. *Neurocomputing* 493 (2022), 28–45. <https://doi.org/10.1016/j.neucom.2022.04.053>
- [56] Geoffrey E. Hinton and Sam Roweis. 2002. Stochastic neighbor embedding. In *Advances in Neural Information Processing Systems*, S. Becker, S. Thrun, and K. Obermayer (Eds.), Vol. 15. MIT Press.
- [57] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Comput.* 9, 8 (1997), 1735–1780.
- [58] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S. Yu, and Xuyun Zhang. 2022. Membership inference attacks on machine learning: A survey. *ACM Comput. Surv.* 54 (Jan. 2022), 1–37. <https://doi.org/10.1145/3523273>
- [59] Zhisheng Huang, Frank van Harmelen, Annette ten Teije, and Kathrin Dentler. 2013. Knowledge-based patient data generation. In *Process Support and Knowledge Representation in Health Care*, David Riaño, Richard Lenz, Silvia Miksch, Mor Peleg, Manfred Reichert, and Annette ten Teije (Eds.). Springer International Publishing, Cham, 83–96.
- [60] Byungduk Jeong, Wonjoon Lee, Deok-Soo Kim, and Hayong Shin. 2016. Copula-based approach to synthetic population generation. *PLoS One* 11, 8 (2016), e0159496.
- [61] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: Past, present and future. *Stroke Vasc. Neurol.* 2, 4 (2017), 230–243. <https://doi.org/10.1136/svn-2017-000101>
- [62] Baoyu Jing, Pengtao Xie, and Eric Xing. 2018. On the automatic generation of medical imaging reports. *arXiv.1711.08195*. Retrieved from <https://arxiv.org/abs/1711.08195>
- [63] Wonkeun Jo and Dongil Kim. 2022. OBGAN: Minority oversampling near borderline with generative adversarial networks. *Expert Syst. Appl.* 197 (2022), 116694.
- [64] Alistair Johnson, Lucas Bulgarelli, Tom Pollard, Steven Horng, Leo Anthony Celi, and Roger Mark. 2022. Mimic-IV. Retrieved from <https://physionet.org/content/mimiciv/2.0/>
- [65] Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. MIMIC-III, a freely accessible critical care database. *Sci. Data* 3, 1 (2016), 1–9.
- [66] James Jordon, Lukasz Szpruch, Florimond Houssiau, Mirko Bottarelli, Giovanni Cherubin, Carsten Maple, Samuel N. Cohen, and Adrian Weller. 2022. Synthetic data—what, why and how? *arXiv:2205.03257*. Retrieved from <https://arxiv.org/abs/2205.03257>
- [67] James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. 2018. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*.
- [68] Anuj Karpatne, Gowtham Atluri, James H. Faghmous, Michael Steinbach, Arindam Banerjee, Auroop Ganguly, Shashi Shekhar, Nagiza Samatova, and Vipin Kumar. 2017. Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Trans. Knowl. Data Eng.* 29, 10 (2017), 2318–2331.
- [69] Uri Kartoun. 2016. A methodology to generate virtual patient repositories. *arXiv:1608.00570*. Retrieved from <https://arxiv.org/abs/1608.00570>
- [70] Dhamanpreet Kaur, Matthew Sobiesk, Shubham Patil, Jin Liu, Puran Bhagat, Amar Gupta, and Natasha Markuzon. 2021. Application of Bayesian networks to generate synthetic health data. *J. Am. Med. Inf. Assoc.* 28, 4 (2021), 801–811.
- [71] Davinder Kaur, Suleyman Uslu, Kaley J. Rittichier, and Arjan Durrresi. 2022. Trustworthy artificial intelligence: A review. *ACM Comput. Surv.* 55, 2 (2022), 1–38.
- [72] Diederik P. Kingma and Max Welling. 2013. Auto-encoding variational Bayes. *arXiv:1312.6114*. Retrieved from <https://arxiv.org/abs/1312.6114>
- [73] Solomon Kullback and Richard A. Leibler. 1951. On information and sufficiency. *Ann. Math. Stat.* 22, 1 (1951), 79–86.
- [74] Rithesh Kumar, Sherjil Ozair, Anirudh Goyal, Aaron Courville, and Yoshua Bengio. 2019. Maximum entropy generators for energy-based models. *arXiv:1901.08508*. Retrieved from <https://arxiv.org/abs/>
- [75] Wai Lam and Fahiem Bacchus. 1994. Learning Bayesian belief networks: An approach based on the MDL principle. *Comput. Intell.* 10, 3 (1994), 269–293.
- [76] Zi-Ching Lan, Guan-Yu Huang, Yun-Pei Li, Seungmin Rho, S. Vimal, and Bo-Wei Chen. 2022. Conquering insufficient/imbalanced data learning for the Internet of Medical Things. *Neural Computing and Applications S.I. : Neural Computing for IOT based Intelligent Healthcare Systems*, 1–10. <https://doi.org/10.1007/s00521-022-06897-z>
- [77] Steffen L. Lauritzen and David J. Spiegelhalter. 1988. Local computations with probabilities on graphical structures and their application to expert systems. *J. Roy. Stat. Soc.: Ser. B (Methodol.)* 50, 2 (1988), 157–194.

- [78] Dongha Lee, Hwanjo Yu, Xiaoqian Jiang, Deevakar Rogith, Meghana Gudala, Mubeen Tejani, Qiuchen Zhang, and Li Xiong. 2020. Generating sequential electronic health records using dual adversarial autoencoder. *J. Am. Med. Inf. Assoc.* 27, 9 (2020), 1411–1419.
- [79] Min Kyung Lee and Katherine Rich. 2021. Who is included in human perceptions of AI?: Trust and perceived fairness around healthcare AI and cultural mistrust. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI'21)*. Association for Computing Machinery, New York, NY. <https://doi.org/10.1145/3411764.3445570>
- [80] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *J. Mach. Learn. Res.* 18, 17 (2017), 1–5.
- [81] Haoran Li, Li Xiong, and Xiaoqian Jiang. 2014. Differentially private synthesis of multi-dimensional data using copula functions. In *Advances in Database Technology: Proceedings. International Conference on Extending Database Technology*, Vol. 2014. NIH Public Access, Bethesda, Maryland, USA, 475.
- [82] Haoran Li, Li Xiong, Lifan Zhang, and Xiaoqian Jiang. 2014. DPSynthesizer: Differentially private data synthesizer for privacy preserving data sharing. In *Proceedings of the VLDB Endowment International Conference on Very Large Data Bases*, Vol. 7. NIH Public Access, NIH Public Access, 1677.
- [83] Weixin Liang, Girmaw Abebe Tadesse, Daniel Ho, Fei-Fei Li, Matei Zaharia, Ce Zhang, and James Zou. 2022. Advances, challenges and opportunities in creating data for trustworthy AI. *Nature Mach. Intell.* 4 (2022), 669–677.
- [84] Martin A. Lindquist, Christian Waugh, and Tor D. Wager. 2007. Modeling state-related fMRI activity using change-point theory. *NeuroImage* 35, 3 (2007), 1125–1141.
- [85] Vladimir Litvak, Jérémie Mattout, Stefan Kiebel, Christophe Phillips, Richard Henson, James Kilner, Gareth Barnes, Robert Oostenveld, Jean Daunizeau, Guillaume Flandin, et al. 2011. EEG and MEG data analysis in SPM8. *Comput. Intell. Neurosci.* 2011 (2011), 852961.
- [86] Hao Luo, Jun Liao, Xuewen Yan, and Li LiU. 2021. Oversampling by a constraint-based causal network in medical imbalanced data classification. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME'21)*. IEEE, 1–6.
- [87] Ashwin Machanavajjhala, Daniel Kifer, John Abowd, Johannes Gehrke, and Lars Vilhuber. 2008. Privacy: Theory meets practice on the map. In *Proceedings of the IEEE 24th International Conference on Data Engineering*. IEEE, IEEE, 277–286.
- [88] Frank J. Massey Jr. 1951. The Kolmogorov-Smirnov test for goodness of fit. *J. Am. Stat. Assoc.* 46, 253 (1951), 68–78.
- [89] Stan Matwin, Jordi Nin, Morvarid Sehatkar, and Tomasz Szapiro. 2015. *A Review of Attribute Disclosure Control*. Springer International Publishing, Cham, 41–61. https://doi.org/10.1007/978-3-319-09885-2_4
- [90] Scott McLachlan, Kudakwashe Dube, and Thomas Gallagher. 2016. Using the caremap with health incidents statistics for generating the realistic synthetic electronic healthcare record. In *Proceedings of the IEEE International Conference on Healthcare Informatics (ICHI'16)*. IEEE, 439–448.
- [91] Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, Providence, Rhode Island, 94–103.
- [92] Linda Moniz, Anna L. Buczak, Lang Hung, Steven Babin, Michael Dorko, and Joseph Lombardo. 2009. Construction and validation of synthetic electronic medical records. *Online J. Publ. Health Inf.* 1, 1 (2009), 1–36.
- [93] George B. Moody and Roger G. Mark. 1996. A database to support development and evaluation of intelligent intensive care monitoring. In *Computers in Cardiology*. IEEE, 657–660.
- [94] Krishnamurthy Muralidhar, Rahul Parsa, and Rathindra Sarathy. 1999. A general additive data perturbation method for database security. *Manage. Sci.* 45, 10 (1999), 1399–1415.
- [95] Hajra Murtaza, Musharif Ahmed, Naurin Farooq Khan, Ghulam Murtaza, Saad Zafar, and Ambreen Bano. 2023. Synthetic data generation: State of the art in health care domain. *Comput. Sci. Rev.* 48 (2023), 100546.
- [96] Karsten Müller, Gabriele Lohmann, Volker Bosch, and D. Yves von Cramon. 2001. On multivariate spectral analysis of fMRI time series. *NeuroImage* 14, 2 (2001), 347–356. <https://doi.org/10.1006/nimg.2001.0804>
- [97] Graham Neubig. 2017. Neural machine translation and sequence-to-sequence models: A tutorial. [arXiv:1703.01619](https://arxiv.org/abs/1703.01619). Retrieved from <https://arxiv.org/abs/1703.01619>
- [98] Diogo Telmo Neves, João Alves, Marcel Ganesh Naik, Alberto José Proença, and Fabian Prasser. 2022. From missing data imputation to data generation. *J. Comput. Sci.* 61 (2022), 101640.
- [99] Sophie J. Nightingale and Hany Farid. 2022. AI-synthesized faces are indistinguishable from real faces and more trustworthy. *Proc. Natl. Acad. Sci. U.S.A.* 119, 8 (2022), e2120481119.
- [100] Marc R. Nuwer. 1988. Quantitative EEG: I. Techniques and problems of frequency analysis and topographic mapping. *J. Clin. Neurophysiol.* 5, 1 (1988), 1–43.
- [101] Augustus Odena, Christopher Olah, and Jonathon Shlens. 2017. Conditional image synthesis with auxiliary classifier GANs. In *Proceedings of the 34th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 70)*. Doina Precup and Yee Whye Teh (Eds.). PMLR, 2642–2651.

- [102] U.S. Department of Health & Human Services. 1996. Health Insurance Portability and Accountability Act (HIPAA). Retrieved from <https://www.hhs.gov/hipaa/index.html>
- [103] Redivis Demo Organization. 2020. CMS Synthetic Patient Data OMOP. Retrieved from <https://redivis.com/datasets/ye2v-6skh7wdr7?v=2.0>
- [104] World Health Organization. 2004. *International Statistical Classification of Diseases and Related Health Problems: Alphabetical Index*. Vol. 3. World Health Organization, USA.
- [105] James Pardey, Stephen Roberts, and Lionel Tarassenko. 1996. A review of parametric modelling techniques for EEG analysis. *Med. Eng. Phys.* 18, 1 (1996), 2–11.
- [106] Noseong Park, Mahmoud Mohammadi, Kshitij Gorde, Sushil Jajodia, Hongkyu Park, and Youngmin Kim. 2018. Data synthesis based on generative adversarial networks. *Proc. VLDB Endow.* 11, 10 (Jun. 2018), 1071–1083. <https://doi.org/10.14778/3231751.3231757>
- [107] Yubin Park and Joydeep Ghosh. 2013. Perturbed Gibbs samplers for synthetic data release. [arXiv:1312.5370](https://arxiv.org/abs/1312.5370). Retrieved from <https://arxiv.org/abs/1312.5370>
- [108] Yubin Park, Joydeep Ghosh, and Mallikarjun Shankar. 2013. Perturbed gibbs samplers for generating large-scale privacy-safe synthetic health data. In *Proceedings of the IEEE International Conference on Healthcare Informatics*. IEEE, 493–498.
- [109] Neha Patki, Roy Wedge, and Kalyan Veeramachaneni. 2016. The synthetic data vault. In *Proceedings of the IEEE International Conference on Data Science and Advanced Analytics (DSAA'16)*. IEEE, 399–410.
- [110] Luiz Pessoa, Eva Gutierrez, Peter A. Bandettini, and Leslie G. Ungerleider. 2002. Neural correlates of visual working memory: fMRI amplitude predicts task performance. *Neuron* 35, 5 (2002), 975–987.
- [111] Haoyue Ping, Julia Stoyanovich, and Bill Howe. 2017. DataSynthesizer: Privacy-preserving synthetic datasets. In *Proceedings of the 29th International Conference on Scientific and Statistical Database Management (SSDBM'17)*. Association for Computing Machinery, New York, NY. <https://doi.org/10.1145/3085504.3091117>
- [112] Kemal Polat. 2019. A hybrid approach to Parkinson disease classification using speech signal: the combination of smote and random forests. In *Proceedings of the Scientific Meeting on Electrical-Electronics & Biomedical Engineering and Computer Science (EBBT'19)*. IEEE, 1–3.
- [113] David Pollard. 2005. *Total Variation Distance Between Measures*. Asymptopia, Virtual.
- [114] Tom J. Pollard, Alistair E. W. Johnson, Jesse D. Raffa, Leo A. Celi, Roger G. Mark, and Omar Badawi. 2018. The eICU collaborative research database, a freely available multi-center database for critical care research. *Sci. Data* 5, 1 (2018), 1–13.
- [115] Pytorch. 2020. RNN Pytorch 1.12 Document. Retrieved August 19, 2022 from <https://pytorch.org/docs/stable/generated/torch.nn.RNN.html>
- [116] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378 (2019), 686–707.
- [117] Sina Rashidian, Fusheng Wang, Richard Moffitt, Victor Garcia, Anurag Dutt, Wei Chang, Vishwam Pandya, Janos Hajagos, Mary Saltz, and Joel Saltz. 2020. SMOOTH-GAN: Towards sharp and smooth synthetic EHR data generation. In *Artificial Intelligence in Medicine*, Martin Michalowski and Robert Moskovitch (Eds.). Springer International Publishing, Cham, 37–48.
- [118] David Riaño and Alberto Fernández-Pérez. 2017. Simulation-based episodes of care data synthetization for chronic disease patients. In *Knowledge Representation for Health Care*, David Riaño, Richard Lenz, and Manfred Reichert (Eds.). Springer International Publishing, Cham, 36–50.
- [119] Donald B. Rubin. 1993. Statistical disclosure limitation. *J. Official Stat.* 9, 2 (1993), 461–468.
- [120] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. 1986. Learning representations by back-propagating errors. *Nature* 323, 6088 (1986), 533–536.
- [121] Mohammed Saeed, Christine Lieu, Greg Raber, and Roger G. Mark. 2002. MIMIC II: A massive temporal ICU patient database to support research in intelligent patient monitoring. In *Computers in Cardiology*. IEEE, 641–644.
- [122] Jacob Schreiber. 2017. Pomegranate: Fast and flexible probabilistic modeling in python. *J. Mach. Learn. Res.* 18, 1 (2017), 5992–5997.
- [123] Gideon Schwarz. 1978. Estimating the dimension of a model. *Ann. Stat.* 6, 2 (1978), 461–464.
- [124] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership inference attacks against machine learning models. In *Proceedings of the IEEE Symposium on Security and Privacy (SP'17)*. IEEE, 3–18.
- [125] Laura Sikstrom, Marta M. Maslej, Katrina Hui, Zoe Findlay, Daniel Z. Buchman, and Sean L. Hill. 2022. Conceptualising fairness: Three pillars for medical algorithms and health equity. *BMJ Health Care Inf.* 29, 1 (2022), e100459. <https://doi.org/10.1136/bmjhci-2021-100459>
- [126] M. Sklar. 1959. Fonctions de repartition an dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris* 8 (1959), 229–231.

- [127] Heithem Sliman, Imen Megdiche, Loay Alajramy, Adel Taweel, Sami Yangui, Aida Drira, and Elyes Lamine. 2023. MedWGAN based synthetic dataset generation for Uveitis pathology. *Intell. Syst. Appl.* 18 (2023), 200223. <https://doi.org/10.1016/j.iswa.2023.200223>
- [128] Jason F. Smith, Kewei Chen, Ajay S. Pillai, and Barry Horwitz. 2013. Identifying effective connectivity parameters in simulated fMRI: A direct comparison of switching linear dynamic system, stochastic dynamic causal, and multivariate autoregressive models. *Front. Neurosci.* 7 (2013), 70.
- [129] S. M. Smith. 2004. Overview of fMRI analysis. *Br. J. Radiol.* 77, suppl_2 (2004), S167–S175. <https://doi.org/10.1259/bjr/33553595>
- [130] Minjae Son, Seungwon Jung, Jihoon Moon, and Eenjun Hwang. 2020. BCGAN-based over-sampling scheme for imbalanced data. In *Proceedings of the IEEE International Conference on Big Data and Smart Computing (BigComp'20)*. IEEE, 155–160.
- [131] Tzu-An Song, Samadrita Roy Chowdhury, Fan Yang, Heidi Jacobs, Georges El Fakhri, Quanzheng Li, Keith Johnson, and Joyita Dutta. 2019. Graph convolutional neural networks for Alzheimer's disease classification. In *Proceedings of the IEEE 16th International Symposium on Biomedical Imaging (ISBI'19)*. IEEE, 414–417.
- [132] Akash Srivastava, Lazar Valkov, Chris Russell, Michael U. Gutmann, and Charles Sutton. 2017. VEEGAN: Reducing mode collapse in GANs using implicit variational learning. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [133] Theresa Stadler, Bristena Oprisanu, and Carmela Troncoso. 2022. Synthetic data—Anonymisation groundhog day. In *Proceedings of the 31st USENIX Security Symposium (USENIX Security'22)*. USENIX Association, 1451–1468.
- [134] Lijun Sun and Alexander Erath. 2015. A Bayesian network approach for population synthesis. *Transport. Res. Part C: Emerg. Technol.* 61 (2015), 49–62.
- [135] Siao Sun, Fusheng Wang, Sina Rashidian, Tahsin Kurc, Kayley Abell-Hart, Janos Hajagos, Wei Zhu, Mary Saltz, and Joel Saltz. 2021. Generating longitudinal synthetic EHR data with recurrent autoencoders and generative adversarial networks. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*. Springer International Publishing, Cham, 153–165.
- [136] Thomas Douglas Victor Swinscow, Michael J. Campbell, et al. 2002. *Statistics at Square One*. Bmj London, London, UK.
- [137] Erdogan Taskesen. 2020. bnlearn—Library for Bayesian Network Learning and Inference. Retrieved from <https://erdogant.github.io/bnlearn>
- [138] Patrick J. Thorat, Jan M. Peppink, Ronald H. Driessen, Eric J. G. Sijbrands, Erwin J. O. Kompanje, Lewis Kaplan, Heatherlee Bailey, Jozef Kesecioglu, Maurizio Cecconi, Matthew Churpek, et al. 2021. Sharing ICU patient data Responsibly under the Society of critical care Medicine/European Society of intensive care medicine joint data science collaboration: The Amsterdam University medical centers database (AmsterdamUMCdb) example. *Crit. Care Med.* 49, 6 (2021), e563.
- [139] Amirsina Torfi and Edward A. Fox. 2020. CorGAN: Correlation-capturing convolutional generative adversarial networks for generating synthetic healthcare records. In *Proceedings of the 33rd International Flairs Conference*. AAAI Press, 1–6.
- [140] Allan Tucker, Zhenchen Wang, Ylenia Rotalinti, and Puja Myles. 2020. Generating high-fidelity synthetic patient data for assessing machine learning healthcare software. *NPJ Digit. Med.* 3, 1 (2020), 1–13.
- [141] uth.edu. 2022. BIG-Arc—Clinical Data Warehouse—Data Dashboard. Retrieved August 19, 2022 from <https://big.uth.edu/bigarc/>
- [142] L. Vivek Harsha Vardhan and Stanley Kok. 2020. Generating privacy-preserving synthetic tabular data using oblivious variational autoencoders. In *Proceedings of the Workshop on Economics of Privacy and Data Labor at the 37th International Conference on Machine Learning*. PMLR, 1–8.
- [143] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc.
- [144] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*. PMLR, Valencia, 1096–1103.
- [145] Paul Voigt and Axel Von dem Bussche. 2017. The eu general data protection regulation (gdpr). In *A Practical Guide, 1st Ed*. Springer International Publishing, Cham.
- [146] Christian Walck. 2007. *Hand-book on Statistical Distributions for Experimentalists*. University of Stockholm, Stockholm.
- [147] Jason Walonoski, Mark Kramer, Joseph Nichols, Andre Quina, Chris Moesel, Dylan Hall, Carlton Duffett, Kudakwashe Dube, Thomas Gallagher, and Scott McLachlan. 2018. Synthea: An approach, method, and software

- mechanism for generating synthetic patients and the synthetic electronic health care record. *J. Am. Med. Inf. Assoc.* 25, 3 (2018), 230–238.
- [148] Lu Wang, Wei Zhang, and Xiaofeng He. 2019. Continuous patient-centric sequence generation via sequentially coupled adversarial learning. In *Database Systems for Advanced Applications*, Guoliang Li, Jun Yang, Joao Gama, Jugapong Natwichai, and Yongxin Tong (Eds.). Springer International Publishing, Cham, 36–52.
 - [149] Zhenchen Wang, Puja Myles, and Allan Tucker. 2021. Generating and evaluating cross-sectional synthetic electronic healthcare data: Preserving data utility and patient privacy. *Comput. Intell.* 37, 2 (2021), 819–851.
 - [150] Marijke Welvaert and Yves Rosseel. 2014. A review of fMRI simulation studies. *PLoS One* 9, 7 (2014), e101953.
 - [151] Jeannette M. Wing. 2021. Trustworthy AI. *Commun. ACM* 64, 10 (Sep. 2021), 64–71. <https://doi.org/10.1145/3448248>
 - [152] Jesper N. Wulff and Linda Ejlskov Jeppesen. 2017. Multiple imputation by chained equations in praxis: Guidelines and review. *Electr. J. Bus. Res. Methods* 15, 1 (2017), 41–56.
 - [153] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. 2018. Differentially private generative adversarial Network. [arXiv:1802.06739](https://arxiv.org/abs/1802.06739). Retrieved from <https://arxiv.org/abs/1802.06739>
 - [154] Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. 2018. Fairgan: Fairness-aware generative adversarial networks. In *Proceedings of the IEEE International Conference on Big Data (Big Data'18)*. IEEE, 570–575.
 - [155] Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. 2019. Fairgan+: Achieving fair data generation and classification through generative adversarial nets. In *Proceedings of the IEEE International Conference on Big Data (Big Data'19)*. IEEE, Los Alamitos, CA, 1401–1406.
 - [156] Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2019. Modeling tabular data using conditional GAN. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc.
 - [157] Lei Xu and Kalyan Veeramachaneni. 2018. Synthesizing tabular data using generative adversarial networks. [arXiv:1811.11264](https://arxiv.org/abs/1811.11264). Retrieved from <https://arxiv.org/abs/1811.11264>
 - [158] Andrew Yale, Saloni Dash, Ritik Dutta, Isabelle Guyon, Adrien Pavao, and Kristin P. Bennett. 2020. Generation and evaluation of privacy preserving synthetic health data. *Neurocomputing* 416 (2020), 244–255.
 - [159] Chao Yan, Yao Yan, Zhiyu Wan, Ziqi Zhang, Larsson Omberg, Justin Guinney, Sean D. Mooney, and Bradley A. Malin. 2022. A multifaceted benchmarking of synthetic electronic health record generation models. *Nat. Commun.* 13, 1 (2022), 7609.
 - [160] Chao Yan, Ziqi Zhang, Steve Nyemba, and Bradley A. Malin. 2020. Generating electronic health records with multiple data types and constraints. In *AMIA Annual Symposium Proceedings*, Vol. 2020. American Medical Informatics Association, American Medical Informatics Association, 1335.
 - [161] Huan Yang and Pengjiang Qian. 2021. GAN-based medical images synthesis: A review. *Int. J. Health Syst. Transl. Med.* 1, 2 (2021), 1–9.
 - [162] Jinsung Yoon, Lydia N. Drumright, and Mihaela Van Der Schaar. 2020. Anonymization through data synthesis using generative adversarial networks (ads-gan). *IEEE J. Biomed. Health Inf.* 24, 8 (2020), 2378–2388.
 - [163] Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. 2019. Time-series generative adversarial networks. In *Advances in Neural Information Processing Systems*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc.
 - [164] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. 2018. GAIN: Missing data imputation using generative adversarial nets. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 5689–5698.
 - [165] Hongyi Yuan, Songchi Zhou, and Sheng Yu. 2023. EHRDiff: Exploring realistic EHR synthesis with diffusion models. [arXiv:2303.05656](https://arxiv.org/abs/2303.05656) [cs.LG]. Retrieved from <https://arxiv.org/abs/2303.05656>
 - [166] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Jiahui Zhang, Shijian Lu, Lingjie Liu, Adam Kortylewski, Christian Theobalt, and Eric Xing. 2021. Multimodal image synthesis and editing: A survey. [arXiv:2112.13592](https://arxiv.org/abs/2112.13592). Retrieved from <https://arxiv.org/abs/2112.13592>
 - [167] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2017. PrivBayes: Private data release via bayesian networks. *ACM Trans. Database Syst.* 42, 4 (2017), 1–41.
 - [168] Zhikun Zhang, Tianhao Wang, Ninghui Li, Jean Honorio, Michael Backes, Shibo He, Jiming Chen, and Yang Zhang. 2020. PrivSyn: Differentially private data synthesis. [arXiv:2012.15128](https://arxiv.org/abs/2012.15128). Retrieved from <https://arxiv.org/abs/2012.15128>
 - [169] Ziqi Zhang, Chao Yan, Thomas A. Lasko, Jimeng Sun, and Bradley A. Malin. 2021. SynTEG: A framework for temporal structured electronic health data simulation. *J. Am. Med. Inf. Assoc.* 28, 3 (2021), 596–604.
 - [170] Ziqi Zhang, Chao Yan, Diego A. Mesa, Jimeng Sun, and Bradley A. Malin. 2019. Ensuring electronic medical record simulation through better training, modeling, and evaluation. *J. Am. Med. Inf. Assoc.* 27, 1 (10 2019), 99–108. <https://doi.org/10.1093/jamia/ocz311>

Received 17 September 2022; revised 13 May 2023; accepted 3 August 2023