
Synthcity: a benchmark framework for diverse use cases of tabular synthetic data

Zhaozhi Qian
University of Cambridge
zq224@cam.ac.uk

Rob Davis
University of Cambridge
rd473@cam.ac.uk

Mihaela van der Schaar
University of Cambridge
mv472@cam.ac.uk

Abstract

1 Accessible high-quality data is the bread and butter of machine learning research,
2 and the demand for data has exploded as larger and more advanced ML models are
3 built across different domains. Yet, real data often contain sensitive information,
4 subject to various biases, and are costly to acquire, which compromise their quality
5 and accessibility. Synthetic data have thus emerged as a complement, sometimes
6 even a replacement, to real data for ML training. However, the landscape of
7 synthetic data research has been fragmented due to the large number of data
8 modalities (e.g., tabular data, time series data, images, etc.) and various use cases
9 (e.g., privacy, fairness, data augmentation, etc.). This poses practical challenges
10 in comparing and selecting synthetic data generators in different problem settings.
11 To this end, we develop Synthcity, an open-source Python library that allows
12 researchers and practitioners to perform one-click benchmarking of synthetic data
13 generators across data modalities and use cases. In addition, Synthcity’s plug-in
14 style API makes it easy to incorporate additional data generators into the framework.
15 Beyond benchmarking, it also offers a single access point to a diverse range of
16 cutting-edge data generators. Through examples on tabular data generation and
17 data augmentation, we illustrate the general applicability of Synthcity, and the
18 insight one can obtain.

19 1 Introduction

20 Access to high quality data is the lifeblood of AI. Although AI holds strong promise in numerous high-
21 stakes domains, the lack of high-quality datasets creates a significant hurdle for the development of
22 AI, leading to missed opportunities. Specifically, three prominent issues contribute to this challenge:
23 *data scarcity*, *privacy*, and *bias* [Mehrabani et al., 2021, Gianfrancesco et al., 2018, Tashea, 2017,
24 Dastin, 2018]. As a result, the dataset may not be available, accessible, or suitable for building
25 performant and socially responsible AI systems [Sambasivan et al., 2021].

26 This challenge is especially prominent for tabular datasets, which are often curated in highly regulated
27 industries including healthcare, finance, manufacturing etc. Synthetic tabular data has the potential
28 to fuel the development of AI by unleashing the information in datasets that are small, sensitive or
29 biased. To achieve this, we need high-performance generative models that both faithfully capture the
30 data distribution and satisfy additional constraints for the desired use cases.

31 To date, the landscape of synthetic data research has been fragmented because the combination of
32 *use cases* (i.e. fairness, privacy, and augmentation) and *data modalities* (e.g. static tabular data, time
33 series data, etc.) creates a plethora of problem settings. In response to the large problem space, the

34 community has taken a divide-and-conquer approach: highly-specialized generative models have
35 been developed to fit in one particular setting. This has led to a proliferation of specialized generative
36 models [Jordon et al., 2018, Yoon et al., 2020, Ho et al., 2021, Mehrabi et al., 2021, van Breugel
37 et al., 2021, Zhu et al., 2017, Yoon et al., 2018, Saxena and Cao, 2021].

38 This fragmented landscape has created four main challenges for benchmarking synthetic data genera-
39 tors, which would hamper the research progress if left unaddressed.

40 **1. Challenge in use case specific evaluation.** Most existing studies in generative model only focus
41 on the fidelity of the synthetic data, i.e. how they resemble the real data in distribution Wang et al.
42 [2019], Tucker et al. [2020], Goncalves et al. [2020], Wang et al. [2021], Kokosi and Harron [2022].
43 However, additional evaluation is needed to assess the specific use cases. For example, the utility to
44 downstream models and the data privacy. This calls for the introduction of new metrics as well as
45 new evaluation pipelines.

46 **2. Challenge in off label uses.** Although specialized generative models are developed for one use
47 case, the practical application often requires them to cover multiple use cases (e.g. data augmentation
48 with privacy). Hence, generative models are often used outside the designed scope. Prior work has
49 shown that this may lead to undesirable and unpredictable consequences [Pereira et al., 2021, Ganev
50 et al., 2022]. As a result, researchers need to comprehensively evaluate the generative model across a
51 variety of use cases to assess the risk of off label uses.

52 **3. Challenge in comparing with a large number of baselines.** In practice, it is often very
53 challenging to systematically compare with a large number of existing baselines because the interfaces
54 (API) of these models are often inconsistent and incompatible (e.g. they may require different formats
55 of input data and conflicting software dependencies). As a result, the researcher usually needs to
56 spend time and effort to harmonize the code rather than focusing on the research question itself.

57 **4. Challenge in understanding the performance gain.** Generative models are complex systems
58 that involve many components, such as the model architecture, the objective function, and the hyper-
59 parameters. These aspects all encode prior assumptions and inductive biases, which would bring
60 unique strengths and weaknesses to the models [Bond-Taylor et al., 2021]. However, it is often
61 difficult to pinpoint the exact component that leads to the performance gain. Most existing studies
62 evaluate the model as a whole and neglect the role of different components.

63 **Contribution.** In this work, we present Synthcity, an open-source Python library available on pip
64 and GitHub, as a solution to these benchmark challenges. Synthcity offers diverse data modalities
65 and supports various use cases. It provides an extensive set of evaluation metrics for assessing
66 dataset fidelity, privacy, and utility, making it a robust tool for evaluating synthetic data across
67 different applications. With a wide array of state-of-the-art generators and customizable architectures,
68 users can perform consistent comparisons with existing models, gaining insights into performance
69 improvements. Accessible through an intuitive interface, Synthcity facilitates tabular data generation
70 and augmentation, demonstrated through two case studies. Researchers can employ Synthcity for
71 benchmarks and guidance in synthetic data research

72 **2 The synthcity library**

73 **2.1 Overview of the synthcity workflow**

74 Despite the fragmented landscape in synthetic data research, Synthcity implements a unified workflow
75 for benchmark studies. We formalize the process as follows. Let X be the random variable of
76 interest (which could be static, temporal or censored). The real data is composed with observations
77 $x_i \sim P(X)$ drawn from the true (but unknown) distribution. For benchmark evaluation, the real data
78 is split into a training (\mathbb{D}_{train}^r) and test (\mathbb{D}_{test}^r) set. The generator is trained using the training set
79 \mathbb{D}_{train}^r . During training, the generator (explicitly or implicitly) learns the distribution $\hat{P}(X)$ in order
80 to sample from it. After training, the generative model generates synthetic data \mathbb{D}^s , which will be
81 evaluated with respect to the test set \mathbb{D}_{test}^r (or in some special cases, the training set \mathbb{D}_{train}^r).

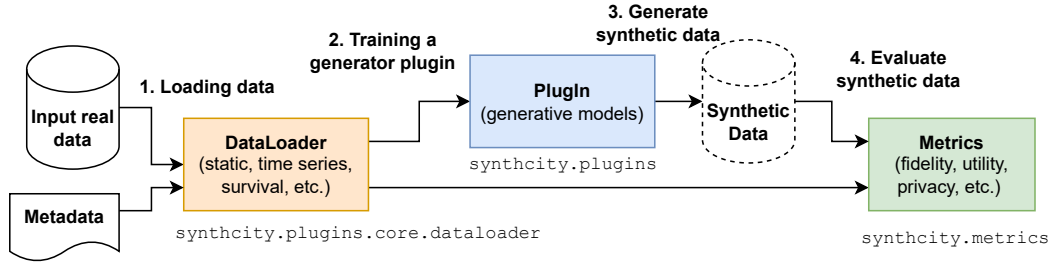


Figure 1: Standard workflow of generating and evaluating synthetic data with synthcity.

82 The synthcity library captures the entire workflow of synthetic data benchmark in four steps (Figure
 83 1). This workflow applies to all use cases, generative models and data modalities.

84 1. *Loading the dataset using a DataLoader.* The DataLoader class provides a consistent interface for
 85 loading, storing, and splitting different types of input data (e.g. tabular, time series, and survival
 86 data). Users can also provide meta-data to inform downstream algorithms, like specifying sensitive
 87 columns for privacy-preserving algorithms.

88 2. *Training the generator using a Plugin.* In synthcity, the users instantiate, train, and apply different
 89 data generators via the Plugin class. Each Plugin represents a specific data generation algorithm.
 90 The generator can be trained using the fit() method of a Plugin.

91 3. *Generating synthetic data.* After the Plugin is trained, the user can use the generate() method to
 92 generate synthetic data. Some plugins also allow for conditional generation.

93 4. *Evaluating synthetic data.* Synthcity provides a large set of metrics for evaluating different aspects
 94 of synthetic data. The Metrics class allows users to perform evaluation.

95 In addition, synthcity also has a Benchmark class that wraps around all four steps. This provides
 96 a one-line interface for comparing and evaluating different generators and produces an evaluation
 97 report at the end of the process.

98 2.2 Evaluation for diverse use cases

99 Synthetic data has many different use cases including fairness, privacy, data augmentation. As a
 100 benchmarking framework, Synthcity provides the users with a comprehensive list of metrics and
 101 routines to evaluate various aspects of synthetic data, including metrics that are specific to these use
 102 cases. In this section, we describe the use cases of synthetic data and how Synthcity performs its
 103 evaluation. A full list of metrics can be found in Appendix.

104 2.2.1 Standard data generation

105 Standard data generation refers to the most basic generation task, where the synthetic data should be
 106 generated as faithfully as possible to the real-data distribution [van Breugel et al., 2023, Hansen et al.,
 107 2023]. This is captured by the fidelity metrics.

108 **Fidelity.** The fidelity of synthetic data captures how much the synthetic data resembles real data. The
 109 fidelity metrics typically evaluate the closeness between the true distribution P and the distribution
 110 learned by the generator \hat{P} using samples from these two distributions. Synthcity supports distribu-
 111 tional divergence measures (e.g. Jensen-Shannon distance, Wasserstein distance, and maximal mean
 112 discrepancy) as well as two sample detection scores (i.e. scores that measure how well a classifier
 113 can distinguish real versus synthetic data) [Gretton et al., 2012, Lopez-Paz and Oquab, 2016, Snoko
 114 et al., 2018].

Use case	Method	Evaluation	Reference
Standard data generation	Generative model	Fidelity	Gretton et al. [2012]
Cross domain augmentation	Domain transfer	Utility	Bing et al. [2022]
ML fairness	Balancing distribution	Minority performance	Lu et al. [2018]
	Causal fairness	Algorithmic fairness	Xu et al. [2018]
Privacy preservation	Differential privacy	Privacy metrics	Abadi et al. [2016]
	Threat model	Attack simulation	Shokri et al. [2017]

Table 1: Synthcity is a unified framework to benchmark diverse use cases of synthetic data. It supports a range of methods and evaluation metrics, and also allows evaluation of off label uses.

115 2.2.2 Cross domain augmentation

116 Here we consider a dataset that is collected from multiple domains or sources (e.g. data from different
117 countries). Often the practitioner is interested in augmenting one particular data source that suffers
118 from data scarcity issues (e.g. it is difficult to collect data from remote areas) by leveraging other
119 related sources. This challenge has been studied in the deep generative model literature [Antoniou
120 et al., 2017, Dina et al., 2022, Das et al., 2022, Bing et al., 2022]. By learning domain-specific and
121 domain-agnostic representations, the generator is able to transfer knowledge across domains, making
122 data augmentation more efficient. Synthcity offers a clean interface so that the user can benchmark
123 the downstream utility of cross-domain generation using only one line of code..

124 **Utility.** Synthcity measures the performance for cross domain augmentation through its utility to
125 downstream tasks. Our approach adapts the common practice of train-on-synthetic evaluate-on-real
126 [Beaulieu-Jones et al., 2019], where a downstream predictive model is trained on fully synthetic
127 training data and then validated on real testing data.

128 For data augmentation, Synthcity augments the data-scarce domain in the training data \mathbb{D}_{train}^r with
129 the synthetic data \mathbb{D}^s . A predictive model is then trained on this augmented dataset and evaluated on
130 the domain of interest in the testing data \mathbb{D}_{test}^r . Synthcity supports various types of predictive tasks,
131 including regression, classification and survival analysis. In addition to linear predictive models,
132 synthcity supports Xgboost and neural nets as downstream models due to their wide adoption in data
133 analytics. In practice, the user may average the performance of several predictive models to reduce
134 the model uncertainty.

135 As a naive baseline, Synthcity reports the predictive performance where no data augmentation is
136 performed. Synthcity provides a pre-configured pipeline to automatically handle this entire procedure,
137 reducing the code and preventing mistakes.

138 2.2.3 Synthetic data for ML fairness

139 Existing research has considered two different ways where Synthetic data could promote fairness.
140 Table 2 shows the corresponding models in synthcity.

141 *1. Balancing distribution.* In this setting, certain groups of people are underrepresented in a dataset
142 that is used for training downstream ML systems. This may lead to a bias being introduced into these
143 ML systems [Lu et al., 2018, de Vassimon Manela et al., 2021, Kadambi, 2021]. As a remedy, one
144 could generate synthetic records for the minority group to augment the real data, thereby achieving
145 balance in distribution. This often requires the data generator to learn the conditional distribution
146 $P(X|G)$, where G is the group label.

147 *2. Causal fairness.* The second approach is to generate fairer synthetic data from a biased real dataset
148 and to use synthetic data alone in downstream tasks [Zemel et al., 2013, Xu et al., 2018, 2019a, van
149 Breugel et al., 2021]. In this setting, it is postulated that the real distribution $P(X)$ reflects existing
150 biases (e.g. unequal access to healthcare). The task for the generator is to learn a distribution $\hat{P}(X)$
151 that is free from such biases but also stay as close to $P(X)$ as possible (to ensure high data fidelity).
152 Typically, notions of causality are employed in the bias removal process.

153 **Fairness.** Synthcity allows users to benchmark both use cases by training a downstream predictive
154 model on the fully synthetic or augmented data and presenting their performance or characteristics.
155 For example, one can evaluate the performance gain on any specified (minority) group as an indicator
156 of the utility of synthetic data. In addition, Synthcity also supports standard algorithmic fairness
157 metrics for the trained predictive model, such as Fairness Through Unawareness, Demographic Parity
158 and Conditional Fairness [van Breugel et al., 2021]

159 **2.2.4 Synthetic data for privacy**

160 Methods for generating privacy-preserving synthetic data mainly fall into two categories: the ones
161 that employ differential privacy, and the ones that are designed for specific threat models.

162 *1. Differential privacy (DP).* DP is a formal way to describe how private a data generator is [Dwork,
163 2008]. Typically, generators with DP property introduce additional noise in the training procedure
164 [Jordon et al., 2022]. For example, adding noise in the gradient or using a noisy discriminator in a
165 GAN architecture [Abadi et al., 2016, Jordon et al., 2018, Long et al., 2019].

166 *2. Threat model (TM).* While DP focuses on giving formal guarantees, the TM approach is designed
167 for specific threat models, such as membership inference, attribute inference, and re-identification
168 [Shokri et al., 2017, Kosinski et al., 2013, Dinur and Nissim, 2003]. These models often involve
169 regularization terms designed to mitigate privacy attack risk [Yoon et al., 2020].

170 **Privacy.** Synthcity evaluates the privacy of synthetic data using a list of well-established privacy met-
171 rics (e.g. k-anonymity [Sweeney, 2002] and l-diversity [Machanavajjhala et al., 2007]). Furthermore,
172 it can measure the privacy of data by performing simulated privacy attacks (e.g. a re-identification
173 attack). The success (or failure) of such an attack quantifies the degree of privacy preservation.

174 **2.2.5 Evaluating off label use cases**

175 Synthcity allows users to conveniently evaluate the off label usage of generative models. For instance,
176 one could evaluate the privacy of synthetic data even if they are not generated by a privacy-enabled
177 generative model. As another example, one could evaluate the fairness for generative models that are
178 differentially private, thereby enabling studies like Ganev et al. [2022].

179 Off-label evaluation is made easy because Synthcity implements generative models and evaluation
180 metrics in two separate modules (PlugIns and Metrics). The consistent interface enables mix and
181 match of models and metrics to empower different benchmark studies.

182 **2.3 Baseline generative models**

183 As a benchmarking framework, Synthcity is a one-stop-shop for state-of-the-art benchmarks with
184 a large collection of baselines covering both deep generative models and other types of generative
185 models. In this way, the user can easily compare with a range of existing methods, without the need
186 to worry about implementation details or interfaces. Table 2 lists the generative models in synthcity
187 for different data modalities.

188 Synthcity covers all major families of deep generative models, including Generative adversarial
189 networks (GAN) [Goodfellow et al., 2020], Variational Autoencoders (VAE) [Kingma et al., 2019],
190 Normalizing flows (NF) [Papamakarios et al., 2021], as well as Diffusion models (DDPM) [Kingma
191 et al., 2021]. In the GAN family, Synthcity currently supports GOGGLE [Liu et al., 2023], CTGAN
192 Xu et al. [2019b], DPGAN [Xie et al., 2018], PATEGAN [Jordon et al., 2019], ADGAN [Yoon et al.,
193 2020], DECAF [van Breugel et al., 2021] for static data, Survival GAN [Norcliffe et al., 2023] for
194 censored data, TimeGAN [Yoon et al., 2019] for time series data, as well as RadialGAN [Yoon et al.,
195 2018] for multi-source data. In the VAE family, it supports TVAE [Xu et al., 2019b], RTVAE for static
196 data [Akrami et al., 2020], Survival VAE [Norcliffe et al., 2023] for censored data, and TimeVAE
197 [Yoon et al., 2019] for time series data. In the NF family, Synthcity implements the standard NF
198 [Papamakarios et al., 2021] for static data, Survival NF [Norcliffe et al., 2023] for censored data,
199 and FourierFlow [Alaa et al., 2021] of time series data. Synthcity also includes the TabDDPM

Data Modality	Model	Standard Gen	Privacy		Fairness	
			DP	TM	Balance	Causal
Static	Bayesian Net	✓				
	NF	✓				
	GREAT	✓				
	ARF	✓				
	GOGGLE	✓				
	TabDDPM	✓				
	TVAE	✓				✓
	RTVAE	✓				✓
	CTGAN	✓				✓
	AIM	✓		✓		
	PrivBayes	✓		✓		
	DPGAN	✓		✓		
	PATEGAN	✓		✓		
	ADSGAN	✓			✓	
DECAF	✓					✓
Static (Censored)	Survival GAN	✓		✓	✓	
	Survival VAE	✓				
	Survival NF	✓				
Time Series (regular, irregular, censored)	TimeGAN	✓			✓	
	TimeVAE	✓				
	FourierFlow*	✓				
	Probabilistic AR*	✓				
Multi-source	RadialGAN	✓			✓	

Table 2: Generative models available in synthcity for different data modalities and use cases. Abbreviations: Differential Privacy (DP), Threat Model (TM). *FourierFlow and Probabilistic AR is compatible with regular time series only while TimeGAN and TimeVAE support both.

200 [Kotelnikov et al., 2022] in the diffusion model family, and GREAT Borisov et al. [2022], which uses
201 auto-regressive generative LLM model.

202 In addition to deep generative models, Synthcity also contains generative models that are not based
203 on neural networks, such as Bayesian networks [Heckerman, 1997], AIM [McKenna et al., 2022],
204 Probabilistic Auto-regressive models [Deodatis and Shinozuka, 1988] and Adversarial random forests
205 (ARF) [Watson et al., 2023].

206 Synthcity implements all generative models using the PlugIn interface. This consistent approach
207 makes it easy to add additional generative models into the benchmark. The GitHub repository
208 includes tutorials and step-by-step instructions on how to add new models.

209 2.4 Architecture and hyper-parameters

210 To help researchers pinpoint the source of performance gain and conduct fair comparison, Synthcity
211 allows the user to incarnate all the deep generative models with different network architectures and
212 hyper-parameters.

213 The architecture can be specified when the user creates a model instance (i.e. a PlugIn). For
214 example, each time-series generative model can be configured using twelve different architectures,
215 including LSTM [Hochreiter and Schmidhuber, 1997], GRU [Dey and Salem, 2017], Transformer
216 [Vaswani et al., 2017], MLSTM-FCN [Karim et al., 2019], TCN [Lea et al., 2017], InceptionTime
217 and InceptionTimePlus [Ismail Fawaz et al., 2020], XceptionTime [Rahimian et al., 2020], ResCNN
218 [Sun et al., 2020], Omni-Scale CNN [Tang et al., 2020], and XCM [Fauvel et al., 2021]. The network
219 architectures compatible with other data modalities are tabulated in the Appendix.

220 Synthcity also has a consistent interface for dealing with hyper-parameters. The library allows
 221 the user to list, set, and sample all relevant hyper-parameters of a generative model. Furthermore,
 222 this interface is compatible with all popular hyper-parameter optimization libraries, such as Optuna
 223 [Akiba et al., 2019]. In this way, Synthcity allows the user to perform hyper-parameter search before
 224 evaluating on the best-performing setting to ensure a like-for-like comparison. Furthermore, Synthcity
 225 also allows the user to configure various early stopping rules to control and compare the training of
 226 generative models.

227 2.5 Data modalities

228 We emphasize that “tabular data” in fact encapsulates many different data modalities, including
 229 static tabular data, time series data, and censored survival data, all of which may contain a mix of
 230 continuous and discrete features (columns). Synthcity can also handle composite datasets composed
 231 of multiple subsets of data. We give a detailed description of the diverse tabular data modalities
 232 Synthcity supports in Figure 2 and further discuss them below. In future versions, we plan to include
 233 more data modalities including relational database-style data, richly annotated images, and texts.

234 2.5.1 Single dataset

235 We start by introducing the most fundamental case where there is a single training dataset (e.g. a
 236 single DataFrame in Pandas). We characterize the data modalities by two axes: the *observation*
 237 *pattern* and the *feature type*. Synthcity supports all combinations.

238 The observation pattern describes whether and how the data are collected over time. There are three
 239 most prominent patterns, static data, regular time series, and irregular time series, which are all
 240 supported by synthcity.

241 The second axis, feature type, describes the do-
 242 main of individual features. Synthcity supports
 243 multivariate tabular data with mixtures of con-
 244 tinuous, categorical, integer, and censored fea-
 245 tures. Censored features are common in survival
 246 analysis applications (e.g. healthcare and insur-
 247 ance). They are represented as a tuple (x, c) ,
 248 where $x \in \mathbb{R}^+$ represents the survival time and
 249 $c \in \{0, 1\}$ is the censoring indicator.

250 2.5.2 Composite dataset

251 A composite dataset involves multiple sub
 252 datasets. Synthcity can handle the benchmark-
 253 ing of different classes of composite datasets.
 254 Currently, it supports (1) static datasets with the
 255 same features, collected from different domains, (2) a static and a time series dataset. The latter
 256 setting is common in applications. For example, a patient’s medical record may contain both static
 257 demographic information and longitudinal follow up data.

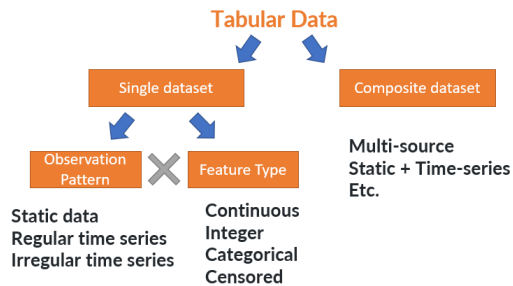


Figure 2: Supported tabular data modalities.

258 3 Comparison with existing libraries

259 In this section, we compare synthcity with other popular open source libraries for synthetic data
 260 generation to demonstrate its suitability as a comprehensive benchmark framework. Here we only
 261 consider the libraries that can generate synthetic data while preserving the statistical properties of
 262 real data, which includes YData Synthetic, Gretel Synthetics, SDV, DataSynthesizer, SmartNoise
 263 and nbsynthetic. Libraries that generate “fake” data for software testing are not considered because
 264 they do not attempt to learn the distribution of real data.

Setting \ Software	Synthcity	YData	Gretel	SDV	DataSynthesizer	SmartNoise	nbsynthetic
Data modalities							
Static data	✓	✓	✓	✓	✓	✓	✓
Regular time series	✓	✓	✓	✓			
Irregular time series	✓						
Censored features	✓						
Composite data	✓			✓			
Use cases							
Generation	✓	✓	✓	✓	✓	✓	✓
Fairness (balance)	✓	✓	✓	✓			✓
Fairness (causal)	✓						
Privacy (DP)	✓		✓			✓	
Privacy (TM)	✓						
Cross domain aug.	✓						

Table 3: The data modalities and use cases supported by synthcity and other open source synthetic data libraries. Comparisons are based on the software versions available at the time of writing.

265 Table 3 shows that synthcity supports much broader use cases and data modalities than the alternatives.
266 The existing libraries often focus on a single data modality or use case because they are intended as a
267 solution to a specific problem rather than a benchmark framework. Furthermore, Synthcity includes
268 many more data generators, including all major flavors of deep generative models as well as traditional
269 generative models. It also contains a built-in evaluation module that assesses various aspects of the
270 generator. A more detailed comparison of the supported data generators and evaluation metrics are
271 available in the Appendix. The broad coverage of data modalities, use cases, data generators and
272 evaluation metrics make Synthcity unique in its capacity for model evaluation and comparison.

273 4 Illustrative case studies

274 In this section, we present two illustrative use cases to show the type of benchmark studies that
275 Synthcity can facilitate. We stress that these examples do not cover the full capability of Synthcity
276 and they are used as illustrations.

277 4.1 Static tabular data generative model benchmark

278 We study which generative model has the strongest performance in generating synthetic tabular data.
279 Synthcity allows us to compare a variety of state-of-the-art algorithms in this study, including ARF,
280 GOGGLE, TabDDPM, CTGAN and TVAE. These algorithms are representative of broader families
281 of generative models such as GANs, VAEs, Diffusion models, and forest-based generative models.

282 Similar to prior tabular data benchmarks [Grinsztajn et al., 2022], we have selected 18 datasets from
283 the OpenML benchmark, which cover common regression and classification datasets encountered in
284 data science projects [Vanschoren et al., 2014]. The datasets cover a range of sample sizes (4,209 to
285 1,025,010) and feature counts (5 to 771).

286 Synthcity can automatically calculate more than 25 supported evaluation metrics in a benchmark. In
287 this study, we focus on evaluating the fidelity of synthetic data. Similar to Liu et al. [2023], we report
288 the average of the three-dimensional metrics (α -precision, β -recall, and authenticity), as proposed in
289 Alaa et al. [2022], as a measure of data quality—whether the synthetic data are realistic, cover the
290 true data distribution, and are generalized. Furthermore, we report the detection score, which reflects
291 how often the synthetic data can be distinguished from the real data. To reduce the variability from
292 the classifiers, we report the average AUROC scores from three different post-hoc data classifiers, as
293 in Liu et al. [2023].

294 Table 4 shows the experimental results averaged across all the datasets. We observe that the ARF
295 model achieves strong performance in the quality score and stands out in terms of the detection score.

296 This suggests that the tree-based generative models are strong competitors to deep generative models
 297 for static tabular data. And this area is a promising avenue for further research.

	Quality	Detection
ARF	0.5475	0.6721
GOGGLE	0.4054	0.9261
TabDDPM	0.5436	0.7074
CTGAN	0.5475	0.7758
TVAE	0.5487	0.7389

Table 4: Benchmark results for static tabular data generation. Quality: the higher the better; Detection: the lower the better.

298 4.2 Tabular data fairness and augmentation benchmark

299 We consider a benchmark study on cross-domain data augmentation for improving predictive perfor-
 300 mance on minority groups. We use the SIVEP-Gripe public dataset as an illustrative example, which
 301 contains anonymized records of COVID-19 patients in Brazil [Baqui et al., 2021]. In this dataset,
 302 'Mixed' and 'White' are the majority ethnicity groups while 'Black', 'East Asian' and 'Indigenous'
 303 are the minority groups (accounting for less than 10% of the total population). The dataset is used for
 304 training a downstream model to predict COVID-19 mortality. Due to the distributional imbalance,
 305 the downstream predictor is likely to under-perform on the minority groups, which may raise fairness
 306 issues (Section 2.2.3). This study aims to benchmark the utility of different generative models for
 307 data augmentation by measuring the AUROC of mortality prediction on the minority groups.

308 Synthcity allows us to easily compare RadialGAN, which was designed for cross-domain data
 309 augmentation, and the conditional generative models (TabDDPM, CTGAN, and TVAE). We use
 310 Synthcity’s pre-configured pipeline for data augmentation benchmark, which reduces the amount of
 311 code and prevents data leakage. Synthcity also allows us to evaluate the performance gain for different
 312 downstream models, and we have selected a multi-layer perceptron classifier and a xgBoost classifier.
 313 The results are listed in Table 5. We observe that data augmentation consistently improves the
 314 accuracy of mortality prediction for minority groups. TabDDPM, a novel diffusion model, achieves
 315 the best overall performance, followed by RadialGAN.

	Neural net	XgBoost
TabDDPM	0.7241	0.7786
RadialGAN	0.7137	0.7627
CTGAN	0.6477	0.7507
TVAE	0.3623	0.7794
Baseline	0.3244	0.7327

Table 5: Benchmark results for cross-domain data augmentation. The metric reported is the AUROC of mortality prediction on the minority groups

316 5 Discussion

317 Synthetic data is an emerging field where many novel algorithms have been proposed; yet there
 318 lacks an easy way to benchmark generative models across different desired or off label use cases,
 319 compare them with diverse baselines, and explain their performance gain. In this work, we present
 320 the open source Synthcity library as a solution to the benchmark challenge. Synthcity contains many
 321 built-in generative models, architectures and evaluation metrics, which are easily accessible through
 322 end-to-end evaluation pipelines. It can help researchers to perform in-depth and comprehensive
 323 benchmark studies with minimal programming effort.

324 **6 Funding Statement**

325 This work is supported by Cancer Research UK. The computational experiments in this work was
326 supported by Azure sponsorship credits granted by Microsoft’s AI for Good Research Lab.

327 **7 Acknowledgment**

328 We thank all the open source contributors to the Synthcity library. The full list of contributors is
329 available on GitHub. Special thanks goes to Bogdan Cebere, who made major contributions to
330 Synthcity while he was working as a Research Engineer at the van der Schaar Lab.

331 **References**

- 332 Martín Abadi, Andy Chu, Ian Goodfellow, H Brendan Mcmahan, Ilya Mironov, Kunal Talwar, and
333 Li Zhang. Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC*
334 *Conference on Computer and Communications Security*, 2016. doi: 10.1145/2976749. URL
335 <http://dx.doi.org/10.1145/2976749.2978318>.
- 336 Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna:
337 A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM*
338 *SIGKDD international conference on knowledge discovery & data mining*, pages 2623–2631,
339 2019.
- 340 Haleh Akrami, Sergul Aydore, Richard M Leahy, and Anand A Joshi. Robust variational autoencoder
341 for tabular data with beta divergence. *arXiv preprint arXiv:2006.08204*, 2020.
- 342 Ahmed Alaa, Alex James Chan, and Mihaela van der Schaar. Generative time-series modeling with
343 fourier flows. In *International Conference on Learning Representations*, 2021.
- 344 Ahmed Alaa, Boris Van Breugel, Evgeny S Saveliev, and Mihaela van der Schaar. How faithful
345 is your synthetic data? sample-level metrics for evaluating and auditing generative models. In
346 *International Conference on Machine Learning*, pages 290–306. PMLR, 2022.
- 347 Antreas Antoniou, Amos Storkey, and Harrison Edwards. Data Augmentation Generative Adversarial
348 Networks. 11 2017. URL <http://arxiv.org/abs/1711.04340>.
- 349 Pedro Baqui, Valerio Marra, Ahmed M Alaa, Ioana Bica, Ari Ercole, and Mihaela van der Schaar.
350 Comparing covid-19 risk factors in brazil using machine learning: the importance of socioeconomic,
351 demographic and structural factors. *Scientific reports*, 11(1):1–10, 2021.
- 352 Brett K Beaulieu-Jones, Zhiwei Steven Wu, Chris Williams, Ran Lee, Sanjeev P Bhavnani,
353 James Brian Byrd, and Casey S Greene. Privacy-preserving generative deep neural networks
354 support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122,
355 2019.
- 356 Simon Bing, Andrea Dittadi, Stefan Bauer, and Patrick Schwab. Conditional Generation of Medical
357 Time Series for Extrapolation to Underrepresented Populations. 2022.
- 358 Sam Bond-Taylor, Adam Leach, Yang Long, and Chris G. Willcocks. Deep Generative Modelling:
359 A Comparative Review of VAEs, GANs, Normalizing Flows, Energy-Based and Autoregressive
360 Models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3 2021. doi: 10.1109/
361 TPAMI.2021.3116668. URL <http://arxiv.org/abs/2103.04922><http://dx.doi.org/10.1109/TPAMI.2021.3116668>.
- 363 Vadim Borisov, Kathrin Seßler, Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. Language
364 models are realistic tabular data generators. *arXiv preprint arXiv:2210.06280*, 2022.
- 365 Hari Prasanna Das, Ryan Tran, Japjot Singh, Xiangyu Yue, Geoffrey Tison, Alberto Sangiovanni-
366 Vincentelli, and Costas J Spanos. Conditional synthetic data generation for robust machine learning
367 applications with limited pandemic data. In *Proceedings of the AAAI Conference on Artificial*
368 *Intelligence*, volume 36, pages 11792–11800, 2022.

- 369 Jeffrey Dastin. Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*,
370 2018.
- 371 Daniel de Vassimon Manela, David Errington, Thomas Fisher, Boris van Breugel, and Pasquale
372 Minervini. Stereotype and Skew: Quantifying Gender Bias in Pre-trained and Fine-tuned Language
373 Models. In *Proceedings of the 16th Conference of the European Chapter of the Association for*
374 *Computational Linguistics: Main Volume*, pages 2232–2242, Online, 4 2021. Association for
375 Computational Linguistics. URL <https://www.aclweb.org/anthology/2021.eacl-main>.
376 190.
- 377 George Deodatis and M Shinozuka. Auto-regressive model for nonstationary stochastic processes.
378 *Journal of engineering mechanics*, 114(11):1995–2012, 1988.
- 379 Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017*
380 *IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600.
381 IEEE, 2017.
- 382 Ayesha S. Dina, A. B. Siddique, and D. Manivannan. Effect of Balancing Data Using Synthetic Data
383 on the Performance of Machine Learning Classifiers for Intrusion Detection in Computer Networks.
384 4 2022. doi: 10.48550/arxiv.2204.00144. URL <https://arxiv.org/abs/2204.00144v1>.
- 385 Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *Proceedings of the*
386 *twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*,
387 pages 202–210, 2003.
- 388 Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and*
389 *applications of models of computation*, pages 1–19. Springer, 2008.
- 390 Kevin Fauvel, Tao Lin, Véronique Masson, Éliisa Fromont, and Alexandre Termier. Xcm: An
391 explainable convolutional neural network for multivariate time series classification. *Mathematics*,
392 9(23):3137, 2021.
- 393 Georgi Ganey, Bristena Oprisanu, and Emiliano De Cristofaro. Robin hood and matthew effects:
394 Differential privacy has disparate impact on synthetic data. In *International Conference on Machine*
395 *Learning*, pages 6944–6959. PMLR, 2022.
- 396 Milena A Gianfrancesco, Suzanne Tamang, Jinoos Yazdany, and Gabriela Schmajuk. Potential biases
397 in machine learning algorithms using electronic health record data. *JAMA internal medicine*, 178
398 (11):1544–1547, 2018.
- 399 Andre Goncalves, Priyadip Ray, Braden Soper, Jennifer Stevens, Linda Coyle, and Ana Paula Sales.
400 Generation and evaluation of synthetic patient data. *BMC medical research methodology*, 20(1):
401 1–40, 2020.
- 402 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
403 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the*
404 *ACM*, 63(11):139–144, 2020.
- 405 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A
406 kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- 407 Leo Grinsztajn, Edouard Oyallon, and Gael Varoquaux. Why do tree-based models still outperform
408 deep learning on typical tabular data? In *Thirty-sixth Conference on Neural Information Processing*
409 *Systems Datasets and Benchmarks Track*, 2022.
- 410 Lasse Hansen, Nabeel Seedat, Mihaela van der Schaar, and Petrovic Andrija. Reimagining syn-
411 thetic data generation through data-centric ai: A comprehensive benchmark. *Advances in neural*
412 *information processing systems*, 30, 2023.
- 413 David Heckerman. Bayesian networks for data mining. *Data mining and knowledge discovery*, 1:
414 79–119, 1997.
- 415 Stella Ho, Youyang Qu, Bruce Gu, Longxiang Gao, Jianxin Li, and Yong Xiang. Dp-gan: Differen-
416 tially private consecutive data publishing using generative adversarial nets. *Journal of Network*
417 *and Computer Applications*, 185:103066, 2021.

418 Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):
419 1735–1780, 1997.

420 Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F Schmidt,
421 Jonathan Weber, Geoffrey I Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petit-
422 jean. Inceptiontime: Finding alexnet for time series classification. *Data Mining and Knowledge*
423 *Discovery*, 34(6):1936–1962, 2020.

424 James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data with
425 differential privacy guarantees. In *International conference on learning representations*, 2018.

426 James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data with
427 differential privacy guarantees. In *International conference on learning representations*, 2019.

428 James Jordon, Lukasz Szpruch, Florimond Houssiau, Mirko Bottarelli, Giovanni Cherubin, Carsten
429 Maple, Samuel N. Cohen, and Adrian Weller. Synthetic Data – what, why and how? 5 2022. doi:
430 10.48550/arxiv.2205.03257. URL <https://arxiv.org/abs/2205.03257v1>.

431 Achuta Kadambi. Achieving fairness in medical devices. *Science*, 372(6537):30–31, 2021.

432 Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Samuel Harford. Multivariate lstm-fcns
433 for time series classification. *Neural networks*, 116:237–245, 2019.

434 Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances*
435 *in neural information processing systems*, 34:21696–21707, 2021.

436 Diederik P Kingma, Max Welling, et al. An introduction to variational autoencoders. *Foundations*
437 *and Trends® in Machine Learning*, 12(4):307–392, 2019.

438 Theodora Kokosi and Katie Harron. Synthetic data in medical research. *BMJ Medicine*, 1(1), 2022.

439 Michal Kosinski, David Stillwell, and Thore Graepel. Private traits and attributes are predictable
440 from digital records of human behavior. *Proceedings of the national academy of sciences*, 110(15):
441 5802–5805, 2013.

442 Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, and Artem Babenko. Tabddpm: Modelling
443 tabular data with diffusion models. *arXiv preprint arXiv:2209.15421*, 2022.

444 Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional
445 networks for action segmentation and detection. In *proceedings of the IEEE Conference on*
446 *Computer Vision and Pattern Recognition*, pages 156–165, 2017.

447 Tennison Liu, Zhaozhi Qian, Jeroen Berrevoets, and Mihaela van der Schaar. Goggle: Generative
448 modelling for tabular data by learning relational structure. In *The Eleventh International Conference*
449 *on Learning Representations*, 2023.

450 Yunhui Long, Boxin Wang, Zhuolin Yang, Bhavya Kailkhura, Aston Zhang, Carl A. Gunter, and
451 Bo Li. G-PATE: Scalable Differentially Private Data Generator via Private Aggregation of Teacher
452 Discriminators. In *Advances in Neural Information Processing Systems*, 6 2019. doi: 10.48550/
453 arxiv.1906.09338. URL <https://arxiv.org/abs/1906.09338v2>.

454 David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. *arXiv preprint*
455 *arXiv:1610.06545*, 2016.

456 Kaiji Lu, Piotr Mardziel, Fangjing Wu, Preetam Amancharla, and Anupam Datta. Gender Bias in
457 Neural Natural Language Processing. *CoRR*, abs/1807.11714, 2018. URL <http://arxiv.org/abs/1807.11714>.

459 Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkitasubrama-
460 niam. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from*
461 *Data (TKDD)*, 1(1):3–es, 2007.

462 Ryan McKenna, Brett Mullins, Daniel Sheldon, and Gerome Miklau. Aim: An adaptive and iterative
463 mechanism for differentially private synthetic data. *arXiv preprint arXiv:2201.12677*, 2022.

464 Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey
465 on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

- 466 Alexander Norcliffe, Bogdan Cebere, Fergus Imrie, Pietro Lio, and Mihaela van der Schaar. Survival-
467 gan: Generating time-to-event data for survival analysis. In *International Conference on Artificial*
468 *Intelligence and Statistics*, pages 10279–10304. PMLR, 2023.
- 469 George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji
470 Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *The Journal of*
471 *Machine Learning Research*, 22(1):2617–2680, 2021.
- 472 Mayana Pereira, Meghana Kshirsagar, Sumit Mukherjee, Rahul Dodhia, and Juan Lavista Ferres.
473 An analysis of the deployment of models trained on private tabular synthetic data: Unexpected
474 surprises. *arXiv preprint arXiv:2106.10241*, 2021.
- 475 Elahe Rahimian, Soheil Zabihi, Seyed Farokh Atashzar, Amir Asif, and Arash Mohammadi. Xcep-
476 tiontime: independent time-window xceptiontime architecture for hand gesture classification. In
477 *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing*
478 *(ICASSP)*, pages 1304–1308. IEEE, 2020.
- 479 Nithya Sambasivan, Shivani Kapania, Hannah Highfill, Diana Akrong, Praveen Paritosh, and Lora M
480 Aroyo. “everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai.
481 In *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15,
482 2021.
- 483 Divya Saxena and Jiannong Cao. Generative adversarial networks (gans) challenges, solutions, and
484 future directions. *ACM Computing Surveys (CSUR)*, 54(3):1–42, 2021.
- 485 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks
486 against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*, pages
487 3–18. IEEE, 2017.
- 488 Joshua Snoke, Gillian M Raab, Beata Nowok, Chris Dibben, and Aleksandra Slavkovic. General
489 and specific utility measures for synthetic data. *Journal of the Royal Statistical Society: Series A*
490 *(Statistics in Society)*, 181(3):663–688, 2018.
- 491 Weitong Sun, Yuping Su, Xia Wu, and Xiaojun Wu. A novel end-to-end 1d-rescnn model to remove
492 artifact from eeg signals. *Neurocomputing*, 404:108–121, 2020.
- 493 Latanya Sweeney. k-anonymity: A model for protecting privacy. *International journal of uncertainty,*
494 *fuzziness and knowledge-based systems*, 10(05):557–570, 2002.
- 495 Wensi Tang, Guodong Long, Lu Liu, Tianyi Zhou, Michael Blumenstein, and Jing Jiang. Omni-scale
496 cnns: a simple and effective kernel size configuration for time series classification. *arXiv preprint*
497 *arXiv:2002.10061*, 2020.
- 498 Jason Tashea. Courts Are Using AI to Sentence Criminals. That Must
499 Stop Now. *WIRED*, 4 2017. URL [https://www.wired.com/2017/04/
500 courts-using-ai-sentence-criminals-must-stop-now/](https://www.wired.com/2017/04/courts-using-ai-sentence-criminals-must-stop-now/).
- 501 Allan Tucker, Zhenchen Wang, Ylenia Rotalinti, and Puja Myles. Generating high-fidelity synthetic
502 patient data for assessing machine learning healthcare software. *NPJ digital medicine*, 3(1):1–13,
503 2020.
- 504 Boris van Breugel, Trent Kyono, Jeroen Berrevoets, and Mihaela van der Schaar. Decaf: Generating
505 fair synthetic data using causally-aware generative networks. *Advances in Neural Information*
506 *Processing Systems*, 34:22221–22233, 2021.
- 507 Boris van Breugel, Zhaozhi Qian, and Mihaela van der Schaar. Synthetic data, real errors: how (not)
508 to publish and use synthetic data. 2023.
- 509 Joaquin Vanschoren, Jan N Van Rijn, Bernd Bischl, and Luis Torgo. Openml: networked science in
510 machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014.
- 511 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz
512 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*
513 *systems*, 30, 2017.

- 514 Zhenchen Wang, Puja Myles, and Allan Tucker. Generating and evaluating synthetic uk primary care
515 data: preserving data utility & patient privacy. In *2019 IEEE 32nd International Symposium on*
516 *Computer-Based Medical Systems (CBMS)*, pages 126–131. IEEE, 2019.
- 517 Zhenchen Wang, Puja Myles, and Allan Tucker. Generating and evaluating cross-sectional synthetic
518 electronic healthcare data: Preserving data utility and patient privacy. *Computational Intelligence*,
519 *37(2):819–851*, 2021.
- 520 David S Watson, Kristin Blesch, Jan Kapar, and Marvin N Wright. Adversarial random forests for
521 density estimation and generative modeling. In *International Conference on Artificial Intelligence*
522 *and Statistics*, pages 5357–5375. PMLR, 2023.
- 523 Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative
524 adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.
- 525 Depeng Xu, Shuhan Yuan, Lu Zhang, and Xintao Wu. Fairgan: Fairness-aware generative adversarial
526 networks. In *2018 IEEE International Conference on Big Data (Big Data)*, pages 570–575, 2018.
- 527 Depeng Xu, Yongkai Wu, Shuhan Yuan, Lu Zhang, and Xintao Wu. Achieving causal fairness
528 through generative adversarial networks. In *IJCAI International Joint Conference on Artificial*
529 *Intelligence*, volume 2019-August, pages 1452–1458. International Joint Conferences on Artificial
530 Intelligence, 2019a. ISBN 9780999241141. doi: 10.24963/ijcai.2019/201.
- 531 Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. Modeling tabular
532 data using conditional gan. *Advances in Neural Information Processing Systems*, 32, 2019b.
- 533 Jinsung Yoon, James Jordon, and Mihaela Schaar. Radialgan: Leveraging multiple datasets to
534 improve target-specific predictive models using generative adversarial networks. In *International*
535 *Conference on Machine Learning*, pages 5699–5707. PMLR, 2018.
- 536 Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. Time-series generative adversarial
537 networks. *Advances in neural information processing systems*, 32, 2019.
- 538 Jinsung Yoon, Lydia N Drumright, and Mihaela Van Der Schaar. Anonymization through data
539 synthesis using generative adversarial networks (ads-gan). *IEEE journal of biomedical and health*
540 *informatics*, 24(8):2378–2388, 2020.
- 541 Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations.
542 In *International conference on machine learning*, pages 325–333, 2013.
- 543 Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation
544 using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference*
545 *on computer vision*, pages 2223–2232, 2017.

546 **A Appendix**

547 **A.1 Code availability**

548 The code for the illustrative use cases are available on GitHub
 549 <https://github.com/vanderschaarlab/synthcity-benchmarking>. The synthcity library is avail-
 550 able on pip and GitHub. The tutorials folder contains additional illustrative examples.

551 **A.2 Supported algorithms and metrics**

Aspect	Evaluation Metric \ Software	Synthcity	YData	Gretel	SDV	DataSynthesizer	SmartNoise	nbsynthetic
Fidelity	Jensen-Shannon distance	✓						
	Wasserstein distance	✓						
	Total variation distance				✓			
	KL divergence	✓						
	Skewness				✓			
	Max-mean discrepancy	✓						✓
	KS test	✓			✓			✓
	PRDC	✓						
	Alpha-precision	✓						
	Survival Kaplan-Meier dist.	✓						
	Detection: linear	✓			✓			
	Detection: NN	✓						
	Detection: XGB	✓						
Detection: Linear					✓			
Utility	Linear model	✓			✓			
	MLP	✓			✓			
	XGBoost	✓			✓			
	Static survival	✓						
	Time-series	✓						
	Survival time-series	✓						
Privacy	Correct attribution prob.	✓			✓			
	K-anonymity	✓						
	K-map	✓						
	Delta-presence	✓						
	L-diversity	✓						
	DOMIAS	✓						
Identifiability score	✓							

Table 6: The evaluation metrics supported by synthcity and other open source synthetic data libraries. Comparisons are based on the software versions available at the time of writing.

Static	Censored	Time Series
Fully connected	Weibull AFT	LSTM
Residual network	Cox PH	GRU
TabNet	Random Survival Forest	RNN
	Survival Xgboost	Transformer
	Deephit	MLSTM_FCN
	Tenn	TCN
	Date	InceptionTime
		InceptionTimePlus
		XceptionTime
		ResCNN
		OmniScaleCNN
		XCM

Table 7: Available network architectures and survival models in synthcity for different data modalities. These components are compatible with multiple algorithms.

Algorithm \Software	Synthcity	YData	Gretel	SDV	DataSynthesizer	SmartNoise	nbsynthetic
AIM	✓						
GREAT	✓						
TabDDPM	✓						
ARF	✓						
GOGGLE	✓						
CTGAN	✓	✓		✓			✓
ACTGAN			✓				
TVAE	✓			✓			
Bayesian Network	✓						
Normalizing Flows	✓						
Survial GAN	✓						
Survival VAE	✓						
DoppelGANger			✓				
TimeGAN	✓	✓					
FourierFlows	✓						
Probabilistic AR	✓			✓			
DECAF	✓						
RadialGAN	✓						
ADSGAN	✓						
DPGAN	✓		✓				
PATEGAN	✓					✓	
PrivBayes	✓				✓	✓	

Table 8: The data generating algorithms supported by synthcity and other open source synthetic data libraries. Comparisons are based on the software versions available at the time of writing.