

---

# GlucoSynth: Generating Differentially-Private Synthetic Glucose Traces

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 We focus on the problem of generating high-quality, private synthetic glucose  
2 traces, a task generalizable to many other time series sources. Existing methods for  
3 time series data synthesis, such as those using Generative Adversarial Networks  
4 (GANs), are not able to capture the innate characteristics of glucose data and cannot  
5 provide any formal privacy guarantees without severely degrading the utility of the  
6 synthetic data. In this paper we present GlucoSynth, a novel privacy-preserving  
7 GAN framework to generate synthetic glucose traces. The core intuition behind our  
8 approach is to conserve relationships amongst motifs (glucose events) within the  
9 traces, in addition to temporal dynamics. Our framework incorporates differential  
10 privacy mechanisms to provide strong formal privacy guarantees. We provide a  
11 comprehensive evaluation on the real-world utility of the data using 1.2 million  
12 glucose traces; GlucoSynth outperforms all previous methods in its ability to  
13 generate high-quality synthetic glucose traces with strong privacy guarantees.

## 14 1 Introduction

15 The sharing of medical time series data can facilitate therapy development. As a motivating example,  
16 sharing glucose traces can contribute to the understanding of diabetes disease mechanisms and the  
17 development of artificial insulin delivery systems that improve people with diabetes' quality of life.  
18 Unsurprisingly, there are serious legal and privacy concerns (e.g., HIPAA, GDPR) with the sharing of  
19 such granular, longitudinal time series data in a medical context [1]. One solution is to generate a set  
20 of synthetic traces from the original traces. In this way, the synthetic data may be shared publicly in  
21 place of the real ones with significantly reduced privacy and legal concerns.

22 This paper focuses on the problem of generating high-quality, privacy-preserving synthetic glucose  
23 traces, a task which generalizes to other time series sources and application domains, including  
24 activity sequences, inpatient events, hormone traces and cyber-physical systems. Specifically, we  
25 focus on long (over 200 timesteps), bounded, univariate time series glucose traces. We assume  
26 that available data does not have any labels or extra information including features or metadata,  
27 which is quite common, especially in diabetes. Continuous Glucose Monitors (CGMs) easily and  
28 automatically send glucose measurements taken subcutaneously at fixed intervals (e.g., every 5  
29 minutes) to data storage facilities, but tracking other sources of diabetes-related data is challenging  
30 [2]. We characterize the quality of the generated traces based on three criteria— synthetic traces  
31 should (1) conserve characteristics of the real data, i.e., glucose dynamics and control-related metrics  
32 (*fidelity*); (2) contain representation of diverse types of realistic traces, without the introduction of  
33 anomalous patterns that do not occur in real traces (*breadth*); and (3) be usable in place of the original  
34 data for real-world use cases (*utility*).

35 Generative Adversarial Networks (GANs) [3] have shown promise in the generation of time series  
36 data. However, previous methods for time series synthesis, e.g., [4, 5, 6], suffer from one or more of

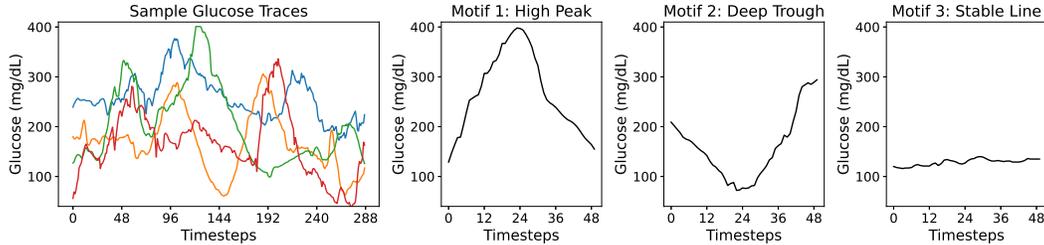


Figure 1: Example Real Glucose Traces and Glucose Motifs from our Dataset.

37 the following issues when applied to glucose traces: 1) surprisingly, they do not generate realistic  
 38 synthetic glucose traces – in particular, they produce human physiologically impossible phenomenon  
 39 in the traces; 2) they require additional information (features, metadata or labels) to guide the model  
 40 learning which are not available for our traces; 3) they do not include any privacy guarantees, or, in  
 41 order to uphold a strong formal privacy guarantee, severely degrade the utility of the synthetic data.

42 Generating high-quality synthetic glucose traces is a difficult task due to the innate characteristics of  
 43 glucose data. Glucose traces can be best understood as sequences of events, which we call *motifs*,  
 44 shown in Figure 1, and they are more event-driven than many other types of time series. As such, a  
 45 current glucose value may be more influenced by an event that occurred in the far past compared to  
 46 values from immediate previous timesteps. For example, a large meal eaten earlier in the day (30-90  
 47 minutes ago) may influence a patient’s glucose more than the glucose values from the past 15 minutes.  
 48 As a result, although there is some degree of temporal dependence within the traces, *only* conserving  
 49 the immediate temporal relationships amongst values at previous timesteps does not adequately  
 50 capture the dynamics of this type of data. In particular, we find that the main reason previous methods  
 51 fail is because they may not sufficiently learn event-related characteristics of glucose traces.

52 **Contributions.** We present *GlucoSynth*, a privacy-preserving GAN framework to generate synthetic  
 53 glucose traces. The core intuition behind our approach is to conserve relationships amongst motifs  
 54 (events) within the traces, in addition to the typical temporal dynamics contained within time series.  
 55 We formalize the concept of motifs and define a notion of *motif causality*, inspired from Granger  
 56 causality [7], which characterizes relationships amongst sequences of motifs within time series traces  
 57 (Section 4). We define a local motif loss to first train a motif causality block that learns the motif  
 58 causal relationships amongst the sequences of motifs in the real traces. The block outputs a motif  
 59 causality matrix, that quantifies the causal value of seeing one particular motif after some other motif.  
 60 Unrealistic motif sequences (such as a peak to an immediate drop in glucose values) will have causal  
 61 relationships close to 0 in the causality matrix. We build a novel GAN framework that is trained  
 62 to optimize motif causality within the traces in addition to temporal dynamics and distributional  
 63 characteristics of the data (Section 5). Explicitly, the generator computes a motif causality matrix  
 64 from each batch of synthetic data it generates, and compares it with the real causality matrix. As  
 65 such, as the generator learns to generate synthetic data that yields a realistic causal matrix (thereby  
 66 identifying appropriate causal relationships from the motifs), it implicitly learns not to generate  
 67 unrealistic motif sequences. We also integrate differential privacy (DP) [8] into the framework  
 68 (Section 6), which provides an intuitive bound on how much information may be disclosed about  
 69 any individual in the dataset, allowing the *GlucoSynth* model to be trained with privacy guarantees.  
 70 Finally, in Section 7, we present a comprehensive evaluation using 1.2 million glucose traces from  
 71 individuals with diabetes collected across 2022, showcasing the suitability of our model to outperform  
 72 all previous models and generate high-quality synthetic glucose traces with strong privacy guarantees.

## 73 2 Related Work

74 We focus the scope of our comparison on current state-of-the-art methods for synthetic time se-  
 75 ries which all build upon Generative Adversarial Networks (GANs) [3] and transformation-based  
 76 approaches [9]. An extended related work is in Appendix A.

77 **Time Series.** Brophy et al. [10] provides a survey of GANs for time series synthesis. TimeGan [4] is a  
 78 popular benchmark that jointly learns an embedding space using supervised and adversarial objectives  
 79 in order to capture the temporal dynamics amongst traces. Esteban et al. [11] develops two time  
 80 series GAN models (RGAN/RCGAN) with RNN architectures, conditioned on auxiliary information

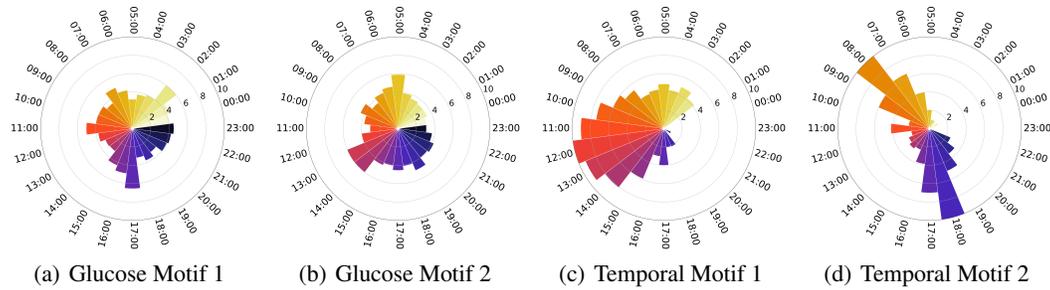


Figure 2: Temporal Distributions of Sample Motifs. Each radial graph displays the temporal distribution of a motif; there are 24 radial bars from 00:00 to 23:00, and each segment displays the % of motif occurrences by each hour. Glucose motifs 1 and 2 are from Fig. 1; they are not temporally-dependent and show up across the day. Temporal motifs 1 and 2 are from a cardiology dataset [15].

81 provided at each timestep during training. TTS-GAN [5] trains a GAN model that uses a transformer  
 82 encoding architecture in order to best preserve temporal dynamics. Transformation-based approaches  
 83 such as real-valued non-volume preserving transformations (NVP) [9] and Fourier Flows (FF) [12],  
 84 have also had success for time series data. These methods model the underlying distribution of the  
 85 real data to transform the input traces into a synthetic data set. Methods that only focus on learning the  
 86 temporal or distributional dynamics in time series are not sufficient for generating realistic synthetic  
 87 glucose traces due to the lack of temporal dependence within sequences of glucose motifs.

88 **Differentially-Private GANs.** To protect sensitive data, several GAN architectures have been  
 89 designed to incorporate privacy-preserving noise needed to satisfy differential privacy guarantees [13].  
 90 Frigerio et al. [14] extends a simple differentially-private architecture (dpGAN) to time-series data  
 91 and RDP-CGAN [6] develops a convolutional GAN architecture specifically for medical data. These  
 92 methods find large gaps in performance between the non-private and private models. Providing strong  
 93 theoretical DP guarantees using these methods often results in synthetic data with too little fidelity  
 94 for use in real-world scenarios. Our framework carefully integrates DP into the motif causality block  
 95 and each network of the GAN, resulting in a better utility-privacy tradeoff than previous methods.

### 96 3 Preliminaries

#### 97 3.1 Motifs

98 Glucose (and many other) traces can be best understood as sequences of events or *motifs*. Motifs  
 99 characterize phenomenon in the traces, such as peaks or troughs. We define a *motif*,  $\mu$ , as a short,  
 100 ordered sequence of values ( $v$ ) of specified length  $\tau$ ,  $\mu = [v_i, v_{i+1}, \dots, v_{i+\tau}]$  and  $\sigma$  is a tolerance  
 101 value to allow approximate matching (within  $\sigma$  for each value). Some examples of glucose traces  
 102 and motifs are shown in Figure 1. We denote a set of  $n$  time series traces as  $X = [x_1, \dots, x_n]$ . Each  
 103 time series may be represented as a sequence of motifs:  $x_i = [\mu_{i_1}, \mu_{i_2}, \dots]$  where each  $i_j$  gives the  
 104 index of the motif in the set that matches  $x_{i_j \cdot \tau}, \dots, x_{(i_j+1) \cdot \tau - 1}$ . Given the motif length  $\tau$ , the motif  
 105 set is the union of all size- $\tau$  chunks in the traces. This definition is chosen for a straightforward  
 106 implementation but motifs can be generated in other ways, such as through the use of rolling windows  
 107 or signal processing techniques [16, 17]. Motifs are pulled from the data such that there is always a  
 108 match from a trace motif to a motif from the set (if multiple matches, the closest one is chosen).

#### 109 3.2 Glucose Dynamics (Why Standard Approaches Fail)

110 We first present a study of the characteristics of glucose data in order to motivate the development  
 111 of our framework. Although there are general patterns in sequences of glucose motifs (e.g., motif  
 112 patterns corresponding to patients that eat 2x vs. 3x a day), individual glucose motifs are typically  
 113 not time-dependent, as illustrated in Figure 2. The radial graphs display the temporal distribution of  
 114 the first two glucose motifs from Figure 1 and two temporally-dependent motifs from a cardiology  
 115 dataset [15]. There are 24 radial bars from 00:00 to 23:00 for each hour of the day, and the bar value  
 116 is the percentage of total motif occurrences at that hour across the entire dataset (i.e., value of 10  
 117 would indicate that 10% of the time that motif occurs during that hour). Note that the glucose motifs

118 show up fairly evenly *across* all hours of the day whereas the motifs from the cardiology dataset have  
 119 shifts in their distribution and show up frequently at *specific* hours of the day. The lack of temporal  
 120 dependence in glucose motifs is likely due to the diverse patient behaviors within a patient population.  
 121 Glucose in particular is highly variable and influenced by many factors including eating, exercise,  
 122 stress levels, and sleep patterns. Moreover, due to innate variability within human physiology, motif  
 123 occurrences can differ even for the *same* patient across weeks or months. These findings indicate that  
 124 only conserving the temporal relationships within glucose traces (as many previous methods do) may  
 125 not be sufficient to properly learn glucose dynamics and output realistic synthetic traces.

### 126 3.3 Granger Causality

127 Granger causality [7] is commonly used to quantify relationships amongst time series without limiting  
 128 the degree to which temporal relationships may be understood as done in other time series models,  
 129 e.g., pure autoregressive ones. In this framework, an entire system (set of traces) is studied *together*,  
 130 allowing for a broader characterization of their relationships, which may be advantageous, especially  
 131 for long time series. We define  $x_t \in \mathbb{R}^n$  as an  $n$ -dimensional vector of time series observed across  $n$   
 132 traces and  $T$  timesteps. To study causality, a vector autoregressive model (VAR) [18] may be used.  
 133 A set of traces at time  $t$  is represented as a linear combination of the previous  $K$  lags in the series:  
 134  $x_t = \sum_{k=1}^K A^{(k)} x_{t-k} + e_t$  where each  $A^{(k)}$  is a  $n \times n$  dimensional matrix that describes how lag  $k$   
 135 affects the future timepoints in the series' and  $e_t$  is a zero mean noise. Given this framework, we state  
 136 that time series  $q$  does not *Granger-cause* time series  $p$ , if and only if for all  $k$ ,  $A_{p,q}^{(k)} = 0$ . To better  
 137 represent nonlinear dynamics amongst traces, a nonlinear autoregressive model (NAR) [19],  $g$ , may  
 138 be defined, in which  $x_t = g(x_{1<t}, \dots, x_{n<t}) + e_t$  where  $x_{p<t} = (x_{p_1}, \dots, x_{p_{t-1}}, x_{p_t})$  describes the  
 139 past of series  $p$ . The NAR nonlinear functions are commonly modeled jointly using neural networks.

## 140 4 Motif Causality

141 Using Granger causality as defined would overwhelm the generator with too much information,  
 142 resulting in convergence issues for the GAN. Instead of looking at traces comprehensively, we  
 143 need a way to *scope* how the generator understands relationships between time series. To this  
 144 end, we aim to use the same intuition developed from Granger causality, namely developing an  
 145 understanding of relationships comprehensively using less stringent temporal constraints, but scope  
 146 these relationships specifically in terms of *motifs*. Therefore, we develop a concept of *motif causality*  
 147 which, by learning causal relationships amongst sequences of motifs, allows the generator to learn  
 148 realistic motif sequences and produce high quality synthetic traces as a result.

### 149 4.1 Extending Granger Causality to Motifs

150 In order to quantify the relationships amongst sequences of motifs to best capture glucose dynamics,  
 151 we extend the idea of Granger causality to work with motifs. Given a motif set with  $m$  motifs,  
 152 we build a separate (component) model, called a *motif network* in our method, for each motif,  
 153 resulting in  $m$  motif networks. For a single motif  $\mu_i$  at time  $t$ ,  $\mu_{i_t}$ , we define a function  $g_i$  specifying  
 154 how motifs in previous timesteps are mapped to that motif:  $\mu_{i_t} = g_i(\mu_{1<t}, \dots, \mu_{m<t}) + e_{i_t}$  where  
 155  $\mu_{j<t} = (\mu_{j_1}, \dots, \mu_{j_{t-1}}, \mu_{j_t})$  describes the past of motif  $\mu_j$ . The output of  $g_i$  is a vector, which is  
 156 added to the noise vector  $e_{i_t}$ . Essentially, we define motif  $\mu_i$  in terms of its relationship to past motifs.  
 157 The  $g_i$  function takes in some *mapping* that describes how motifs in previous timesteps are mapped  
 158 to the current motif  $\mu_{i_t}$ . The mapping is not specified in this notation, and could be defined in many  
 159 different ways. In our case, we instantiate  $g_i$  using a single-layer LSTM, described next.

160 A  $g_i$  function for each motif  $\mu_i$  in the motif set is modeled using a motif network with a single-  
 161 layer RNN architecture. For a RNN predicting a single component motif, let  $h_t \in \mathbb{R}^m$  represent  
 162 the  $m$ -dimensional hidden state at time  $t$ . This represents the historical context of the motifs in  
 163 the series for predicting a component motif at time  $t$ ,  $\mu_{i_t}$ . At time  $t$ , the hidden state is updated:  
 164  $h_t = g_i(h_{t-1}) + e_{i_t}$ .  $g_i$  here is the function describing how motifs in previous timesteps are mapped  
 165 to the current motif, and is modeled (instantiated) as a single-layer LSTM as they are good at modeling  
 166 long, nonlinear dependencies amongst traces [20]. The output for a motif  $\mu_i$  at time  $t$ ,  $\mu_{i_t}$  can be  
 167 obtained by a linear decoding of the hidden state,  $\mu_{i_t} = W^o h_t + e_{i_t}$ , where  $W^o$  is a matrix of the  
 168 output weights. These weights control the update of the hidden state and thereby control the influence

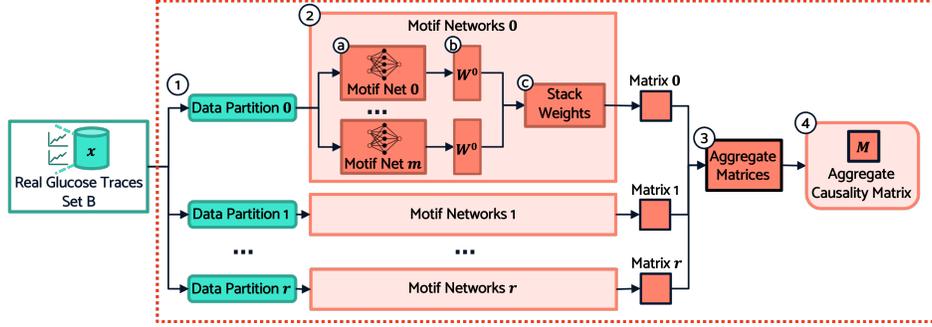


Figure 3: Motif Causality Block.

169 of past motifs on this component motif. Essentially, this function learns a weighting that quantifies  
 170 how helpful motifs in previous timesteps are for predicting the specified motif  $\mu_i$  at time  $t$ . We note  
 171 that we define causality in this way based on how Granger causality models such relationships, which  
 172 is different from traditional causality models.

173 If all elements in the  $j$ th column of  $W^o$  are zero ( $W^o_{:j} = 0$ ), this is a sufficient condition for an  
 174 input motif  $\mu_j$  being motif non-causal on an output  $\mu_i$ . Therefore, we can find the motifs that are  
 175 motif-causal for motif  $\mu_i$  using a group lasso penalty optimization across the columns of  $W^o$ :

$$\min_W \sum_{t=2}^T (\mu_{i_t} - g_i(\mu_{0<t}, \dots, \mu_{m<t}))^2 + \sum_{j=1}^m \|W^o_{:j}\|_2$$

176 We define this as the *local motif loss*,  $\mathcal{L}_{ml}$ , which is optimized in each motif network using proximal  
 177 gradient descent.

## 178 4.2 Training the Motif Causality Block

179 We next describe how the motif causality block is trained to learn motif causal relationships amongst  
 180 traces, displayed in Figure 3. The block is structured in this way to accommodate the privacy  
 181 integration (Section 6.2); here, we present its implementation without any privacy noise.

182 **Partition data.** First, the data is partitioned into  $r$  partitions (Step 1, Figure 3) such that no models  
 183 are trained on overlapping data. The number of partitions,  $r$ , is a user-specified hyperparameter.

184 **Build motif network for each motif.** Next, within each data partition a set of motif networks is  
 185 trained. As a pre-processing step, we assume each trace has been chunked into a sequence of motifs  
 186 of size  $\tau$  (Section 3.1).  $\tau$  is a hyperparameter, which we suggest chosen based on the longest effect  
 187 time of a trace event. We use  $\tau = 48$ , corresponding to 4 hours of time, because large glucose events  
 188 (from behaviors like eating) are encompassed within that time frame. We assume a tolerance of  $\sigma = 2$   
 189 mg/dL, chosen to allow for reasonable variations in glucose. To model motif causality for an entire  
 190 set of data, a  $g_i$  function is implemented for each motif via a separate RNN motif net following the  
 191 description provided previously, resulting in  $m$  total networks (Step 2a, Figure 3).

192 **Combine outputs of individual motif networks.** Each motif network outputs a vector of weights  
 193  $W^o$  of dimensionality  $1 \times m$ , corresponding to the learned causal relationships (Step 2b, Figure 3).  
 194 Values in the vector are between 0 (no causal relationship) and 1 (strongest causal relationship) and  
 195 give the degree to which every other motif is motif causal of the particular motif  $\mu_i$  the RNN was  
 196 specialized for. To return a complete matrix that summarizes causal relationships amongst *all* motifs,  
 197 we stack the weights (Step 2c). The output of each data partition is a complete motif causality matrix,  
 198 resulting in  $r$  total matrices, each of dimensionality  $m \times m$ .

199 **Aggregate matrices and integrate with GAN.** After motif causality matrices have been outputted  
 200 from each data partition, the weights in the matrices are aggregated (Step 3, Figure 3) to return the final  
 201 aggregate causality matrix,  $M$  (Step 4). In the nonprivate version, the weights are averaged. Finally,  
 202  $M$  is sent to the generator to help it learn how to conserve motif relationships within sequences of  
 203 motifs in the synthetically generated data. Details are described next in the subsequent section.

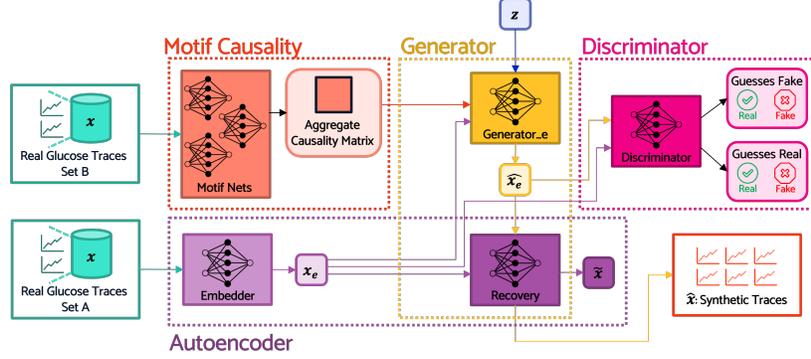


Figure 4: Overview of GlucoSynth Architecture.

## 204 5 GlucoSynth

205 The complete GlucoSynth framework, shown in Figure 4, comprises four key blocks: the motif  
 206 causality block (explained previously in Section 4), an autoencoder, a generator and a discriminator.  
 207 We walk through the remaining components of the framework surrounding the GAN next.

### 208 5.1 GAN Architecture Components

209 **Autoencoder.** We use an autoencoder (AE) with an RNN architecture to learn a lower dimensional  
 210 representation of the traces, allowing the generator to better preserve underlying temporal dynamics  
 211 of the traces. The autoencoder consists of two networks: an *embedder* and a *recovery network*.  
 212 The embedder uses an encoding function to map the real data into a lower dimensional space:  
 213  $Enc(x) : x \in \mathbb{R}^n \rightarrow x_e \in \mathbb{R}^e$  while the recovery network reverses this process, mapping the  
 214 embedded data back to the original dimensional space:  $Dec(x_e) : x_e \in \mathbb{R}^e \rightarrow \tilde{x} \in \mathbb{R}^n$ . A foolproof  
 215 autoencoder perfectly reconstructs the original input data, such that  $x = \tilde{x} \equiv Dec(Enc(x))$ . This  
 216 process yields the Reconstruction Loss,  $\mathcal{L}_R$ , the Mean Square Error (MSE) between the original data  
 217  $x$  and the recovered data,  $\tilde{x}$ :  $MSE(x, \tilde{x})$ .

218 **Generator.** We implement the generator via an RNN or LSTM. Importantly, the generator works in  
 219 the embedded space, by receiving the input traces passed through the embedder ( $x_e$ ). To generate  
 220 synthetic data, a random vector of noise,  $z$  is passed through the generator and then the recovery  
 221 network to return the synthetic traces in the original dimensional space. To learn how to produce  
 222 high-quality synthetic data, the generator receives three key pieces of information:

223 *1 – Stepwise.* The generator receives batches of real data to guide the generation of realistic next step  
 224 vectors. To do this, a Stepwise Loss,  $\mathcal{L}_S$ , is computed at time  $t$  using the MSE between the batch of  
 225 embedded real data,  $x_{et}$ , and the batch of embedded synthetic data,  $\hat{x}_{et}$ :  $MSE(x_{et}, \hat{x}_{et})$ . This allows  
 226 the generator to compare (and learn to correct) the discrepancies in stepwise data distributions.

227 *2 – Motif Causality.* The generator needs to preserve sequences of motifs in addition to temporal  
 228 dynamics. Using the aggregate causality matrix  $M$  returned from the Motif Causality Block, the  
 229 generator computes a motif causality matrix,  $M_{\hat{x}}$ , on the set of synthetic data  $\hat{x}$ . Because the original  
 230 causality matrix was not trained on data in the embedded space, we first run the set of embedded  
 231 synthetic data through the recovery network  $\hat{x}_e \rightarrow \hat{x}$ . From there, the Motif Causality Loss,  $\mathcal{L}_M$ ,  
 232 is computed as the MSE error between the two matrices:  $MSE(M, M_{\hat{x}})$ . These matrices give a causal  
 233 value of seeing a motif  $\mu_i$  in the future after some motif  $\mu_j$ — unrealistic motif sequences will have  
 234 causal values close to 0. As the generator learns to generate synthetic data that yields a realistic causal  
 235 matrix (thereby identifying appropriate causal relationships from the motifs), it implicitly learns to  
 236 not generate unrealistic motif sequences.

237 *3 – Distributional.* To guide the generator to produce a diverse set of traces, the generator computes a  
 238 Distributional Loss,  $\mathcal{L}_D$ , the moments loss (MML), between the overall distribution of the real data  
 239  $x_e$  and the distribution of the synthetic data  $\hat{x}_e$ :  $MML(x_e, \hat{x}_e)$ . The MML is the difference in the  
 240 mean and variance of two matrices.

241 **Discriminator.** The discriminator is a traditional discriminator model using an RNN, the only change  
242 being it also works in the embedded space. The discriminator yields the Adversarial Loss Real,  $\mathcal{L}_{Ar}$ ,  
243 the Binary Cross Entropy (BCE) between the discriminator guesses on the real data  $y_{x_e}$  and the  
244 ground truth  $y$ , a vector of 0's,  $\text{BCE}(y_{x_e}, y)$  and the Adversarial Loss Fake,  $\mathcal{L}_{Af}$ , the BCE between  
245 the discriminator guesses on the fake data  $y_{\hat{x}_e}$  and the ground truth  $y$ , a vector of 1's,  $\text{BCE}(y_{\hat{x}_e}, y)$ .

## 246 5.2 Training Procedure

247 First, the motif causality block is trained following the procedure described in Section 4.2, and then  
248 the rest of the GAN is trained. The autoencoder is optimized to minimize  $\mathcal{L}_R + \alpha\mathcal{L}_S$ , where  $\alpha$  is  
249 a hyperparameter that balances the two loss functions. If the AE only receives  $\mathcal{L}_R$  (as is typically  
250 done), it becomes overspecialized, i.e., it becomes too good at learning the best lower dimensional  
251 representation of the data such that the embedded data are no longer helpful to the generator. For  
252 this reason, the AE also receives  $\mathcal{L}_S$ , enabling the dual training of the generator and embedder. The  
253 generator is optimized using  $\min(1 - \mathcal{L}_{Af}) + \eta(\mathcal{L}_S + \mathcal{L}_D) + \mathcal{L}_M$ , where  $\eta$  is a hyperparameter that  
254 balances the effect of the stepwise and distributional loss. Finally the discriminator is optimized using  
255 the traditional adversarial feedback  $\min \mathcal{L}_{Af} + \mathcal{L}_{Ar}$ . The networks are trained in sequence (within  
256 each epoch) in the following order: autoencoder, generator, then discriminator. In our experiments  
257 we set  $\alpha = 0.1$  and  $\eta = 10$  as they enable GlucoSynth to converge fastest, i.e., in the fewest epochs.

## 258 6 Providing Differential Privacy

259 There are two components to our privacy architecture, described in the following two subsections:  
260 (1) each network in the GAN (Embedder, Recovery, Generator and Discriminator networks) is  
261 trained in a differentially private manner using the Differentially-Private Stochastic Gradient Descent  
262 (DP-SGD) algorithm from Abadi et al. [21]; and (2) the motif causality block is trained using  
263 the PATE framework from Papernot et al. [22]. Importantly, two completely separate datasets are  
264 used for the training of the motif causality block (dataset B in Figure 4) and the GAN (dataset  
265 A in Figure 4). We structure the privacy integration in this way to allow for better privacy-utility  
266 trade-offs. Our design satisfies the formal differential privacy notion introduced by Dwork et al.  
267 [23]. Differential Privacy (DP) provides an intuitive bound on the amount of information that can  
268 be learned about any individual in a dataset. A randomized algorithm  $\mathcal{M}$  satisfies  $(\epsilon, \delta)$ -differential  
269 privacy if, for all datasets  $D_1$  and  $D_2$  differing by at most a single unit, and all  $S \subseteq \text{Range}(\mathcal{M})$ ,  
270  $\Pr[\mathcal{M}(D_1) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D_2) \in S] + \delta$ . The parameters  $\epsilon$  and  $\delta$  determine the *privacy loss*  
271 *budget*, which provide a way to tradeoff privacy and utility; smaller values have stronger privacy.

### 272 6.1 Training the GAN Networks with DP

273 To add privacy to the GAN components, each of the networks (Embedder, Recovery, Generator and  
274 Discriminator) is trained in a differentially private manner using DP-SGD [21]. Although the overall  
275 GAN framework is complicated, the individual networks all use simple RNN or LSTM architectures  
276 with Adam optimizers. As such, adding DP noise to their network weights is straightforward. We  
277 employ the following procedure using Tensorflow Privacy functions [24]. Since there are four  
278 networks being trained with DP, we divide the privacy loss budget evenly to get the budget per  
279 network,  $\epsilon_{net} = \epsilon/4$ . Then, we use Tensorflow's built-in DP accountant to determine how much noise  
280 must be added to the weights of each network based on the number of epochs, batch size, number  
281 of traces and  $\epsilon_{net}$ . This function returns a noise multiplier, which we use when we instantiate a  
282 Tensorflow DP Keras Adam Optimizer for each network. Finally, we train each of the networks using  
283 their respective DP Keras Adam Optimizer, which automatically trains the network using DP-SGD.

### 284 6.2 Training the Motif Causality Block with DP

285 We train the motif causality block using the PATE framework [22]. PATE provides a way to return  
286 aggregated votes about the class a data point belongs to. First, the data is partitioned into  $r$  partitions,  
287 where  $r$  is determined based on the size of the dataset and the privacy loss budget. Then, a class  
288 membership model is trained independently for each partition. The class membership votes from  
289 each partition are aggregated by adding noise to the vote matrix and the noisiest votes are returned  
290 using the max-of-Laplacian mechanism (LNMax), tuned based on the privacy budget and  $r$ .

291 We use PATE to train the motif causality block: instead of predicting the degree of class membership  
 292 we predict *causal* membership, e.g., does motif  $\mu_i$  have a causal relationship to  $\mu_j$ . The motif  
 293 causality block is trained in the same procedure described in Section 4.2 with two changes: (1) the  
 294 number of data partitions,  $r$ , is determined based on the privacy budget, instead of a user-specified  
 295 value; (2) the final causality matrix  $M$  is aggregated using DP across the partitions. In normal PATE,  
 296 carefully calibrated noise is added to a matrix of votes for each class, such that the classes with the  
 297 noisiest votes are outputted. In our use, each value in a motif causality matrix may be likened to a  
 298 class (i.e., causal “class” prediction between motif  $\mu_i$  and  $\mu_j$ ). Thus, we use the LNMax mechanism  
 299 (from predefined Tensorflow Privacy functions [24]) to aggregate the matrices weights and return  $M$ .

300 We use PATE instead of training each motif network using DP-SGD for better privacy-utility trade-  
 301 offs. With DP-SGD, we would need to add noise to *every* motif net, eating up our privacy budget  
 302 quickly and severely impacting the quality of the returned causality matrices. PATE allows us to train  
 303 each of the motif networks without any noise on the gradients, but then aggregates their returned  
 304 causality matrices in a privacy-preserving manner, resulting in a better privacy-utility trade-off.

## 305 7 Evaluation

306 Evaluating synthetic data is notoriously difficult [25], so we provide an extensive evaluation across  
 307 three criteria. Synthetic data should: 1) conserve characteristics of the real data (*fidelity*, Section 7.1);  
 308 2) contain diverse patterns from the real data without the introduction of anomalous patterns (*breadth*,  
 309 Section 7.2); and 3) be usable in place of the original for real-world use cases (*utility*, Section 7.3).

310 **Data and Benchmarks.** We use 100,000 single-day glucose traces randomly sampled across each  
 311 month from January to December 2022, for a total of 1.2 million traces, collected from Dexcom’s  
 312 G6 Continuous Glucose Monitors (CGMs) [26]. Data was recorded every 5 minutes ( $T = 288$ ) and  
 313 each trace was aligned temporally from 00:00 to 23:59. We restrict our comparison to the five most  
 314 closely related state-of-the-art models for generating synthetic univariate time series with no labels or  
 315 auxiliary data: Three nonprivate—TimeGAN [4], Fourier Flows (FF) [12], non-volume preserving  
 316 transformations (NVP) [9]; and two private—RGAN [11] and dpGAN [14]. Additional experimental  
 317 details and all hyperparameter settings are available in Appendix B.

### 318 7.1 Fidelity

319 **Population Statistics.** To evaluate fidelity on a population scale, we compute a common set of glucose  
 320 metrics and test if the difference between the synthetic and real data is statistically significant. Table 1  
 321 provides an abbreviated summary of the results; Appendix C.2 has complete results. GlucoSynth  
 322 performs the best, with few statistical differences between the real and synthetic data for  $\epsilon \geq 0.1$ .

323 **Distributional Comparisons.** We visualize differences in distributions between the real and synthetic  
 324 data by plotting the distribution of variances and using PCA [27]. Figure 5 shows the variance  
 325 distribution for the nonprivate models. Additional comparisons across privacy budgets are available  
 326 in Appendix C.3. In both nonprivate and private settings, GlucoSynth produces synthetic distributions  
 327 closest to the real ones, better than all other models.

### 328 7.2 Breadth

329 We quantify breadth in terms of glucose motifs. For each model’s synthetic traces, we build a motif  
 330 set (see Section 3.1). Given a real motif set from the validation traces  $S_x$ , for each synthetic motif set  
 331  $S_{\hat{x}}$ , we compute “Validation Motifs”, (VM), the fraction of motifs found in the validation motif set  
 332 that are present in the synthetic motif set,  $VM/|S_{\hat{x}}|$ . This metric quantifies how good our synthetic  
 333 motif set is (e.g., are its motifs mostly similar to motifs found in real traces). We also compute metrics  
 334 related to *coverage*, the fraction of motifs in the validation motif set that are found in our synthetic  
 335 data, defined as  $VM/|S_x|$ . This gives a sense of the breadth in a more traditional manner. To compare  
 336 actual *distributions* of motifs (not just counts), we compute the MSE between the distribution of  
 337 real motifs  $S_x$  and the distribution of synthetic motifs  $S_{\hat{x}}$ . This gives a measure about how close the  
 338 synthetic motif distribution is to the real one. We want high VM and coverage, and low MSE. Results  
 339 are in Table 1 with additional analysis in Appendix D; overall our model provides the best breadth.

Table 1: Fidelity, Breadth and Utility Evaluation. Fidelity: bolded values do not have a statistically significant difference from the real data (what we want). Breadth and Utility: VM = fraction found validation motifs; We want high VM, Coverage and low MSE, RMSE; Bolded values indicate the best ones at each privacy budget (nonprivate compared with private models when  $\epsilon = \infty$ ).

Model	$\epsilon$	Fidelity (metric, p-val)		Breadth			Utility
		Variance	Time-in-Range	VM	Coverage	MSE	RMSE
GlucoSynth	0.01	2576, $<1e-5$	61.8, $2e-5$	<b>1.000</b>	0.010	<b>99.0</b>	<b>0.038 <math>\pm</math> 3e-4</b>
	0.1	<b>2809, 0.356</b>	<b>60.1, 0.532</b>	<b>1.000</b>	0.083	<b>11.2</b>	<b>0.036 <math>\pm</math> 3e-4</b>
	1	2761, 0.022	<b>60.6, 0.410</b>	<b>0.992</b>	0.145	<b>6.7</b>	<b>0.030 <math>\pm</math> 1e-4</b>
	10	<b>2801, 0.316</b>	<b>60.2, 0.845</b>	<b>1.000</b>	0.167	<b>5.0</b>	<b>0.029 <math>\pm</math> 1e-4</b>
	$\infty$	<b>2812, 0.503</b>	<b>60.2, 0.682</b>	<b>0.987</b>	<b>0.534</b>	<b>1.6</b>	<b>7e-3 <math>\pm</math> 2e-4</b>
TimeGAN	$\infty$	2235, $8e-3$	<b>62.3, 0.420</b>	0.625	$6e-3$	107.7	0.061 $\pm$ 3e-4
FF	$\infty$	<b>2836, 0.902</b>	46.6, $<1e-5$	0.642	0.405	2.0	0.038 $\pm$ 3e-4
NVP	$\infty$	1789, $<1e-5$	65.5, $<1e-5$	0.482	0.328	1.9	0.029 $\pm$ 3e-5
RGAN	0.01	57, $<1e-5$	78.8, $<1e-5$	0.013	$1e-3$	108.6	0.819 $\pm$ 0.010
	0.1	53, $<1e-5$	71.6, $3e-5$	0.015	0.031	107.3	0.688 $\pm$ 6e-3
	1	67, $<1e-5$	78.2, $<1e-5$	0.015	0.033	103.3	0.651 $\pm$ 0.018
	10	77, $<1e-5$	83.7, $<1e-5$	0.017	0.053	100.3	0.619 $\pm$ 0.016
	$\infty$	90, $<1e-5$	78.0, $<1e-5$	0.026	0.091	79.6	0.460 $\pm$ 0.013
dpGAN	0.01	451, $<1e-5$	95.3, $<1e-5$	0.094	<b>0.054</b>	180.1	0.205 $\pm$ 5e-3
	0.1	1057, $<1e-5$	86.4, $<1e-5$	0.390	<b>0.195</b>	28.9	0.045 $\pm$ 2e-4
	1	875, $<1e-5$	86.6, $<1e-5$	0.480	<b>0.239</b>	23.2	0.030 $\pm$ 2e-5
	10	1030, $<1e-5$	88.1, $<1e-5$	0.743	<b>0.251</b>	16.1	0.035 $\pm$ 8e-5
	$\infty$	1121, $<1e-5$	81.8, $<1e-5$	0.855	0.293	10.9	0.028 $\pm$ 5e-5

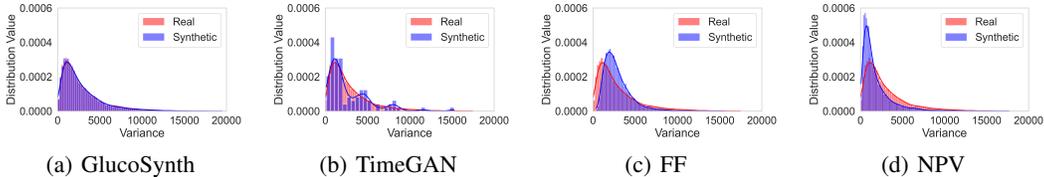


Figure 5: Distributional Variance for Nonprivate Models

### 340 7.3 Utility

341 We evaluate our synthetic glucose traces for use in a glucose forecasting task using the common  
 342 paradigm TSTR (Train on Synthetic, Test on Real), in which the synthetic data is used to train  
 343 the model and then tested on the real validation data. We train an LSTM network optimized for  
 344 glucose forecasting tasks [28] and report the Root Mean Square Error (RMSE) in Table 1. Since  
 345 RMSE provides a limited view about the model’s predictions, we also plot the Clarke Error Grid [29],  
 346 which visualizes the differences between a predictive and reference measurement, and is a basis for  
 347 evaluating the safety of diabetes-related medical devices. More details are in Appendix E. GlucoSynth  
 348 provides the best forecasting results compared to all other models across all privacy budgets.

### 349 8 Limitations & Conclusion

350 **Limitations.** In order to train on a huge set of glucose traces, we used a private dataset, not publicly  
 351 available (one of the motivations for this project was actually to share a synthetic version of these  
 352 traces). That being said, smaller samples of glucose traces with similar patient populations are  
 353 available at OpenHumans [30] and T1D Exchange Registry [31]. In addition, one of the reasons our  
 354 privacy results perform well is because we use two *separate* datasets for the training of the motif  
 355 causality block and the GAN. However, this may be a limiting factor for others that do not have a  
 356 large enough set of traces available to be able to train adequately on partitioned data.

357 **Conclusion.** In this paper we have presented GlucoSynth, a novel GAN framework with integrated  
 358 differential privacy to generate synthetic glucose traces. GlucoSynth conserves motif relationships  
 359 within the traces, in addition to the typical temporal dynamics contained within time series. We pre-  
 360 sented a comprehensive evaluation using 1.2 million glucose traces wherein our model outperformed  
 361 all previous models across three criteria of fidelity, breadth and utility.

## References

- [1] Katherine E Britton and Jennifer D Britton-Colonnese. Privacy and security issues surrounding the protection of data generated by continuous glucose monitors. *Journal of diabetes science and technology*, 11(2):216–219, 2017.
- [2] Deborah Young-Hyman, Mary De Groot, Felicia Hill-Briggs, Jeffrey S Gonzalez, Korey Hood, and Mark Peyrot. Psychosocial care for people with diabetes: a position statement of the american diabetes association. *Diabetes care*, 39(12):2126–2140, 2016.
- [3] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [4] Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. Time-series generative adversarial networks. *Advances in neural information processing systems*, 32, 2019.
- [5] Xiaomin Li, Vangelis Metsis, Huangyingrui Wang, and Anne Hee Hiong Ngu. Tts-gan: A transformer-based time-series generative adversarial network. In Martin Michalowski, Syed Sibte Raza Abidi, and Samina Abidi, editors, *Artificial Intelligence in Medicine*, pages 133–143, Cham, 2022. Springer International Publishing.
- [6] Amirsina Torfi, Edward A Fox, and Chandan K Reddy. Differentially private synthetic medical data generation using convolutional gans. *Information Sciences*, 586:485–500, 2022.
- [7] Clive WJ Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438, 1969.
- [8] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.
- [9] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [10] Eoin Brophy, Zhengwei Wang, Qi She, and Tomas Ward. Generative adversarial networks in time series: A survey and taxonomy. *arXiv preprint arXiv:2107.11098*, 2021.
- [11] Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- [12] Ahmed Alaa, Alex James Chan, and Mihaela van der Schaar. Generative time-series modeling with fourier flows. In *International Conference on Learning Representations*, 2021.
- [13] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.
- [14] Lorenzo Frigerio, Anderson Santana de Oliveira, Laurent Gomez, and Patrick Duverger. Differentially private generative adversarial networks for time series, continuous, and discrete open data. In *IFIP International Conference on ICT Systems Security and Privacy Protection*, pages 151–164. Springer, 2019.
- [15] Cynthia Binanay, Robert M Califf, Vic Hasselblad, Christopher M O’Connor, Monica R Shah, George Sopko, Lynne W Stevenson, Gary S Francis, Carl V Leier, Leslie W Miller, et al. Evaluation study of congestive heart failure and pulmonary artery catheterization effectiveness: the escape trial. *Jama*, 294(13):1625–1633, 2005.
- [16] George EP Box, Gwilym M Jenkins, Gregory C Reinsel, and Greta M Ljung. *Time series analysis: forecasting and control*. John Wiley & Sons, 2015.
- [17] Lexiang Ye and Eamonn Keogh. Time series shapelets: a new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 947–956, 2009.
- [18] Helmut Lütkepohl. *New introduction to multiple time series analysis*. Springer Science & Business Media, 2005.

- 409 [19] Stephen A Billings. *Nonlinear system identification: NARMAX methods in the time, frequency,*  
410 *and spatio-temporal domains*. John Wiley & Sons, 2013.
- 411 [20] Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A review of recurrent neural  
412 networks: Lstm cells and network architectures. *Neural computation*, 31(7):1235–1270, 2019.
- 413 [21] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar,  
414 and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC*  
415 *conference on computer and communications security*, pages 308–318, 2016.
- 416 [22] Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Úlfar  
417 Erlingsson. Scalable private learning with pate. *arXiv preprint arXiv:1802.08908*, 2018.
- 418 [23] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to  
419 sensitivity in private data analysis. In *Theory of Cryptography Conference*, 2006.
- 420 [24] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro,  
421 Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow,  
422 Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser,  
423 Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek  
424 Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal  
425 Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete  
426 Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-  
427 scale machine learning on heterogeneous systems, 2015. URL [https://www.tensorflow.](https://www.tensorflow.org/)  
428 [org/](https://www.tensorflow.org/). Software available from tensorflow.org.
- 429 [25] Ali Borji. Pros and cons of gan evaluation measures: New developments. *Computer Vision and*  
430 *Image Understanding*, 215:103329, 2022.
- 431 [26] Halis Kaan Akturk, Robert Dowd, Kaushik Shankar, and Mark Derdzinski. Real-world evidence  
432 and glycemic improvement using dexcom g6 features. *Diabetes Technology & Therapeutics*, 23  
433 (S1):S–21, 2021.
- 434 [27] Fred B Bryant and Paul R Yarnold. Principal-components analysis and exploratory and confir-  
435 matory factor analysis. *American Psychological Association*, 1995.
- 436 [28] Taisa Kushner, Marc D Breton, and Sriram Sankaranarayanan. Multi-hour blood glucose  
437 prediction in type 1 diabetes: A patient-specific approach using shallow neural network models.  
438 *Diabetes Technology & Therapeutics*, 22(12):883–891, 2020.
- 439 [29] William L Clarke. The original clarke error grid analysis (ega). *Diabetes technology &*  
440 *therapeutics*, 7(5):776–779, 2005.
- 441 [30] Open humans. <https://www.openhumans.org/>.
- 442 [31] T1d exchange registry. <https://t1dexchange.org/registry/>.
- 443 [32] Hao Ni, Lukasz Szpruch, Magnus Wiese, Shujian Liao, and Baoren Xiao. Conditional sig-  
444 wasserstein gans for time series generation. *arXiv preprint arXiv:2006.05421*, 2020.
- 445 [33] Zinan Lin, Alankar Jain, Chen Wang, Giulia Fanti, and Vyas Sekar. Using gans for sharing  
446 networked time series data: Challenges, initial promise, and open questions. In *Proceedings of*  
447 *the ACM Internet Measurement Conference*, pages 464–483, 2020.
- 448 [34] Juan Miguel Lopez Alcaraz and Nils Strodthoff. Diffusion-based time series imputation and  
449 forecasting with structured state space models. *arXiv preprint arXiv:2208.09399*, 2022.
- 450 [35] Mihai Dogariu, Liviu-Daniel Ștefan, Bogdan Andrei Boteanu, Claudiu Lamba, Bomi Kim, and  
451 Bogdan Ionescu. Generation of realistic synthetic financial time-series. *ACM Transactions on*  
452 *Multimedia Computing, Communications, and Applications (TOMM)*, 18(4):1–27, 2022.
- 453 [36] Mina Razghandi, Hao Zhou, Melike Erol-Kantarci, and Damla Turgut. Variational autoencoder  
454 generative adversarial network for synthetic data generation in smart home. *arXiv preprint*  
455 *arXiv:2201.07387*, 2022.

- 456 [37] Debapriya Hazra and Yung-Cheol Byun. Synsiggan: Generative adversarial networks for  
457 synthetic biomedical signal generation. *Biology*, 9(12):441, 2020.
- 458 [38] James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. Pate-gan: Generating synthetic data  
459 with differential privacy guarantees. In *International conference on learning representations*,  
460 2018.

Table 2: Summary of Previous Methods for Time Series Synthesis. \*CI = conditional information or extra features

Name	Private?	No Labels Required?	No CI*?	Length
TimeGAN [4]	x	✓	✓	24 - 58
TTS-GAN [5]	x	x	✓	24 - 150
SigCWGAN [32]	x	✓	x	80,000
RGAN [11]	✓	✓	✓	16 - 30
RCGAN [11]	✓	✓	x	16 - 30
dpGAN [14]	✓	✓	✓	96
RDP-CGAN [6]	✓	✓	x	2 - 4097
DoppelGANger [33]	✓	✓	x	50 - 600
Glucosynth (Ours)	✓	✓	✓	288+

## 461 A Extended Related Work

462 We overview related work in three lines of research: time series, conditional time series, and time  
 463 series methods that employ differential privacy. Table 2 summarizes previous time series synthesis  
 464 methods. We note that there have been exciting developments for adjacent research tasks (data  
 465 augmentation, forecasting) such as diffusion models [34], but there are not yet any publicly available  
 466 models specifically for the generation of complete synthetic time series datasets. As such, we focus  
 467 the scope of our comparison on the current state-of-the-art methods for synthetic time series which all  
 468 build upon Generative Adversarial Networks (GANs) [3] and transformation-based approaches [9].  
 469 In particular TimeGAN [4], RGAN [11] and dpGAN [14] are most similar to ours and used as  
 470 benchmarks in the evaluation in Section 7.

471 **Time Series.** There have been promising models to generate synthetic time series across a variety of  
 472 domains such as financial data [35], cyber-physical systems (e.g., smart homes [36]), and medical  
 473 signals [37]. Brophy et al. [10] provides a survey of GANs for time series synthesis. TimeGan [4]  
 474 is a popular benchmark that jointly learns an embedding space using supervised and adversarial  
 475 objectives in order to capture the temporal dynamics amongst traces. TTS-GAN [5], trains a GAN  
 476 model that uses a transformer encoding architecture in order to best preserve temporal dynamics.  
 477 Transformation-based approaches have also had success for time series data. Real-valued non-  
 478 volume preserving transformations (NVP) [9] model the underlying distribution of the real data using  
 479 generative probabilistic modeling and use this model to output a set of synthetic data. Similarly,  
 480 Fourier Flows (FF) [12] transform input traces into the frequency domain and output a set of synthetic  
 481 data from the learned spectral representation of the original data. Methods that only focus on learning  
 482 the temporal or distributional dynamics in time series are not sufficient for generating *realistic*  
 483 synthetic glucose traces due to the lack of temporal dependence within sequences of glucose motifs.

484 **Conditional Time Series.** Many works have developed time series models that supplement their  
 485 training using extra features or conditional data. Esteban et al. [11] develops two GAN models  
 486 (RGAN/RCGAN) with RNN architectures, conditioned on auxiliary information provided at each  
 487 timestep during training. SigCWGAN [32] uses a mathematical conditional metric ( $Sig - W_1$ )  
 488 characterizing the signature of a path to capture temporal dependence of joint probability distributions  
 489 in long time series data. However, our glucose traces do not have any additional information available  
 490 so these methods cannot be used<sup>1</sup>.

491 **Differentially-Private GANs.** To protect sensitive data, several GAN architectures (DP GANs)  
 492 have been designed to incorporate privacy-preserving noise needed to satisfy differential privacy  
 493 guarantees [13]. Although DP GANs such as PateGAN [38] have had great success for other data  
 494 types and learning tasks (e.g., tabular data, supervised classification tasks), results have been less  
 495 satisfactory in DP GANs developed for time series.

496 RGAN/RCGAN [11] also includes a DP implementation, but the authors find large gaps in perfor-  
 497 mance between the nonprivate and private models. Frigerio et al. [14] extends a simple DP GAN  
 498 architecture (denoted dpGAN) to time-series data. The synthetic data from their private model

<sup>1</sup>There is a caveat here that RGAN does not use auxiliary information, hence why we compare with it in our benchmarks.

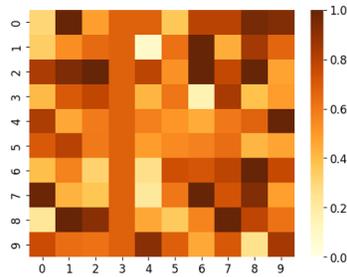


Figure 6: Example motif causality matrix for a small motif set ( $m = 10$ ). Each value in the grid is between 0 and 1. 0 indicates no motif-causal relationship, and 1 indicates the strongest motif causal relationship.

499 conserves the distribution of the real data but loses some of the variability (diversity) from the original  
500 samples. RDP-CGAN [6] develops a convolutional GAN architecture that uses Rényi differential  
501 privacy specifically for medical data. Across different datasets, they find that reasonable privacy  
502 budgets result in major drops in the performance of the synthetic data. Finally, DoppelGANger [33]  
503 develops a temporal GAN framework for time series with metadata and perform an in-depth privacy  
504 evaluation. Notably, they find that providing strong theoretical DP guarantees results in destroying  
505 the fidelity of the synthetic data, beyond anything feasible for use in real-world scenarios. Each  
506 of these methods touches on the innate challenge of generating DP synthetic time series due to  
507 very high tradeoffs between utility and privacy. Our DP framework uses two different methods  
508 to integrate privacy into our GAN architecture, resulting in a better utility-privacy trade-off than  
509 previous methods.

## 510 B Additional Experimental Details

511 **Note on Data Use.** As explained in the approach (Section 5), our model uses two *separate* datasets  
512 for the training of the motif causality block and the rest of the GAN. As such, we used two different  
513 samples of glucose traces with no overlap between patients for the training of each section (meaning  
514 we actually used a total of 2.4 million traces across the entire model).

515 **Extra Benchmark Details.** TimeGAN [4] is implemented from [www.github.com/jsyoon0823/](https://www.github.com/jsyoon0823/TimeGAN)  
516 TimeGAN; Fourier Flows (FF) [12] are implemented from [www.github.com/ahmedmalaa/](https://www.github.com/ahmedmalaa/Fourier-flows)  
517 Fourier-flows; RGAN [11] is implemented from [www.github.com/ratschlab/RGAN](https://www.github.com/ratschlab/RGAN);  
518 and DPGAN [14] is adapted from [www.github.com/SAP-samples/security-research-](https://www.github.com/SAP-samples/security-research-differentially-private-generative-models)  
519 [differentially-private-generative-models](https://www.github.com/SAP-samples/security-research-differentially-private-generative-models).

520 **Hyperparameters.** We use a separate validation dataset (not the set of original training traces) for  
521 all experimental results. Throughout all our experiments we use a motif tolerance  $\sigma = 2$  mg/dL,  
522 motif length  $\tau = 48$ , and GlucoSynth model parameters of  $\alpha = 0.1$  and  $\eta = 10$ . Motif length of  
523 48 timesteps is equivalent to 4 hours of time; this threshold was chosen because the effect of any  
524 behaviors on glucose occur within 4 hours of the event (e.g., the effect from eating a meal – a rise in  
525 glucose – will occur within 4 hours after eating.) There are  $m = 5,977,610$  total motifs in the motif  
526 set. We vary  $\epsilon$  in our privacy experiments, but keep  $\delta$  the same at  $5e-4$ . All the benchmarks were  
527 trained according to their suggested parameters, with most models trained for 10,000 epochs. We  
528 note that we trained for more than the suggested epochs (50,000 instead of 10,000) and tried many  
529 additional hyperparameter settings for RGAN to attempt to improve its performance and provide the  
530 fairest comparison possible. Our experiments were completed in the Google Cloud platform on an  
531 Intel Skylake 96-core cpu with 360 GB of memory.

## 532 C Additional Evaluation: Fidelity

### 533 C.1 Visualizations

534 **Traces.** We provide visualizations of sample real and synthetic glucose traces from all the models.  
535 Although this is not a comprehensive way to evaluate trace quality, it does give a snapshot view about  
536 what synthetic traces may look like. Figure 7 shows randomly sampled individual traces across the

537 nonprivate models, and Figure 8 shows traces across different privacy budgets for the private models.  
 538 As evidenced by the figures, GlucoSynth produces highly realistic synthetic glucose traces, even at  
 539 small privacy budgets.

540 **Heatmaps.** We also provide a heatmap visualization of the traces, to give a slightly larger snapshot  
 541 view of the outputted synthetic vs real traces. Each heatmap contains 100 randomly sampled glucose  
 542 traces. Each row is a single trace from timestep 0 to 288. The values (coloring) in each row indicate  
 543 the glucose value (between 40 mg/dL and 400 mg/dL). Figure 9 shows the nonprivate models, and  
 544 Figures 10, 11, 12 show the private models with different privacy budgets. Upon examining the  
 545 heatmaps, we notice that GlucoSynth consistently generates realistic looking glucose traces, even at  
 546 very small privacy budgets.

Table 3: Glycemic Metric Explanations

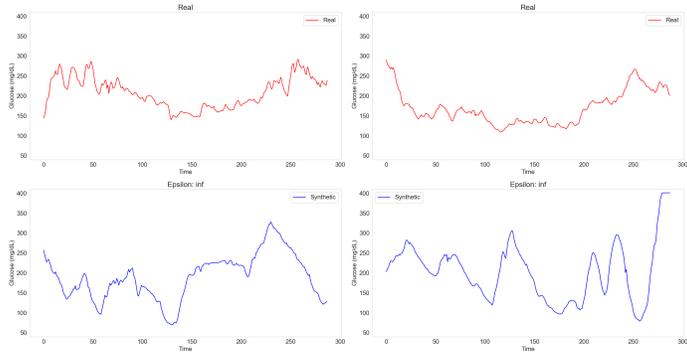
Metric	Name	Explanation
VAR	Signal Variance	average trace variability
TIR	Time in Range	% of time glucose $\geq 70$ & $\leq 180$
Hypo	Time Hypoglycemic	% of time glucose $< 70$
Hyper	Time Hyperglycemic	% of time glucose $> 180$
GVI	Glycaemic Variability Index	more detailed measure of glucose variability
PGS	Patient Glycaemic Status	metric combining GVI and TIR

Table 4: Population Data Statistics. Each cell value for the synthetic data shows the (metric, p-value) using a 0.05 testing threshold. Bolded values do not have a statistically significant difference from the real data (what we want).

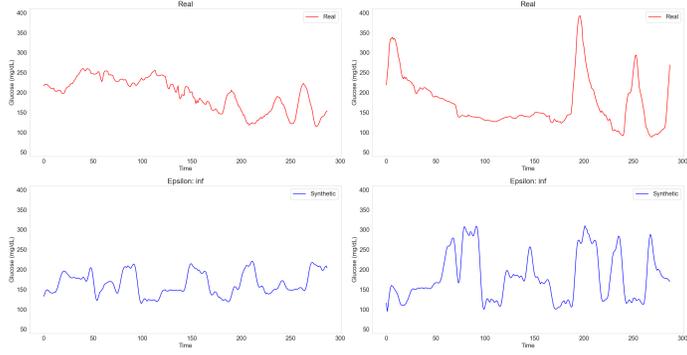
Model	$\epsilon$	VAR	TIR	Hypo	Hyper	GVI	PGS
Real Data	N/A	2832.76	60.31	1.58	38.11	4.03	349.23
GlucoSynth	0.01	2575.501, 0.0	61.759, $2.0e-5$	1.331, 0.0	36.91, $5.66e-4$	<b>4.002, 0.085</b>	323.056, 0.0
	0.1	<b>2803.513, 0.356</b>	<b>60.088, 0.532</b>	1.264, 0.0	<b>38.648, 0.137</b>	$3.969, 2.74e-4$	<b>347.562, 0.712</b>
	1	2760.853, 0.022	<b>60.597, 0.41</b>	<b>1.512, 0.163</b>	<b>37.892, 0.537</b>	<b>4.019, 0.577</b>	<b>345.159, 0.368</b>
	10	<b>2800.805, 0.316</b>	<b>60.24, 0.845</b>	<b>1.538, 0.395</b>	<b>38.222, 0.76</b>	$3.963, 6.7e-5$	<b>344.376, 0.28</b>
	100	<b>2796.424, 0.244</b>	<b>60.138, 0.625</b>	<b>1.567, 0.808</b>	<b>38.295, 0.609</b>	<b>4.044, 0.32</b>	<b>352.679, 0.449</b>
$\infty$	<b>2811.622, 0.503</b>	<b>60.165, 0.682</b>	<b>1.54, 0.416</b>	<b>38.295, 0.61</b>	<b>4.056, 0.083</b>	<b>353.584, 0.339</b>	
TimeGAN	$\infty$	2234.576, $8.08e-3$	<b>62.315, 0.42</b>	$0.657, 8.233e-3$	<b>37.028, 0.669</b>	5.482, 0.0	503.148, $0.2e-5$
FF	$\infty$	<b>2836.067, 0.902</b>	46.578, 0.0	5.627, 0.0	47.795, 0.0	4.931, 0.0	528.773, 0.0
NVP	$\infty$	1789.430, 0.0	65.499, 0.0	<b>1.507, 0.154</b>	32.994, 0.0	6.607, 0.0	589.473, 0.0
RGAN	0.01	56.96, 0.0	78.756, 0.0	0.0, $1.78e-4$	21.244, 0.0	2.52, 0.0	93.409, 0.0
	0.1	52.553, 0.0	71.617, $3.7e-5$	0.0, $1.78e-4$	25.715, 0.0	2.208, 0.0	98.944, 0.0
	1	67.346, 0.0	78.154, 0.0	0.0, $1.78e-4$	21.846, 0.0	2.251, 0.0	85.417, 0.0
	10	76.632, 0.0	83.681, 0.0	0.0, $1.78e-4$	16.319, 0.0	2.23, 0.0	64.562, 0.0
	100	84.918, 0.0	74.285, 0.0	0.0, $1.78e-4$	$25.715, 0.6e-5$	2.208, 0.0	98.944, 0.0
$\infty$	89.702, 0.0	78.044, 0.0	0.0, $1.78e-4$	21.956, 0.0	2.184, 0.0	82.923, 0.0	
dpGAN	0.01	451.098, 0.0	95.275, 0.0	4.60, 0.0	0.124, 0.0	7.718, 0.0	41.549, 0.0
	0.1	1057.205, 0.0	86.43, 0.0	0.837, 0.0	12.732, 0.0	6.349, 0.0	148.412, 0.0
	1	874.663, 0.0	86.631, 0.0	1.135, 0.0	12.234, 0.0	4.794, 0.0	118.286, 0.0
	10	1029.971, 0.0	88.122, 0.0	2.002, 0.0	9.876, 0.0	4.759, 0.0	93.632, 0.0
	100	821.636, 0.0	89.354, 0.0	0.664, 0.0	9.982, 0.0	4.613, 0.0	82.561, 0.0
$\infty$	1120.553, 0.0	81.773, 0.0	$1.359, 0.3e-5$	16.868, 0.0	6.248, 0.0	188.991, 0.0	

## 547 C.2 Population Statistics

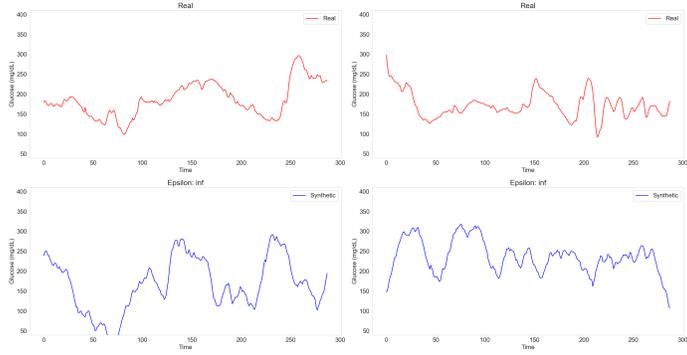
548 In order to evaluate fidelity on a population scale, we compute a common set of glucose metrics used  
 549 to evaluate patient glycemic control on the real and synthetic data, including average trace variability  
 550 (VAR), Time in Range (TIR), the percentage of time glucose is within the clinical guided range of  
 551 70-180mg/dL; and time hypo- and hyper- glycemic (time below and above range, respectively) in  
 552 Table 4. More details on each of the metrics are included in Table 3. We test if the difference in  
 553 metrics between the synthetic and real data is statistically significant, using a p-value of 0.05. A  
 554 p-value  $< 0.05$  indicates the difference is statistically significant. We want synthetic data that has  
 555 similar population statistics to the real data: p-values  $> 0.05$  such that the differences in statistics  
 556 between real and synthetic data are not significant. GlucoSynth outperforms all other models, with no  
 557 statistically significant difference in all metrics for privacy budgets of  $\epsilon \geq 100$  and only one metric  
 558 with a statistically significant difference for budgets  $\epsilon = 1 - 10$ .



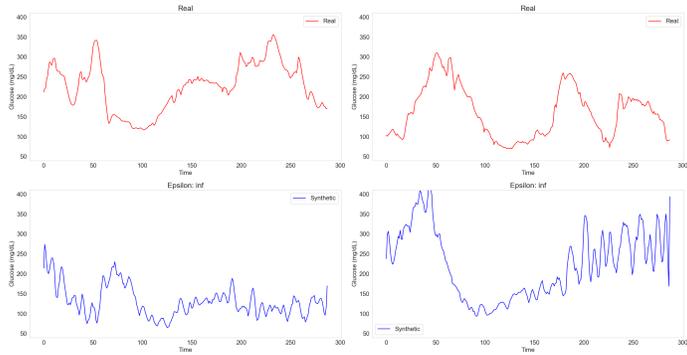
(a) GlucoSynth



(b) TimeGAN



(c) FF



(d) NPV

Figure 7: Sample Traces for Nonprivate Models

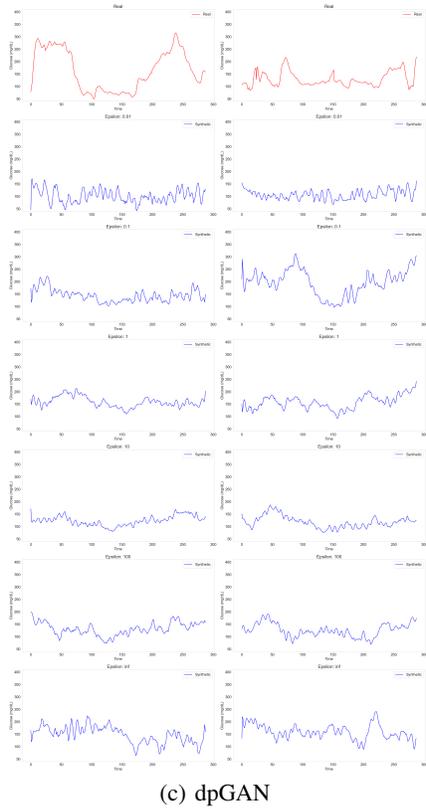
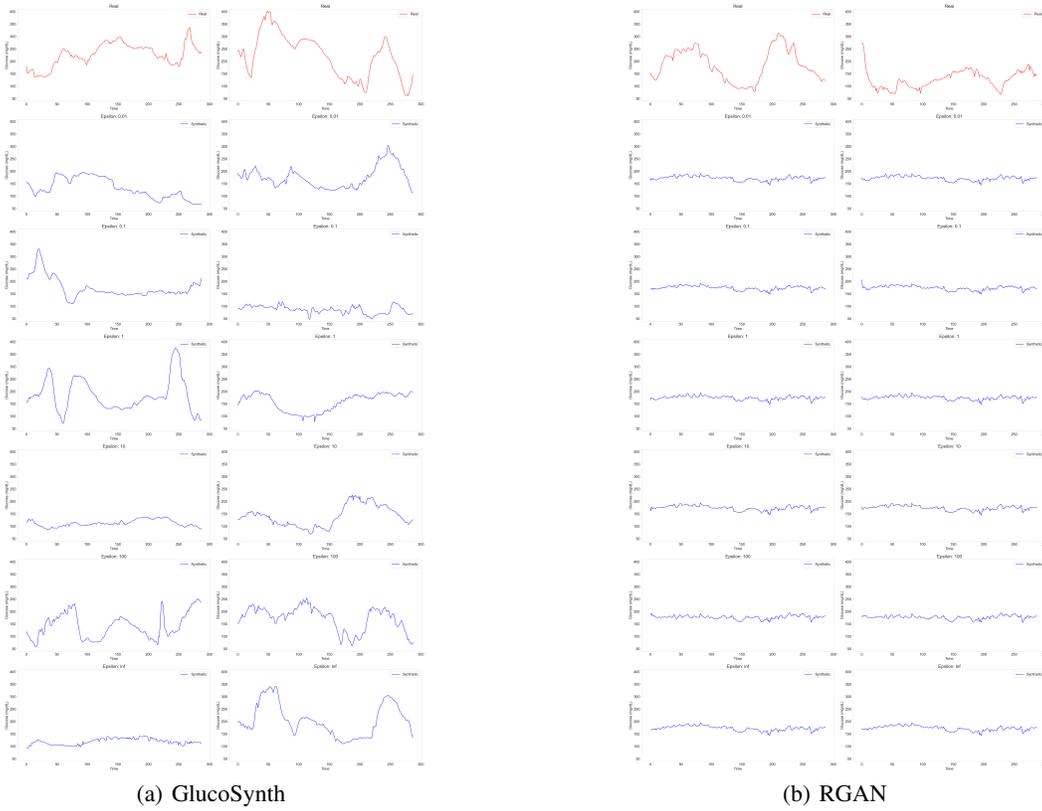


Figure 8: Sample Traces for Private Models Across Privacy Budgets

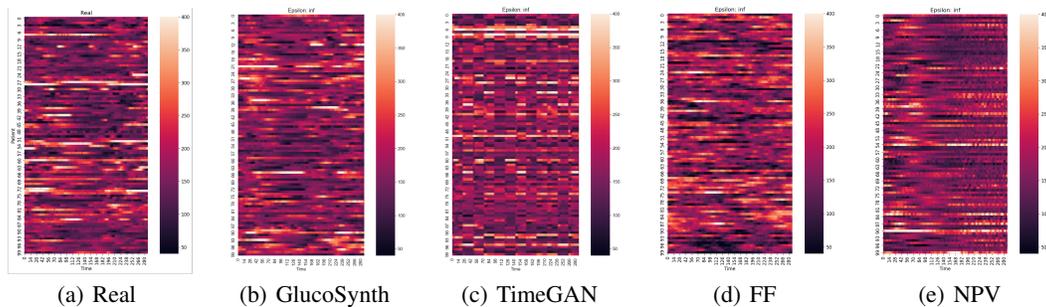


Figure 9: Heatmaps for Nonprivate Models

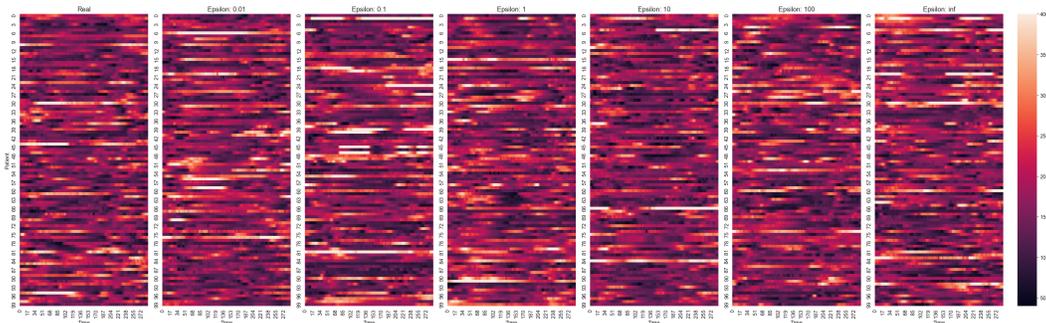


Figure 10: Heatmaps for GlucoSynth Across Different Privacy Budgets

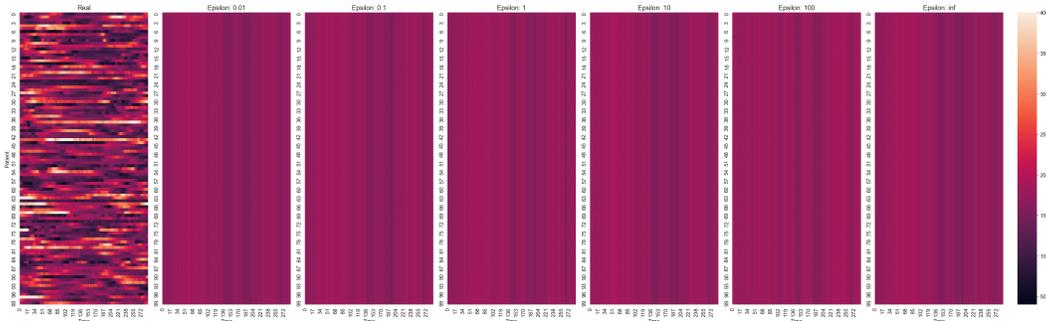


Figure 11: Heatmaps for RGAN Across Different Privacy Budgets

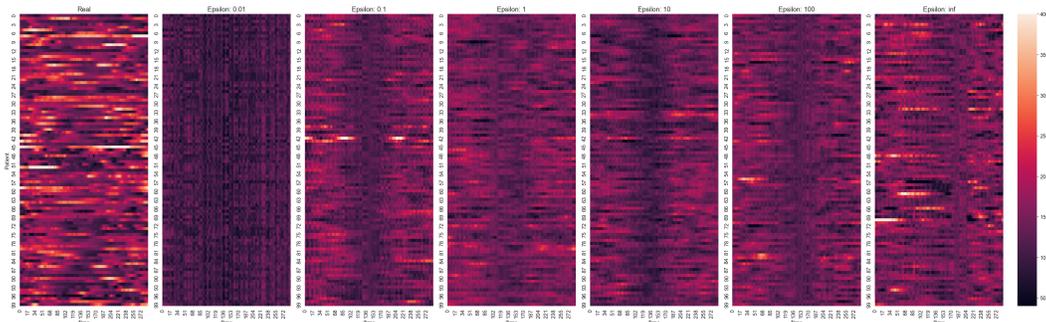


Figure 12: Heatmaps for dpGAN Across Different Privacy Budgets

### 559 C.3 Distributional Comparisons

560 We visualize differences in distributions between the real and synthetic data by plotting the distribution  
561 of variances and using PCA [27]. Figure 5 and Figure 13 show the variance distribution and PCA  
562 plots, respectively for the nonprivate models. We also compare distributional changes across privacy  
563 budgets: Figures 14 and 15 show GlucoSynth, Figures 16 and 17 show RGAN and Figures 18 and 19  
564 show dpGAN.

565 Looking at the figures, GlucoSynth better captures the distribution of the real data compared to all of  
566 the nonprivate models. As evidenced in the PCA plot, (Fig. 13), FF comes the closest to capturing  
567 the real distribution in its synthetic data, but ours does a better job of representing the more rare types  
568 of traces. GlucoSynth also outperforms all of the private models across all privacy budgets. Even at  
569 small budgets ( $\epsilon < 1$ ), the general shape of the overall distribution is conserved (e.g., see Figure 14).

### 570 D Additional Evaluation: Breadth

571 Compared to all other models across all privacy budgets, our model has the best ratio of found  
572 validation motifs, with close to 1.0 for VM and the lowest MSEs. It also has the best coverage for  
573 nonprivate settings and an  $\epsilon$  of 100. Interestingly, dpGAN has the best coverage compared to all other  
574 models for privacy budgets  $\epsilon \leq 10$  but worse MSEs across all budgets than GlucoSynth. This means  
575 that although it finds a broader *number* of motifs contained in the real data, the overall distributions  
576 of motifs it creates in the synthetic data have much higher error rates. We argue that the tradeoff  
577 found by our model is better because although it does miss some of the *types* of motifs from the real  
578 data (misses some breadth), from the ones it does find it constructs realistic distributions of the motifs  
579 and generates very few anomalous ones.

### 580 E Additional Evaluation: Utility

581 Since RMSE may provide a limited view about the predictions from the glucose forecasting model,  
582 we also plot the Clarke Error Grid [29], which visualizes the differences between a predictive  
583 measurement and a reference measurement, and is the basis used for evaluation of the safety of  
584 diabetes-related medical devices (for example, used for evaluating glucose outputs from predictive  
585 models integrated into artificial insulin delivery systems). The Clarke Error Grid is implemented  
586 using [www.github.com/suetAndTie/ClarkeErrorGrid](http://www.github.com/suetAndTie/ClarkeErrorGrid). The grids are shown in Figure 20.

587 In the figures, the x-axis is the reference value and the y-axis is the prediction. A diagonal line means  
588 the predicted value is exactly the same as the reference value (the best case). There are 5 total zones  
589 that make up the grid, listed in order from best to worst:

- 590 • Zone A – Clinically Accurate: Predictions differ from actual values by no more than 20%  
591 and lead to clinically correct treatment decisions.
- 592 • Zone B – Clinically Acceptable: Predictions differ from actual values by more than 20% but  
593 would not lead to any treatment decisions.
- 594 • Zone C – Overcorrections: Acceptable glucose levels would be corrected (overcorrection).
- 595 • Zone D – Failure to Detect: Predictions lie within the acceptable range but the actual values  
596 are outside the acceptable range, resulting in a failure to detect and treat errors in glucose.
- 597 • Zone E – Erroneous Treatment: Predictions are opposite the actual values, resulting in  
598 erroneous treatment, opposite of what is clinically recommended.

599 We show Clarke Error grids for all models (and the private models with no privacy included,  $\epsilon = \infty$ ).  
600 This is because comparing the models at different privacy budgets is not very informative – it can be  
601 hard to tell exactly where changes between different budgets may occur. We also present a table with  
602 the percentages of predicted datapoints in each category in Table 5. This table includes a comparison  
603 among different privacy budgets for the private models (much more effective than the figures by  
604 themselves.)

605 Looking at the grids, we can see that GlucoSynth performs the best, with most of the values along  
606 the diagonal axis (Zone A and B) and less around the other zones (Zones C-E) as compared to the

Table 5: Clarke Error Grid Zones. Value is the percentage of predicted datapoints. Categories go from A to E, best to worst. Bolded rows indicate the best results on the synthetic data at each privacy budget (nonprivate models compared with private models when  $\epsilon = \infty$ )

Model	$\epsilon$	A: Accurate	B: Acceptable	C: Overcorrection	D: Failure to Detect	E: Error
GlucoSynth	0.01	<b>0.858 ± 1.057e-3</b>	<b>0.131 ± 1.172e-3</b>	<b>3.271e-3 ± 0.0</b>	<b>0.017 ± 1.158e-4</b>	<b>5.79e-6 ± 1.2e-6</b>
	0.1	<b>0.863 ± 6.947e-3</b>	<b>0.126 ± 7.526e-4</b>	<b>3.054e-3 ± 1.45e-5</b>	<b>0.018 ± 4.34e-5</b>	<b>5.79e-6 ± 0.0</b>
	1	<b>0.862 ± 1.578e-3</b>	<b>0.128 ± 1.259e-3</b>	<b>3.343e-3 ± 1.45e-5</b>	<b>0.016 ± 3.329e-4</b>	<b>5.79e-6 ± 0.0</b>
	10	<b>0.864 ± 6.947e-3</b>	<b>0.125 ± 6.513e-4</b>	<b>3.039e-3 ± 5.79e-5</b>	<b>0.017 ± 4.34e-5</b>	<b>8.68e-6 ± 2.89e-5</b>
	100	<b>0.864 ± 1.74e-3</b>	<b>0.126 ± 1.447e-3</b>	<b>3.387e-3 ± 0.0</b>	<b>0.017 ± 2.895e-4</b>	<b>5.79e-6 ± 0.0</b>
	$\infty$	<b>0.964 ± 1.201e-3</b>	<b>0.035 ± 1.158e-3</b>	<b>3.039e-4 ± 2.89e-5</b>	<b>1.732e-4 ± 1.158e-4</b>	<b>8.68e-6 ± 1.45e-5</b>
TimeGAN	$\infty$	0.741 ± 0.012	0.233 ± 0.012	2.240e-3 ± 9.8e-5	0.024 ± 8.44e-4	2.19e-4 ± 1.9e-5
FF	$\infty$	0.824 ± 6.624e-3	0.156 ± 6.148e-3	3.547e-3 ± 9.0e-5	0.017 ± 3.940e-4	3.57e-4 ± 8.0e-6
NVP	$\infty$	0.79 ± 3.03e-4	0.186 ± 3.87e-4	3.49e-3 ± 1.5e-5	0.02 ± 1.04e-4	3.58e-4 ± 5.0e-6
RGAN	0.01	0.54 ± 0.014	0.435 ± 0.014	3.389e-4 ± 1.197e-4	0.024 ± 2.71e-4	2.429e-4 ± 3.43e-5
	0.1	0.594 ± 1.998e-3	0.38 ± 1.74e-3	1.326e-3 ± 1.429e-4	0.025 ± 1.069e-4	2.873e-4 ± 8.68e-6
	1	0.637 ± 6.785e-3	0.336 ± 6.128e-3	2.661e-3 ± 1.87e-5	0.024 ± 6.464e-4	2.792e-4 ± 2.95e-5
	10	0.634 ± 3.452e-3	0.338 ± 3.247e-3	2.253e-3 ± 1.004e-4	0.025 ± 2.894e-4	3.027e-4 ± 1.71e-5
	100	0.638 ± 4.709e-3	0.335 ± 4.219e-3	1.991e-3 ± 2.17e-5	0.025 ± 4.884e-4	2.949e-4 ± 2.26e-5
	$\infty$	0.646 ± 6.89e-4	0.326 ± 7.19e-4	2.613e-3 ± 2.852e-4	0.024 ± 3.006e-4	2.859e-4 ± 1.5e-5
dpGAN	0.01	0.308 ± 3.482e-3	0.509 ± 3.71e-3	2.894e-7 ± 0.0	0.183 ± 2.33e-4	1.114e-5 ± 4.196e-6
	0.1	0.781 ± 6.35e-4	0.191 ± 5.37e-4	3.226e-3 ± 5.715e-5	0.024 ± 3.8e-5	2.533e-4 ± 1.881e-6
	1	0.786 ± 5.44e-4	0.187 ± 5.81e-4	2.409e-3 ± 2.894e-7	0.024 ± 3.6e-5	2.078e-4 ± 5.787e-7
	10	0.806 ± 7.34e-4	0.169 ± 6.09e-4	2.386e-3 ± 1.476e-5	0.023 ± 1.113e-4	2.146e-4 ± 2.749e-6
	100	0.813 ± 3.18e-4	0.161 ± 2.86e-4	2.266e-3 ± 2.083e-5	0.023 ± 5.4e-5	1.889e-4 ± 1.013e-6
	$\infty$	0.819 ± 1.487e-3	0.16 ± 1.306e-3	3.193e-3 ± 2.677e-5	0.018 ± 1.60e-4	3.166e-4 ± 5.208e-6

607 other models. This means that most of the predicted glucose values from the model trained on our  
608 synthetic data are in the Clinically Accurate and Acceptable ranges, with less in the erroneous zones.  
609 Moreover, by examining the table we see that GlucoSynth outperforms all other models across all  
610 privacy budgets as well.

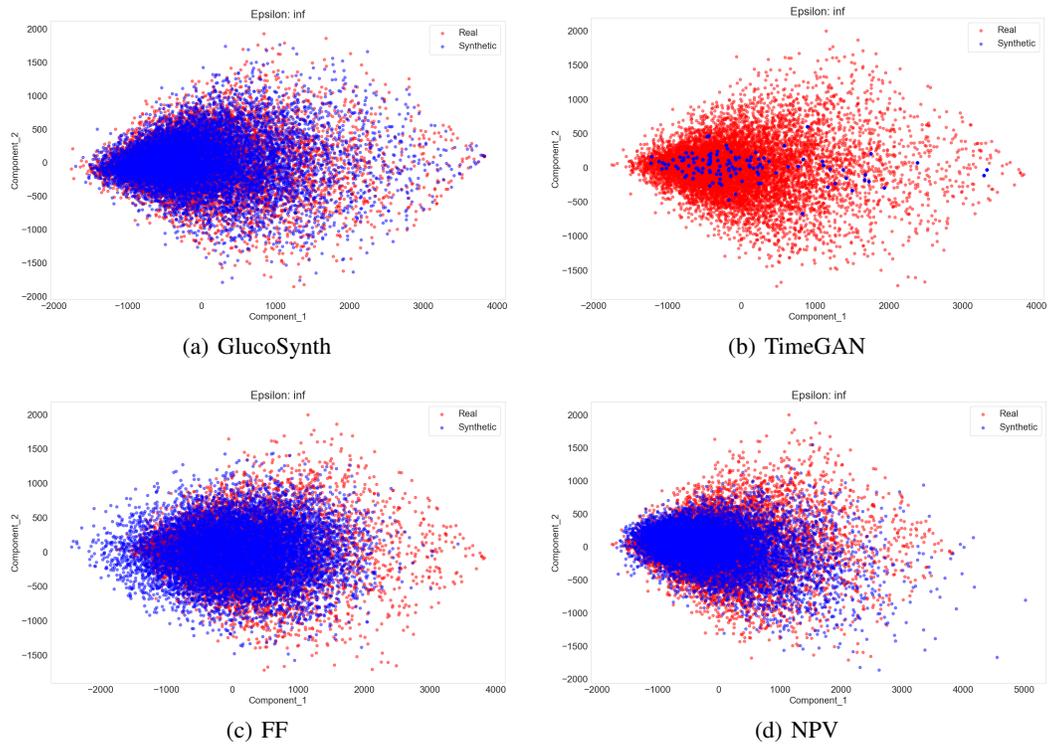


Figure 13: PCA Comparison for Nonprivate Models

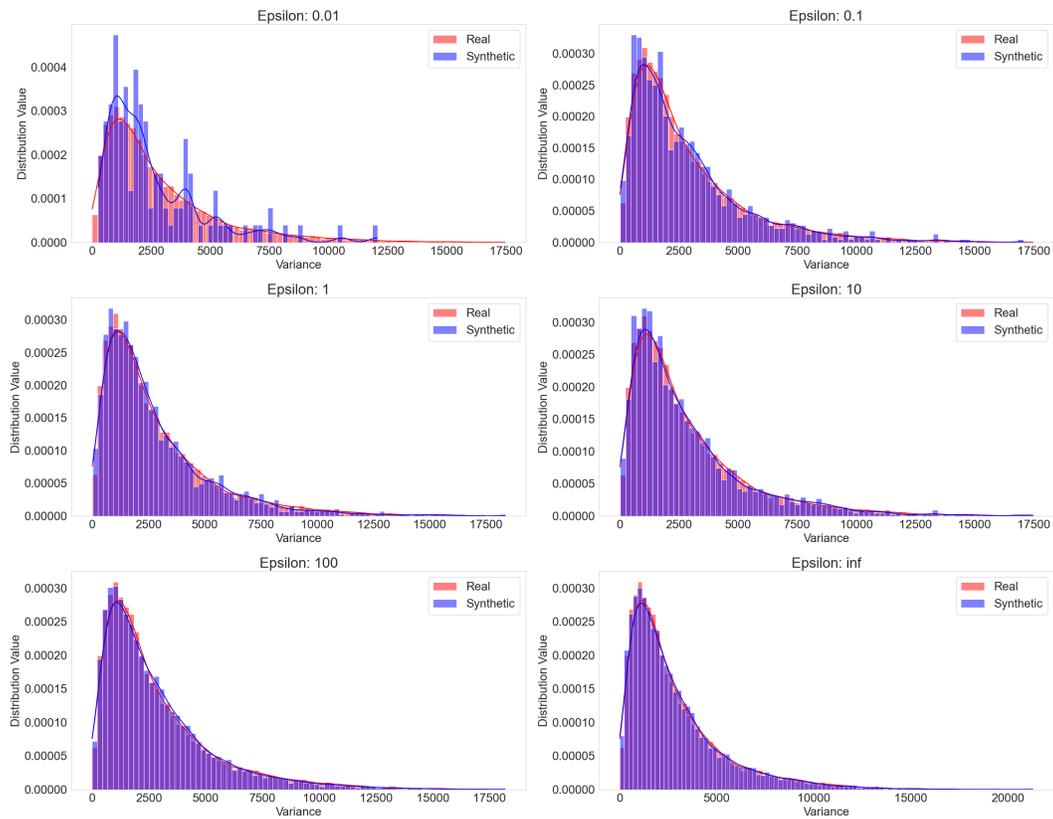


Figure 14: GlucoSynth Distributional Variance Comparison Across Privacy Budgets

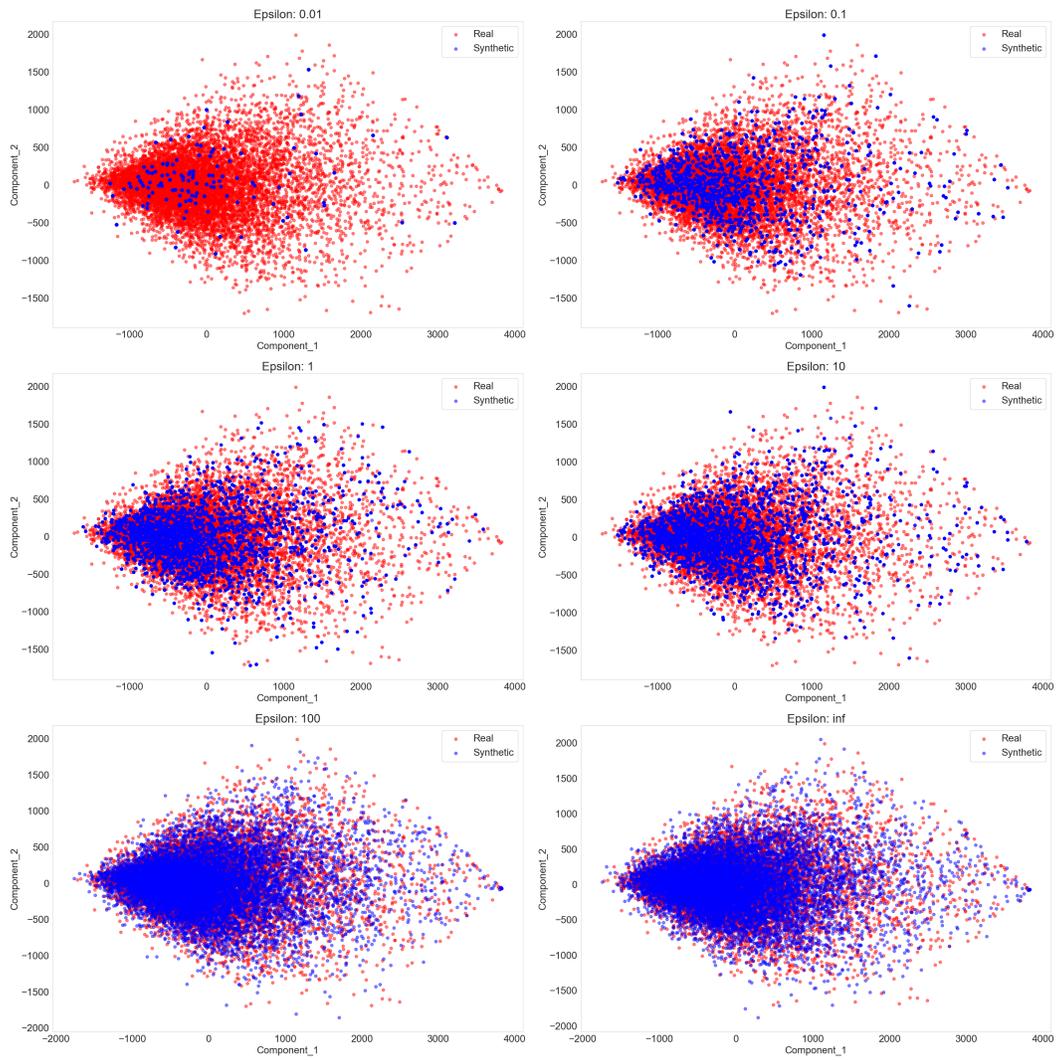


Figure 15: GlucoSynth PCA Comparison Across Privacy Budgets

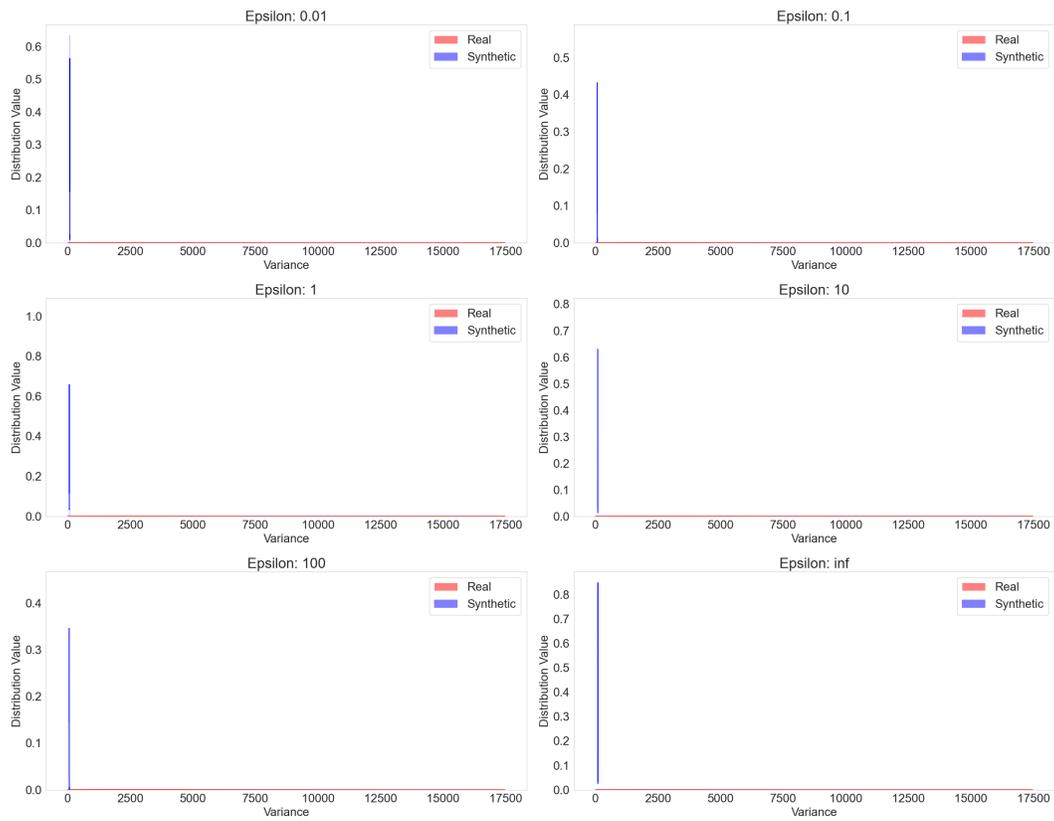


Figure 16: RGAN distributional Variance Comparison Across Privacy Budgets

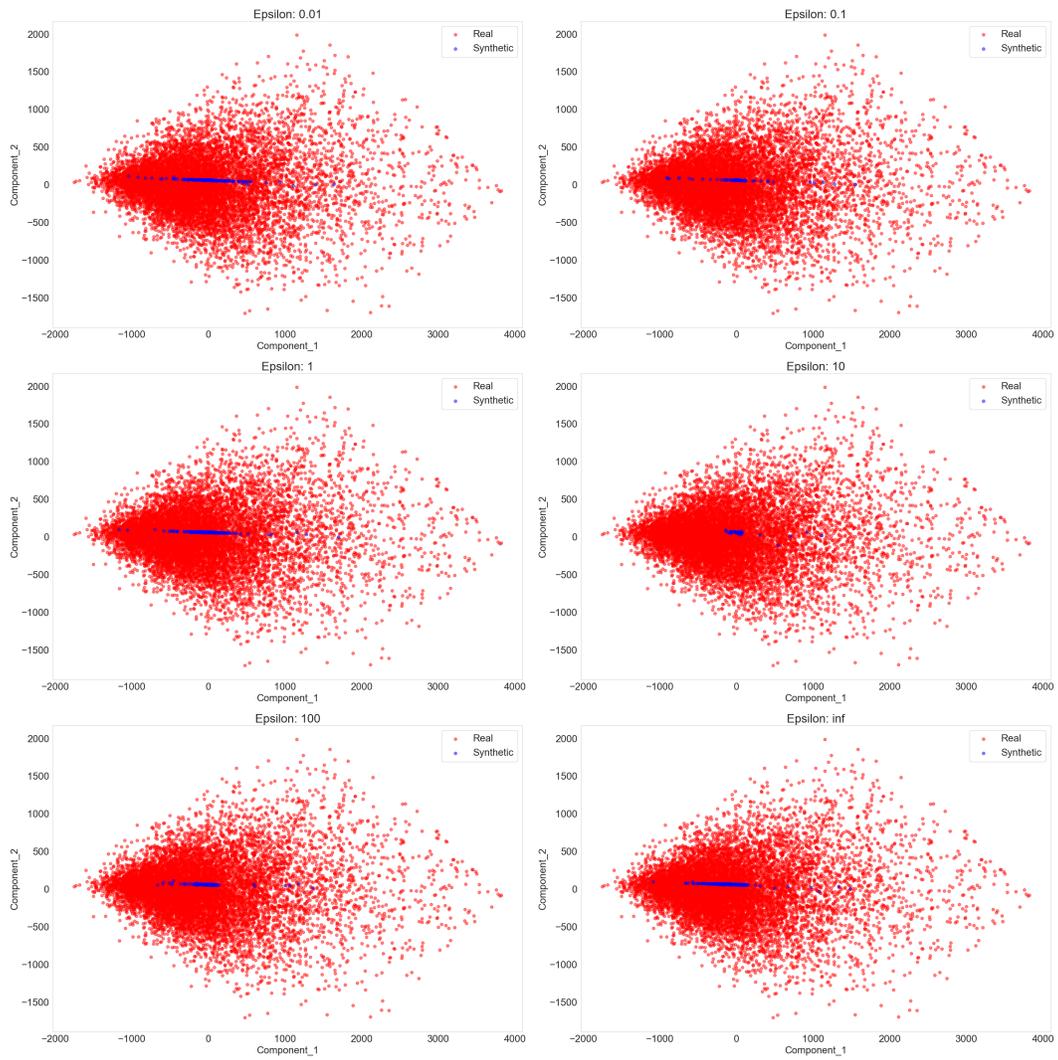


Figure 17: RGAN PCA Comparison Across Privacy Budgets

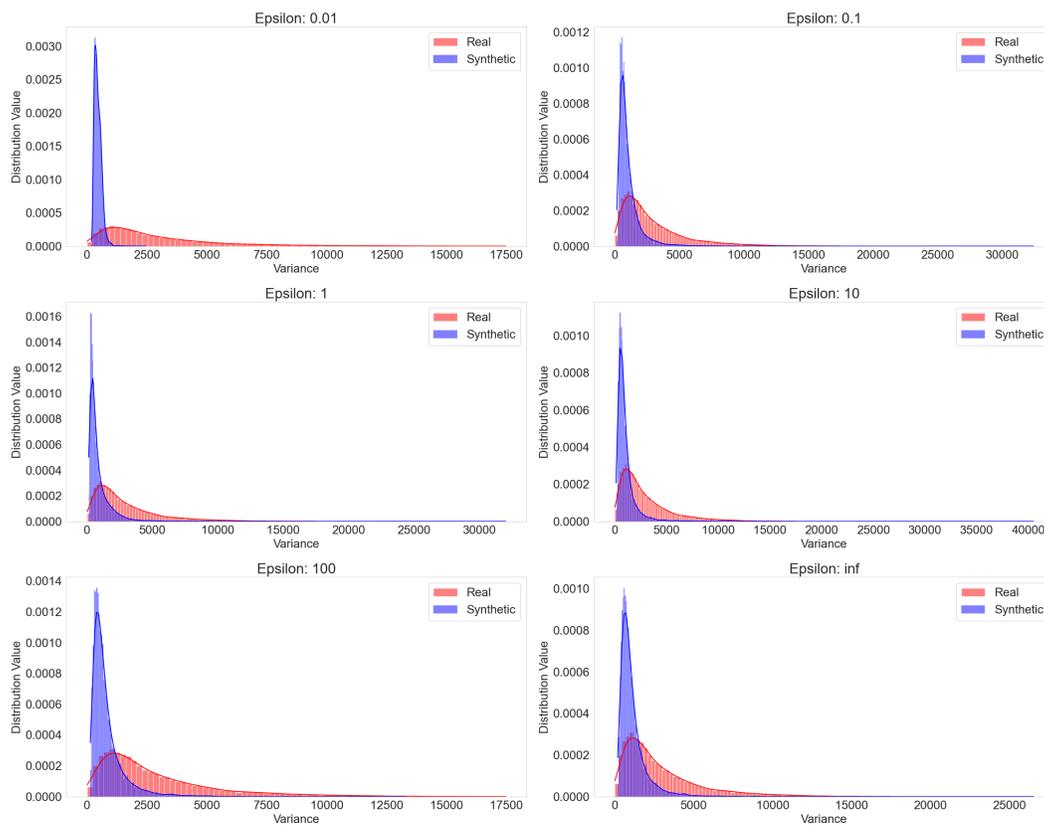


Figure 18: dpGAN distributional Variance Comparison Across Privacy Budgets

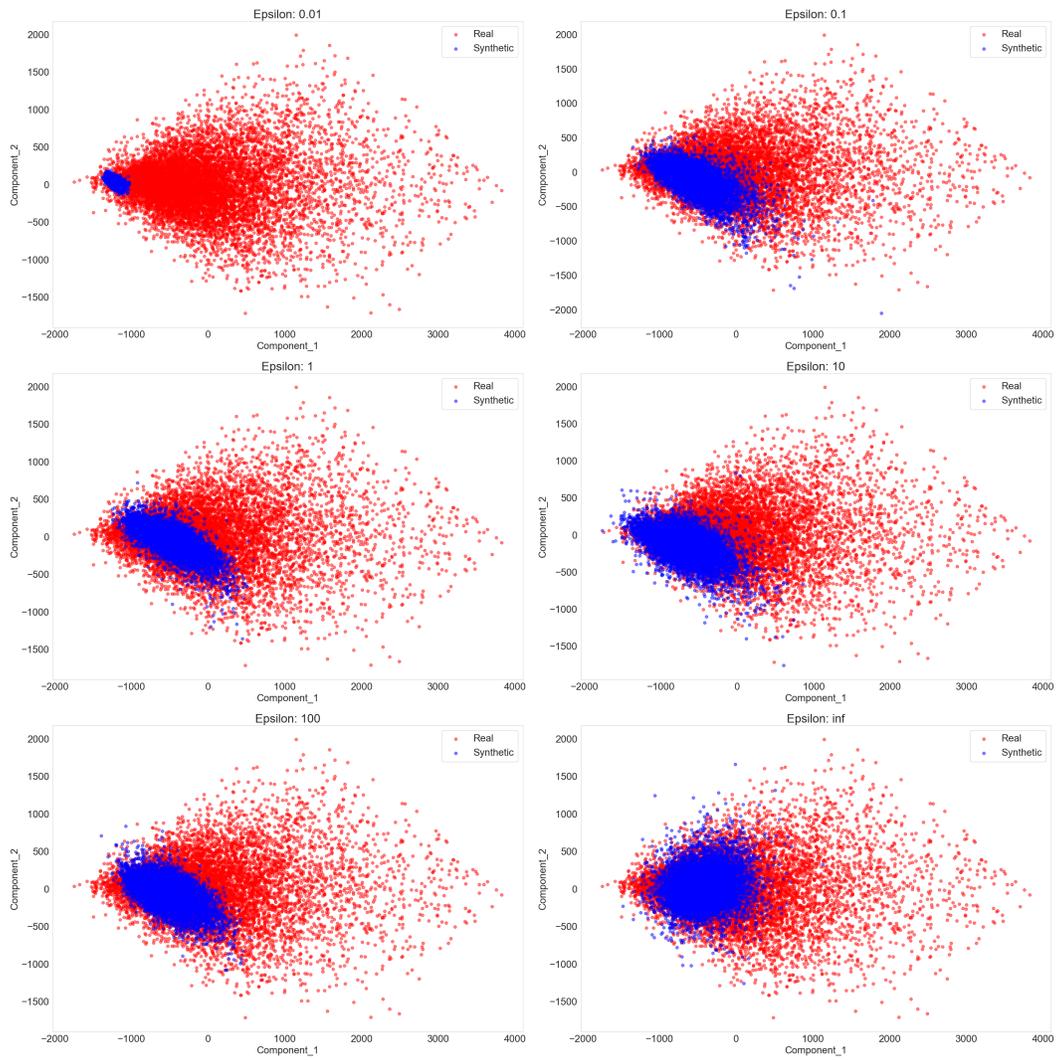


Figure 19: dpGAN PCA Comparison Across Privacy Budgets

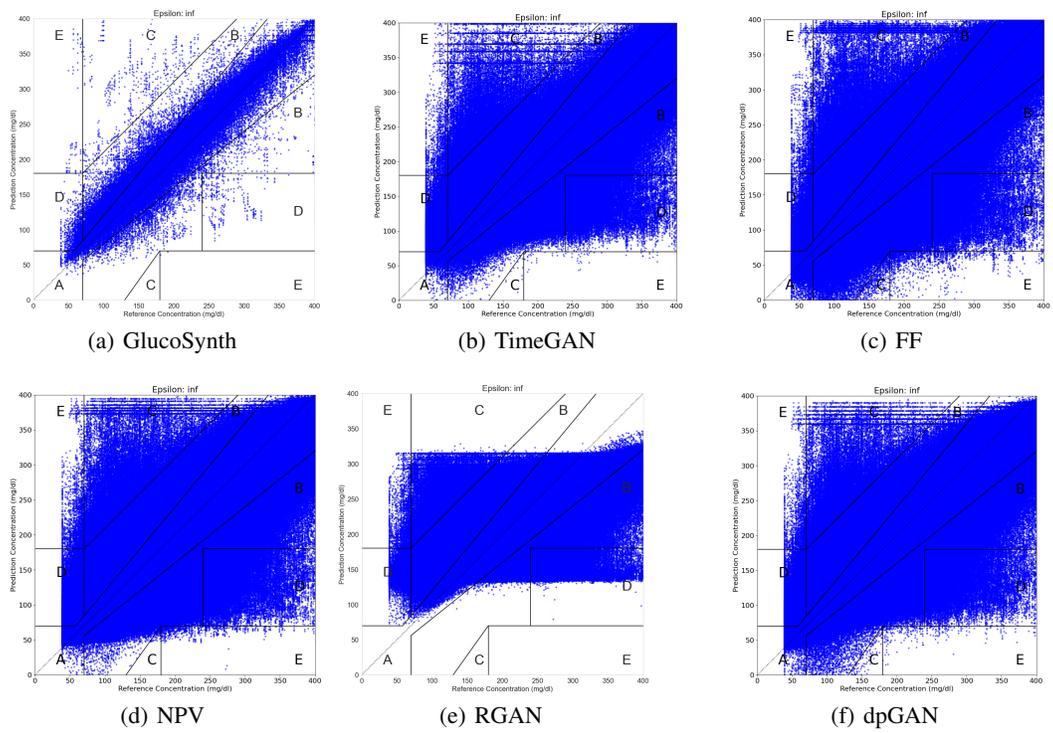


Figure 20: Clarke Error Zone Figures for All Models