

On Utility and Privacy in Synthetic Genomic Data

Bristena Oprisanu
University College London
bristena.oprisanu.10@ucl.ac.uk

Georgi Ganev
University College London and Hazy
georgi.ganev.16@ucl.ac.uk

Emiliano De Cristofaro
University College London
e.decrisofaro@ucl.ac.uk

Abstract—The availability of genomic data is essential to progress in biomedical research, personalized medicine, etc. However, its extreme sensitivity makes it problematic, if not outright impossible, to publish or share it. As a result, several initiatives have been launched to experiment with *synthetic genomic data*, e.g., using generative models to learn the underlying distribution of the real data and generate artificial datasets that preserve its salient characteristics without exposing it.

This paper provides the first evaluation of both utility and privacy protection of six state-of-the-art models for generating synthetic genomic data. We assess the performance of the synthetic data on several common tasks, such as allele population statistics and linkage disequilibrium. We then measure privacy through the lens of *membership inference attacks*, i.e., inferring whether a record was part of the training data. Our experiments show that no single approach to generate synthetic genomic data yields both high utility and strong privacy across the board. Also, the size and nature of the training dataset matter. Moreover, while some combinations of datasets and models produce synthetic data with distributions close to the real data, there often are target data points that are vulnerable to membership inference. Looking forward, our techniques can be used by practitioners to assess the risks of deploying synthetic genomic data in the wild and serve as a benchmark for future work.

I. INTRODUCTION

Progress in genome sequencing has paved the way towards preventing, diagnosing, and treating several diseases and conditions. Much of this progress is dependent on the availability of genomic data. Consequently, numerous initiatives have been established to support and encourage genomic data sharing. For example, the International HapMap Project [42] helped identify common genetic variations and study their involvement in human health and disease, while the 1000 Genomes Project [1] created a catalog of human variation and genotype data. Funding agencies, e.g., the National Institutes of Health (NIH), increasingly often make data sharing a requirement to fund grant applications [46].

Data sharing in genomics is crucial to enable progress in Precision Medicine [21]. However, this inherently conflicts with the need to protect individuals' privacy. Genomic data contains sensitive information related to heritage, predisposition to diseases, phenotype traits, etc., making it hard to

anonymize [23]. Hiding “sensitive” portions of the genome is not effective, as they can still be inferred via high-order correlation models [54].¹

As a result, genomics researchers have begun to investigate the possibility of releasing *synthetic* datasets rather than real/anonymized data [52]. This follows a general trend in healthcare; for instance, the National Health Service (NHS) in England has recently concluded a project focused on releasing synthetic Emergency Room (“A&E”) records [45]. Basically, the idea is to use generative models to learn to generate samples with the same characteristics—more precisely, with very close distributions—as the real data. In other words, rather than releasing data of actual individuals, entities share artificially generated data so that the statistical properties of the original data are preserved, but minimizing the risk of malicious inference of sensitive information [15].

Generative Models and Genomics. In particular, previous work has experimented with both statistical and generative models. Samani et al. [54] propose an inference model based on the recombination rate, which can also be used to generate new synthetic genomic samples. Yelmen et al. [71] use Generative Adversarial Models (GANs) and Restricted Boltzmann Machines (RMBs) to mimic the distribution of real genomes and capture population structures. Finally, Killoran et al. [36] use ad-hoc training techniques for GANs and architectures for computer vision tasks.

Technical Roadmap. Prior work on synthetic genomic data (Section III) has not evaluated the privacy guarantees they provide and only evaluated utility from the shallow point of view of the statistical fidelity provided by the generative model. To address this gap, we introduce a novel evaluation methodology and perform a series of experiments geared to quantify both the utility and the privacy of six state-of-the-art models used to generate human genomic synthetic data.

Our analysis unfolds along two main axes:

- 1) *Utility.* We focus on several, very common computational tasks on genomic data, measuring how well generative models preserve summary statistics (e.g., allele frequencies, population statistics) or linkage disequilibrium (see Section IV). We also compute the Kologorov-Smirnov(KS) test on the main metrics discussed as a non-parametric test.
- 2) *Privacy.* We mount *membership inference attacks* [29], having an attacker infer whether a target record was part

¹For a thorough review of privacy threats in genomics, please see [4, 44, 67].

of the real data used to train the model producing the synthetic dataset, as opposed to releasing the real dataset (see Section V). In the process, we also introduce a novel attack where the adversary only has *partial* information for a target individual (Section V-B). We choose to study MIA as it is one of the most studied attacks mounted for both genomics data and machine learning, however, it has never been studied in the context of synthetic genomic data.

Main Findings. Overall, our evaluation shows that there is no single approach for generating genomic synthetic data that performs well across the board, both in terms of utility and privacy. Among other things, we find that:

- A high-order correlation model specifically build for genomic data (Recomb) has the best utility metrics for small datasets but does so at the cost of privacy, even against weaker adversaries who only have partial information available.
- The RBM model has a better performance with increasing dataset sizes, both in terms of utility and privacy, as targets get stronger privacy protection when synthetic data is generated using a larger training set.
- Releasing synthetic datasets does not provide robust protection against membership inference attacks. We find cases where releasing the synthetic dataset sometimes offers better protection against membership inference attacks. However, because of the randomness introduced by the generative models, one cannot meaningfully predict a target’s susceptibility to privacy attacks without fixing the training set and quantifying the respective privacy loss/gain for all targets in the set.

II. PRELIMINARIES

A. Genomic Primer

Genomes and Genes. The genome represents the entirety of an organism’s hereditary information. It is encoded in DNA: each DNA strand comprises four chemical units, called nucleotides, represented by the letters A, C, G, and T. The human genome consists of approximately 3 billion nucleotides packaged into thread-like structures called chromosomes. The genome includes both the genes and the non-coding sequences of the DNA. The former determine specific traits, characteristics, or control activity within an organism. We refer to the group of genes inherited from a single parent as a *haplotype*. An *allele* is a different variation of a gene; any individual inherits two alleles for each gene, one from each of their parents. The genotype consists of the alleles that an organism has for a particular characteristic.

SNPs and SNVs. About 99.5% of the genome is shared among all humans; the rest differs due to genetic variations. Single nucleotide polymorphisms (SNPs) are the most common type of genetic variation. They occur at a single position in the genome and at least 1% of the population. SNPs are usually biallelic and can be encoded by $\{0, 1, 2\}$, with 0 denoting a

combination of two major (i.e., common) alleles, 2 a combination of two minor alleles, and 1 a combination of a major and a minor allele (which is also referred to as a heterozygous SNP). Single nucleotide variants (SNVs) are single nucleotide positions in the genomic DNA at which different sequence alternative exists [17].

Recombination Rate (RR). Recombination is the process of determining the frequency with which characteristics are inherited together. The RR is the probability that a transmitted haplotype constitutes a new combination of alleles different from that of either parental haplotype [14].

Genome-Wide Association Studies (GWAS). GWAS are hypothesis-free methods for identifying associations between genetic regions and traits. A typical GWAS looks for common variants in several individuals, both with and without a trait, using genome-wide SNP arrays [16, 43].

B. Membership Inference Attacks

Genomics. A well-understood privacy threat in genomics is determining whether the data of a target individual is part of an aggregate genomic dataset or mixture. This is known as a *membership inference attack* (MIA) [29, 66, 74]. The ability to infer the presence of an individual’s data in a dataset constitutes an inherent privacy leak whenever the dataset has some sensitive attributes. For instance, if a mixture includes DNA from patients with a specific disease, learning that a person is part of that mixture exposes their health status.

In general, genomic data contains extremely sensitive information about individuals; hence, MIAs against genomic data prompt severe privacy threats, including denial of life/health insurance, revealing predisposition to diseases, ancestry, etc.

Machine Learning. MIAs have also been studied in the context of machine learning to infer whether or not a target data point was used to train a target model. This has been done both for discriminative [11, 30, 38, 41, 51, 53, 55] and generative models [11, 24, 26]. Inferring training set membership might yield serious privacy violations. For instance, if a model for drug dose prediction is trained using data from patients with a certain disease, or synthetic health images are produced by a generative model trained on patients’ images, learning that data of a particular individual was part of the training set leaks information about that person’s health. Overall, MIAs are also used as signals that access to a target model is “leaky” and can be a gateway to additional attacks [24]. This is the main reason we choose MIA as our leading privacy metric while deferring the quantification of other attacks (e.g., attribute inference or reconstruction attacks) to future work.

C. Privacy Gain (PG) from Synthetic Data Release

Our main privacy evaluation metric is denoted as *Privacy Gain (PG)*, first proposed in [59]. The PG quantifies the privacy advantage obtained by a target t , vis-à-vis an MIA adversary, when a synthetic dataset is published instead of the real data.

MIA Training (Algorithm 1). The adversary is trained as follows. First, we take a reference dataset (which may or may

not overlap with the real data used to generate the synthetic dataset) and use it to generate a synthetic dataset (lines 1–6). This dataset is labeled as 0, i.e., it does not include the target record (line 7). The target record is added to the dataset (line 8), and synthetic data is generated from the new dataset, which includes the target record; this dataset is labeled as 1 (lines 10–13). These models are referred to as *generative shadow* models. Finally, the adversary uses the synthetic datasets to train a classifier (line 14), which distinguishes whether or not the target was used to train a generative model.

PG Estimation (Algorithm 2). We estimate the PG for a fixed target record and input dataset as follows. First, the algorithm takes the target training set and generates n_s synthetic datasets without the target record (lines 1–4). Then, the target record is added to the training set, and n_s of synthetic datasets are generated from this dataset (lines 5–9). Then, the adversary trains their attack model MIATrain (line 10). Finally, the PG for a target t is computed as $PG_t = \frac{1 - \overline{MIA}_t(S_{test})}{2}$, where $\overline{MIA}_t(S_{test}) = \sum_{S_i \in S_{test}} \frac{\Pr[MIA_t(S_i)=1]}{2^{n_s}}$ (lines 11–12).

Put simply, the PG quantifies the difference between the probability that an attacker correctly identifies that the target record belongs to the real dataset and that the attacker correctly identifies that the target record was used in training a generative model that outputs a synthetic dataset.

PG Values. The PG ranges between 0, when publishing the synthetic dataset leads to the same privacy loss as publishing the real dataset (i.e., $\overline{MIA}_t(S_{test}) = 1$) and 0.5, when publishing the synthetic dataset perfectly protects the target from MIA (i.e., $\overline{MIA}_t(S_{test})=0$). This means that $PG = 0.25$ when the probability of the adversary inferring whether or not a target is part of the training set used to generate the synthetic dataset is the same as random guessing (i.e., $\overline{MIA}_t(S_{test}) = 0.5$).

Dimensionality Reduction. To reduce the effects of high dimensionality, the attacker first maps the synthetic data to a lower feature space. This helps detect the influence of the target record on the training dataset. We experiment with four different feature sets, as done in [59]: namely, a *naive* feature set, which encodes the number of distinct categories plus the most and least frequent category for each attribute, a *histogram*, which computes the frequency counts for each attribute, a *correlation*, which encodes pairwise correlations between attributes, and an *ensemble* feature set, which combines all the previously mentioned feature sets.

D. Datasets

Our evaluation uses data from two projects: HapMap [42] and the 1000 Genome Project [1]. More specifically, we use 1,000 SNPs from chromosome 13 from the following datasets:

- 1) *CEU Population (HapMap)*. Samples from 117 Utah residents with Northern and Western European ancestry, released in phase 2 of the HapMap project.
- 2) *CHB Population (HapMap)*. Samples from 120 Han Chinese individuals from Beijing, China.

- 3) *1,000 Genomes*. Samples from 2,504 individuals from 26 different populations released from phase 3 of the 1000 Genomes project.

III. SYNTHETIC DATA APPROACHES IN GENOMICS

In this section, we provide an overview of the state-of-the-art models for generating synthetic genomic data. In particular, we discuss the Recombination model presented by Samani et al. [54], the RBM and GAN models proposed by Yelmen et al. [71], and the WGAN model from Killoran et al. [36]. We also introduce and consider two other “hybrid” models.

Recombination Model (Recomb). Samani et al. [54] propose the use of a *recombination model* as an inference method for quantifying individuals’ genomic privacy. This statistical model is based on a high-order SNV correlation that relates linkage disequilibrium patterns to the underlying recombination rate. [54] shows how to use this method to generate synthetic samples and perform Principal Component Analysis (PCA). The recombination model yields a distribution closer to the real data than models using only linkage disequilibrium and allele frequencies. The model uses a “genetic map,” which includes the recombination rate. This is provided with the dataset for the HapMap datasets, but not for the 1000 genomes data. For the latter, we use the scripts from [47] to generate the genomic map.

Restricted Boltzmann Machines (RBMs). RBMs [58] are generative models geared to learn a probability distribution over a set of inputs. RBMs are shallow, two-layer neural nets: the first is known as the “visible” (on input) layer and the second as the hidden layer. The two layers are connected via a bipartite graph, i.e., every node in the visible layer is connected to every node in the hidden one, but no two nodes in the same group are connected, allowing for more efficient training algorithms. The learning procedure consists of maximizing the likelihood function over the visible variables of the model. The RBM models re-create data in an unsupervised manner through many forward and backward passes between the layers, corresponding to sampling from the learned distribution. The output of the hidden layer passes through an activation function, which becomes the input for the former.

As mentioned, Yelmen et al. [71] use RBMs to generate synthetic genomic data. In our evaluation, we follow the same RBM settings as [71]. More specifically, we use a ReLU activation function, with the visible layer having the same size as the input we considered (1,000 features) and with the number of hidden nodes set to 100. The learning rate is set to 0.01, the batch size to 32, and we iterate over 2,000 epochs.

Generative Adversarial Networks (GANs). A GAN is an unsupervised deep learning model consisting of two neural networks, a generator and a discriminator, which “compete” against each other. During training, the generator’s goal is to produce synthetic data, and the discriminator evaluates them against real data samples to distinguish the synthetic from the real samples. The training objective is to learn the data distribution so that the data samples produced by the generator cannot be distinguished from real data by the discriminator.

Algorithm 1 MIATrain [59]

Input: A generative model $GM()$, the target record t , a reference dataset R of size n , the number of synthetic test sets n_s of size m , and the number k of shadow models.

Output: $MIA_t()$

```

1: for  $i = 1, \dots, k$  do
2:    $R_i \sim R^n$ 
3:    $f_i \sim GM(R_i)$ 
4:   for  $j = 1, \dots, n_s$  do
5:      $S_j^m \sim f_i$ 
6:      $S_{train} \leftarrow S_{train} \cup S_j^m$ 
7:      $l_{train} \leftarrow l_{train} \cup 0$ 
8:    $R'_i \leftarrow R_i \cup t$ 
9:    $f'_i \sim GM(R'_i)$ 
10:  for  $j = 1, \dots, n_s$  do
11:     $S_j^m \sim f'_i$ 
12:     $S_{train} \leftarrow S_{train} \cup S_j^m$ 
13:     $l_{train} \leftarrow l_{train} \cup 1$ 
14:  $MIA_t() \leftarrow \text{Classifier}(S_{train}, l_{train})$ 

```

We use the GAN approach also proposed by Yelmen et al. [71], mirroring their experimental settings. That is, the generator model consists of an input layer with latent dimension set to 600 and two hidden layers, of sizes 512 and 1,024, respectively. The discriminator consists of an input layer with a size equal to the number of SNPs evaluated (1,000) and two hidden layers of sizes 512 and 256, respectively, and an output layer of size 1. The output layer for the generator uses tanh as an activation function, and the output layer for the discriminator uses the sigmoid activation function. We compile both the generator and discriminator using the Adam optimization and binary cross-entropy as the loss function.

Recombination RBM (Rec-RBM). To overcome issues caused by low numbers of training samples, we propose a hybrid approach between the Recomb and the RBM models. We use the former to generate extra samples, which we then use, together with the real data samples, to train the RBM model with the same parameters as before. We do so to explore whether having more data points to train the model improves the utility of the synthetic data.

Recombination GAN (Rec-GAN). Like Rec-RBM, we use the Recomb model to generate extra training samples for the GAN model, using the same parameters as before. Again, we want to study whether having a larger dataset available for training improves the synthetic data output’s overall utility.

Wasserstein GAN (WGAN). Killoran et al. [36] propose an alternative GAN model by treating DNA sequences as a hybrid between natural language and computer vision data. The sequences are one-hot encoded, the GAN is based on a WGAN architecture trained with a gradient penalty [22], and both the generator and discriminator use convolutional neural networks [37] and a residual architecture [25], which includes skip connections that jump over some layers. The authors also propose a joint method combining the GAN model with an activation maximization design [40, 57, 73] to tune the sequences to have desired properties. However, we do not

Algorithm 2 MIAGain [59]

Input: A generative model $GM()$, the target record t , the target training set R_{out}^t of size n , the size m of the synthetic dataset, the number n_s of synthetic test sets, a reference dataset R^a , the number k of shadow models.

Output: PG_t

```

1:  $f_{out} \sim GM(R_{out}^t)$ 
2: for  $i = 1, \dots, n_s$  do
3:    $S_i \sim f_{out}^m$ 
4:    $S_{test} \leftarrow S_{test} \cup S_i$ 
5:  $R_{in}^t \leftarrow R_{out}^t \cup t$ 
6:  $f_{in} \sim GM(R_{in}^t)$ 
7: for  $i = 1, \dots, n_s$  do
8:    $S_i \sim f_{in}^m$ 
9:    $S_{test} \leftarrow S_{test} \cup S_i$ 
10:  $MIA_t() \leftarrow \text{MIATrain}(GM(), t, R^a, n, m, n_s, k)$ 
11:  $\overline{MIA_t}(S_{test}) = \sum_{S_i \in S_{test}} \Pr[MIA_t(S_i) = 1] / 2n_s$ 
12:  $PG_t \leftarrow (1 - \overline{MIA_t}(S_{test})) / 2$ 

```

include the joint model in our evaluation, as we focus on a range of statistics instead of a single desired property.

In our evaluation, we use the WGAN model with the default parameters from the implementation in [35]. The generator consists of an input layer with a dimension of the latent space set to 100, followed by a hidden layer with a size 100 times the length of the sequence (1,000), which is then reshaped to (length of the sequence, 100), followed by 5 resblocks. Finally, there is a 1-D convolutional layer followed by the output layer, which uses softmax. The discriminator has a very similar architecture but in a different order – i.e., it starts with the input layer to which the one-hot sequences are fed, that is followed by the 1-D convolutional layer, then the 5 resblocks, followed by the reshape layer and the output layer of size 1. We perform 5 discriminator updates for every generator update. Both the generator and discriminator use Adam optimization, and their learning rates are set to 0.0001, while a gradient penalty adjusts the loss as mentioned. We use a batch size of 64. For more details please refer to Appendix A.

IV. UTILITY EVALUATION

This section presents a comprehensive utility evaluation of the synthetic data generated by the models introduced in Section III. We look at common summary statistics used in genome-wide association studies, aiming to assess the accuracy loss due to the use of synthetic datasets. We choose summary statistics as our main utility evaluation metrics since prior work shows that they provide a good starting point for assessing the utility of synthetic genomic data [7, 48]. While high(er) granularity may be required in some biomedical applications, if most generative models do not accurately replicate the high-level statistics of the real data, then it is extremely unlikely they would provide reasonable accuracy on single sequences. More specifically, we analyze how well data generated by the generative models preserve allele frequencies, population statistics, and linkage disequilibrium. We also compute the KS two-sample test [27] for the goodness of fit

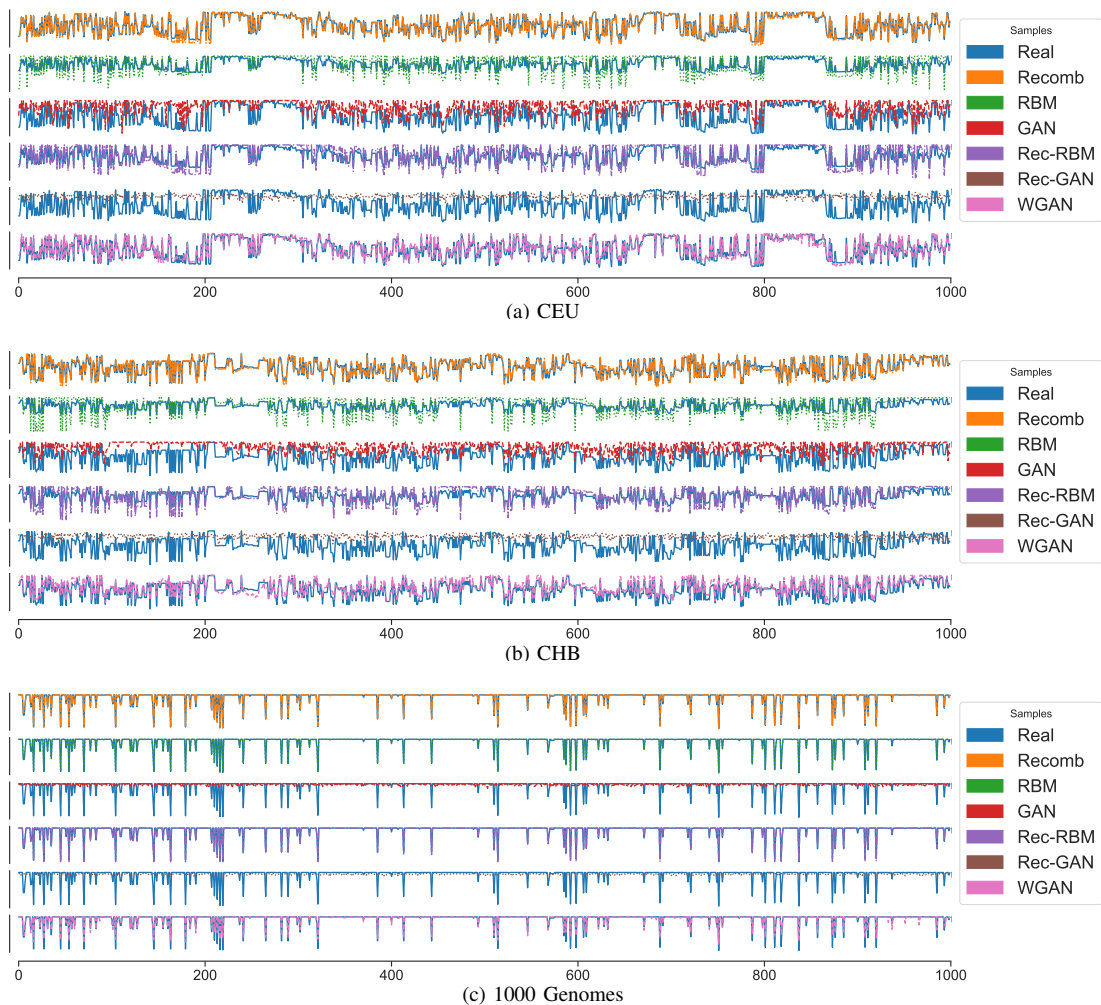


Fig. 1: Major allele frequencies for synthetic data generated by the models, plotted against the real data.

on the most relevant metrics in each subsection for each real dataset vs. the synthetic data.

A. Allele Statistics

Major Allele Frequency (MAF). In population genetics, the *major allele frequency* (MAF) is routinely used to provide helpful information to differentiate between common and rare variants in the population, as it quantifies the frequency at which the most common allele occurs in a given population. We start our utility analysis by comparing MAFs in the synthetic data vs. the real data.

In Fig. 1, we plot the MAF at each position for the real datasets and the synthetic samples, over the CEU and CHB populations, and the 1000 Genomes dataset. For CEU/CHB (Fig. 1a–1b), the Recomb and WGAN replicate best the allele frequencies in the real data. On the other hand, GAN and Rec-GAN fail to do so, and in fact, the generated samples seem random. Even though not as close to the real frequencies as Recomb, the RBM model performs better than the GAN and Rec-GAN models. In fact, RBM further improves when combined with Recomb (see Rec-RBM).

For 1000 Genomes (Fig. 1c), Recomb’s MAF distribution is also similar to the real data’s. However, RBM and Rec-RBM both display MAFs close to the real data, whereas, even with more training samples available, the GAN and Rec-GAN models still seem to produce random results. Moreover, WGAN does not match the MAF distribution for this population as closely. Overall, the difference in the MAF distributions across datasets is likely due to fewer samples available for the HapMap populations than the 1000 Genomes. To further support our findings, we also compute the KS two-sample test [27] on the statistics presented in this section. (see Table I). The test compares the agreement between the cumulative distributions of two independent samples. If the resulting D value is low, and the p-value is high, we cannot reject the null hypothesis (i.e., there is no difference between the distributions). For every two-samples test, the 95% critical value is approximately 0.195 (as we have 100 samples in each dataset), so we can reject the null hypothesis (that there is no difference between the distributions) for all synthetic data above this value. For the CEU and CHB populations, we cannot reject the null hypothesis for samples generated by the

Models	CEU		CHB		1000 Geno	
	D	p-value	D	p-value	D	p-value
MAF						
Recomb	0.04	0.36	0.03	0.53	0.008	0.99
RBM	0.47	<0.001	0.33	<0.001	0.06	0.03
GAN	0.37	<0.001	0.52	<0.001	0.58	<0.001
Rec-RBM	0.31	<0.01	0.23	<0.01	0.08	<0.1
Rec-GAN	0.54	<0.001	0.61	<0.001	0.863	0
WGAN	0.10	<0.001	0.14	<0.1	0.28	0.82
SFS						
Recomb	0.07	0.89	0.07	0.89	0.18	<0.1
RBM	0.34	<0.001	0.28	<0.001	0.12	0.44
GAN	0.40	<0.001	0.41	<0.001	0.23	<0.01
Rec-RBM	0.26	<0.01	0.23	<0.01	0.06	0.99
Rec-GAN	0.64	<0.001	0.58	<0.001	0.25	<0.01
WGAN	0.09	0.79	0.18	<0.1	0.19	<0.1
% Heterozygous Samples						
Recomb	0.19	0.47	0.29	<0.001	0.32	<0.001
RBM	0.64	<0.001	0.70	<0.001	0.13	0.34
GAN	0.90	<0.001	1.00	<0.001	0.57	<0.001
Rec-RBM	0.99	<0.001	1.00	<0.001	0.40	<0.001
Rec-GAN	0.55	<0.001	0.68	<0.001	0.52	<0.001
WGAN	0.17	<0.1	0.39	<0.001	0.45	<0.001
Linkage Disequilibrium						
Recomb	0.11	<0.001	0.12	<0.001	0.31	0
RBM	0.61	0	0.69	0	0.08	<0.001
GAN	0.07	<0.001	0.09	<0.001	0.86	0
Rec-RBM	0.40	0	0.41	0	0.38	0
Rec-GAN	0.23	0	0.21	0	0.91	0
WGAN	0.10	<0.001	0.09	<0.001	0.53	0

TABLE I: Two-sample (real vs. synthetic data) Kolmogorov-Smirnov test performed on the the main utility summary statistics presented in Section IV.

Recomb and WGAN., whereas for the 1000 Genomes dataset we cannot reject the null hypothesis for Recomb, RBM and Rec-RBM.

Alternate Allele Correlation (AAC). To evaluate whether the real and synthetic data are *genetically* different, in Appendix Fig. 12, we plot the *alternate allele correlation* (AAC). The more similar the two populations are, the closer the SNPs should be to the diagonal, as in the leftmost plots, where we have the real data against itself. The strongest AAC is with the synthetic data generated by Recomb. On the opposite side of the spectrum, the synthetic data generated by GAN and Rec-GAN have weak correlations. For the CEU and CHB populations, Rec-RBM yields stronger AACs than simple RBM and the WGAN. For the 1000 genomes dataset (Fig. 12c), there is a strong correlation between the alternate alleles for the real data and Recomb, RBM, Rec-RBM, and WGAN.

Site Frequency Spectrum (SFS). Another summary statistic that captures essential information about the underlying distribution of the allele frequencies of a given set of SNPs in a population or sample is the SFS [18, 19]. Basically, it provides a histogram whose size depends on the number of sequenced individuals. In Fig. 13, we plot the scaled folded SFS, which is the distribution of counts of minor alleles in a sample calculated over all segregating sites. We scale this value so that

a constant value is expected across the spectrum for neutral variation and constant population size, which yields the best visual comparisons. If the distribution of allele frequencies for the synthetic samples matches that of the real data, we would expect to see the two spectra aligned.

With the HapMap populations (Fig. 13a–13b), Rec-GAN suggests an excess of rare variants for a minor allele frequency around 0.1. Whereas GAN seems to generate data closer to a neutral expectation, i.e., the synthetic dataset describes a more stable population. Similarly, for the 1000 Genomes (Fig. 13c), Rec-GAN has an excess of rare variants for a minor allele frequency less than 0.1, and this is also displayed, at a lower scale, by the GAN-generated data. After computing the KS-two sample test, we cannot reject the null hypothesis for the samples generated by the Recomb and the WGAN models, for both CEU and CHB populations. We reject the null hypothesis for the 1000 Genomes dataset for synthetic data generated by the GAN and Rec-GAN.

B. Population Statistics

Next, we study population statistics to determine how close to the real dataset is the synthetic data. In particular, we look at the percentage of heterozygous variants, the fixation index, and the Euclidean Genetic Distance.

Heterozygosity. The condition of having two different alleles at a locus is denoted as heterozygosity. The percentage of heterozygous variants is commonly used in population studies, as a low percentage of heterozygous variants implies less diversity in the population. In Fig. 2a–2b, we plot the percentage of heterozygous variants in each sample for the CEU/CHB populations, comparing the real statistics (blue/leftmost bars) vs. those computed on the synthetic data. From a visual inspection, both CEU and CHB datasets, the Recomb and WGAN samples yield a similar distribution to the real data. However, when computing the KS statistic, we find that we cannot reject the null hypothesis only for those two models for the CEU datasets, and we reject the null hypothesis for all synthetic datasets generated from the CHB dataset.

For the GAN and RBM, the percentage of heterozygous samples decreases, suggesting that both models produce more homozygous variants. Moreover, even though for the major allele frequencies, Rec-RBM produces variants with statistics closer to the real data, the percentage of heterozygous variants turns out to be the lowest for both populations. By contrast, Rec-GAN produces a higher percentage of heterozygous variants than GAN, even though the major allele frequencies are not aligned with the original samples. With the 1000 Genomes (Fig. 2c), the % of heterozygous samples in the real data is lower across all samples. Once again, and in line with previous results, GAN and the Rec-GAN significantly deviate from the %s of heterozygous samples found in the real data. In fact, in Fig. 2d, we remove the GAN and the Rec-GAN models in order to take a closer look at the other models; with more data samples, the samples generated yield percentages of heterozygous samples relatively similar to the real data. For

this dataset, the only model for which we cannot reject the null hypothesis is the RBM.

Fixation Index (F_{ST}). Another way to assess how *different* are groups of populations from each other is to use the fixation index [28]. It provides a comparison of differences in allele frequency, with values ranging from 0 (not different) to 1 (completely different/no alleles in common). In Fig. 3, we compare the F_{ST} values for the real data against the synthetic samples. For illustration purposes, we also include F_{ST} of the real data against itself, which obviously yields $F_{ST} = 0$.

Recomb is once again the closest to the real data, which confirms the alignment from Fig. 1 of the allele frequencies of the synthetic recombination data with the real data. The F_{ST} value for the synthetic data produced by RBM is, for both CEU and CHB populations, less than 0.10; however, the hybrid Rec-RBM model further reduces this value to less than 0.04, and so does WGAN. For both populations, data generated by GAN and Rec-GAN has the highest F_{ST} , although, for the CHB population, the latter increases it, and for the CEU population reduces it. Finally, for the 1000 genomes, Recomb, RBM, and Rec-RBM all have F_{ST} close to the real data. While still having a low F_{ST} , WGAN has a slightly higher value. Whereas, with GAN and Rec-GAN, F_{ST} significantly deviates from the real data, even with the increased number of samples of this dataset.

Euclidean Genetic Distance (EGD). Since the fixation index does not easily allow for pairwise comparisons among populations, in Appendix Fig. 14, we plot the Euclidean Genetic Distance (EGD) between the samples in each dataset. EGD is routinely used as a measure of divergence between populations and shows the number of differences, or mutations, between two populations; a distance of 0 means there is no difference in the results, i.e., there is an exact match. From Fig. 14a–14b, where the EGD on the diagonal is 0, we observe that, for both CEU and CHB populations, the synthetic samples generated by GAN are closer to each other than by the other models. Rec-GAN generates samples with EGD close to 0, suggesting that there are very few differences between them and samples with a distance of around 30. As for the other population statistics, Recomb generates samples that match the differences observed in the real data the closest for both populations. For RBM, the samples generated have fewer differences than the real data. Perhaps more interestingly, Rec-RBM yields samples with a higher divergence than the real data; this can be a consequence of the low percentage of heterozygous samples found in the synthetic samples generated by this model (recall Fig. 2). The samples from WGAN match some of the differences observed in the real data, but the model also yields a few samples with a higher divergence.

Finally, for the 1000 Genomes (Fig. 14c), all samples in the real data have closer EGDs between each other. In fact, the samples generated by RBM yield a similar pattern in the EGD distances. Although Recomb, Rec-RBM, and WGAN do too, they exhibit a lower distance, on average, between samples. As for CEU/CHB populations, GAN and Rec-GAN models

fail to capture the differences between samples.

C. Linkage Disequilibrium (LD) Analysis

Linkage disequilibrium (LD) captures the non-random association of alleles at two or more positions in a general population – i.e., those alleles do not occur randomly with respect to each other. In Genome-Wide Association Studies, LD allows researchers to optimize genetic studies, e.g., by preventing genotyping SNPs that provide redundant information [9]. In Fig. 4, we plot the r^2 value for LD based on the Rogers-Huff method [50]. This ranges from 0 (there is no LD between the 2 SNPs) to 1 (the SNPs are in complete LD, i.e., the two SNPs have not been separated by recombination and have the same allele frequencies).

For CEU and CHB populations, RBM generates samples that display a stronger LD than the real data. With more training samples, Rec-RBM yields a weaker LD but still stronger than the real data for both models. On the other side of the spectrum, for Rec-GAN, the LD for the synthetic data is the weakest. For the 1000 Genomes, we find a stronger LD between the real samples than with the other two datasets. RBM generates samples that are almost indistinguishable from the real data in terms of LD. The LD in the synthetic datasets generated by Recomb, Rec-RBM, and WGAN have lower correlations than RBM, with GAN and Rec-GAN both failing to preserve the LD. In this case, the non-parametric KS test yields more surprising results for both the CEU and CHB populations. We find that we cannot reject the null hypothesis for samples generated by Recomb, GAN and WGAN. This implies that, even though the GAN model does not preserve well the allele frequencies or the population statistics, it manages to preserve the correlations between alleles.

For the 1000 Genomes data, the only generative models that yields samples which cannot be rejected under the null hypothesis is the RBM.

D. Takeaways

Our utility evaluation shows that there are only a handful of cases where generative models produce synthetic genomic data with high utility on popular tasks.

The Recomb model, which is based on high-order SNV correlations, generates synthetic data preserving most statistical properties displayed by the real data, even when few samples are available. We get better utility when the genetic map is included with the data rather than generated from the existing data. This conclusion is also supported by the fact that we cannot reject the null hypothesis for this model for most statistics analyzed (with the exception of the percentage of heterozygous samples for the CHB).

With RBM, more training samples improve the quality of the synthetic data, as evidenced by the difference between the HapMap populations and the 1000 Genomes dataset.

We also find that when few samples are available for training, the hybrid Rec-RBM model approach helps improve the quality of samples compared to just RBM, but the samples are not as close to the real data as for Recomb. This is clear from

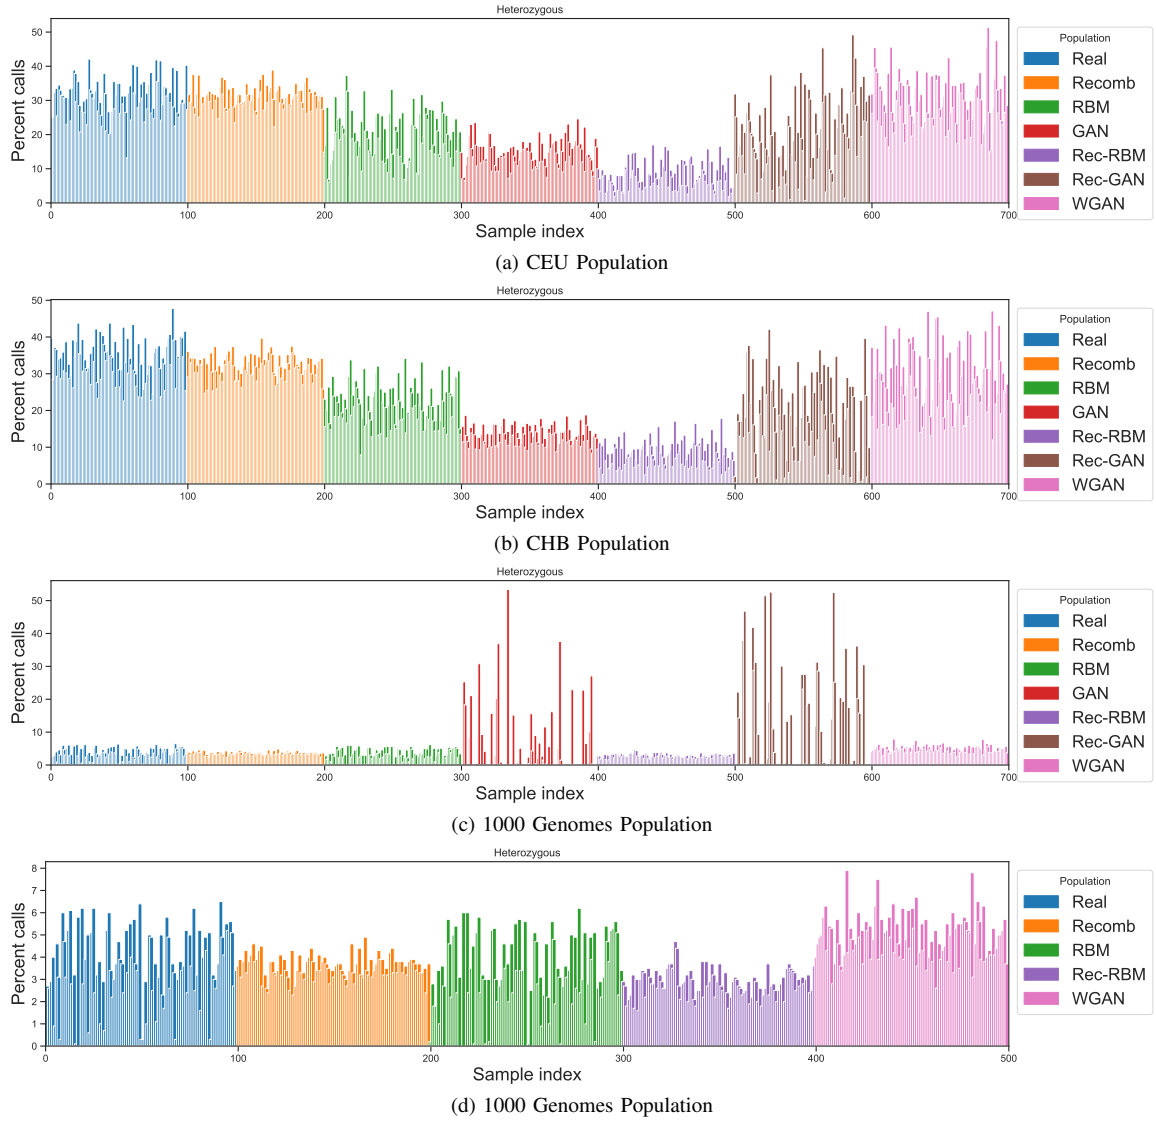


Fig. 2: Percentage of heterozygous variants in each sample in the dataset for CEU, CHB populations, and the 1000 Genomes dataset.

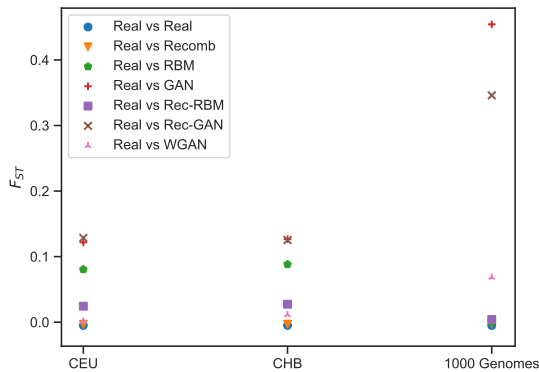


Fig. 3: Fixation index values (F_{ST}) for the CEU and CHB populations, and the 1000 Genomes dataset.

the utility of the synthetic data on the two smaller HapMap datasets. For the 1000 genomes, it is not surprising that Rec-RBM’s performance is worse than RBM since Recomb does not generate as “useful” samples as for the other two

datasets. Finally, the GAN and the Rec-GAN models generate samples with the lowest utility, regardless of the number of samples available for training. However, even though the allele and population statistics are now well preserved, our non-parametric test showcases that the correlations between alleles are preserved for the LD. In contrast to the other GAN models studied, the data generated by WGAN preserves most statistical properties of the real data.

Overall, our analysis exposes the limitations, in terms of statistical utility, of using generative machine learning models to produce synthetic genomic data.

V. PRIVACY EVALUATION

Next, we quantify the privacy provided by synthetic data by evaluating its vulnerability to Membership Inference Attacks (MIAs). More precisely, we compute the Privacy Gain, PG, obtained by releasing a synthetic dataset instead of the real data (see Section II-C). Recall that if the synthetic data does not hide nor give any additional information to an MIA

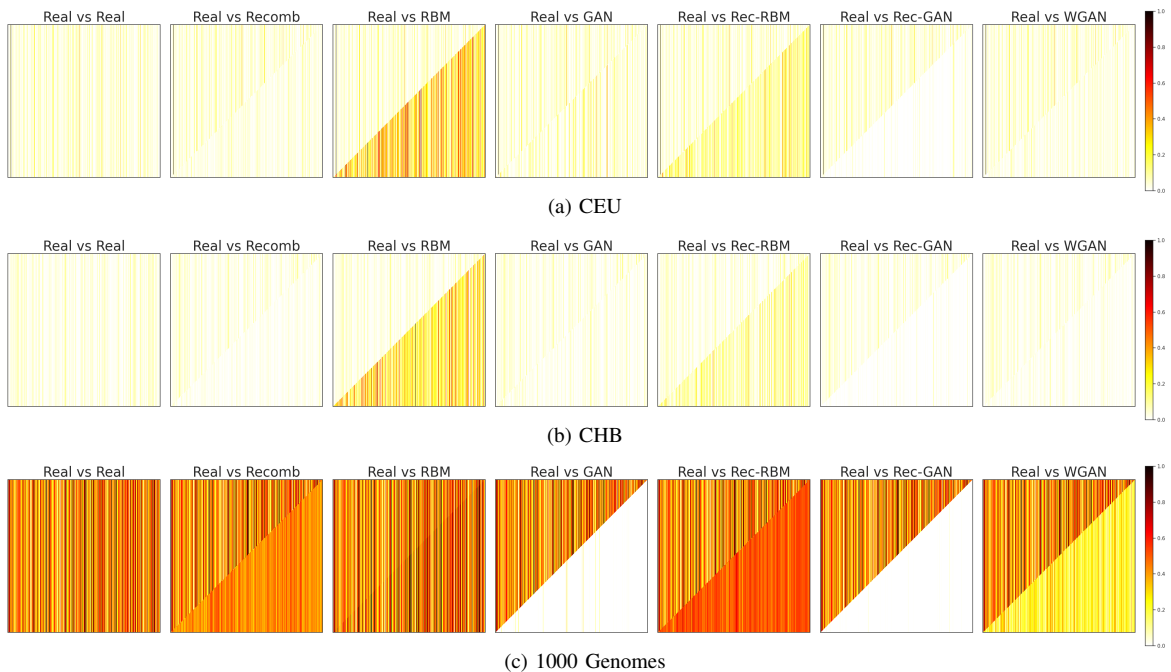


Fig. 4: Pairwise Linkage Disequilibrium for Real vs. Synthetic Samples.

attacker, PG_t , for a target record t , should have a value of around 0.25.

We present experiments for both a “standard” MIA and a novel attack, which we denote as *MIA with partial information*. The latter essentially assumes that the adversary only has access to partial data from the target sequence. We exclude GAN and Rec-GAN from the evaluation since they yield poor utility performance, so there is not really any point in evaluating their privacy.

Throughout our evaluation, we randomly choose 10 targets from each dataset across 10 test runs. In each run, we fix the target and sample a new training cohort. We train the attack classifier using 5 shadow models, using 100 synthetic training sets for each of them. We then compute the PG on 100 synthetic datasets, with a split of 50 sets generated from a training set including the target and 50 sets generated without. Finally, we report the PG for each test and each target as the average PG across all synthetic datasets tested.

A. Privacy Gain Under Membership Inference Attack

Threat Model. We assume that the adversary has access to the target, to the synthetic datasets, as well as public datasets with similar distribution to the original data (which may or may not include the target) for training his shadow models. The goal of the adversary is to identify whether the target sequence was used for training the model that generated the synthetic dataset. This could lead to the leakage of additional information about the target, as for example if a synthetic dataset of cancer patients is released and the adversary can identify that the target is part of the training set, it would lead to exposing the cancer status of the target.

We use three adversarial classifiers: K-Nearest Neighbor (KNN), Logistic Regression (LogReg), and Random Forest

(RandForest). We use four feature sets, as described in Section II-C: Naive ($\bullet F_{\text{Naive}}$), Histogram ($\bullet F_{\text{Hist}}$), Correlations ($\bullet F_{\text{Corr}}$), and an Ensemble feature set ($\bullet F_{\text{Ens}}$).

1) *HapMap Populations:* In Fig. 5, we report the PG value for targets randomly chosen from the two HapMap populations.

KNN. For CEU, using KNN (Fig. 5a, left), we find that over 74% of the targets in the synthetic dataset generated by Recomb have a PG lower than the random baseline (0.25) for the Ensemble, Correlations, and Histogram feature sets. With RBM, there are between 84% and 88% of the targets, depending on the feature set, that have a PG of 0.25; in other words, for these targets, the adversary’s probability of inferring their presence in the training set is the same as random guessing. However, between 10% and 15% of the targets have no PG at all, whereas there are between 1% and 4% of the targets, depending on the feature set, for which the synthetic data perfectly protects the target from MIA ($PG=0.5$). With Rec-RBM, at least 59% of the targets have a PG of 0.25 under the four feature sets. With WGAN, at least 48% of the targets have a PG of 0.25, depending on the feature set.

For the CHB population (Fig. 5b, left), we find that over 60% of the targets generated by the Recomb, with all features, have PG below the random guess baseline. With RBM, between 89% and 97% of targets have PG of exactly 0.25 across all feature sets, i.e., the synthetic dataset generated by the RBM for these targets does not hide nor give new information to the attacker about their membership to the synthetic dataset. Interestingly, for the Correlations feature set, no target has PG lower than 0.25. As for the CEU population, about 50% of the targets across all feature sets have a PG of 0.25 for data from Rec-RBM and at least 47% of targets from WGAN.

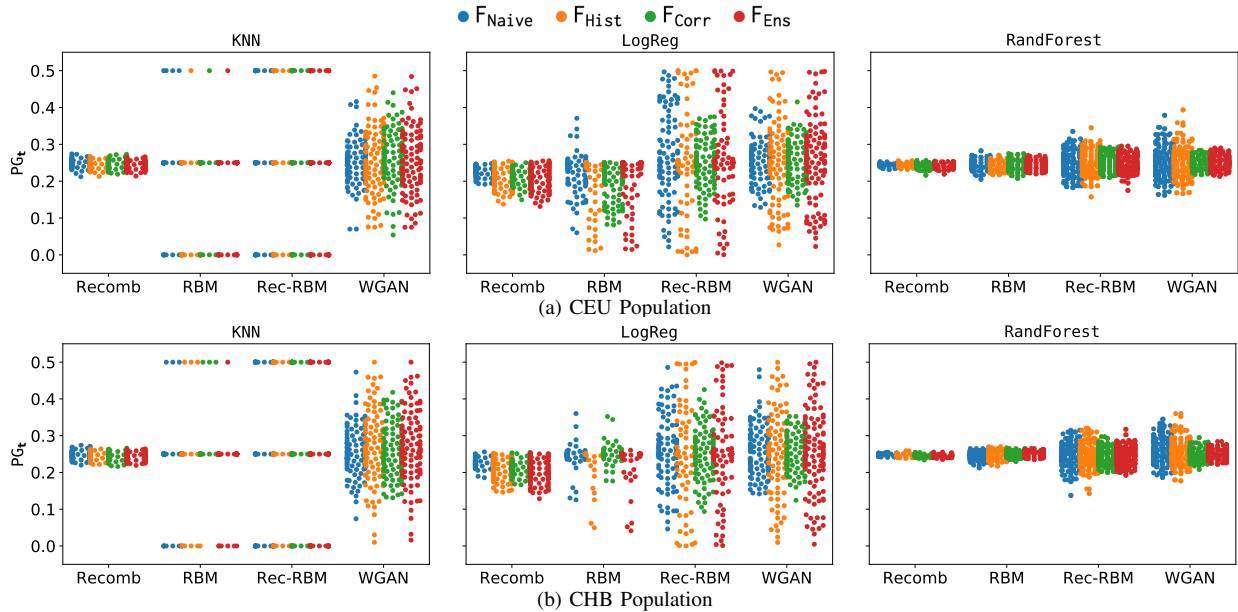


Fig. 5: Privacy Gain (PG) of different models over the two HapMap populations.

LogReg. Using LogReg, both Recomb and RBM have the lowest PG among all attack classifiers for both HapMap populations. For CEU (Fig. 5a, middle), using the Histogram feature set, 94% (resp., 96%) of the targets from Recomb (resp., RBM) have PG below 0.25, which is the random guess baseline. Under Correlations, 99% (resp., 97%) of the targets in Recomb (resp., RBM) have PG below 0.25, while, for the Ensemble feature set, 96% (resp., 98%) of the targets from Recomb (resp., RBM) have PG below 0.25. With Rec-RBM, we find that between 52% and 56% of the targets across all feature sets have PG above 0.25, and with WGAN, between 50% and 57% of the targets across all feature sets have PG above 0.25. Moreover, for the Rec-RBM and WGAN-generated data, no target consistently has a lower PG than the random guess baseline across all test runs.

For CHB (Fig. 5b, middle), with synthetic data generated by Recomb, the average PG is below the random baseline (0.25) for 99% of the targets in the Histogram feature set, 97% for Correlations, and 96% for Ensemble. For RBM, 79% of the targets in the Histogram feature set have a PG below 0.25. Under the Ensemble feature set, 84% of the targets have PG below 0.25. For synthetic data generated by Rec-RBM, we find that 54% of the targets from the Histogram and Ensemble feature sets have PG lower than 0.25, and 46% have PG over 0.25. For the Naive and Correlations feature sets, respectively, 45% and 47% of the targets have PG lower than the random guess baseline. For WGAN-generated data, the Correlations feature set yields most targets (55%) with $PG < 0.25$.

RandForest. When using RandForest as the attack classifier on data from the CEU population (Fig. 5a, right), with RBM-generated data, 73% of the targets from both Correlation and Histogram feature sets have lower PG than the random baseline. This is for 69% and 62% of the targets with, respectively, Ensemble and Naive feature sets. For the synthetic data

generated by Rec-RBM, about 51% of targets have a PG of over 0.25, and 49% of the targets have PG of less than 0.25 across all feature sets. For WGAN, between 46% and 59% of the targets from all feature sets have a PG less than the random guess baseline (0.25), with the Correlations feature set having the least percentage of vulnerable targets (59%).

For CHB (Fig. 5b, right), the lowest privacy gain for the samples generated by Recomb: over 79% of all targets for each of the four feature sets have PG lower than the random baseline. For the synthetic samples from RBM, 71%, 61%, and 53% of the targets under the Naive, Histogram, respectively, Ensemble feature sets have PG lower than 0.25. However, for the Correlations feature sets, we find that 55% of the targets have a PG of 0.25, meaning that those targets are protected from MIA. For the synthetic samples generated by Rec-RBM, 54% of the targets from the Histogram and Ensemble feature sets and 47 and 45% of the targets from the Correlations and Naive, respectively, have PG lower than 0.25. Finally, for the WGAN, we find that between 44% and 53% of the targets have $PG > 0.25$ across all four feature sets. To give a better idea of different trade-offs between utility and privacy, we also plot the PG vs. the F_{ST} (see Section IV-B) for all the models presented in Appendix D.

2) *1000 Genomes Population:* For this dataset, we focus our analysis on the RBM model, as it generated synthetic data closest to the real data across all utility metrics evaluated.

Random Target. In Figure 6a, we plot the PG for the synthetic data generated by RBM for randomly chosen targets. With the RandForest MIA classifier, we observe that 56%, 75%, 92%, and 50% of the targets have PG higher than the random guess baseline ($PG \geq 0.25$) for, respectively, Naive, Histogram, Correlations, and Ensemble feature sets. The high percentage of targets that have a $PG \geq 0.25$ for the Correlations suggests that the impact of a single target in the training dataset of the

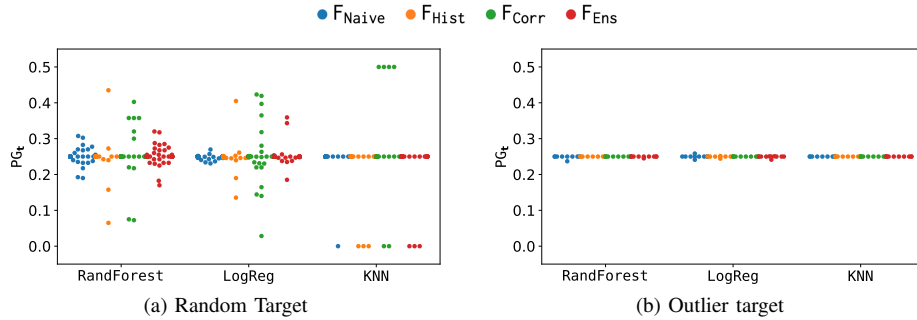


Fig. 6: Privacy Gain (PG) of RBM over the 1000 Genomes dataset.

RBM on the correlations of the synthetic dataset is minimal.

Then, with the LogReg classifier, we find a high variation in PG, similar to the HapMap populations in the case of Rec-RBM. Under the Naive feature set, half of the targets have PG gain below the random guess baseline. Similarly, 43%, 45%, and 43% of targets have a PG lower than 0.25 for, respectively, the Correlations, Histogram, and Ensemble feature sets. Finally, for the KNN classifier, over 94% of the targets have a PG of 0.25 across all feature sets.

Outlier Target. To better understand whether or not, with more training data, a target’s signal in the synthetic dataset is diluted, we also test an “extreme” outlier case. That is, we craft an outlier target that has only minor alleles at all positions. While we are aware that this case would be sporadic in a real-world scenario, our goal is to observe whether and how much this impacts PG. To this end, in Figure 6b, we plot the PG of this outlier case across 10 test runs.

With RandForest, we find that PG is below 0.25 for 8 of the 10 test runs under the Ensemble feature set. In fact, this is the only combination between attack classifier and feature set for which a greater percentage of the targets have a lower privacy gain than in the random target case. For the Naive feature set, PG is below the random baseline in only 3 of the test runs. For the Correlations and Histogram feature sets, all test runs yield PG of 0.25 or above. With LogReg, 4 out of 10 of the test runs for the Naive, Histogram, and Correlations feature sets yield PG below 0.25. For Ensemble, this happens for 6 test runs. Finally, with KNN, across all feature sets, PG for all test runs is 0.25, i.e., the synthetic data does not disclose any membership information regarding the outlier.

While there are differences across classifiers/feature sets, the PG is centered around 0.25 for all test runs in the outlier target setting. This is evident from Figure 6b, which implies that, across all test runs, the accuracy of the MIA is not much better than random guessing.

3) *Takeaways:* The different combinations of datasets, attack classifiers, and feature sets yield varied results with respect to privacy. This is due to two main reasons: first, not all classifiers have the same accuracy on tasks for the same dataset, as shown in previous work [3]. Second, the features that the generative model preserves after training will “reflect” in the synthetic data; thus, this will impact the PG based on the feature extraction method. On the HapMap populations, while the utility experiments show that Recomb-generated synthetic

data is “closest” to the real data, it does so with a significant privacy loss. The RBM-generated synthetic data is the most vulnerable under the LogReg classifier, with at least 70% of the targets across both populations and all feature sets having PG below the random guess baseline. This suggests that, with few data samples available for training, the RBM model is likely to overfit and is thus susceptible to MIAs.

For Rec-RBM and WGAN, the attacker cannot reliably predict membership, i.e., extra samples from the Recomb model in the training of the Rec-RBM dilute the target’s signal in the training data. However, for both models, we still find combinations of targets and training sets for the attack classifier for which PG is significantly lower than the random guess baseline; i.e., the synthetic data will still expose membership information about the respective targets. On the 1000 Genomes, PG values have a higher variation overall when the target is chosen randomly from the dataset than for the two smaller HapMap datasets. The results for RBM confirm our hypothesis that the target’s influence is diluted within larger datasets. Once again, this does not mean that MIAs are not possible for Rec-RBM and WGAN; it depends on the combination of the target, training set, attack classifier, and feature set.

However, with an “extreme” outlier (i.e., a target with minor alleles at all positions), the synthetic data generated by RBM does not have a big impact on PG. In this case, the PG is actually close to the random guess baseline across all test runs.

We find that targets get very different levels of privacy protection, based on the combination of target and training set used to generate the synthetic data. For a privacy mechanism to provide good privacy protection, it needs to be predictable, i.e., all targets should have a privacy gain above the random guess baseline, which our evaluation shows it is not the case for the data generated by the models we have evaluated.

B. Privacy Gain under Membership Inference Attack with Partial Information

Next, we introduce a novel attack, which we denote as MIA with Partial Information (MIA-PI). Basically, we only give the attacker access to a fraction of SNVs from the target sequence, chosen at random. The attacker then uses the Recombination model from [54] as an inference method to predict the rest of the sequence. Compared to the previous attack, the adversary trains their (attack) classifier using the sequence *inferred* from

the partial data. Thus, the PG formula also needs to be adjusted to account for how likely an adversary is to identify a target from partial information.

Threat Model. We assume that, similar to the MIA attack, the adversary has access to the synthetic dataset, as well as public datasets with similar distribution to the original data (which may or may not include the target) for training his shadow models. In contrast with the previous MIA, the adversary does not obtain access to the full sequence of the target, but only to a percentage of the SNPs from the target sequence (as for example, the victim undertakes some specific genomic tests and the adversary learns that information about the victim from the outcome of those tests), as well as public knowledge about genomic data. The goal of the adversary is to identify if the partial target sequence was part of the training set which generated the synthetic data.

PG for MIA-PI. Assuming the attacker has partial information t' as a fraction of the SNVs from t , they first use the Recomb model, as an inference algorithm, to predict the rest of the SNVs from the target sequence, which we denote by t_p . The privacy gain is computed as:

$$PG_t = \frac{\overline{MIA}_{t_p}(R_t) - \overline{MIA}_{t_p}(S_{test})}{2}, \text{ where}$$

$$\overline{MIA}_{t_p}(S_{test}) = \sum_{S_i \in S_{test}} \frac{\Pr[MIA_{t_p}(S_i) = 1]}{2 * n_s}, \text{ and}$$

$$\overline{MIA}_{t_p}(R_t) = \sum_{R_i \in R_t} \frac{\Pr[MIA_{t_p}(R_i) = 1]}{2 * n_s}.$$

That is, the privacy gain, in this case, is computed as the difference between the probability that the attacker, who has partial information about the target record, correctly identifies that the target is part of the real dataset versus the target being part of the training set used to generate the synthetic dataset.

As a result, PG now ranges between -0.5 and 0.5, where 0.5 means that having the real dataset R and the partial information t' about the target allows the adversary to infer the membership of t in R , while the synthetic dataset reduces the adversary's chance of success (i.e. $\overline{MIA}_{t_p}(R_t) = 1$ and $\overline{MIA}_{t_p}(S_{test}) = 0$). A negative PG value means that publishing the synthetic data, instead of the real data, improves the adversary's chance to correctly infer membership of the target t (i.e. $\overline{MIA}_{t_p}(R_t) < \overline{MIA}_{t_p}(S_{test})$). If publishing the synthetic data does not increase nor decrease the adversary's inference powers, we should have PG=0 (i.e. $\overline{MIA}_{t_p}(R_t) = \overline{MIA}_{t_p}(S_{test})$).

Experiments. From the experiments presented in Section V-A, we find that the attack classifier that yields the lowest PG is Logistic Regression; thus, we only experiment with that one to ease presentation. In the following, we present the results of the MIA-PI experiments for the *CEU population*, focusing on the *Recomb and RBM models* (as mentioned, with a LogReg attack classifier). We do so as these two models yield the lowest PG in Section V-A.




Fig. 7: Accuracy of the Membership Inference Attack with access to full and partial information (0.05, 0.1, and 0.2 ratio) for *Recomb*.




Fig. 8: Privacy Gain (PG) for synthetic samples from *Recomb*.

Recomb. In Figure 7, we plot the Cumulative Distribution Function (CDF) of the accuracy of the attack for Recomb when the adversary has access to the full sequence vs. partial information, specifically, a ratio of 0.05, 0.1, and 0.2 of the total SNVs from the target sequence. Interestingly, even when only 0.05 of the target SNVs are available to the attacker, for 90% and 91% of the targets from the Histogram and respectively Ensemble feature sets, the attacker's accuracy is still above the random guess baseline (50% accuracy). Our intuition is that many targets are vulnerable to the attack, even with little partial information, since we use the Recomb model not only for the attack but also as an inference method to predict the rest of the sequence.

To explore how much of the MIA-PI vulnerability is due to the release of synthetic datasets, and not only by how much information the attacker has available, in Figure 8, we plot the CDF of the PG with MIA-PI. In line with the accuracy results, we find that the PG is greater than 0 for at least 88% of the targets for all ratios of partial information tested in the case of the Correlations feature set. However, for the other three feature sets, releasing the synthetic dataset instead of the real data decreases the privacy gain (i.e., $PG > 0$) for the majority of targets. When the adversary has access to just 5% of the SNVs from the target, there is a negative PG for 61% of the targets under the Histogram and 62% of the targets under the Ensemble feature sets. With 10% of the target sequence available, 54% and 60% of the targets under, respectively, the Histogram and the Ensemble feature sets have negative PG, and with 20%, these numbers go up to 64% and 67%. There are more targets with negative PG with increasing partial information available to the attacker for the Naive feature set, i.e., 59%, 70%, and 84% with, respectively, 5%, 10%, and 20% of the target sequence available. Overall, this shows that releasing the synthetic dataset instead of the synthetic data does not mitigate privacy, even when the attacker does not have access to the full sequence.

RBM. In Figure 9, we plot the CDF for the accuracy of the attack for the RBM model for both full and partial information about the target record available to the attacker. Across all feature sets, there is an increase in the accuracy of the attack with




Fig. 9: Accuracy of the Membership Inference Attack with access to full and partial information (0.05, 0.1, and 0.2 ratio) for *RBM*.




Fig. 10: Privacy Gain (PG) for synthetic samples generated by *RBM*.

more information available to the attacker, as is expected. We also look at CDF for the PG in the case of partial information available to the attacker in Figure 10. Once again, under the Naive feature set, increasing the partial information available to the attacker negatively correlates with the percentage of targets with a negative PG. Under all other feature sets, for most targets, releasing the synthetic dataset instead of the real data yields a positive PG, meaning that releasing the synthetic dataset instead of the real dataset improves the PG.

Takeaways. We find that not even decreasing the attacker’s power by only giving him partial information from the target sequence mitigates privacy for the Recomb-generated synthetic data. This is likely because, using the Recomb model as both generative and inference model, the adversary’s power is increased since the feature set extracted from the synthetic data will be closer to the feature set for the predicted target.

However, in this case, for *RBM*, we see an increase in the privacy gained by releasing synthetic data instead of real data. This implies that, even if the *RBM* is likely to overfit when few samples are available for training, it does so on the predicted sequence of the target rather than on the full sequence and thus decreases the accuracy of the MIA.

Overall, not even with partial data from the target sequence, we obtain privacy gain values constantly better than random guessing, which, as mentioned before, indicates that synthetic data is not really a reliable privacy defense.

VI. RELATED WORK

In this section, we review related work on synthetic data, genomic privacy, and MIAs against machine learning models.

Synthetic Data Initiatives. In recent years, researchers have focused on generating synthetic electronic health records (EHR), aiming to facilitate research in and adoption of machine learning in medicine. Choi et al. [13] use a combination of an autoencoder with a GAN model, called medGAN, to generate high-dimensional multi-label discrete data. ADS-GAN [72] uses a quantifiable definition for “identifiability” combined with the discriminator’s loss to minimize the probability of patient’s re-identification, while CorGAN [61] combines convolutional GANs and convolutional autoencoders

to capture the correlations between adjacent medical features. Biswal et al. [8] use variational autoencoder to synthesize sequences of discrete EHR encounters and encounter features.

Other initiatives focus on synthetic data modeled on primary care data [45, 63, 64, 68]. Researchers have also explored generating synthetic health patient data to detect cancer and other diseases, e.g., RDP-CGAN [62] combines convolutional GANs and convolutional autoencoders, both trained with Rényi differential privacy [39].

Specific to genomes is the work presented in Section III [36, 54, 71], which we have evaluated in terms of their utility–privacy tradeoffs.

Privacy in Genomics. Researchers have focused on studying and mitigating privacy risks in genomics. One of the first attacks on genomic data is the Membership Inference Attack proposed by Homer et al. [29], showing that an adversary can infer the presence of an individual’s genotype within a complex DNA mixture. This attack has been improved by Wang et al. [66] using correlation statistics of a few hundred SNPs. Then, Im et al. [32] show that the summary information from genome-wide association studies, such as regression coefficients, can also reveal an individual’s participation within the respective study. Membership inference has also been shown possible in the context of the Beacon network [49, 56, 65], a federated service that answers queries of the form “does your data have a specific nucleotide at a specific genomic coordinate?”

Chen et al. [12] study the effects of differential privacy protection against membership inference attack on machine learning for genomic data. However, their study is focused on privacy leakage via providing access to trained classification models, whereas we study the privacy leakage from sharing synthetic datasets.

MIAs against Machine Learning Models. MIAs have long been studied in the context of machine learning. Shokri et al. [55] present the first attack against discriminative models, aiming to identify whether a data record was used in training, using an approach based on shadow models. Hui et al. [30] also propose an attack against discriminative models, which probes the target model and infers membership directly from the probes instead of shadow models. Hayes et al. [24] present the first MIA against generative models like GANs; they use a discriminator to output the data with the highest confidence values as the original training data. Hilprecht et al. [26] study MIAs against both GANs and Variational AutoEncoders (VAEs), while Chen et al. [11] propose a generic MIA model against GANs.

Stadler et al. [59] recently evaluate MIAs in the context of synthetic data and show that even access to a single synthetic dataset output by the target model can lead to privacy leakage. We re-use their framework for quantifying the privacy gain when a synthetic dataset is released instead of the real dataset. However, not only do we do so for a specific context (namely, genomics), but we also measure utility (while they only study privacy). In fact, there are several distinguishing characteristics

between “generic” synthetic data and *genomic* data, which makes the evaluation significantly different, if not harder; paramount among these is the fact that all features found in genomic sequences are correlated with each other, unlike with generic datasets. Finally, we introduce and measure a novel attack whereby the attacker only has access to partial genomic information about the target.

VII. CONCLUSION

This paper presented an in-depth measurement of state-of-the-art methods to generate synthetic genomic data. We did so vis-à-vis their utility, with respect to several common analytical tasks performed by researchers and the privacy protection they provide compared to releasing real data.

High-quality synthetic data must accurately capture the relations between data points; however, this can enable attackers to infer sensitive information about the training data used to generate the synthetic data. This was illustrated by the performance of the Recomb model on the HapMap datasets: while it achieves the best utility, it does so at the cost of significantly reducing privacy.

Discussion. Overall, we found that there is no single method that outperforms the others for all metrics and all datasets. While this is perhaps not surprising, we are the first to present a systematic, re-usable methodology to analyze all kinds of methods to generate synthetic genomic data. For instance, this allows us to be the first to show that, unlike what previously suggested in previous work, models based on a simple GAN architecture (i.e., GAN and Rec-GAN) are not a good fit for genomic data. In fact, these provide low (the lowest) utility across the board. Our methodology also allowed us to shed light on the influence of the size of the training dataset, especially in the case of generative models. For example, utility improves with more samples in the hybrid Rec-RBM for the smaller HapMap datasets, and the RBM for the 1000 Genomes dataset; also, we measured a decrease in the number of targets exposed to membership inference.

We also introduced a new “MIA with partial information,” which shed light on the fact that not even decreasing the adversary’s power by limiting their knowledge of the target to partial information fully mitigates the privacy loss. Finally, our modular evaluation framework paves the way for practitioners, scientists, and researchers to easily build on our work and assess the risks of deploying synthetic genomic data in the wild, for a wide range of applications, and serve as a benchmark for techniques proposed in the future.

Limitations & Future Work. Our evaluation focuses on existing generative methods for synthetic genomic data; thus, we have not engaged in fine-tuning the (hyper-)parameters of the models evaluated. Moreover, one might argue that the techniques we evaluate were not designed with privacy in mind, unlike previous work on differentially private generative models for images or clinical data [2, 6, 10]. That is, it is not entirely surprising that they yield small privacy gains. However, to the best of our knowledge, *no differentially*

private generative model has been proposed for genomic data, which is the focus of our study.

In fact, prior work has shown that, for precision medicine applications, the high dimensionality of the data tends to be a major limitation, resulting in poor utility for differentially private mechanisms [5, 20, 34, 69, 70]. Differentially private techniques for GWAS are also known to yield poor accuracy as the number of features is large, relative to the number of patients in a study [33]. Nonetheless, we plan to experiment with the possible adaptation of differentially private models to genomic data and evaluate them in future work.

Finally, future work should search for, and experiment with, genomics use cases that have more data points, possibly relying on further collaborations with biomedical researchers. We also intend to extend our privacy evaluation to understand how much the privacy loss stemming from releasing (synthetic) datasets affects the relatives of those included in the training set of the corresponding generative models [31, 60].

Acknowledgments. This work was supported by a Google Faculty Award on “Enabling Progress in Genomic Research via Privacy Preserving Data Sharing.”

REFERENCES

- [1] 1000 Genomes Project Consortium. A global reference for human genetic variation. *Nature*, 526(7571), 2015.
- [2] G. Acs, L. Melis, C. Castelluccia, and E. De Cristofaro. Differentially private mixture of generative neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 31(6):1109–1121, 2018.
- [3] D. R. Amancio, C. H. Comin, D. Casanova, G. Travieso, O. M. Bruno, F. A. Rodrigues, and L. da Fontoura Costa. A systematic comparison of supervised classifiers. *PLoS one*, 9(4), 2014.
- [4] E. Ayday, E. De Cristofaro, J.-P. Hubaux, and G. Tsudik. Whole genome sequencing: Revolutionary medicine or privacy nightmare? *Computer*, 48(2), 2015.
- [5] C.-A. Azencott. Machine learning and genomics: precision medicine versus patient privacy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128), 2018.
- [6] B. K. Beaulieu-Jones, Z. S. Wu, C. Williams, R. Lee, S. P. Bhavnani, J. B. Byrd, and C. S. Greene. Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7):e005122, 2019.
- [7] C. Benner, A. S. Havulinna, M.-R. Järvelin, V. Salomaa, S. Ripatti, and M. Pirinen. Prospects of fine-mapping trait-associated genomic regions by using summary statistics from genome-wide association studies. *The American Journal of Human Genetics*, 101(4):539–551, 2017.
- [8] S. Biswal, S. Ghosh, J. Duke, B. Malin, W. Stewart, and J. Sun. EVA: Generating Longitudinal Electronic Health Records Using Conditional Variational Autoencoders. *arXiv:2012.10020*, 2020.
- [9] W. S. Bush and J. H. Moore. Genome-wide association studies. *PLoS Computational Biology*, 8(12), 2012.
- [10] D. Chen, T. Orekondy, and M. Fritz. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. *arXiv preprint arXiv:2006.08265*, 2020.
- [11] D. Chen, N. Yu, Y. Zhang, and M. Fritz. GAN-Leaks: A taxonomy of membership inference attacks against GANs. In *ACM CCS*, 2020.
- [12] J. Chen, W. H. Wang, and X. Shi. Differential Privacy Protection Against Membership Inference Attack on Machine Learning for

- Genomic Data. <https://www.biorxiv.org/content/early/2020/08/04/2020.08.03.235416.full.pdf>, 2020.
- [13] E. Choi, S. Biswal, B. Malin, J. Duke, W. F. Stewart, and J. Sun. Generating multi-label discrete patient records using generative adversarial networks. In *Machine learning for Healthcare Conference*, 2017.
- [14] G. M. Clarke and L. R. Cardon. Disentangling linkage disequilibrium and linkage from dense single-nucleotide polymorphism trio data. *Genetics*, 171(4), 2005.
- [15] H. Elias. Synthetic data – All the perks without the risk? <https://techhq.com/2020/10/synthetic-data-all-the-perks-without-the-risk/>, 2020.
- [16] EMBL-EBI Training. What are genome wide association studies (GWAS)? <https://bit.ly/3wP0zA8>, 2021.
- [17] Ensembl Variation. Variant classification. <https://m.ensembl.org/info/genome/variation/prediction/classification.html>, 2020.
- [18] S. N. Evans, Y. Shvets, and M. Slatkin. Non-equilibrium theory of the allele frequency spectrum. *Theoretical Population Biology*, 71(1), 2007.
- [19] R. A. Fisher. The distribution of gene ratios for rare mutations. *Proceedings of the Royal Society of Edinburgh*, 50, 1931.
- [20] M. Fredrikson, E. Lantz, S. Jha, S. Lin, D. Page, and T. Ristenpart. Privacy in pharmacogenetics: An end-to-end case study of personalized warfarin dosing. In *USENIX Security*, 2014.
- [21] Genetics Home Reference. What is Precision Medicine? <https://medlineplus.gov/genetics/understanding/precisionmedicine/definition/>, 2020.
- [22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved Training of Wasserstein GANs. In *NeurIPS*, 2017.
- [23] M. Gymrek, A. L. McGuire, D. Golan, E. Halperin, and Y. Erlich. Identifying personal genomes by surname inference. *Science*, 339(6117), 2013.
- [24] J. Hayes, L. Melis, G. Danezis, and E. De Cristofaro. LOGAN: Membership inference attacks against generative models. *Proceedings on Privacy Enhancing Technologies*, 2019.
- [25] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *IEEE CVPR*, 2016.
- [26] B. Hilprecht, M. Härterich, and D. Bernau. Monte Carlo and Reconstruction Membership Inference Attacks against Generative Models. *PETS*, 2019.
- [27] J. L. Hodges. The significance probability of the Smirnov two-sample test. *Arkiv för Matematik*, 3(5), 1958.
- [28] K. E. Holsinger and B. S. Weir. Genetics in geographically structured populations: defining, estimating and interpreting FST. *Nature Reviews Genetics*, 10(9), 2009.
- [29] N. Homer, S. Szelinger, M. Redman, D. Duggan, W. Tembe, J. Muehling, J. V. Pearson, D. A. Stephan, S. F. Nelson, and D. W. Craig. Resolving individuals contributing trace amounts of DNA to highly complex mixtures using high-density SNP genotyping microarrays. *PLoS Genetics*, 4(8), 2008.
- [30] B. Hui, Y. Yang, H. Yuan, P. Burlina, N. Z. Gong, and Y. Cao. Practical blind membership inference attack via differential comparisons. arXiv preprint arXiv:2101.01341, 2021.
- [31] M. Humbert, E. Ayday, J.-P. Hubaux, and A. Telenti. Addressing the concerns of the lacks family: quantification of kin genomic privacy. In *ACM CCS*, 2013.
- [32] H. K. Im, E. R. Gamazon, D. L. Nicolae, and N. J. Cox. On sharing quantitative trait GWAS results in an era of multiple-omics data and the limits of genomic privacy. *The American Journal of Human Genetics*, 90(4), 2012.
- [33] A. Johnson and V. Shmatikov. Privacy-preserving data exploration in genome-wide association studies. In *ACM KDD*, 2013.
- [34] J. Jordon, J. Yoon, and M. Van Der Schaar. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *ICLR*, 2018.
- [35] N. Killoran. Generating and designing DNA. <https://github.com/co9olgu/Generating-and-designing-DNA>, 2021.
- [36] N. Killoran, L. Lee, A. Delong, D. Duvenaud, and B. Frey. Generating and designing DNA with deep generative models. In *NIPS Computational Biology Workshop*, 2017.
- [37] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Comp.*, 1(4), 1989.
- [38] K. Leino and M. Fredrikson. Stolen Memories: Leveraging Model Memorization for Calibrated White-Box Membership Inference. In *USENIX Security Symposium*. USENIX, 2020.
- [39] I. Mironov. Rényi differential privacy. In *CSF*, 2017.
- [40] A. Mordvintsev, C. Olah, and M. Tyka. Inceptionism: Going Deeper into Neural Networks. <https://ai.googleblog.com/2015/06/inceptionism-going-deeper-into-neural.html>, 2015.
- [41] M. Nasr, R. Shokri, and A. Houmansadr. Machine Learning with Membership Privacy using Adversarial Regularization. In *ACM CCS*. ACM, 2018.
- [42] National Human Genome Research Institute. International HapMap Project. <https://www.genome.gov/10001688/international-hapmap-project>, 2012.
- [43] National Human Genome Research Institute. Genome-Wide Association Studies Fact Sheet. <https://www.genome.gov/about-genomics/fact-sheets/Genome-Wide-Association-Studies-Fact-Sheet>, 2021.
- [44] M. Naveed, E. Ayday, E. W. Clayton, J. Fellay, C. A. Gunter, J.-P. Hubaux, B. A. Malin, and X. Wang. Privacy In The Genomic Era. *ACM Computing Surveys*, 48(1), 2015.
- [45] NHS England. A&E Synthetic Data. <https://data.england.nhs.uk/dataset/a-e-synthetic-data>, 2021.
- [46] NIH. Genomic Data Sharing. <https://osp.od.nih.gov/scientific-sharing/genomic-data-sharing-faqs/>, 2021.
- [47] J. Pickrell. Genetic maps for the 1000 Genomes Project variants. <https://github.com/joepickrell/1000-genomes-genetic-maps>, 2021.
- [48] F. Privé, Z. Zhu, and B. J. Vilhjalmsson. Finding hidden treasures in summary statistics from genome-wide association studies. *Nature Genetics*, 53(4):431–432, 2021.
- [49] J. L. Raisaro, F. Tramer, Z. Ji, D. Bu, Y. Zhao, K. Carey, D. Lloyd, H. Sofia, D. Baker, P. Flicek, et al. Addressing Beacon re-identification attacks: quantification and mitigation of privacy risks. *Journal of the American Medical Informatics Association*, 24(4), 2017.
- [50] A. R. Rogers and C. Huff. Linkage disequilibrium between loci with unknown phase. *Genetics*, 182, 2009.
- [51] A. Sablayrolles, M. Douze, C. Schmid, Y. Ollivier, and H. Jégou. White-box vs Black-box: Bayes Optimal Strategies for Membership Inference. In *ICML*, 2019.
- [52] P. Sadrach. How Artificial Intelligence Can Revolutionize Healthcare. <https://builtin.com/healthcare-technology/how-ai-revolutionize-healthcare>, 2020.
- [53] A. Salem, Y. Zhang, M. Humbert, P. Berrang, M. Fritz, and M. Backes. ML-Leaks: Model and Data Independent Membership Inference Attacks and Defenses on Machine Learning Models. In *NDSS*, 2019.
- [54] S. S. Samani, Z. Huang, E. Ayday, M. Elliot, J. Fellay, J.-P. Hubaux, and Z. Kutalik. Quantifying genomic privacy via inference attack with high-order SNV correlations. In *IEEE Security and Privacy Workshops*, 2015.
- [55] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning models. In *IEEE S&P*, 2017.
- [56] S. S. Shringarpure and C. D. Bustamante. Privacy risks from genomic data-sharing beacons. *The American Journal of Human Genetics*, 97(5), 2015.
- [57] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. arXiv:1312.6034, 2013.

- [58] P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. Technical report, Colorado Univ at Boulder Dept of Computer Science, 1986.
- [59] T. Stadler, B. Oprisanu, and C. Troncoso. Synthetic Data – A Privacy Mirage. arXiv:2011.07018, 2020.
- [60] A. Telenti, E. Ayday, and J. P. Hubaux. On genomics, kin, and privacy. *F1000Research*, 3, 2014.
- [61] A. Torfi and E. A. Fox. CorGAN: Correlation-capturing convolutional generative adversarial networks for generating synthetic healthcare records. arXiv:2001.09346, 2020.
- [62] A. Torfi, E. A. Fox, and C. K. Reddy. Differentially Private Synthetic Medical Data Generation using Convolutional GANs. arXiv:2012.11774, 2020.
- [63] A. Tucker, Z. Wang, Y. Rotalinti, and P. Myles. Generating high-fidelity synthetic patient data for assessing machine learning healthcare software. *NPJ Digital Medicine*, 3(1), 2020.
- [64] UK Government. Press Release – New synthetic datasets to assist COVID-19 and cardiovascular research. <https://bit.ly/3eFCFB4>, 2020.
- [65] N. Von Thenen, E. Ayday, and A. E. Cicek. Re-identification of individuals in genomic data-sharing beacons via allele inference. *Bioinformatics*, 35(3), 2019.
- [66] R. Wang, Y. F. Li, X. Wang, H. Tang, and X. Zhou. Learning your identity and disease from research papers: information leaks in genome wide association study. In *ACM CCS*, 2009.
- [67] S. Wang, X. Jiang, S. Singh, R. Marmor, L. Bonomi, D. Fox, M. Dow, and L. Ohno-Machado. Genome Privacy: Challenges, Technical Approaches to Mitigate Risk, and Ethical Considerations in the United States. *Annals of the New York Academy of Sciences*, 1387(1), 2017.
- [68] Z. Wang, P. Myles, and A. Tucker. Generating and Evaluating Synthetic UK Primary Care Data: Preserving Data Utility & Patient Privacy. In *IEEE International Symposium on Computer-Based Medical Systems*, 2019.
- [69] L. Xie, K. Lin, S. Wang, F. Wang, and J. Zhou. Differentially private generative adversarial network. arXiv preprint arXiv:1802.06739, 2018.
- [70] A. Yale, S. Dash, R. Dutta, I. Guyon, A. Pavao, and K. Bennett. Privacy preserving synthetic health data. 2019.
- [71] B. Yelmen, A. Decelle, L. Ongaro, D. Marnetto, F. Montinaro, C. Furtlehner, L. Pagani, and F. Jay. Creating Artificial Human Genomes Using Generative Models. <https://www.biorxiv.org/content/10.1101/769091v2>, 2019.
- [72] J. Yoon, L. Drumright, and M. Schaar. Anonymization Through Data Synthesis Using Generative Adversarial Networks (ADS-GAN). *IEEE Journal of Biomedical and Health Informatics*, PP, 2020.
- [73] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding Neural Networks Through Deep Visualization. arXiv:1506.06579, 2015.
- [74] X. Zhou, B. Peng, Y. F. Li, Y. Chen, H. Tang, and X. Wang. To release or not to release: evaluating information leaks in aggregate human-genome data. In *ESORICS*, 2011.

APPENDIX A WGAN SETTINGS

We experimented with training the WGAN models for various iterations (up to 10,000) and observed that the utility metrics stabilized after around the 100th iteration. Specifically, it can be seen that for MAF in the CEU dataset in Fig. 11a training the WGAN model for more than 100 iterations does not result in better metrics. The discriminator training and validation losses converge to the same values around 100th iteration for all datasets as can be seen in Fig. 11b. Therefore, in our evaluation we train the WGAN model for 100

iterations, as opposed to the 100,000 proposed in the original implementations. Overall, we believe this could be due to the characteristics of the datasets as, for example, the sequences in our evaluation all have a length of 1,000. Killoran et al. [36] experiment with DNA sequences of length around 50.

APPENDIX B ADDITIONAL FIGURES FROM EXPERIMENTS

In this section, we include some additional experiments complementing our utility evaluation presented in Section IV. More precisely, Figure 12 presents the alternate allele correlation (AAC) between the real and the synthetic datasets, while in Figure 13, we plot the scaled folded Site Frequency Spectrum (SFS), which is the distribution of counts of minor alleles in a sample calculated over all segregating sites. Finally, Figure 14 shows the Euclidean Genetic Distance (EGD) between the samples in each dataset.

APPENDIX C PRINCIPAL COMPONENT ANALYSIS (PCA)

Here, we study the difference between synthetic and real data vis-à-vis a principal component analysis on the corresponding samples. We extract the first two principal components and project the real and synthetic datasets on these two components to show how the synthetic samples are distributed compared to the real data.

Fig. 15 presents this 2D visualization. Recomb has a close distribution to the real data for both HapMap populations, which, according to [54], is because the genetic recombination model considers all the correlations between SNPs and builds a higher-order model. For the 1000 Genomes, as for the other utility metrics studied in Section IV, the GAN and RecGAN models perform quite poorly, generating samples with a different distribution than the real samples. Therefore, in Figure 15d, we exclude them in order to take a closer look at the non-GAN-based models. Here, we can better observe that the RBM-generated samples have the closest distribution to the real data. In contrast to the HapMap populations, the samples from Recomb are all centered around 0 and fail to simulate the distribution given by the real data, and similar results are in the case of samples generated by Rec-RBM.

APPENDIX D PRIVACY GAIN VS. F_{ST}

To provide a quick visualization of the *trade-offs* between privacy and utility, we also plot the F_{ST} (see Section IV-B) against the PG in Figure 16 for the two smaller HapMap Populations. For this set of experiments, we randomly sample 10 targets from each dataset. We train a Logistic Regression attack classifier using 5 shadow models, using 100 synthetic datasets for each of them. We then compute the PG and F_{ST} on 100 synthetic datasets, with a split of 50 sets generated from a training set including the target and 50 sets generated without. We report the PG and F_{ST} as the average across all targets and synthetic datasets.




Fig. 11: WGANs experiments for 100/1,000/10,000 iterations.

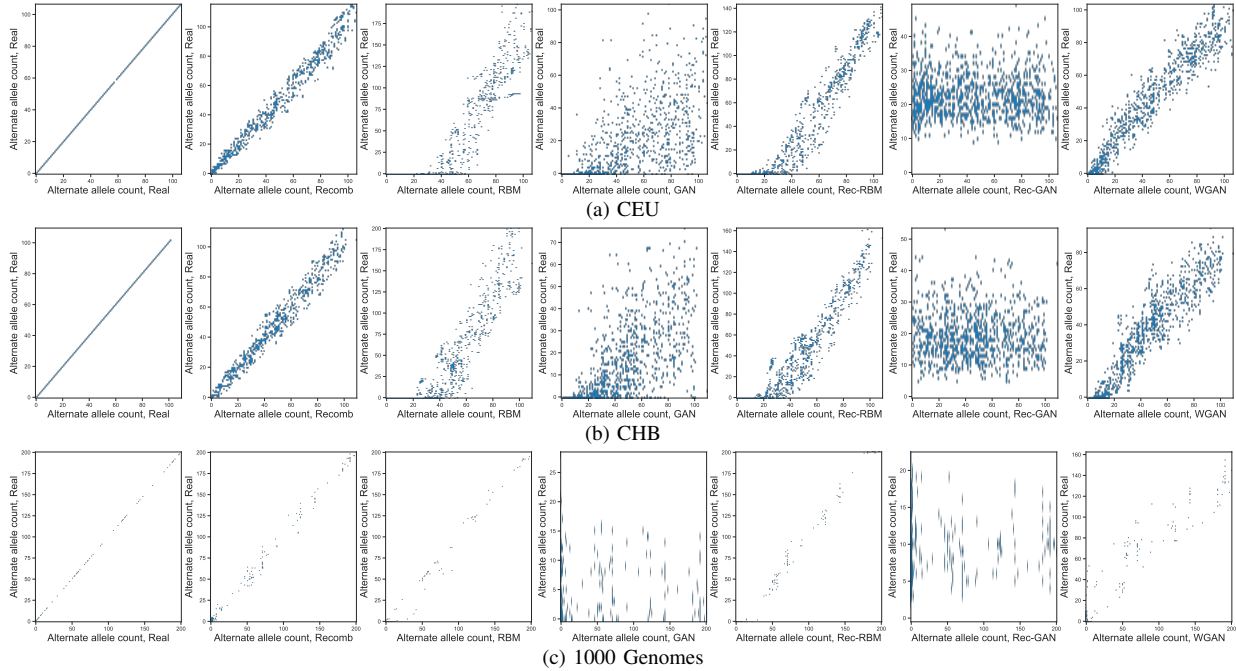


Fig. 12: Alternate allele correlation for the CEU population, the CHB population, and the 1000 Genomes dataset.




Fig. 13: Frequency spectrum analysis for the CEU population, the CHB population, and the 1000 Genomes dataset.

Recall that the lower the F_{ST} value the closer the synthetic samples are to the real data, and PG should be above the random guess baseline of 0.25 for the synthetic data to offer better privacy protection than releasing the real. Put simply, the ideal “place” in the plots in Figure 16 is the top left.

For the CEU population (Figure 16a), the synthetic Recomb samples have the best overall utility; however, the average PG is below 0.25. In line with our utility evaluation, we find that the hybrid model Rec-RBM model generates samples with a F_{ST} value closer to the real data than the RBM model. Interestingly, while for Histogram, Correlation, and Ensemble feature sets, the PG for the Rec-RBM is close to 0.25, there is

a significant privacy loss for Naive. In contrast, for the samples generated by the RBM, the Naive yields the highest PG across all feature sets. The samples generated by the WGAN also have a F_{ST} value close to 0; however, for the Histogram, Correlation and Ensemble, the PG values are below 0.25.

For CHB (Figure 16b), the utility results are similar to the CEU population. However, all targets apart from the ones generated by RecRBM have a lower PG than the CEU population. This reiterates the fact that one cannot reliably use synthetic data as a good privacy mechanism, as the value of PG is unpredictable and can fluctuate based on target and training set combinations chosen.

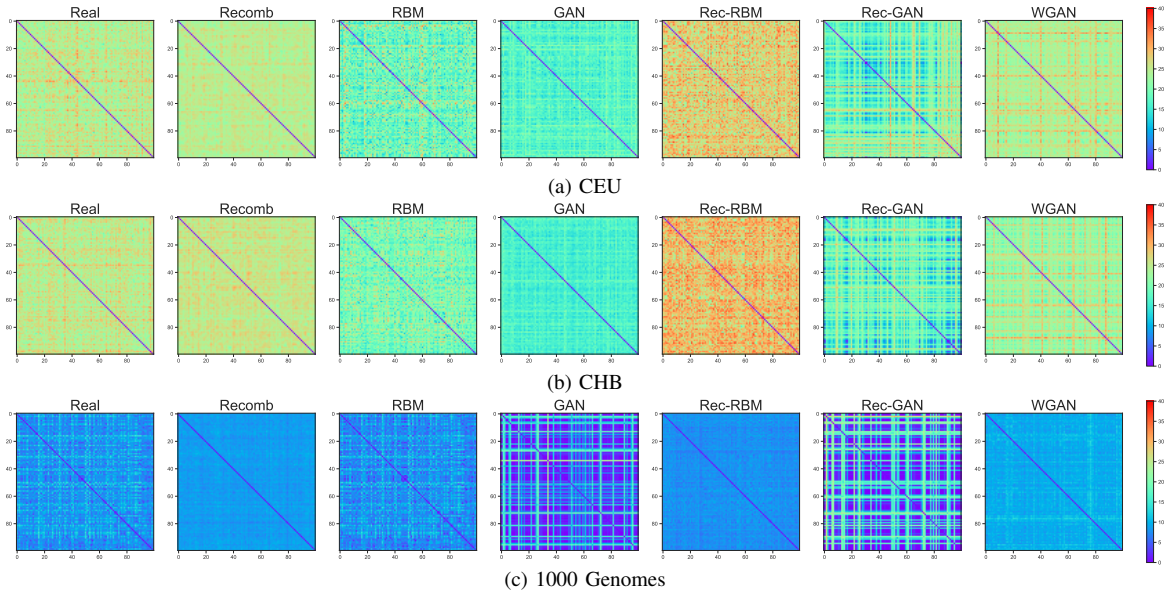


Fig. 14: Pairwise Euclidean Genetic Distance (EGD) between individuals.

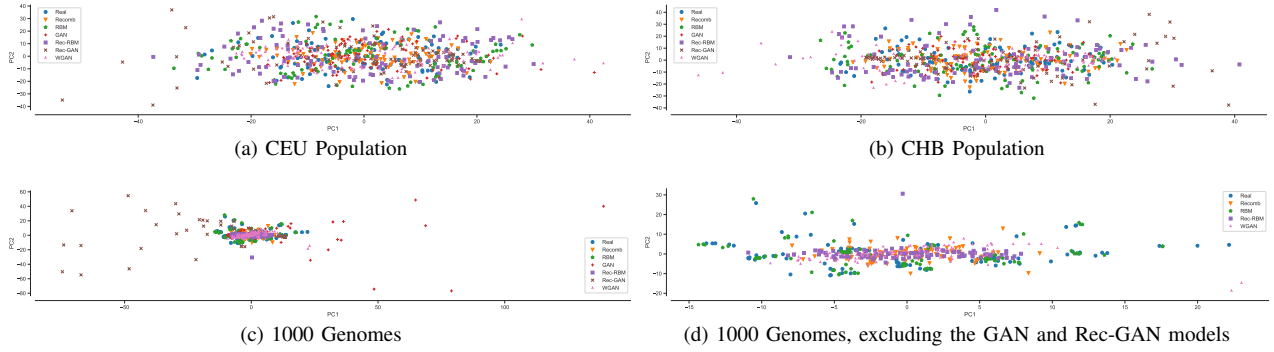


Fig. 15: 2D Principal Component Analysis (PCA) visualization of the real and synthetic sequences.

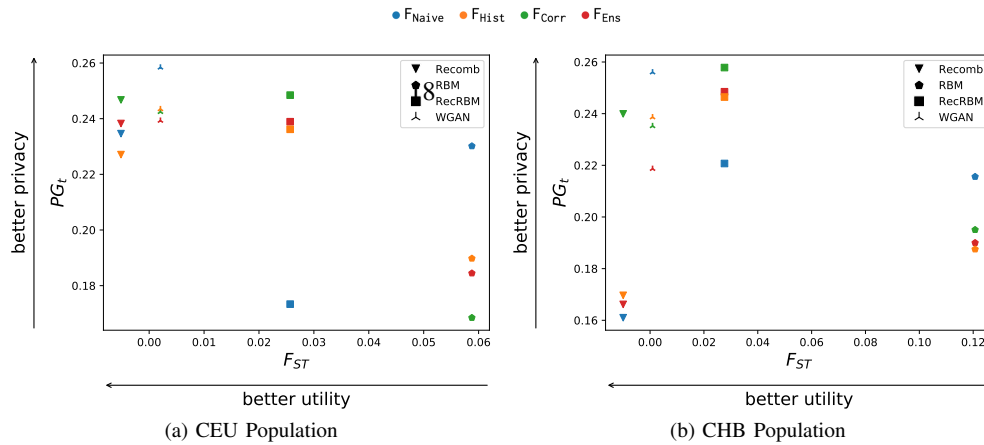


Fig. 16: Fixation index values vs Privacy gain for the HapMap populations.