

Importance Weighted Generative Networks

Maurice Diesendruck¹ [✉], Ethan R. Elenberg², Rajat Sen³, Guy W. Cole¹,
Sanjay Shakkottai¹, and Sinead A. Williamson¹⁴

¹ The University of Texas at Austin, USA*

² ASAPP, Inc.**

³ Amazon, Inc.

⁴ CognitiveScale

{momod, rajat.sen, guywcole}@utexas.edu, elenberg@asapp.com,
shakkott@austin.utexas.edu, sinead.williamson@mcombs.utexas.edu

Abstract. While deep generative networks can simulate from complex data distributions, their utility can be hindered by limitations on the data available for training. Specifically, the training data distribution may differ from the target sampling distribution due to sample selection bias, or because the training data comes from a different but related distribution. We present methods to accommodate this difference via *importance weighting*, which allow us to estimate a loss function with respect to a target distribution *even if we cannot access that distribution directly*. These estimators, which differentially weight the contribution of data to the loss function, offer theoretical guarantees that heuristic approaches lack, while giving impressive empirical performance in a variety of settings.

Keywords: importance weights · generative networks · bias correction

1 Introduction

Deep generative models have important applications in many fields: we can automatically generate illustrations for text [48]; simulate video streams [46] or molecular fingerprints [27]; and create privacy-preserving versions of medical time-series data [14]. Such models use a neural network to parametrize a function $G(Z)$, which maps random noise Z to a target probability distribution \mathbb{P} . This is achieved by minimizing a loss function between simulations and data, which is equivalent to learning a distribution over simulations that is indistinguishable from \mathbb{P} under an appropriate two-sample test. In this paper we focus on Generative Adversarial Networks (GANs) [17, 2, 4, 30], which incorporate an adversarially learned neural network in the loss function; however the results are also applicable to non-adversarial networks [13, 31].

An interesting challenge arises when we do not have direct access to i.i.d. samples from \mathbb{P} . This could arise either because observations are obtained via

* R. Sen and S. Shakkottai were partially supported by ARO grant W911NF-17-1-0359

** Work done primarily while at UT Austin.

a biased sampling mechanism [6, 49], or in a transfer learning setting where our target distribution differs from our training distribution. As an example of the former, a dataset of faces generated as part of a university project may contain disproportionately many young adult faces relative to the population. As an example of the latter, a Canadian hospital system might want to customize simulations to its population while still leveraging a training set of patients from the United States (which has a different statistical distribution of medical records). In both cases, and more generally, we want to generate data from a target distribution \mathbb{P} but only have access to representative samples from a *modified* distribution $M\mathbb{P}$. We give a pictorial example of this setting in Figure 1.

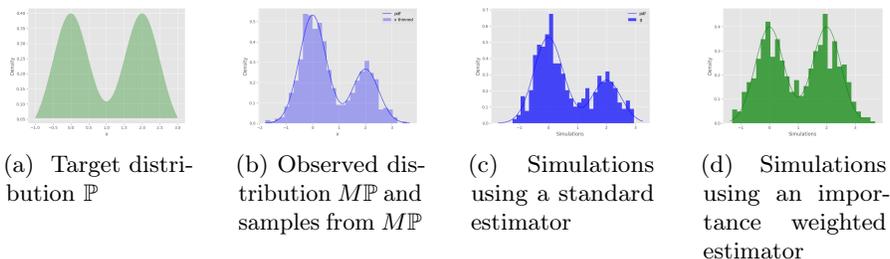


Fig. 1: If our target distribution \mathbb{P} differs from our observed distribution $M\mathbb{P}$, using the standard estimator will replicate $M\mathbb{P}$, while an importance weighted estimator can replicate the target \mathbb{P} .

In some cases, we can approach this problem using existing methods. For example, if we can reduce our problem to a conditional data-generating mechanism, we can employ Conditional Generative Adversarial Networks (C-GANs) or related models [33, 36], which enable conditional sampling given one or more latent variables. However, this requires that M can be described on a low-dimensional space, and that we can sample from our target distribution over that latent space. Further, C-GANs rely on a large, labeled dataset of training samples with diversity over the conditioning variable (within each batch), which becomes a challenge when conditioning on a high-dimensional variable. For example, if we wish to modify a distribution over faces with respect to age, gender and hair length, there may be few exemplars of 80-year-old men with long hair with which to learn the corresponding conditional distribution.

In this paper, we propose an alternate approach based on importance sampling [37]. Our method modifies an existing GAN by rescaling the observed data distribution $M\mathbb{P}$ during training, or equivalently by reweighting the contribution of each data point to the loss function. When training a GAN with samples from $M\mathbb{P}$, the standard estimator equally weights the contribution of each point, yielding an estimator of the loss with respect to $M\mathbb{P}$ and corresponding simulations, as shown in Fig. 1b and Fig. 1c. This is not ideal.

In order to yield the desired estimator with respect to our target distribution \mathbb{P} , we modify the estimator by reweighting the loss function evaluation for each sample. When the Radon-Nikodym derivative between the target and observed distributions (aka the modifier function M) is known, we inversely scale each evaluation by that derivative, yielding the finite-sample importance sampling transform on the estimate, which we call the *importance weighted* estimator. This reweighting asymptotically ensures that discrimination, and the corresponding GAN update, occurs with respect to \mathbb{P} instead of $M\mathbb{P}$, as shown in Fig. 1a and Fig. 1d.

This approach has multiple advantages and extensions. First, if M is known, we can estimate importance weighted losses using robust estimators like the median-of-means estimator, which is crucial for controlling variance in settings where the modifier function M has a large dynamic range. Second, even when the modifier function is only known up to a scaling factor, we can construct an alternative estimator using self-normalized sampling [41, 37] to use this partial information, while still maintaining asymptotic correctness. Finally and importantly, for the common case of an unknown modifier function, we demonstrate techniques for estimating it from partially labeled data.

Our contributions are as follows: 1) We provide a novel application of traditional importance weighting to deep generative models. This has connections to many types of GAN loss functions through the theory of U-statistics. 2) We propose several variants of our importance weighting framework for different practical scenarios. When dealing with particularly difficult functions M , we propose to use robust median-of-means estimation and show that it has similar theoretical guarantees under weaker assumptions, *i.e.* bounded second moment. When M is not known fully (only up to a scaling factor), we propose a self-normalized estimator. 3) We conduct an extensive experimental evaluation of the proposed methods on both synthetic and real-world datasets. This includes estimating M when less than 4% of the data is labeled with user-provided exemplars.

1.1 Related Work

Our method aims to generate samples from a distribution \mathbb{P} , given access to samples from $M\mathbb{P}$. While to the best of our knowledge this has not been explicitly addressed in the GAN literature, several approaches have related goals.

Domain adaptation: Our formulation is related to but distinct from the problem of Domain Adaptation (DA). The challenge of DA is, “If I train on one distribution and test on another, how do I maximize performance on test data?” Critically, the test data is available and extensively used. Instead, our method solves the problem, “Given only a training data distribution, how do I generate from arbitrarily modified versions of it?” The former uses two datasets – one source and one target – while the latter uses one dataset and accommodates an arbitrary number of targets. The methodologies are inherently different because the information available is different.

Typical approaches to DA involve finding domain-invariant feature representations for both source and target data. Blitzer, Pereira, Ben-David, and Daume

[5, 3, 9] write extensively on techniques involving feature correlation and mutual information within classification settings. Pan, Huang, and Gong [38, 39, 24, 16] propose methods with similar goals that find kernel representations under which source and target distributions are close. The work of [24] and [43] address covariate shift using kernel-based and importance-weighted techniques, but still inhabit a different setting from our problem since they perform estimation on specific source and target datasets.

Recently, the term DA has been used in the context of adversarially-trained image-to-image translation and downstream transfer learning tasks [25, 45, 50, 22]. Typically the goal is to produce representations of the same image in both source and target domains. Such problems begin with datasets from both domains, whereas our setting presents only one source dataset and seeks to generate samples from a hypothetical, user-described target domain.

Inverse probability weighting: Inverse probability weighting (IPW), originally proposed by [23] and still in wide use in the field of survey statistics [32], can be seen as a special case of importance sampling. IPW is a weighting scheme used to correct for biased treatment assignment methods in survey sampling. In such settings, the target distribution is known and the sampling distribution is typically finite and discrete, and can easily be estimated from data.

Conditional GANs: Conditional GANs (C-GANs) are an extension of GANs that aim to simulate from a conditional distribution, given some covariate. In the case where our modifier function M can be represented in terms of a low-dimensional covariate space, and if we can generate samples from the marginal distribution of $M\mathbb{P}$ on that space, then we can, in theory, use a C-GAN to generate samples from \mathbb{P} , by conditioning on the sampled covariates. This strategy suffers from two limitations. First, it assumes we can express M in terms of a sampleable distribution on a low-dimensional covariate space. For settings where M varies across many data dimensions or across a high-dimensional latent embedding, this ability to sample becomes untenable. Second, learning a family of conditional distributions is typically more difficult than learning a single joint distribution. As we show in our experiments, C-GANs often fail if there are too few real exemplars for a given covariate setting.

Related to C-GANs, [8] proposes conditional generation and a classifier for assigning samples to specific discriminators. While not mentioned, such a structure could feasibly be used to preferentially sample certain modes, if a correspondence between latent features and numbered modes were known.

Weighted loss: In the context of domain adaptation for data with discrete class labels, the strategy of reweighting the Maximum Mean Discrepancy (MMD) [18] based on class probabilities has been proposed by [47]. This approach, however, differs from ours in several ways: It is limited to class imbalance problems, as opposed to changes in continuous-valued latent features; it requires access to the non-conforming target dataset; it provides no theoretical guarantees about the weighted estimator; and it is not in the generative model setting.

Other uses of importance weights in GANs: The language and use of importance weights is not unique to this application, and has been used for other purposes within the GAN context. In [19], for example, importance weights are used to provide policy gradients for GANs in a discrete-data setting. Our application is different in that our target distribution is not that of our data, as it is in [19]. Instead we view our data as having been modified, and use importance weights to simulate closer to the hypothetical and desired *unmodified* distribution.

2 Problem Formulation and Technical Approach

The problem: Given training samples from a distribution $M\mathbb{P}$, our goal is to construct (train) a generator function $G(\cdot)$ that produces i.i.d. samples from a distribution \mathbb{P} .

To train $G(\cdot)$, we follow the methodology of a Generative Adversarial Network (GAN) [17]. In brief, a GAN consists of a pair of interacting and evolving neural networks – a generator neural network with outputs that approximate the desired distribution, and a discriminator neural network that distinguishes between increasingly realistic outputs from the generator and samples from a training dataset.

The loss function is a critical feature of the GAN discriminator, and evaluates the closeness between the samples of the generator and those of the training data. Designing good loss functions remains an active area of research [2, 30]. One popular loss function is the Maximum Mean Discrepancy (MMD) [18], a distributional distance that is zero if and only if the two distributions are the same. As such, MMD can be used to prevent mode collapse [42, 7] during training.

Our approach: We are able to train a GAN to generate samples from \mathbb{P} using a simple reweighting modification to the MMD loss function. Reweighting forces the loss function to apply greater penalties in areas of the support where the target and observed distributions differ most.

Below, we formally describe the MMD loss function, and describe its importance weighted variants.

Remark 1 (Extension to other losses). While this paper focuses on the MMD loss, we note that the above estimators can be extended to any estimator that can be expressed as the expectation of some function with respect to one or more distributions. This class includes losses such as squared mean difference between two distributions, cross entropy loss, and autoencoder losses [44, 20, 34]. Such losses can be estimated from data using a combination of U-statistics, V-statistics and sample averages. Each of these statistics can be reweighted, in a manner analogous to the treatment described above. We provide more comprehensive details in Table 1, and in Section 3.1 we evaluate all three importance weighting techniques as applied to the standard cross entropy GAN objective.

2.1 Maximum Mean Discrepancy between Two Distributions

The MMD projects two distributions \mathbb{P} and \mathbb{Q} into a reproducing kernel Hilbert space (RKHS) \mathcal{H} , and evaluates the maximum mean distance between the two projections, *i.e.*

$$\text{MMD}(\mathbb{P}, \mathbb{Q}) := \sup_{f \in \mathcal{H}} (\mathbf{E}_{X \sim \mathbb{P}}[f(X)] - \mathbf{E}_{Y \sim \mathbb{Q}}[f(Y)]).$$

If we specify the *kernel mean embedding* $\mu_{\mathbb{P}}$ of \mathbb{P} as $\mu_{\mathbb{P}} = \int k(x, \cdot) d\mathbb{P}(x)$, where $k(\cdot, \cdot)$ is the characteristic kernel defining the RKHS, then we can write the square of this distance as

$$\begin{aligned} \text{MMD}^2(\mathbb{P}, \mathbb{Q}) &= \|\mu_{\mathbb{P}} - \mu_{\mathbb{Q}}\|_{\mathcal{H}}^2 \\ &= \mathbb{E}_{X, X' \sim \mathbb{P}}[k(X, X')] + \mathbb{E}_{Y, Y' \sim \mathbb{Q}}[k(Y, Y')] \\ &\quad - 2\mathbb{E}_{X \sim \mathbb{P}, Y \sim \mathbb{Q}}[k(X, Y)]. \end{aligned} \tag{1}$$

In order to be a useful loss function for training a neural network, we must be able to estimate $\text{MMD}^2(\mathbb{P}, \mathbb{Q})$ from data, and compute gradients of this estimate with respect to the network parameters. Let $\{x_i\}_n$ be a sample $\{X_1 = x_1, \dots, X_n = x_n\} : X_i \sim \mathbb{P}$, and $\{y_i\}_m$ be a sample $\{Y_1 = y_1, \dots, Y_m = y_m\} : Y_i \sim \mathbb{Q}$. We can construct an unbiased estimator $\widehat{\text{MMD}}^2(\mathbb{P}, \mathbb{Q})$ of $\text{MMD}^2(\mathbb{P}, \mathbb{Q})$ [18] using these samples as

$$\begin{aligned} \widehat{\text{MMD}}^2(\mathbb{P}, \mathbb{Q}) &= \frac{1}{n(n-1)} \sum_{i \neq j}^n k(x_i, x_j) \\ &\quad + \frac{1}{m(m-1)} \sum_{i \neq j}^m k(y_i, y_j) \\ &\quad - \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m k(x_i, y_j). \end{aligned} \tag{2}$$

2.2 Importance Weighted Estimator for Known M

We begin with the case where M (which relates the distribution of the samples and the desired distribution; formally the Radon-Nikodym derivative) is known. Here, the reweighting of our loss function can be framed as an *importance sampling* problem: we want to estimate $\text{MMD}^2(\mathbb{P}, \mathbb{Q})$, which is in terms of the target distribution \mathbb{P} and the distribution \mathbb{Q} implied by our generator, but we have samples from the modified $M\mathbb{P}$. Importance sampling [37] provides a method for constructing an estimator for the expectation of a function $\phi(X)$ with respect to a distribution \mathbb{P} , by taking an appropriately weighted sum of evaluations of ϕ at values sampled from a different distribution. We can therefore modify the estimator in (2) by weighting each term in the estimator involving data point x_i using the likelihood ratio $\mathbb{P}(x_i)/M(x_i)\mathbb{P}(x_i) = 1/M(x_i)$, yielding an unbiased

importance weighted estimator that takes the form

$$\begin{aligned} \widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q}) &= \frac{1}{n(n-1)} \sum_{i \neq j}^n \frac{k(x_i, x_j)}{M(x_i)M(x_j)} \\ &+ \frac{1}{m(m-1)} \sum_{i \neq j}^m k(y_i, y_j) \\ &- \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m \frac{k(x_i, y_j)}{M(x_i)}. \end{aligned} \quad (3)$$

While importance weighting using the likelihood ratio yields an unbiased estimator (3), the estimator may not concentrate well because the weights $\{1/M(x_i)\}_n$ may be large or even unbounded. We now provide a concentration bound for the estimator in (3) for the case where weights $\{1/M(x_i)\}_n$ are upper-bounded by some maximum value.

Theorem 1. *Let $\widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})$ be the unbiased, importance weighted estimator for $\text{MMD}^2(\mathbb{P}, \mathbb{Q})$ defined in (3), given m i.i.d samples from $M\mathbb{P}$ and \mathbb{Q} , and maximum kernel value K . Further assume that $1 \leq 1/M(x) \leq W$ for all $x \in \mathcal{X}$. Then*

$$\begin{aligned} \mathbb{P}\left(\widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}^2(\mathbb{P}, \mathbb{Q}) > t\right) &\leq C, \\ \text{where } C &= \exp((-2t^2 m_2)/(K^2(W+1)^4)) \\ m_2 &:= \lfloor m/2 \rfloor \end{aligned}$$

These guarantees are based on estimator guarantees in [18], which in turn build on classical results by Hoeffding [21, 20]. We defer the proof of this theorem to the extended version of this work [12].

2.3 Robust Importance Weighted Estimator for Known M

Theorem 1 is sufficient to guarantee good concentration of our importance weighted estimator only when $1/M(x)$ is uniformly bounded by some constant W , which is not too large. Many class imbalance problems fall into this setting. However, $1/M(x)$ may be unbounded in practice. Therefore, we now introduce a different estimator, which enjoys good concentration even when only $\mathbb{E}_{X \sim M\mathbb{P}}[1/M(X)^2]$ is bounded, while $1/M(x)$ may be unbounded for many values of x .

The estimator is based on the classical idea of median of means [35, 26, 1, 29]⁵. Given m samples from $M\mathbb{P}$ and \mathbb{Q} , we divide these samples uniformly at random into k equal sized groups, indexed $\{(1), \dots, (k)\}$. Let $\widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})^{(i)}$ be the value obtained when the estimator in (3) is applied on the i -th group of samples. Then our median of means based estimator is given by

$$\widehat{\text{MMD}}_{MIW}^2(\mathbb{P}, \mathbb{Q}) = \text{median} \left\{ \widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})^{(1)}, \dots, \widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})^{(k)} \right\}. \quad (4)$$

⁵ [29] appeared concurrently and contains a different approach for the unweighted estimator. Comparisons are left for future work.

Theorem 2. Let $\widehat{\text{MMD}}_{MIW}^2(\mathbb{P}, \mathbb{Q})$ be the asymptotically unbiased median of means estimator defined in (4) using $k = mt^2/(8K^2\sigma^2)$ groups. Further assume that $n = m$ and let $W_2 = \mathbb{E}_{X \sim M\mathbb{P}}[1/M(X)^2]$ be bounded. Then

$$\mathbb{P}\left(\left|\widehat{\text{MMD}}_{MIW}^2(\mathbb{P}, \mathbb{Q}) - \text{MMD}^2(\mathbb{P}, \mathbb{Q})\right| > t\right) \leq C,$$

$$\text{where } C = \exp((-mt^2)/(64K^2\sigma^2))$$

$$\sigma^2 = O(W_2^2 + \text{MMD}^4(\mathbb{P}, \mathbb{Q})).$$

We defer the proof of this theorem to the extended version of this work [12]. Note that the confidence bound in Theorem 2 depends on the term W_2 being bounded. This is the second moment of $1/M(X)$ where $X \sim M\mathbb{P}$. Thus, unlike in Theorem 1, this confidence bound may still hold even if $1/M(x)$ is *not uniformly bounded*. When $1/M(X)$ is heavy-tailed with finite variance, *e.g.* Pareto ($\alpha > 2$) or log-normal, then Theorem 2 is valid but Theorem 1 does not apply.

In addition to increased robustness, the median of means MMD estimator is more computationally efficient: since calculating $\widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})$ scales quadratically in the batch size, using the median of means estimator introduces a speed-up that is linear in the number of groups.

2.4 Self-normalized Importance Weights for Unknown M

To specify M , we must know the forms of our target and observed distributions along any marginals where the two differ. In some settings this is available: consider for example a class rebalancing setting where we have class labels and a desired class ratio, and can estimate the observed class ratio from data. This, however, may be infeasible if M is continuous and/or varies over several dimensions, particularly if data are arriving in a streaming manner. In such a setting it may be easier to specify a *thinning function* T that is proportional to M , *i.e.* $M\mathbb{P} = \frac{T\mathbb{P}}{Z}$ for some unknown Z , than to estimate M directly. This is because T can be directly obtained from an estimate of how much a given location is underestimated, without any knowledge of the underlying distribution.

This setting—where the $1/M$ weights used in Section 2.2 are only known up to a normalizing constant—motivates the use of a *self-normalized* importance sampling scheme, where the weights $w_i \propto \frac{\mathbb{P}(x_i)}{M(x_i)\mathbb{P}(x_i)} = \frac{Z}{T(x_i)}$ are normalized to sum to one [41, 37]. For example, by letting $w_i = \frac{1}{T(x_i)}$, the resulting self-normalized estimator for the squared MMD takes the form

$$\begin{aligned} \widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q}) &= \frac{\sum_{i \neq j}^n w_i w_j k(x_i, x_j)}{\sum_{i \neq j}^n w_i w_j} \\ &+ \sum_{i \neq j}^m \frac{k(y_i, y_j)}{m(m-1)} \\ &- 2 \frac{\sum_{i=1}^n \sum_{j=1}^m w_i k(x_i, y_j)}{m \sum_{i=1}^n w_i}. \end{aligned} \tag{5}$$

While use of self-normalized weights means this self-normalized estimator is biased, it is asymptotically unbiased, with the bias decreasing at a rate of $1/n$

[28]. Although we have motivated self-normalized weights out of necessity, in practice they often trade off bias for reduced variance, making them preferable in some practical applications [37].

More generally, in addition to not knowing the normalizing constant Z , we might also not know the thinning function T . For example, T might vary along some latent dimension—perhaps we want to have more images of people fitting a certain aesthetic, rather than corresponding to a certain observed covariate or class. In this setting, a practitioner may be able to estimate $T(x_i)$, or equivalently w_i , for a small number of training points x_i , by considering how much those training points are under- or over-represented. Continuous-valued latent preferences can therefore be expressed by applying higher weights to points deemed more appealing. From here, we can use function estimation techniques, such as neural network regression, to estimate T from a small number of labeled data points.

2.5 Approximate Importance Weighting by Data Duplication

In the importance weighting scheme described above, each data point is assigned a weight $1/M(x_i)$. We can obtain an approximation to this method by including $\lceil 1/M(x_i) \rceil$ duplicates of data point x_i in our training set. We refer to this approach as *importance duplication*. Importance duplication obviously introduces discretization errors, and if our estimator is a U-statistic it will introduce bias (e.g. in the MMD example, if two or more copies of the data point x_i appear in a minibatch, then $k(x_i, x_i)$ will appear in the first term of (2)). However, as we show in the experimental setting, even though this approach lacks theoretical guarantees it provides generally good performance.

Data duplication can be done as a pre-processing step, making it an appealing choice if we have an existing GAN implementation that we do not wish to modify. In other settings, it is less appealing, since duplicating data adds an additional step and increases the amount of data the algorithm must process. Further, if we were to use this approximation in a setting where M is unknown, we would have to perform this data duplication on the fly as our estimate of M changes.

3 Evaluation

In this section, we show that our estimators, in conjunction with an appropriate generator network, allow us to generate simulations that are close in distribution to our target distribution, even when we only have access to this distribution via a biased sampling mechanism. Further, we show that our method performs comparably with, or better than, conditional GAN baselines.

Most of our weighted GAN models are based on the MMD-GAN of [30], replacing the original MMD loss with either our importance weighted loss $\widehat{\text{MMD}}_{IW}^2(\mathbb{P}, \mathbb{Q})$ (IW-MMD), our median of means loss $\widehat{\text{MMD}}_{MIW}^2(\mathbb{P}, \mathbb{Q})$ (MIW-MMD), or our self-normalized loss $\widehat{\text{MMD}}_{SNIW}^2(\mathbb{P}, \mathbb{Q})$ (SNIW-MMD). We also use a standard MMD loss with an importance duplicated dataset (ID-MMD). Other losses used

Table 1: Constructing importance weighted estimators for losses involving U-statistics, V-statistics and sample averages. Here, \mathcal{U} is the set of all r -tuples of numbers from 1 to n without repeats, and \mathcal{V} is the set of r -tuples allowing repeats. Below, let $X_{u,*} = X_{u_1}, \dots, X_{u_r}$.

	$\widehat{D}(\mathbb{P}, \mathbb{Q})$	$\widehat{D}_{IW}(\mathbb{P}, \mathbb{Q})$	$\widehat{D}_{SNIW}(\mathbb{P}, \mathbb{Q})$
U-statistic	$\frac{1}{nP_r} \sum_{u \in \mathcal{U}} g(X_{u,*})$	$\frac{1}{nP_r} \sum_{u \in \mathcal{U}} \frac{g(X_{u,*})}{M(X_{u_1}) \cdots M(X_{u_r})}$	$\frac{\sum_{u \in \mathcal{U}} w_{u_1} \cdots w_{u_r} g(X_{u,*})}{\sum_{u \in \mathcal{U}} w_{u_1} \cdots w_{u_r}}$
V-statistic	$\frac{1}{n^r} \sum_{v \in \mathcal{V}} g(X_{v,*})$	$\frac{1}{n^r} \sum_{v \in \mathcal{V}} \frac{g(X_{v,*})}{M(X_{v_1}) \cdots M(X_{v_r})}$	$\frac{\sum_{v \in \mathcal{V}} w_{v_1} \cdots w_{v_r} g(X_{v,*})}{\sum_{v_r=1}^n w_{v_1} \cdots w_{v_r}}$
Average	$\frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m f(X_i, Y_j)$	$\frac{1}{nm} \sum_{i=1}^n \sum_{j=1}^m \frac{f(X_i, Y_j)}{M(X_i)}$	$\frac{\sum_{i=1}^n w_i \sum_{j=1}^m f(X_i, Y_j)}{m \sum_{i=1}^n w_i}$

in [30] are also appropriately weighted, following the form in Table 1. In the synthetic data examples of Section 3.1, the kernel is a fixed radial basis function, while in all other sections it is adversarially trained using a discriminator network as in [30].

To demonstrate that our method is applicable to other losses, in Section 3.1 we also create models that use the standard cross entropy GAN loss, replacing this loss with either an importance weighted estimator (IW-CE), a median of means estimator (MIW-CE) or a self-normalized estimator (SNIW-CE). We also combine a standard cross entropy loss with an importance duplicated dataset (ID-CE). These models used a two-layer feedforward neural network with ten nodes per layer.

Where appropriate, we compare against a conditional GAN (C-GAN). If M is known exactly and expressible in terms of a lower-dimensional covariate space, a conditional GAN (C-GAN) offers an alternative method to sample from \mathbb{P} : learn the appropriate conditional distributions given each covariate value, sample new covariate values, and then sample from \mathbb{P} using each conditional distribution.

3.1 Can GANs with Importance Weighted Estimators Recover Target Distributions, Given M ?

To evaluate whether using importance weighted estimators can recover target distributions, we consider a synthetically generated distribution that has been manipulated along a latent dimension. Under the target distribution, a latent representation θ_i of each data point lives in a ten-dimensional space, with each dimension independently Uniform(0,1). The observed data points x_i are then obtained as $\theta_i^T F$, where $F_{ij} \sim \mathcal{N}(0, 1)$ represents a fixed mapping between the latent space and D -dimensional observation space. In the training data, the first dimension of θ_i has distribution $p(\theta) = 2\theta, 0 < \theta \leq 1$. We assume that the modifying function $M(x_i) = 2\theta_{i,1}$ is observed, but that the remaining latent dimensions are unobserved.

In our experiments, we generate samples from the target distribution using each of the methods described above, and include weighted versions of the cross entropy GAN to demonstrate that importance weighting can be generalized to other losses.

To compare methods, we report the empirically estimated KL divergence between the target and generated samples in Table 2. Similar results using squared MMD and energy distance are shown in the extended version of this work [12]. For varying real dimensions D , importance weighted methods outperform C-GAN under a variety of measures.

In some instances C-GAN performs well in two dimensions, but deteriorates quickly as the problem becomes more challenging with higher dimensions. We also note that many runs of C-GAN either ran into numerical issues or diverged; in these cases we report the best score among runs, before training failure.

Table 2: Estimated KL divergence between generated and target samples (mean \pm standard deviation over 20 runs).

Model	2D	4D	10D
IW-CE	0.1768 \pm 0.0635	0.4934 \pm 0.1238	2.7945 \pm 0.5966
MIW-CE	0.3265 \pm 0.1071	0.6251 \pm 0.1343	3.3093 \pm 0.7179
SNIW-CE	0.0925 \pm 0.0272	0.3864 \pm 0.1478	2.3060 \pm 0.6915
ID-CE	0.1526 \pm 0.0332	0.3444 \pm 0.0766	1.4128 \pm 0.3288
IW-MMD	0.0343 \pm 0.0230	0.0037 \pm 0.0489	0.5133 \pm 0.1718
MIW-MMD	0.2698 \pm 0.0618	0.0939 \pm 0.0522	0.8501 \pm 0.3271
SNIW-MMD	0.0451 \pm 0.0132	0.1435 \pm 0.0377	0.6623 \pm 0.0918
C-GAN	0.0879 \pm 0.0405	0.3108 \pm 0.0982	6.9016 \pm 2.8406

While the above experiment can be evaluated numerically and provide good results for thinning on a continuous-valued variable, it is difficult to visualize the outcome. In order to better visualize whether the target distribution is correctly achieved, we also run experiments with explicit and easily measurable class distributions. In Figure 2, we show a class rebalancing problem on MNIST digits, where an initial uneven distribution between three classes can be accurately rebalanced. In the extended version of this work [12], we also show good performance modifying a balanced distribution to specific boosted levels. Together, these experiments provide evidence that importance weighting controls the simulated distribution in the desired way.

3.2 In a High-dimensional Image Setting, How Does Importance Weighting Compare with Conditional Generation?

Next we evaluate performance of importance weighted MMD on high-dimensional image generation. In this section we address two questions: Can our estimators generate simulations from \mathbb{P} in such a setting, and how do the resulting images

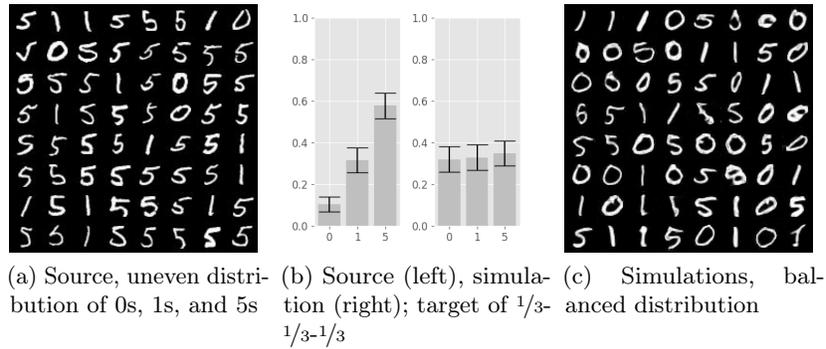


Fig. 2: Importance weights are used to accurately rebalance an uneven class distribution.

compare with those obtained using a C-GAN? To do so, we evaluate several generative models on the Yearbook dataset [15], which contains over 37,000 high school yearbook photos across over 100 years and demonstrates evolving styles and demographics. The goal is to produce images uniformly across each half decade. Each GAN, however, is trained on the original dataset, which contains many more photos from recent decades.

Since we have specified M in terms of a single covariate (time), we can compare with C-GANs. For the C-GAN, we use a conditional version of the standard DCGAN architecture (C-DCGAN) [40].

Figure 3 shows generated images from each network. All networks were trained until convergence. The images show a diversity across hairstyles, demographics and facial expressions, indicating the successful temporal rebalancing. Even while importance duplication introduces approximations and lacks the theoretical guarantees of the other two methods, all three importance-based methods achieve comparable quality. Since some covariates have fewer than 65 images, C-DCGAN cannot learn the conditional distributions, and is unstable across a variety of training parameters. Implementation details and additional experiments are shown in the extended version of this work [12].

3.3 When M Is Unknown, but Can Be Estimated Up to a Normalizing Constant on a Subset of Data, Are We Able to Sample from our Target Distribution?

In many settings, especially those with high-dimensional latent features, we will not know the functional form of M , or even the corresponding thinning function T . We would still, however, like to be able to express a preference for certain areas of the latent space. To do so, we propose labeling a small subset of data using weights that correspond to preference. To expand those weights to the entire dataset, we train a neural network called the estimated weighting function. This

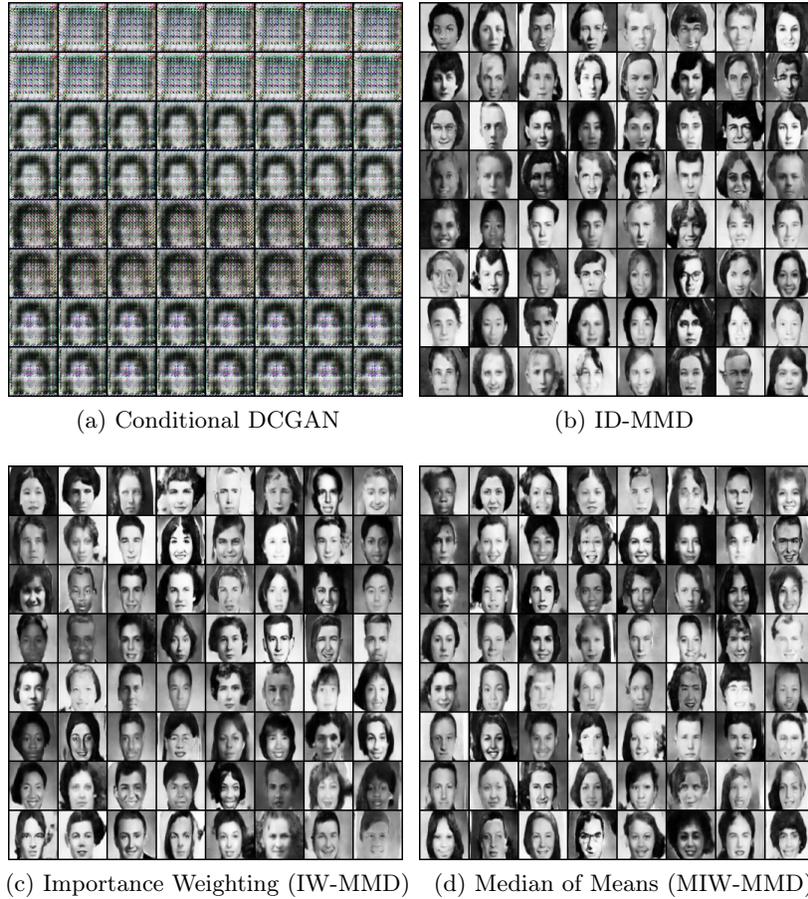


Fig. 3: Example generated images for all example networks, Yearbook dataset [15]. Target distribution is uniform across half-decades, while the training set is unbalanced.

weighting function takes encoded images as input, and outputs continuous-valued weights. Since this function exists in a high-dimensional space that changes as the encoder is updated, and since we do not know the full observed distribution on this space, we are in a setting unsuitable for conditional methods, and therefore use self-normalized estimators (SNIW-MMD).

We evaluate using a collection of sevens from the MNIST dataset, where the goal is to generate more European-style sevens with horizontal bars. Out of 5915 images, 200 were manually labeled with a weight (reciprocal of a thinning function value), where sevens with no horizontal bar were assigned a 1, and sevens with horizontal bars were assigned weights between 2 and 9 based on the width of the bar.

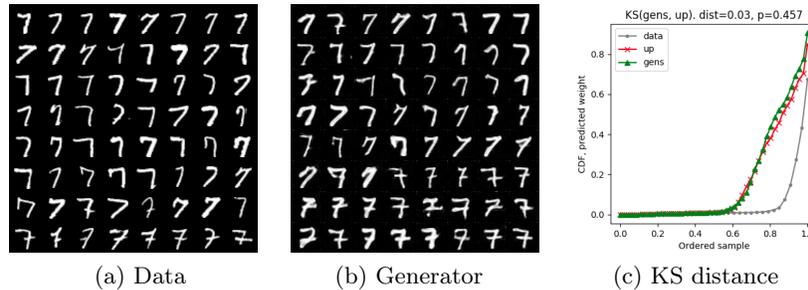


Fig. 4: Partial labeling and an importance weighted estimator boost the presence of sevens with horizontal bars. In 4a and 4b, samples are sorted by predicted weight, and in 4c, the empirical CDFs of data, generated, and importance duplicated draws, are shown, where the latter serves as a theoretical target. The generated distribution is close in distance to the target.

Fig. 4a shows 64 real images, sorted in terms of their predicted weights – note that the majority have no horizontal bar. Fig. 4b shows 64 generated simulations, sorted in the same manner, clearly showing an increase in the number of horizontal-bar sevens.

To test the quantitative performance, we display and compare the empirical CDFs of weights from simulations, data, and importance duplicated data. For example, if a batch of data $[A, B, C]$ has weights $[1, 3, 2]$, this implies that we expected three times as many B -like points and two times as many C -like points as A -like points. A simulator that achieves this target produces simulations like $[A, B, B, B, C, C]$ with weights $[1, 3, 3, 3, 2, 2]$, equivalent to an importance duplication of data weights. Using importance duplicated weights as a theoretical target, we measure our model’s performance by computing the Kolmogorov-Smirnov (KS) distance between CDFs of simulated and importance duplicated weights. Fig. 4c shows a small distributional distance between simulations and their theoretical target, with $d_{KS} = 0.03$, $p = 0.457$.

4 Conclusions and Future Work

We present three estimators for the MMD (and a wide class of other loss functions) between target distribution \mathbb{P} and the distribution \mathbb{Q} implied by our generator. These estimators can be used to train a GAN to simulate from the target distribution \mathbb{P} , given samples from a modified distribution $M\mathbb{P}$. We present solutions for when M is potentially unbounded, is unknown, or is known only up to a scaling factor.

We demonstrate that importance weighted estimators allow deep generative models to match target distributions for common and challenging cases with continuous-valued, multivariate latent features. This method avoids heuristics while providing good empirical performance and theoretical guarantees.

Though the median of means estimator offers a more robust estimate of the MMD, we may still experience high variance in our estimates, for example if we rarely see data points from a class we want to boost. An interesting future line of research is exploring how variance-reduction techniques [11] or adaptive batch sizes [10] could be used to overcome this problem.

Bibliography

- [1] Alon, N., Matias, Y., Szegedy, M.: The space complexity of approximating the frequency moments. In: ACM symposium on Theory of Computing. pp. 20–29. ACM (1996)
- [2] Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. In: ICML (2017)
- [3] Ben-David, S., Blitzer, J., Crammer, K., Pereira, F.: Analysis of representations for domain adaptation. In: Advances in neural information processing systems. pp. 137–144 (2007)
- [4] Bińkowski, M., Sutherland, D., Arbel, M., Gretton, A.: Demystifying MMD GANs. In: ICLR (2018)
- [5] Blitzer, J., Dredze, M., Pereira, F.: Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In: Proceedings of the 45th annual meeting of the association of computational linguistics. pp. 440–447 (2007)
- [6] Bolukbasi, T., Chang, K.W., Zou, J., Saligrama, V., Kalai, A.: Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: NIPS (2016)
- [7] Che, T., Li, Y., Jacob, A., Bengio, Y., Li, W.: Mode regularized generative adversarial networks. In: ICLR (2017)
- [8] Csaba, B., Boukhayma, A., Kulharia, V., Horváth, A., Torr, P.H.: Domain partitioning network. arXiv preprint arXiv:1902.08134 (2019)
- [9] Daumé III, H., Kumar, A., Saha, A.: Frustratingly easy semi-supervised domain adaptation. In: Proceedings of the 2010 Workshop on Domain Adaptation for Natural Language Processing. pp. 53–59. Association for Computational Linguistics (2010)
- [10] De, S., Yadav, A., Jacobs, D., Goldstein, T.: Automated inference with adaptive batches. In: AISTATS. pp. 1504–1513 (2017)
- [11] Defazio, A., Bach, F., Lacoste-Julien, S.: SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. In: NIPS. pp. 1646–1654 (2014)
- [12] Diesendruck, M., Elenberg, E.R., Sen, R., Cole, G.W., Shakkottai, S., Williamson, S.A.: Importance weighted generative networks. arXiv preprint arXiv:1806.02512 (2018)
- [13] Dziugaite, G., Roy, D., Ghahramani, Z.: Training generative neural networks via maximum mean discrepancy optimization. In: UAI (2015)
- [14] Esteban, C., Hyland, S., Rätsch, G.: Real-valued (medical) time series generation with recurrent conditional GANs. arXiv:1706.02633 (2017)

- [15] Ginosar, S., Rakelly, K., Sachs, S.M., Yin, B., Lee, C., Krähenbühl, P., Efros, A.A.: A century of portraits: A visual historical record of American high school yearbooks. *IEEE Transactions on Computational Imaging* **3**(3), 421–431 (Sept 2017). <https://doi.org/10.1109/TCI.2017.2699865>
- [16] Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. pp. 2066–2073. IEEE (2012)
- [17] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *NIPS* (2014)
- [18] Gretton, A., Borgwardt, K., Rasch, M., Schölkopf, B., Smola, A.: A kernel two-sample test. *JMLR* **13**(Mar), 723–773 (2012)
- [19] Hjelm, R.D., Jacob, A.P., Che, T., Trischler, A., Cho, K., Bengio, Y.: Boundary-seeking generative adversarial networks. *arXiv preprint arXiv:1702.08431* (2017)
- [20] Hoeffding, W.: A class of statistics with asymptotically normal distribution. *The annals of mathematical statistics* pp. 293–325 (1948)
- [21] Hoeffding, W.: Probability inequalities for sums of bounded random variables. *JASA* **58**(301), 13–30 (1963)
- [22] Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213* (2017)
- [23] Horvitz, D., Thompson, D.: A generalization of sampling without replacement from a finite universe. *JASA* **47**(260), 663–685 (1952)
- [24] Huang, J., Gretton, A., Borgwardt, K.M., Schölkopf, B., Smola, A.J.: Correcting sample selection bias by unlabeled data. In: *Advances in neural information processing systems*. pp. 601–608 (2007)
- [25] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. *arXiv preprint* (2017)
- [26] Jerrum, M.R., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science* **43**, 169–188 (1986)
- [27] Kadurin, A., Aliper, A., Kazennov, A., Mamoshina, P., Vanhaelen, Q., Khrabrov, K., Zhavoronkov, A.: The cornucopia of meaningful leads: Applying deep adversarial autoencoders for new molecule development in oncology. *Oncotarget* **8**(7), 10883 (2017)
- [28] Kong, A.: A note on importance sampling using standardized weights. *University of Chicago, Dept. of Statistics, Tech. Rep* **348** (1992)
- [29] Lerasle, M., Szabó, Z., Mathieu, T., Lecué, G.: Monk-outlier-robust mean embedding estimation by median-of-means. *arXiv preprint arXiv:1802.04784* (2018)
- [30] Li, C.L., Chang, W.C., Cheng, Y., Yang, Y., Póczos, B.: MMD GAN: Towards deeper understanding of moment matching network. In: *NIPS* (2017)
- [31] Li, Y., Swersky, K., Zemel, R.: Generative moment matching networks. In: *ICML* (2015)

- [32] Mansournia, M., Altman, D.: Inverse probability weighting. *BMJ* **352**, i189 (2016)
- [33] Mehdi, M., Osindero, S.: Conditional generative adversarial nets. arXiv:1411.1784 (2014)
- [34] Mises, R.v.: On the asymptotic distribution of differentiable statistical functions. *The Annals of Mathematical Statistics* **18**(3), 309–348 (1947)
- [35] Nemirovskii, A., Yudin, D.B., Dawson, E.R.: Problem complexity and method efficiency in optimization. Wiley (1983)
- [36] Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In: ICML (2017)
- [37] Owen, A.: Monte Carlo theory, methods and examples. Book draft (2013)
- [38] Pan, S.J., Kwok, J.T., Yang, Q.: Transfer learning via dimensionality reduction. In: AAAI. vol. 8, pp. 677–682 (2008)
- [39] Pan, S.J., Tsang, I.W., Kwok, J.T., Yang, Q.: Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* **22**(2), 199–210 (2011)
- [40] Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434 (2015)
- [41] Robert, C., Casella, G.: Monte Carlo Statistical Methods. Springer Texts in Statistics, Springer, 2 edn. (2004)
- [42] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training GANs. In: NIPS (2016)
- [43] Sugiyama, M., Nakajima, S., Kashima, H., Buenau, P.V., Kawanabe, M.: Direct importance estimation with model selection and its application to covariate shift adaptation. In: Advances in neural information processing systems. pp. 1433–1440 (2008)
- [44] Székely, G.J., Rizzo, M.L.: Energy statistics: A class of statistics based on distances. *Journal of statistical planning and inference* **143**(8), 1249–1272 (2013)
- [45] Taigman, Y., Polyak, A., Wolf, L.: Unsupervised cross-domain image generation. arXiv preprint arXiv:1611.02200 (2016)
- [46] Vondrick, C., Pirsaviash, H., Torralba, A.: Generating videos with scene dynamics. In: NIPS (2016)
- [47] Yan, H., Ding, Y., Li, P., Wang, Q., Xu, Y., Zuo, W.: Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. arXiv preprint arXiv:1705.00609 (2017)
- [48] Zhang, H., Xu, T., Li, H., Zhang, S., Huang, X., Wang, X., Metaxas, D.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: ICCV (2017)
- [49] Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In: EMNLP (2017)
- [50] Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint (2017)