
Entropic GANs meet VAEs: A Statistical Approach to Compute Sample Likelihoods in GANs

Yogesh Balaji¹ Hamed Hassani² Rama Chellappa³ Soheil Feizi¹

Abstract

Building on the success of deep learning, two modern approaches to learn a probability model from the data are Generative Adversarial Networks (GANs) and Variational AutoEncoders (VAEs). VAEs consider an explicit probability model for the data and compute a generative distribution by maximizing a variational lower-bound on the log-likelihood function. GANs, however, compute a generative model by minimizing a distance between observed and generated probability distributions without considering an explicit model for the observed data. The lack of having explicit probability models in GANs prohibits computation of sample likelihoods in their frameworks and limits their use in statistical inference problems. In this work, we resolve this issue by constructing an explicit probability model that can be used to compute sample likelihood statistics in GANs. In particular, we prove that under this probability model, a family of Wasserstein GANs with an entropy regularization can be viewed as a generative model that maximizes a variational lower-bound on average sample log likelihoods, an approach that VAEs are based on. This result makes a principled connection between two modern generative models, namely GANs and VAEs. In addition to the aforementioned theoretical results, we compute likelihood statistics for GANs trained on Gaussian, MNIST, SVHN, CIFAR-10 and LSUN datasets. Our numerical results validate the proposed theory.

1. Introduction

Learning generative models is an important problem in machine learning and statistics with a wide range of applications in self-driving cars (Santana & Hotz, 2016), robotics (Hirose et al., 2017), natural language processing (Lee & Tsao, 2018), domain-transfer (Sankaranarayanan et al., 2018), computational biology (Ghahramani et al., 2018), etc. Two modern approaches to deal with this problem are Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) and Variational AutoEncoders (VAEs) (Kingma & Welling, 2013; Makhzani et al., 2015; Rosca et al., 2017; Tolstikhin et al., 2017; Mescheder et al., 2017b).

VAEs (Kingma & Welling, 2013) compute a generative model by maximizing a variational lower-bound on average sample log-likelihoods using an explicit probability distribution for the data. GANs, however, learn a generative model by minimizing a distance between observed and generated distributions without considering an explicit probability model for the data. Empirically, GANs have been shown to produce higher-quality generative samples than that of VAEs (Karras et al., 2017). However, since GANs do not consider an explicit probability model for the data, we are unable to compute sample likelihoods using their generative models. Obtaining sample likelihoods and *posterior* distributions of latent variables are critical in several statistical inference applications. Inability to obtain such statistics within GAN’s framework severely limits their applications in statistical inference problems.

In this paper, we resolve this issue for a general formulation of GANs by providing a theoretically-justified approach to compute sample likelihoods using GAN’s generative model. Our results can facilitate the use of GANs in massive-data applications such as model selection, sample selection, hypothesis-testing, etc.

We first state our main results *informally* without going into technical details while precise statements of our results are presented in Section 2. Let Y and $\hat{Y} := \mathbf{G}(X)$ represent the observed (i.e. real) and generative (i.e. fake or synthetic) variables, respectively. X (i.e. the latent variable) is the random vector used as the input to the generator $\mathbf{G}(\cdot)$. Consider the following explicit probability model of the data

¹Department of Computer Science, University of Maryland, College Park ²Department of Electrical and Systems Engineering, University of Pennsylvania ³Department of Electrical and Computer Engineering, University of Maryland, College Park. Correspondence to: Yogesh Balaji <yogesh@cs.umd.edu>, Soheil Feizi <sfeizi@cs.umd.edu>.

given a latent sample $X = \mathbf{x}$:

$$f_{Y|X=\mathbf{x}}(\mathbf{y}) \propto \exp(-\ell(\mathbf{y}, \mathbf{G}(\mathbf{x}))), \quad (1.1)$$

where $\ell(\cdot, \cdot)$ is a loss function. $f_{Y|X=\mathbf{x}}(\mathbf{y})$ is the model that we consider for the underlying data distribution. This is a reasonable model for the data as the function \mathbf{G} can be a complex function. Similar data models have been used in VAEs. Under this explicit probability model, we show that minimizing the objective of an *optimal transport* GAN (e.g. Wasserstein GAN (Arjovsky et al., 2017)) with the cost function $\ell(\cdot, \cdot)$ and an entropy regularization (Cuturi, 2013; Seguy et al., 2017) maximizes a variational lower-bound on average sample likelihoods. That is

$$\underbrace{\mathbb{E}_{\mathbb{P}_Y} [\log f_Y(Y)]}_{\text{ave. sample log likelihoods}} \geq -\frac{1}{\lambda} \underbrace{\left\{ \mathbb{E}_{\mathbb{P}_{Y, \hat{Y}}} [\ell(Y, \hat{Y})] - \lambda H(\mathbb{P}_{Y, \hat{Y}}) \right\}}_{\text{entropic GAN objective}} + \text{constants.}$$

If $\ell(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2$, the optimal transport (OT) GAN simplifies to WGAN (Arjovsky et al., 2017) while if $\ell(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2$, the OT GAN simplifies to the quadratic GAN (or, W2GAN) (Feizi et al., 2017). The precise statement of this result can be found in Theorem 1.

This result provides a statistical justification for GAN’s optimization and puts it in par with VAEs whose goal is to maximize a lower bound on sample likelihoods. We note that entropy regularization has been proposed primarily to improve computational aspects of GANs (Genevay et al., 2018). Our results provide an additional statistical justification for this regularization term. Moreover, using the GAN’s training, we obtain a coupling between the observed variable Y and the latent variable X . This coupling provides the conditional distribution of the latent variable X given an observed sample $Y = \mathbf{y}$. The explicit model of (1.1) acts similar to the *decoder* in the VAE framework, while the coupling computed using GANs acts as an *encoder*.

Another key question is how to estimate the likelihood of a new sample \mathbf{y}^{test} given the generative model trained using GANs. For instance, if we train a GAN on *stop-sign* images, upon receiving a new image, one may wish to compute the likelihood of the new sample \mathbf{y}^{test} according to the trained generative model. In standard GAN formulations, the support of the generative distribution lies on the range of the optimal generator function. Thus, if the observed sample \mathbf{y}^{test} does not lie in that range (which is very likely in practice), there is no way to assign a sensible likelihood score to the sample. Below, we show that using the explicit probability model of (1.1), we can *lower-bound* the likelihood of this sample \mathbf{y}^{test} . This is similar to the variational lower-bound on sample likelihoods used in VAEs. Our numerical results

in Section 4 show that this lower-bound well-reflects the expected trends of the true sample likelihoods.

Let \mathbf{G}^* and $\mathbb{P}_{Y, X}^*$ be the optimal generator and the optimal coupling between real and latent variables, respectively. The optimal coupling $\mathbb{P}_{Y, X}^*$ can be computed efficiently for entropic GANs as we explain in Section 3. For other GAN architectures, one may approximate such couplings as we explain in Section 4. The log likelihood of a new test sample \mathbf{y}^{test} can be lower-bounded as

$$\underbrace{\log f_Y(\mathbf{y}^{\text{test}})}_{\text{log likelihood}} \geq -\underbrace{\mathbb{E}_{\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*} [\ell(\mathbf{y}^{\text{test}}, \mathbf{G}^*(\mathbf{x}))]}_{\text{distance to the generative model}} + \underbrace{H(\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*)}_{\text{coupling entropy}} + \underbrace{\mathbb{E}_{\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*} \left[-\frac{\|\mathbf{x}\|^2}{2} \right]}_{\text{likelihood of latent variable}}. \quad (1.2)$$

We present the precise statement of this result in Corollary 2. This result combines three components in order to approximate the likelihood of a sample given a trained generative model: (1) if the distance between \mathbf{y}^{test} to the generative model is large, the likelihood of observing \mathbf{y}^{test} from the generative model is small, (2) if the entropy of the coupled latent variable is large, the coupled latent variable has large randomness, thus, this contributes positively to the sample likelihood, and (3) if the likelihood of the coupled latent variable is large, the likelihood of the observed test sample will be large as well. Figure 1 provides a pictorial illustration of these components.

To summarize, we make the following **theoretical contributions** in this paper:

- We construct an explicit probability model for a family of optimal transport GANs (such as the Wasserstein GAN) that can be used to compute likelihood statistics within GAN’s framework (eq. (2.4) and Corollary 2).
- We prove that, under this probability model, the objective of an entropic GAN is a variational lower bound for average sample log likelihoods (Theorem 1). This result makes a principled connection between two modern generative models, namely GANs and VAEs.

Moreover, we conduct the following **empirical experiments** in this paper:

- We compute likelihood statistics for GANs trained on Gaussian, MNIST, SVHN, CIFAR-10 and LSUN datasets and shown the consistency of these empirical results with the proposed theory (Section 4 and Supplementary Material Section 4).

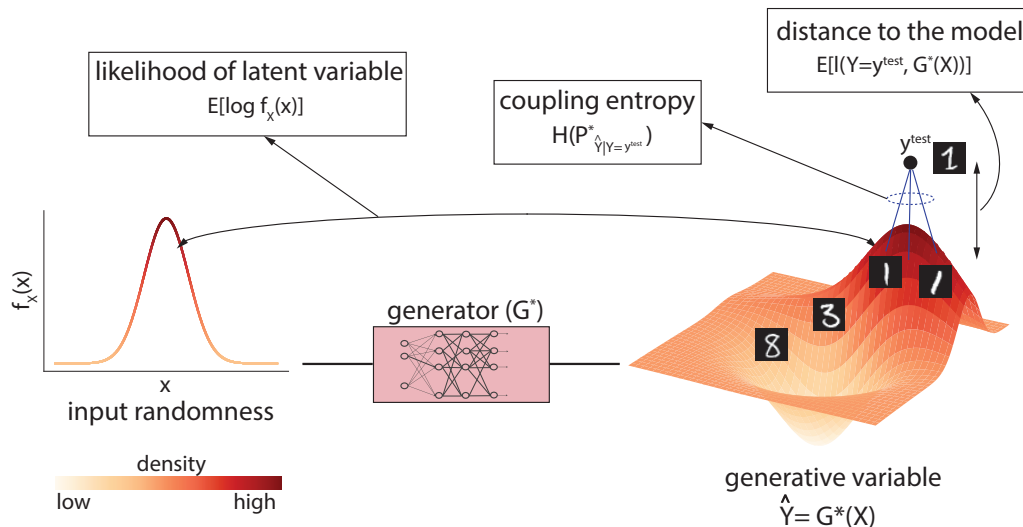


Figure 1: A statistical framework for GANs. By training a GAN model, we first compute optimal generator G^* and optimal coupling between the observed variable Y and the latent variable X . The likelihood of a test sample y^{test} can then be lower-bounded using a combination of three terms: (1) the expected distance of y^{test} to the distribution learnt by the generative model, (2) the entropy of the coupled latent variable given y^{test} and (3) the likelihood of the coupled latent variable with y^{test} .

- We demonstrate the tightness of the variational lower bound of entropic GANs for both linear and non-linear generators (Section 4.3).

1.1. Related Work

Connections between GANs and VAEs have been investigated in some of the recent works as well (Hu et al., 2018; Mescheder et al., 2017a). In (Hu et al., 2018), GANs are interpreted as methods performing variational inference on a generative model in the label space. In their framework, observed data samples are treated as latent variables while the generative variable is the indicator of whether data is real or fake. The method in (Mescheder et al., 2017a), on the other hand, uses an auxiliary discriminator network to rephrase the maximum-likelihood objective of a VAE as a two-player game similar to the objective of a GAN. Our method is different from both of these approaches as we consider an explicit probability model for the data, and show that the entropic GAN objective maximizes a variational lower bound under this probability model, thus allowing sample likelihood computation in GANs similar to VAEs.

Of relevance to our work is (Wu et al., 2016), in which annealed importance sampling (AIS) is used to evaluate the approximate likelihood of decoder-based generative models. More specifically, a Gaussian observation model with a fixed variance is used as the generative distribution for GAN-based models on which the AIS is computed. Gaussian

observation models may not be appropriate specially in high-dimensional spaces. Our approach, on the other hand, makes a connection between GANs and VAEs by constructing a theoretically-motivated model for the data distribution in GANs, and uses this model to compute sample likelihoods.

In an independent recent work, the authors of (Rigollet & Weed, 2018) show that entropic optimal transport corresponds to the objective function in maximum-likelihood estimation for deconvolution problems involving additive Gaussian noise. Our result is similar in spirit, however, our focus is in providing statistical interpretations to GANs by constructing a probability model for data distribution using GANs. We show that under this model, the objective of the entropic GAN is a variational lower bound to the average log-likelihood function. This lower-bound can then be used for computing sample likelihood estimates in GANs.

2. A Variational Bound for GANs

Let $Y \in \mathbb{R}^d$ represent the real-data random variable with a probability density function $f_Y(y)$. GAN's goal is to find a generator function $G: \mathbb{R}^r \rightarrow \mathbb{R}^d$ such that $\hat{Y} := G(X)$ has a similar distribution to Y . Let X be an r -dimensional random variable with a fixed probability density function $f_X(x)$. Here, we assume $f_X(\cdot)$ is the density of a normal distribution. In practice, we observe m samples $\{y_1, \dots, y_m\}$ from Y and generate m' samples from \hat{Y} , i.e., $\{\hat{y}_1, \dots, \hat{y}_{m'}\}$ where $\hat{y}_i = G(x_i)$ for $1 \leq i \leq m'$. We rep-

represent these empirical distributions by \mathbb{P}_Y and $\mathbb{P}_{\hat{Y}}$, respectively. Note that the number of generative samples m' can be arbitrarily large.

GAN computes the optimal generator \mathbf{G}^* by minimizing a distance between the observed distribution \mathbb{P}_Y and the generative one $\mathbb{P}_{\hat{Y}}$. Common distance measures include *optimal transport* measures (e.g. Wasserstein GAN (Arjovsky et al., 2017), WGAN+Gradient Penalty (Gulrajani et al., 2017), GAN+Spectral Normalization (Miyato et al., 2018), WGAN+Truncated Gradient Penalty (Petzka et al., 2017), relaxed WGAN (Guo et al., 2017)), and *divergence* measures (e.g. the original GAN's formulation (Goodfellow et al., 2014), f -GAN (Nowozin et al., 2016)), etc.

In this paper, we focus on GANs based on optimal transport (OT) distance (Villani, 2008; Arjovsky et al., 2017) defined for a general loss function $\ell(\cdot, \cdot)$ as follows

$$W^{(\ell)}(\mathbb{P}_Y, \mathbb{P}_{\hat{Y}}) := \min_{\mathbb{P}_{Y, \hat{Y}}} \mathbb{E}[\ell(Y, \hat{Y})]. \quad (2.1)$$

$\mathbb{P}_{Y, \hat{Y}}$ is the joint distribution whose marginal distributions are equal to \mathbb{P}_Y and $\mathbb{P}_{\hat{Y}}$, respectively. If $\ell(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2$, this distance is called the first-order Wasserstein distance and is referred to by $W_1(\cdot, \cdot)$, while if $\ell(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2^2$, this measure is referred to by $W_2(\cdot, \cdot)$ where W_2 is the second-order Wasserstein distance (Villani, 2008).

The optimal transport (OT) GAN is formulated using the following optimization problem (Arjovsky et al., 2017):

$$\min_{\mathbf{G} \in \mathcal{G}} W^{(\ell)}(\mathbb{P}_Y, \mathbb{P}_{\hat{Y}}), \quad (2.2)$$

where \mathcal{G} is the set of generator functions. Examples of the OT GAN are WGAN (Arjovsky et al., 2017) corresponding to the first-order Wasserstein distance¹ and the quadratic GAN (or, the W2GAN) (Feizi et al., 2017) corresponding to the second-order Wasserstein distance.

Note that optimization (2.2) is a *min-min* optimization. The objective of this optimization is not smooth in \mathbf{G} and it is often computationally expensive to obtain a solution for it (Sanjabi et al., 2018). One approach to improve computational aspects of this optimization problem is to add a regularization term to make its objective *strongly* convex (Cuturi, 2013; Seguy et al., 2017). The Shannon entropy function is defined as $H(\mathbb{P}_{Y, \hat{Y}}) := -\mathbb{E}[\log \mathbb{P}_{Y, \hat{Y}}]$. The negative Shannon entropy is a common strongly-convex regularization term. This leads to the following optimal transport GAN formulation with the entropy regularization,

or for simplicity, the *entropic GAN* formulation:

$$\min_{\mathbf{G} \in \mathcal{G}} \min_{\mathbb{P}_{Y, \hat{Y}}} \mathbb{E}[\ell(Y, \hat{Y})] - \lambda H(\mathbb{P}_{Y, \hat{Y}}), \quad (2.3)$$

where λ is the regularization parameter.

There are two approaches to solve the optimization problem (2.3). The first approach uses an iterative method to solve the *min-min* formulation (Genevay et al., 2017). Another approach is to solve an equivalent *min-max* formulation by writing the dual of the inner minimization (Seguy et al., 2017; Sanjabi et al., 2018). The latter is often referred to as a GAN formulation since the min-max optimization is over a set of generator functions and a set of discriminator functions. The details of this approach are further explained in Section 3.

In the following, we present an explicit probability model for entropic GANs under which their objective can be viewed as maximizing a lower bound on average sample likelihoods.

Theorem 1 *Let the loss function be shift invariant, i.e., $\ell(\mathbf{y}, \hat{\mathbf{y}}) = h(\mathbf{y} - \hat{\mathbf{y}})$. Let*

$$f_{Y|X=\mathbf{x}}(\mathbf{y}) = C \exp(-\ell(\mathbf{y}, \mathbf{G}(\mathbf{x}))/\lambda), \quad (2.4)$$

be an explicit probability model for Y given $X = \mathbf{x}$ for a well-defined normalization

$$C := \frac{1}{\int_{\mathbf{y} \in \mathbb{R}^d} \exp(-\ell(\mathbf{y}, \mathbf{G}(\mathbf{x}))/\lambda)}. \quad (2.5)$$

Then, we have

$$\underbrace{\mathbb{E}_{\mathbb{P}_Y} [\log f_Y(Y)]}_{\text{ave. sample likelihoods}} \geq -\frac{1}{\lambda} \underbrace{\left\{ \mathbb{E}_{\mathbb{P}_{Y, \hat{Y}}} [\ell(Y, \hat{Y})] - \lambda H(\mathbb{P}_{Y, \hat{Y}}) \right\}}_{\text{entropic GAN objective}} + \text{constants}. \quad (2.6)$$

In words, the entropic GAN maximizes a lower bound on sample likelihoods according to the explicit probability model of (2.4).

The proof of this theorem is presented in supplementary material. This result has a similar flavor to that of VAEs (Kingma & Welling, 2013; Rosca et al., 2017) where a generative model is computed by maximizing a lower bound on sample likelihoods.

Having a shift invariant loss function is critical for Theorem 1 as this makes the normalization term C to be independent from \mathbf{G} and \mathbf{x} (to see this, one can define $\mathbf{y}' := \mathbf{y} - \mathbf{G}(\mathbf{x})$ in (2.6)). The most standard OT GAN loss functions such as the L_2 for WGAN (Arjovsky et al., 2017) and the quadratic loss for W2GAN (Feizi et al., 2017) satisfy this property.

One can further simplify this result by considering specific loss functions. For example, we have the following result for the entropic GAN with the quadratic loss function.

¹Note that some references (e.g. (Arjovsky et al., 2017)) refer to the first-order Wasserstein distance simply as the Wasserstein distance. In this paper, we explicitly distinguish between different Wasserstein distances.

Corollary 1 Let $\ell(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|^2/2$. Then, $f_{Y|X=\mathbf{x}}(\cdot)$ of (2.4) corresponds to the multivariate Gaussian density function and $C = \frac{1}{\sqrt{(2\pi\lambda)^d}}$. In this case, the constant term in (2.6) is equal to $-\log(m) - d\log(2\pi\lambda)/2 - r/2 - \log(2\pi)/2$.

Let \mathbf{G}^* and $\mathbb{P}_{Y,X}^*$ be optimal solutions of an entropic GAN optimization (2.3) (note that the optimal coupling can be computed efficiently using (3.7)). Let \mathbf{y}^{test} be a newly observed sample. An important question is what the likelihood of this sample is given the trained generative model. Using the explicit probability model of (2.4) and the result of Theorem 1, we can (approximately) compute sample likelihoods as explained in the following corollary.

Corollary 2 Let \mathbf{G}^* and $\mathbb{P}_{Y,\hat{Y}}^*$ (or, alternatively $\mathbb{P}_{Y,X}^*$) be optimal solutions of the entropic GAN (2.3). Let \mathbf{y}^{test} be a new observed sample. We have

$$\begin{aligned} \log f_Y(\mathbf{y}^{\text{test}}) &\geq -\frac{1}{\lambda} \left\{ \mathbb{E}_{\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*} [\ell(\mathbf{y}^{\text{test}}, \mathbf{G}^*(\mathbf{x}))] \right. \\ &\quad \left. - \lambda H(\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*) \right\} + \mathbb{E}_{\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*} \left[-\frac{\|\mathbf{x}\|^2}{2} \right] \\ &\quad + \text{constants}. \end{aligned} \quad (2.7)$$

The inequality becomes tight iff $KL(\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^* \| f_{X|Y=\mathbf{y}^{\text{test}}}) = 0$, where $KL(\cdot, \|\cdot)$ is the Kullback-Leibler divergence between two distributions. The r.h.s of equation (2.7) will be denoted as "surrogate likelihood" in the rest of the paper.

3. Review of GAN's Dual Formulations

In this section, we discuss dual formulations for OT GAN (2.2) and entropic GAN (2.3) optimizations. These dual formulations are *min-max* optimizations over two function classes, namely the generator and the discriminator. Often local search methods such as alternating gradient descent (GD) are used to compute a solution for these min-max optimizations.

First, we discuss the dual formulation of OT GAN optimization (2.2). Using the duality of the inner minimization, which is a linear program, we can re-write optimization (2.2) as follows (Villani, 2008):

$$\min_{\mathbf{G} \in \mathcal{G}} \max_{\mathbf{D}_1, \mathbf{D}_2} \mathbb{E}[\mathbf{D}_1(Y)] - \mathbb{E}[\mathbf{D}_2(\mathbf{G}(X))], \quad (3.1)$$

where $\mathbf{D}_1(\mathbf{y}) - \mathbf{D}_2(\hat{\mathbf{y}}) \leq \ell(\mathbf{y}, \hat{\mathbf{y}})$ for all $(\mathbf{y}, \hat{\mathbf{y}})$. The maximization is over two sets of functions \mathbf{D}_1 and \mathbf{D}_2 which are coupled using the loss function. Using the Kantorovich duality (Villani, 2008), we can further simplify this optimization as follows:

$$\min_{\mathbf{G} \in \mathcal{G}} \max_{\mathbf{D}: \ell\text{-convex}} \mathbb{E}[\mathbf{D}(Y)] - \mathbb{E}[\mathbf{D}^{(\ell)}(\mathbf{G}(X))], \quad (3.2)$$

where $\mathbf{D}^{(\ell)}(\hat{Y}) := \inf_Y \ell(Y, \hat{Y}) + \mathbf{D}(Y)$ is the ℓ -conjugate function of $\mathbf{D}(\cdot)$ and \mathbf{D} is restricted to ℓ -convex functions (Villani, 2008). The above optimization provides a general formulation for OT GANs. If the loss function is $\|\cdot\|_2$, then the optimal transport distance is referred to as the first order Wasserstein distance. In this case, the min-max optimization (3.2) simplifies to the following optimization (Arjovsky et al., 2017):

$$\min_{\mathbf{G} \in \mathcal{G}} \max_{\mathbf{D}: 1\text{-Lip}} \mathbb{E}[\mathbf{D}(Y)] - \mathbb{E}[\mathbf{D}(\mathbf{G}(X))]. \quad (3.3)$$

This is often referred to as Wasserstein GAN, or WGAN (Arjovsky et al., 2017). If the loss function is quadratic, then the OT GAN is referred to as the quadratic GAN (or, W2GAN) (Feizi et al., 2017).

Similarly, the dual formulation of the entropic GAN (2.3) can be written as the following optimization (Cuturi, 2013; Seguy et al., 2017)²:

$$\begin{aligned} \min_{\mathbf{G} \in \mathcal{G}} \max_{\mathbf{D}_1, \mathbf{D}_2} \mathbb{E}[\mathbf{D}_1(Y)] - \mathbb{E}[\mathbf{D}_2(\mathbf{G}(X))] \\ - \lambda \mathbb{E}_{\mathbb{P}_Y \times \mathbb{P}_{\hat{Y}}} [\exp(v(\mathbf{y}, \hat{\mathbf{y}})/\lambda)], \end{aligned} \quad (3.4)$$

where

$$v(\mathbf{y}, \hat{\mathbf{y}}) := \mathbf{D}_1(\mathbf{y}) - \mathbf{D}_2(\hat{\mathbf{y}}) - \ell(\mathbf{y}, \hat{\mathbf{y}}). \quad (3.5)$$

Note that the hard constraint of optimization (3.1) is being replaced by a soft constraint in optimization (3.2). In this case, optimal primal variables $\mathbb{P}_{Y,\hat{Y}}^*$ can be computed according to the following lemma (Seguy et al., 2017):

Lemma 1 Let \mathbf{D}_1^* and \mathbf{D}_2^* be the optimal discriminator functions for a given generator function \mathbf{G} according to optimization (3.4). Let

$$v^*(\mathbf{y}, \hat{\mathbf{y}}) := \mathbf{D}_1^*(\mathbf{y}) - \mathbf{D}_2^*(\hat{\mathbf{y}}) - \ell(\mathbf{y}, \hat{\mathbf{y}}). \quad (3.6)$$

Then,

$$\mathbb{P}_{Y,\hat{Y}}^*(\mathbf{y}, \hat{\mathbf{y}}) = \mathbb{P}_Y(\mathbf{y})\mathbb{P}_{\hat{Y}}(\hat{\mathbf{y}}) \exp(v^*(\mathbf{y}, \hat{\mathbf{y}})/\lambda). \quad (3.7)$$

This lemma is important for our results since it provides an efficient way to compute the optimal coupling between real and generative variables (i.e. $\mathbb{P}_{Y,\hat{Y}}^*$) using the optimal generator (\mathbf{G}^*) and discriminators (\mathbf{D}_1^* and \mathbf{D}_2^*) of optimization (3.4). It is worth noting that without the entropy regularization term, computing the optimal coupling using the optimal generator and discriminator functions is not straightforward in general (unless in some special cases such as W2GAN (Villani, 2008; Feizi et al., 2017)). This is another additional computational benefit of using entropic GAN.

²Note that optimization (3.4) is dual of optimization (2.3) when the terms $\lambda H(\mathbb{P}_Y) + \lambda H(\mathbb{P}_{\hat{Y}})$ have been added to its objective. Since for a fixed \mathbf{G} (fixed marginals), these terms are constants, they can be ignored from the optimization objective without loss of generality.

4. Experimental Results

In this section, we supplement our theoretical results with experimental validations. One of the main objectives of our work is to provide a framework to compute sample likelihoods in GANs. Such likelihood statistics can then be used in several statistical inference applications that we discuss in Section 5. With a trained entropic WGAN, the likelihood of a test sample can be lower-bounded using Corollary 2. Note that this likelihood estimate requires the discriminators \mathbf{D}_1 and \mathbf{D}_2 to be solved to optimality. In our implementation, we use the algorithm presented in (Sanjabi et al., 2018) to train the Entropic GAN. It has been proven (Sanjabi et al., 2018) that this algorithm leads to a good approximation of stationary solutions of Entropic GAN. We also discuss an approximate likelihood computation approach for un-regularized GANs in SM Section 4.

To obtain the surrogate likelihood estimates using Corollary 2, we need to compute the density $\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x})$. As shown in Lemma 1, WGAN with entropy regularization provides a closed-form solution to the conditional density of the latent variable ((3.7)). When \mathbf{G}^* is injective, $\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x})$ can be obtained from (3.7) by change of variables. In general case, $\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x})$ is not well defined as multiple \mathbf{x} can produce the same \mathbf{y}^{test} . In this case,

$$\mathbb{P}_{\hat{Y}|Y=\mathbf{y}^{\text{test}}}^*(\hat{\mathbf{y}}) = \sum_{\mathbf{x}|\mathbf{G}^*(\mathbf{x})=\hat{\mathbf{y}}} \mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x}). \quad (4.1)$$

Also, from (3.7), we have

$$\mathbb{P}_{\hat{Y}|Y=\mathbf{y}^{\text{test}}}^*(\hat{\mathbf{y}}) = \sum_{\mathbf{x}|\mathbf{G}^*(\mathbf{x})=\hat{\mathbf{y}}} \mathbb{P}_X(\mathbf{x}) \exp(v^*(\mathbf{y}^{\text{test}}, \mathbf{G}^*(\mathbf{x}))/\lambda). \quad (4.2)$$

One solution (which may not be unique) that satisfies both (4.1) and (4.2) is

$$\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x}) = \mathbb{P}_X(\mathbf{x}) \exp(v^*(\mathbf{y}^{\text{test}}, \mathbf{G}^*(\mathbf{x}))/\lambda). \quad (4.3)$$

Ideally, we would like to choose $\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x})$, satisfying (4.1) and (4.2) that maximizes the lower bound of Corollary 2. But finding such a solution can be difficult in general. Instead we use (4.3) to evaluate the surrogate likelihoods of Corollary 2 (note that our results still hold in this case). In order to compute our proposed surrogate likelihood, we need to draw samples from the distribution $\mathbb{P}_{X|Y=\mathbf{y}^{\text{test}}}^*(\mathbf{x})$. One approach is to use a Markov chain Monte Carlo (MCMC) method to sample from this distribution. In our experiments, however, we found that MCMC demonstrates poor performance owing to the high dimensional nature of X . A similar issue with MCMC has been reported for VAEs in (Kingma & Welling, 2013). Thus, we use a different estimator to compute the likelihood surrogate which provides a better exploration of the latent space. We present our sampling procedure in Algorithm 1 of Supplementary material.

4.1. Likelihood Evolution in GAN’s Training

In the experiments of this section, we study how sample likelihoods vary during GAN’s training. An entropic WGAN is first trained on MNIST training set. Then, we randomly choose 1,000 samples from MNIST test set to compute the surrogate likelihoods using Algorithm 1 at different training iterations. Surrogate likelihood computation requires solving \mathbf{D}_1 and \mathbf{D}_2 to optimality for a given \mathbf{G} (refer to Lemma 1), which might not be satisfied at the intermediate iterations of the training process. Therefore, before computing the surrogate likelihoods, discriminators \mathbf{D}_1 and \mathbf{D}_2 are updated for 100 steps for a fixed \mathbf{G} . We expect sample likelihoods to increase over training iterations as the quality of the generative model improves.

Fig. 2a demonstrates the evolution of sample likelihood distributions at different training iterations of the entropic WGAN. Note that the constant in surrogate likelihood (Corollary 2) was ignored for obtaining the plot since its inclusion would have only shifted every curve by the same offset. At iteration 1, surrogate likelihood values are very low as GAN’s generated images are merely random noise. The likelihood distribution shifts towards high values during the training and saturates beyond a point. The depicted likelihood values (after convergence) is roughly in the ballpark of the values reported by direct likelihood optimization methods (e.g., DRAW (Gregor et al., 2015), PixelRNN (Oord et al., 2016)).

4.2. Likelihood Comparison Across Different Datasets

In this section, we perform experiments across different datasets. An entropic WGAN is first trained on a subset of samples from the MNIST dataset containing digit 1 (which we call the MNIST-1 dataset). With this trained model, likelihood estimates are computed for (1) samples from the entire MNIST dataset, and (2) samples from the Street View House Numbers (SVHN) dataset (Netzer et al., 2011) (Fig. 2b). In each experiment, the likelihood estimates are computed for 1000 samples. We note that highest likelihood estimates are obtained for samples from MNIST-1 dataset, the same dataset on which the GAN was trained. The likelihood distribution for the MNIST dataset is bimodal with one mode peaking inline with the MNIST-1 mode. Samples from this mode correspond to digit 1 in the MNIST dataset. The other mode, which is the dominant one, contains the rest of the digits and has relatively low likelihood estimates. The SVHN dataset, on the other hand, has much smaller likelihoods as its distribution is significantly different than that of MNIST. Furthermore, we observe that the likelihood distribution of SVHN samples has a large spread (variance). This is because samples of the SVHN dataset is more diverse with varying backgrounds and styles than samples from MNIST. We note that SVHN samples

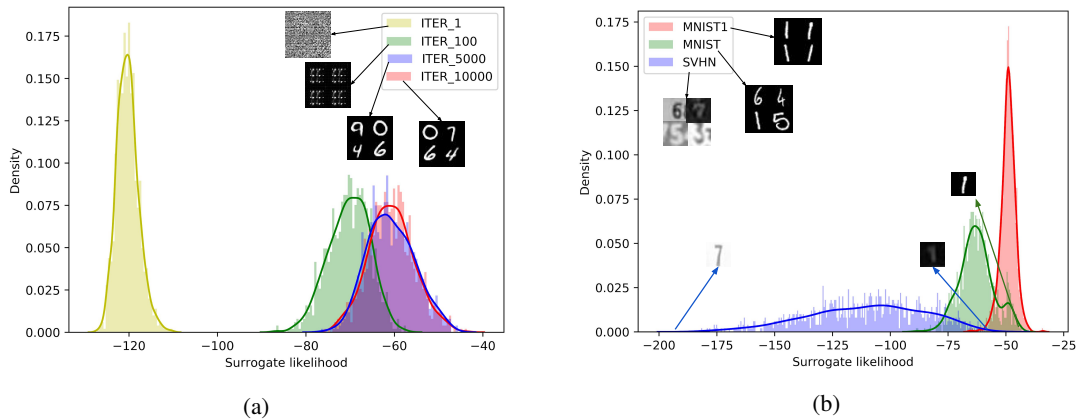


Figure 2: (a) Distributions of surrogate sample likelihoods at different iterations of entropic WGAN’s training using MNIST dataset. (b) Distributions of surrogate sample likelihoods of MNIST, MNIST-1 and SVHM datasets using a GAN trained on MNIST-1.




Figure 3: A visualization of density functions of $\mathbb{P}_{X|Y=\mathbf{y}^{test}}$ and $f_{X|Y=\mathbf{y}^{test}}$ for a random two-dimensional \mathbf{y}^{test} . Both distributions are very similar to one another making the approximation gap (i.e. $\text{KL}(\mathbb{P}_{X|Y=\mathbf{y}^{test}} || f_{X|Y=\mathbf{y}^{test}})$) very small. Our other experimental results presented in Table 1 are consistent with this result.

with high likelihood estimates correspond to images that are similar to MNIST digits, while samples with low scores are different than MNIST samples. Details of this experiment are presented in Supplementary material.³

Table 1: The tightness of the entropic GAN lower bound. Approximation gaps are orders of magnitudes smaller than the surrogate log-likelihoods. Results are averaged over 100 samples drawn from the underlying data distribution.

Data dimension	Approximation gap	Surrogate Log-Likelihood
2	9.3×10^{-4}	-4.15
5	4.7×10^{-2}	-15.35
10	6.2×10^{-2}	-46.3

³Training code available at https://github.com/yogeshbalaji/EntropicGANs_meet_VAEs

4.3. Tightness of the Variational Bound

In Theorem 1, we have shown that the Entropic GAN objective maximizes a lower-bound on the average sample log-likelihoods. This result has the same flavor as variational lower bounds used in VAEs, thus providing a connection between these two areas. One drawback of VAEs in general

Table 2: The tightness of the entropic GAN lower bound for non-linear generators.

Data dimension	Exact Log-Likelihood	Surrogate Log-Likelihood
5	-16.38	-17.94
10	-35.15	-43.6
20	-58.04	-66.58
30	-91.80	-100.69
64	-203.46	-217.52

is the lack of tightness analysis of the employed variational lower bounds. In this section, we aim to understand the tightness of the entropic GAN’s variational lower bound for some generative models.

4.3.1. LINEAR GENERATORS

Corollary 2 shows that the entropic GAN lower bound is tight when $\text{KL}(\mathbb{P}_{X|Y=y} \| f_{X|Y=y})$ approaches 0. Quantifying this term can be useful for assessing the quality of the proposed likelihood surrogate function. We refer to this term as the approximation gap.

Computing the approximation gap can be difficult in general as it requires evaluating $f_{X|Y=y}$. Here we perform an experiment for linear generative models and a quadratic loss function (same setting of Corollary 1). Let the real data Y be generated from the following underlying model

$$f_{Y|X=\mathbf{x}} \sim \mathcal{N}(\mathbf{G}\mathbf{x}, \lambda\mathbf{I})$$

where $X \sim \mathcal{N}(0, \mathbf{I})$

Using the Bayes rule, we have,

$$f_{X|Y^{test}} \sim \mathcal{N}(\mathbf{R}\mathbf{y}^{test}, \mathbf{I} - \mathbf{R}\mathbf{G})$$

where $\mathbf{R} = \mathbf{G}^T(\mathbf{G}\mathbf{G}^T + \lambda\mathbf{I})^{-1}$.

Since we have a closed-form expression for $f_{X|Y}$, $\text{KL}(\mathbb{P}_{X|Y=y} \| f_{X|Y=y})$ can be computed efficiently.

The matrix \mathbf{G} to generate Y is chosen randomly. Then, an entropic GAN with a linear generator and non-linear discriminators are trained on this dataset. $\mathbb{P}_{X|Y=y}$ is then computed using (4.3). Table 1 reports the average surrogate log-likelihood values and the average approximation gaps computed over 100 samples drawn from the underlying data distribution. We observe that the approximation gap is orders of magnitudes smaller than the log-likelihood values.

Additionally, in Figure 3, we demonstrate the density functions of $\mathbb{P}_{X|Y=y}$ and $f_{X|Y=y}$ for a random \mathbf{y} and a two-dimensional case ($r = 2$). In this figure, one can observe that both distributions are very similar to one another making the approximation gap very small.

Architecture and hyper-parameter details: For the generator network, we used 3 linear layers without any non-linearities ($2 \rightarrow 128 \rightarrow 128 \rightarrow 2$). Thus, it is an over-parameterized linear system. Over-parameterization was needed to improve convergence of the EntropicGAN training. The discriminator architecture (both D_1 and D_2) is a 2-layer MLP with ReLU non-linearities ($2 \rightarrow 128 \rightarrow 128 \rightarrow 1$). $\lambda = 0.1$ was used in all the experiments. Both generator and discriminator were trained using the Adam optimizer with a learning rate 10^{-6} and momentum 0.5. The discriminators were trained for 10 steps per generator iteration. Batch size of 512 was used.

4.3.2. NON-LINEAR GENERATORS

In this part, we consider the case of non-linear generators. The approximation gap $\text{KL}(\mathbb{P}_{X|Y=y} \| f_{X|Y=y})$ cannot be computed efficiently for non-linear generators as computing the optimal coupling $\mathbb{P}_{X|Y=y}$ is intractable. Instead, we demonstrate the tightness of the variational lower bound by comparing the exact data log-likelihood and the estimated lower-bound. As before, a d -dimensional Gaussian data distribution is used as the data distribution. The use of Gaussian distribution enables us to compute the exact data likelihood in closed-form. A table showing exact likelihood and the estimated lower-bound is shown in Table 2. We observe that the computed likelihood surrogate provides a good estimate to the exact data likelihood.

5. Conclusion

In this paper, we have provided a statistical framework for a family of GANs. Our main result shows that the entropic GAN optimization can be viewed as maximization of a variational lower-bound on average sample log-likelihoods, an approach that VAEs are based upon. This result makes a connection between two most-popular generative models, namely GANs and VAEs. More importantly, our result constructs an explicit probability model for GANs that can be used to compute a lower-bound on sample likelihoods. Our experimental results on various datasets demonstrate that this likelihood surrogate can be a good approximation of the true likelihood function. Although in this paper we mainly focus on understanding the behavior of the sample likelihood surrogate in different datasets, the proposed statistical framework of GANs can be used in various statistical inference applications. For example, our proposed likelihood surrogate can be used as a quantitative measure to evaluate the performance of different GAN architectures, it can be used to quantify the domain shifts, it can be used to select a proper generator class by balancing the bias term vs. variance, it can be used to detect outlier samples, it can be used in statistical tests such as hypothesis testing, etc. We leave exploring these directions for future work.

Acknowledgement

The Authors acknowledge support of the following organizations for sponsoring this work: (1) MURI from the Army Research Office under the Grant No. W911NF-17-1-0304, (2) Capital One Services LLC.

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- Cuturi, M. Sinkhorn distances: Lightspeed computation

- of optimal transport. In *Advances in neural information processing systems*, pp. 2292–2300, 2013.
- Feizi, S., Suh, C., Xia, F., and Tse, D. Understanding GANs: the LQG setting. *arXiv preprint arXiv:1710.10793*, 2017.
- Genevay, A., Peyré, G., and Cuturi, M. Sinkhorn-autodiff: Tractable wasserstein learning of generative models. *arXiv preprint arXiv:1706.00292*, 2017.
- Genevay, A., Peyre, G., and Cuturi, M. Learning generative models with sinkhorn divergences. In Storkey, A. and Perez-Cruz, F. (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1608–1617, Playa Blanca, Lanzarote, Canary Islands, 09–11 Apr 2018. PMLR.
- Ghahramani, A., Watt, F. M., and Luscombe, N. M. Generative adversarial networks uncover epidermal regulators and predict single cell perturbations. *bioRxiv*, pp. 262501, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Gregor, K., Danihelka, I., Graves, A., Rezende, D., and Wierstra, D. Draw: A recurrent neural network for image generation. In Bach, F. and Blei, D. (eds.), *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1462–1471, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/gregor15.html>.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of Wasserstein GANs. *arXiv preprint arXiv:1704.00028*, 2017.
- Guo, X., Hong, J., Lin, T., and Yang, N. Relaxed wasserstein with applications to GANs. *arXiv preprint arXiv:1705.07164*, 2017.
- Hirose, N., Sadeghian, A., Goebel, P., and Savarese, S. To go or not to go? a near unsupervised learning approach for robot navigation. *arXiv preprint arXiv:1709.05439*, 2017.
- Hu, Z., Yang, Z., Salakhutdinov, R., and Xing, E. P. On unifying deep generative models. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rylSzl-R->.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of gans for improved quality, stability, and variation. *CoRR*, abs/1710.10196, 2017. URL <http://arxiv.org/abs/1710.10196>.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Lee, H.-y. and Tsao, Y. Generative adversarial network and its applications to speech signal and natural language processing. *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018.
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*. PMLR, August 2017a.
- Mescheder, L., Nowozin, S., and Geiger, A. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. *arXiv preprint arXiv:1701.04722*, 2017b.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., and Ng, A. Y. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, pp. 5, 2011.
- Nowozin, S., Cseke, B., and Tomioka, R. f-GAN: training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*, pp. 271–279, 2016.
- Oord, A. V., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1747–1756, New York, New York, USA, 20–22 Jun 2016. PMLR. URL <http://proceedings.mlr.press/v48/oord16.html>.
- Petzka, H., Fischer, A., and Lukovnicov, D. On the regularization of Wasserstein GANs. *arXiv preprint arXiv:1709.08894*, 2017.
- Rigollet, P. and Weed, J. Entropic optimal transport is maximum-likelihood deconvolution. *Comptes Rendus Mathematique*, 356(11-12), 2018.

- Rosca, M., Lakshminarayanan, B., Warde-Farley, D., and Mohamed, S. Variational approaches for auto-encoding generative adversarial networks. *arXiv preprint arXiv:1706.04987*, 2017.
- Sanjabi, M., Ba, J., Razaviyayn, M., and Lee, J. D. Solving approximate Wasserstein GANs to stationarity. *Neural Information Processing Systems (NIPS)*, 2018.
- Sankaranarayanan, S., Balaji, Y., Castillo, C. D., and Chellappa, R. Generate to adapt: Aligning domains using generative adversarial networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Santana, E. and Hotz, G. Learning a driving simulator. *arXiv preprint arXiv:1608.01230*, 2016.
- Seguy, V., Damodaran, B. B., Flamary, R., Courty, N., Rolet, A., and Blondel, M. Large-scale optimal transport and mapping estimation. *arXiv preprint arXiv:1711.02283*, 2017.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- Villani, C. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- Wu, Y., Burda, Y., Salakhutdinov, R., and Grosse, R. B. On the quantitative analysis of decoder-based generative models. *CoRR*, abs/1611.04273, 2016.