
Learning Implicit Generative Models with the Method of Learned Moments

Suman Ravuri¹ Shakir Mohamed¹ Mihaela Rosca¹ Oriol Vinyals¹

Abstract

We propose a method of moments (MoM) algorithm for training large-scale implicit generative models. Moment estimation in this setting encounters two problems: it is often difficult to define the millions of moments needed to learn the model parameters, and it is hard to determine which properties are useful when specifying moments. To address the first issue, we introduce a *moment network*, and define the moments as the network’s hidden units and the gradient of the network’s output with respect to its parameters. To tackle the second problem, we use asymptotic theory to highlight desiderata for moments – namely they should minimize the asymptotic variance of estimated model parameters – and introduce an objective to *learn* better moments. The sequence of objectives created by this *Method of Learned Moments* (MoLM) can train high-quality neural image samplers. On CIFAR-10, we demonstrate that MoLM-trained generators achieve significantly higher Inception Scores and lower Fréchet Inception Distances than those trained with gradient penalty-regularized and spectrally-normalized adversarial objectives. These generators also achieve nearly perfect Multi-Scale Structural Similarity Scores on CelebA, and can create high-quality samples of 128×128 images.

1. Introduction

The method of moments (MoM) is an ancient principle of learning (Pearson, 1893; 1936). At its heart lies a simple procedure: given a model with parameters θ , estimate θ such that the moments — or more generally feature averages — of the model match those of the data. While the technique is simple and yields consistent estimators under weak conditions, other properties of the moment estimator are less desirable. Moment estimators are often biased, sometimes

lie outside the parameter space (such as negative probabilities), and, unless the model is in an exponential family, are less statistically efficient than maximum likelihood estimators. It is perhaps this last property that has relegated the method of moments to a niche technique.

There are, however, situations in which moment estimation is preferable to maximum likelihood estimation (MLE). One is when MLE is more computationally challenging than MoM. For example, one can scale training of Latent Dirichlet Allocation to large datasets by moment estimation using the first three order moments (Anandkumar et al., 2012a). Second, for latent variable models more generally, maximum likelihood estimation using the EM algorithm results in a local optimum, while MoM enjoys stronger formal guarantees (Anandkumar et al., 2012b). Finally, one can use moment estimation to determine model parameters in settings where likelihoods are unnatural. Instrumental variable estimation, an example of MoM, is used to learn parameters in supply and demand models (Wright, 1928).

We study another scenario: when data come from unknown or difficult-to-capture likelihood models. Data such as images, speech, or music often arise from complicated distributions, and for image data in particular, models based on likelihoods often yield low-quality samples. Researchers interested in generating more realistic samples have shifted their efforts to training neural network samplers with alternative losses. They have studied training these implicit generative models with the Wasserstein distance (Arjovsky et al., 2017), Maximum Mean Discrepancy (MMD) (Li et al., 2015; Dziugaite et al., 2015; Sutherland et al., 2016), and other divergences (Nowozin et al., 2016; Mao et al., 2016). While direct minimization of these distances or divergences – namely Wasserstein (Salimans et al., 2018), and MMD (Li et al., 2015; Dziugaite et al., 2015), Cramér (Bellemare et al., 2017) – in pixel space has led to poor sample quality, indirect minimization using adversarial training (Goodfellow et al., 2014) dramatically improves samples.

We pursue an alternative strategy: we explicitly define our moments and train them so that the moment estimators are statistically efficient. This choice creates two practical problems. The first is that traditional neural samplers, such as Deep Convolutional Generative Adversarial Networks (DCGANs) (Radford et al., 2015), have millions of parameters, and typically one needs at least one moment per parameter

¹DeepMind London, NIC 4AG, UK . Correspondence to: Suman Ravuri <ravuris@google.com>.

to train the model. Second, it is not obvious how to choose or train moments so that they are statistically efficient.

To address the first issue, we introduce a *moment network*, whose activations and gradients constitute the set of moments with which we train the generator. To tackle the second problem, we appeal to asymptotic theory to determine desiderata for moments (namely that they minimize the asymptotic variance of estimated model parameters) and explicitly specify and learn moments to train the generator. It is this theory that is used in the literature of the generalized method of moments in econometrics to reweight moments to make estimators more statistically efficient (Hansen, 1982; Hall, 2005).

We make the following contributions:

- We demonstrate that method of moments can scale to neural network models with tens of millions of parameters.
- We highlight the importance of statistical efficiency in method of moments and provide a method for learning moments such that moment estimators minimize asymptotic variance.
- We show that implicit generative models trained with our algorithm, the *Method of Learned Moments*, generate samples that are as good as, or better than, models that use adversarial learning, as measured by standard metrics.

2. The Method of Learned Moments

2.1. A Review of the Method of Moments

Suppose our data are drawn i.i.d. from $x_i \sim p^*$, and our samples s are drawn i.i.d. from a model p_θ , whose parameters $\theta \in \mathbb{R}^n$ we wish to learn. Method of Moments (MoM) estimation requires us to define feature functions $\Phi(x) \in \mathbb{R}^k$ and an associated moment function $m(\theta) := \mathbb{E}_{p_\theta(s)}[\Phi(s)]$. The moment estimator $\hat{\theta}_N$ matches the moment function with the feature average over the data:

$$m(\theta) = \frac{1}{N} \sum_{i=1}^N \Phi(x_i) \quad \theta = \hat{\theta}_N$$

If $p_\theta \stackrel{d}{=} p^*$ for some θ^* , the moments exist, and $m(\theta) \neq m(\theta^*) \quad \forall \theta \neq \theta^*$, this feature matching will yield a consistent estimator of θ^* . When we have access to the likelihood, we can recover the maximum likelihood estimate by setting $\Phi(x) = \nabla_\theta \log p_\theta(x)$ and noting that the expected value of the score function is zero at θ^* . Since maximum likelihood estimators are asymptotically efficient, they are generally preferable to their moment counterparts (Van der Vaart, 1998, Ch. 8).

With implicit generative models (Mohamed and Lakshminarayanan, 2016), we no longer have explicit access to the likelihood. This precludes straightforward application of

maximum likelihood. Instead, we have indirect access to p_θ through a parametric sampler $g_\theta(z) \sim p_\theta$. MoM estimation, however, is still applicable by replacing the moment function $m(\theta) := \mathbb{E}_{p_\theta(s)}[\Phi(s)]$ with $m(\theta) := \mathbb{E}_{p(z)}[\Phi(g_\theta(z))]$, where z is a draw from a prior distribution $p(z)$, such as a Gaussian or uniform distribution. If the generative model is sufficiently expressive to model the data distribution, the same regularity conditions on Φ will ensure a consistent estimator of generator parameters θ .

While consistency is guaranteed, the asymptotic efficiency argument of maximum likelihood implies that the specification of features $\Phi(x)$ affects the quality of the learned model parameters. One desirable aspect of the moment-matching framework is a developed asymptotic theory that provides large-sample behavior of the moment estimator. Intuitively, it tells us that our estimated generator parameters after N datapoints is roughly distributed as a Gaussian with mean θ^* and variance V/N . Minimizing V makes estimation of generator parameters more data-efficient and depends on the quality of $\Phi(x)$. The following theorem allows us to connect our choice of moments with its asymptotic variance.

Theorem 1 (Asymptotic Normality of Invertible Moment Functions). *Let $m(\theta) = \mathbb{E}_{p(z)}[\Phi(g_\theta(z))]$ be a one-to-one function on an open set $\Theta \subset \mathbb{R}^d$ and continuously differentiable at θ^* with nonsingular derivative $G = \nabla_\theta \mathbb{E}_{p(z)}[\Phi(g_{\theta^*}(z))]$. Then assuming $\mathbb{E}_{p(z)}[\|\Phi(g_{\theta^*}(z))\|^2] < \infty$, moment estimators $\hat{\theta}_N$ exist with probability tending to one and satisfy*

$$\begin{aligned} \sqrt{N}(\hat{\theta}_N - \theta^*) &\rightarrow \mathcal{N}(0, G^{-1}\Sigma G^{-T}) \\ \Sigma &:= \text{cov}(\Phi(g_{\theta^*}(z))) \end{aligned}$$

Proof. See Theorem 4.1 in Van der Vaart (1998) □

Minimizing this asymptotic covariance requires balancing G and Σ . G asserts that one should maximize the difference in features between the optimal generator parameters and those a small distance away, while Σ expresses that one should minimize the covariance of the features for the optimal generator parameters.

While this theorem requires the restrictive condition of invertible moment functions, the theorem in this ideal setting allows us to design statistically efficient moments. Moreover, in section 2.2.1 we later relax the assumptions of invertibility while showing the design choices still hold.

To obtain better moments, we will explicitly create parametric moments and optimize those moments to be more statistically efficient. This approach introduces two hurdles: 1) defining millions of sufficiently different moments and 2) creating an objective to learn desirable moments. Our practical contribution comprises how we solve these two issues.




Figure 1. Illustration of method of learned moments architecture.

2.2. From Theory to Practice

2.2.1. MOMENT SPECIFICATION

Moment specification is an exercise in ensuring consistency of the moment estimator. The assumption in Theorem 1 that the moment function is invertible is rather restrictive. One can instead weaken it to require identifiability: $m(\theta) = \mathbb{E}_{p^*(x)}[\Phi(x)]$ iff $\theta = \theta^*$. Even this condition, however, is difficult to verify in practice. We thus resort to a local identifiability assumption used in econometrics: that there are more moments than model parameters, and that G is full rank (Hall, 2005, Ch. 3). For linear models, this ensures global identifiability, and for nonlinear ones, this heuristic tends to work well in the literature.

Explicitly specifying moments for neural samplers seems especially daunting, however, since generally one needs at least as many moments as sampler parameters for moment matching to yield a consistent estimator of the model parameters.¹ For the DCGAN architecture with batch normalization, depending on the resolution of the dataset, the number of feature maps, and kernel size, one would need up to 20 million moments.

Our solution is to create a *moment network* $f_\phi(x)$ and define moments as:

$$\Phi(x) = [\nabla_\phi f_\phi(x), x, h_1(x) \dots h_{L-1}(x)]^\top$$

where $h_i(x)$ are the activations for layer i . $h_L(x)$ is implicitly included in the gradient. As long as the moment network has as many parameters as the generator, there are enough moments to train the model. Although we cannot ensure that G is full rank, we typically scale the moment network to produce between 1.5 and 5 times as many moments as generator parameters. We use both the gradients and activations since they encode different inductive biases. We discuss those biases in Section 2.5.

¹Moreover, neural samplers also suffer from identifiability issues. For example, one can express the same function by permuting nodes and associated parameters. We assume that, given an initialization θ_0 , only one θ^* is achievable using gradient descent, but we leave a more rigorous argument for future work.

Since the moment function is not invertible, we replace the moment-matching objective with the squared error loss between moments of the data and samples:

$$\mathcal{L}^G(\theta) = \frac{1}{2} \left\| \frac{1}{N} \sum_{i=1}^N \Phi(x_i) - \mathbb{E}_{p(z)}[\Phi(g_\theta(z))] \right\|_2^2 \quad (1)$$

Figure 1 illustrates our setup. In practice, we take Monte Carlo estimates of the sampler expectation. The change in objective function and Monte Carlo estimate induce a change in the asymptotics of this modified moment estimator. We address the effect of these changes in Section 2.3.

2.2.2. LEARNING EFFICIENT MOMENTS

It may seem plausible that given enough moments parameterized by random ϕ , the underlying objective is sufficiently good in practice to train neural samplers. Unfortunately, as we show in section 4, sample quality is poor. Theorem 1 allows us to diagnose the problem. G , the Jacobian of the moment estimator with respect to the optimal generator parameters is not sufficiently “large”. From the definition of the Jacobian:

$$G(\theta - \theta^*) = \mathbb{E}_{p(z)}[\Phi(g_\theta(z))] - \mathbb{E}_{p(z)}[\Phi(g_{\theta^*}(z))] + o(\|\theta - \theta^*\|)$$

For θ near the optimum, maximizing the difference in expected features also maximizes the “directional Jacobian” G . One does not expect, however, that moments produced for random ϕ to maximize this difference. This motivates learning ϕ .

Of course, since we do not have access to θ^* , exact computation of the asymptotic variance components G and Σ is impossible. Under the assumption that the generator is sufficiently expressive to represent the data distribution, then $p(x) \stackrel{d}{=} p(g_{\theta^*}(z))$, $\mathbb{E}_{p(z)}[\Phi(g_{\theta^*}(z))] = \mathbb{E}_{p(x)}[\Phi(x)]$ and $\Sigma = \text{cov}[\Phi(x)]$. We make a first approximation of maximizing this directional Jacobian:

$$G(\theta - \theta^*) \approx \mathbb{E}_{p(z)}[\Phi(g_\theta(z))] - \mathbb{E}_{p(x)}[\Phi(x)]$$

In early experiments optimizing this difference, moments became correlated and as a result the estimator of θ was no longer consistent. Thus, we make a second approximation. Inspired by the work of Jaakkola and Haussler (1999) in extracting feature vectors from auxiliary models for use in a linear SVM classifier, Tsuda et al. (2002) proposed the gradient of the log-odds ratio of a probabilistic binary classifier as a model-dependent feature. The authors empirically and theoretically show improved binary classification (and thus separability).

Applying this idea to our method, we perform binary classification on real images and our samples. Denoting

$D_\phi(x) = P(y = 1|x)$ as the probability that x is a real image, then the tangent of posterior odds is:

$$\nabla_\phi \log \frac{P(y = 1|x)}{P(y = -1|x)} = \nabla_\phi \log \frac{\sigma(f_\phi(x))}{1 - \sigma(f_\phi(x))} = \nabla_\phi f_\phi(x)$$

where σ is the logistic sigmoid function. When the log-odds ratio is included in the feature, the resulting kernel $K(x, y) = \Phi(x)^\top \Phi(y)$ is known as Tangent of Posterior Odds Kernel.

To help control Σ , we add a quadratic penalty on the squared norm of the minibatch gradient, so that the average squared moment is close to 1. The moment objective is now:

$$\mathcal{L}^M(\phi) = \mathbb{E}_{p(x)}[\log D_\phi(x)] + \mathbb{E}_{p(z)}[\log(1 - D_\phi(g_\theta(z)))] \\ + \lambda \left(\frac{\|\mathbb{E}_{p(x)}[\nabla_\phi f_\phi(x)]\|^2}{k} - 1 \right)^2$$

We do not include regularization on the hidden units as they represent a small percentage of our moments and its regularization yielded no difference in performance. A more correct penalty term is $(\frac{1}{k} \mathbb{E}_{p(x)}[\|\nabla_\phi f_\phi(x)\|^2] - 1)^2$ (to make sure the second moments are close to 1), but we found no performance benefit from the extra computational cost.

The upshot of learning moments is that each parameterization of f_ϕ , subject to regularity conditions, produces a consistent estimator of θ . On the other hand, each set of moments is not necessarily asymptotically efficient, as they rely on poor estimates of G . So, we learn better Φ iteratively (usually every 1,000-2,000 generator steps), and then match those moments. Algorithm 2.5 describes the proposed method.

2.3. Refinement of Asymptotic Theory

Note that while we defined the asymptotics for the feature-matching objective and used that to learn moments, our loss is actually the squared error objective in Equation 1. Although the same tradeoff applies to that objective, its asymptotic variance is somewhat more complicated. To develop the asymptotics of that expression, for clarity let us define the moment function:

$$m_N(x_{1,\dots,N}, \Phi, \theta) = \frac{1}{N} \sum_{i=1}^N \Phi(x_i) - \mathbb{E}_{p(z)}[\Phi(g_\theta(z))]$$

The asymptotics of the weighted squared loss function:

$$\mathcal{L}^G(\theta) = m_N(x_{1,\dots,N}, \Phi, \theta)^\top W m_N(x_{1,\dots,N}, \Phi, \theta)$$

are:

Theorem 2 (Asymptotic Normality of Squared Error Functions). *Under the consistency and asymptotic normality conditions in Appendix B.1, the estimator satisfies:*

$$\sqrt{N}(\hat{\theta}_N - \theta^*) \rightarrow \mathcal{N}(0, V_{SE}) \\ V_{SE} := (G^\top W G)^{-1} G^\top W \Sigma W G (G^\top W G)^{-1}$$

where $\Sigma := \text{cov}(\Phi(g_{\theta^*}(z)))$.

Proof. Theorem 3.2 of Hall (2005) □

When $W \propto \Sigma^{-1}$, then

$$\sqrt{N}(\hat{\theta} - \theta^*) \rightarrow \mathcal{N}(0, (G^\top \Sigma^{-1} G)^{-1})$$

It can be shown that this is the optimal weighting matrix. The inverse of this matrix is known as the Godambe Information Matrix (Godambe, 1960) and serves as a generalization of the Fisher Information Matrix.

Of course, the above theorem presupposes that we can analytically calculate $\mathbb{E}_{p(z)}[\Phi(g_\theta(z))]$. In implicit generative modeling, however, we only have access to $\frac{1}{K} \sum_{k=1}^K \Phi(g_\theta(z_k))$. It turns out we only pay a constant factor penalty for sampling. More specifically, suppose our moment function is now:

$$\hat{m}_N(x_{1,\dots,N}, \Phi, \theta) = \frac{1}{N} \sum_{i=1}^N \Phi(x_i) - \frac{1}{\mathcal{T}(N)} \sum_{k=1}^{\mathcal{T}(N)} \Phi(g_\theta(z_{i,k}))$$

where $\mathcal{T}(N)$ is the number of samples used to estimate generator moments for N points in the dataset. Then the asymptotic variance of this method, known as the simulated method of moments (Hall, 2005, Ch. 10), is:

Theorem 3 (Asymptotic Normality of Simulated Method of Moments). *Suppose that $\frac{\mathcal{T}(N)}{N} \rightarrow K$ as $N \rightarrow \infty$. Assuming the conditions in Appendix B.2, then $\hat{\theta}_N$ satisfies.*

$$\sqrt{N}(\hat{\theta}_N - \theta^*) \rightarrow \mathcal{N}\left(0, \left(1 + \frac{1}{K}\right) V_{SE}\right)$$

Proof. See Duffie and Singleton (1993) □

2.4. Computational Considerations

The gradient for the squared-error objective is:

$$\nabla_\theta \mathcal{L}^G(\theta) = \frac{1}{K} \sum_i J_i^\top \bar{m}(x, \Phi, \theta)$$

where $J_i := \nabla_\theta \Phi(g_\theta(z_i))$ is a Jacobian matrix and $\bar{m}(x, \Phi, \theta) = \frac{1}{N} \sum_{n=1}^N \Phi(x_n) - \frac{1}{K} \sum_{i=1}^K \Phi(g_\theta(z_i))$.

When using only gradient features, one can speed up gradient computation by $\sim 20\%$ by using a Hessian-vector product-like trick (Pearlmutter, 1994; Schraudolph, 2002).

Note that Hessian-vector products are defined as:

$$Hv = \left[\begin{array}{c|c} H_{\theta\theta} & J^\top \\ \hline J & H_{\phi\phi} \end{array} \right] [v]$$

Algorithm 1 Method of Learned Moments

Input: Learning rate α ; number of objectives N_o ; number of moment training steps N_m ; number of generator training steps N_g ; norm penalty parameter λ

Initialize generator and moment network parameters to θ and ϕ respectively

for $n = 0, \dots, N_o$ **do**

for $n = 1, \dots, N_m$ **do**

$d_\phi \leftarrow \nabla_\phi \mathcal{L}^M(\phi)$

$\phi_{t+1} \leftarrow \phi_t - \alpha \cdot \text{AdamOptimizer}(\phi_t, d_\phi)$

Calculate $\frac{1}{N} \sum_i \Phi(x_i)$ over the entire dataset.

for $n = 1, \dots, N_g$ **do**

$d_\theta \leftarrow \nabla_\theta \mathcal{L}^G(\theta)$

$\theta_{t+1} \leftarrow \theta_t - \alpha \cdot \text{AdamOptimizer}(\theta_t, d_\theta)$

Let v be defined as a partitioned vector as follows:

$$v = \begin{bmatrix} 0 \\ \bar{m}(x, \Phi, \theta) \end{bmatrix} \quad \text{Then: } Hv = \begin{bmatrix} J^\top \bar{m}(x, \Phi, \theta) \\ H_{\phi\phi} \bar{m}(x, \Phi, \theta) \end{bmatrix}$$

Performing a Hessian-vector computation only through the moment network provides the desired gradients.

2.5. Moment Architectures and Inductive Biases

When choosing a particular moment architecture, we implicitly specify the inductive biases of our features. Researchers typically design such architectures such that the forward pass encodes properties such as translational invariance and local receptive fields. These are properties of the hidden units. Gradients, however, often encode far different properties. For the convolutional moment architectures used in our experiments, gradient features encode global properties of the data. To see this, note that the output of a convolution is $y_{l,m,n} = \sum_a \sum_b \sum_c \phi_{a,b,c,n} x_{l+a,m+b,c}$. Its partial derivatives are $\frac{\partial y_{l,m,n}}{\partial \phi_{a,b,c,n}} = x_{l+a,m+b,c}$. Applying chain rule gives us the partial derivative with respect to $o = f_\phi(x)$

$$\frac{\partial f_\phi(x)}{\partial \phi_{a,b,c,n}} = \sum_l \sum_m \frac{\partial f_\phi(x)}{\partial y_{l,m,n}} x_{l+a,m+b,c}$$

Tying weights, which helps us encode local properties of the image in the forward pass, instead gives us global properties in the backward pass. Hence, we augment gradient features with hidden units to balance both local and global structure.

3. Connection to Other Methods

3.1. Maximum Mean Discrepancy

The method of moments probably bears the closest relationship Maximum Mean Discrepancy (MMD)². One can con-

²We assume the Reproducing Kernel Hilbert Space (RKHS) version of MMD; i.e., the function class is $\mathcal{F} = \{f \mid \|f\|_{\mathcal{H}} \leq 1\}$.

sider method of moments as embedding a probability distribution into a finite-dimensional vector. MMD, on the other hand, embeds a distribution into an infinite-dimensional vector. By enforcing $\phi \in \mathbb{L}^2$, one can calculate this ‘‘infinite-moment’’ matching loss as sum of expectations of kernels:

$$\begin{aligned} \text{MMD}^2(\theta) &= \sum_{i=1}^{\infty} (\mathbb{E}_{p(x)}[\phi_i(x)] - \mathbb{E}_{p(z)}[\phi_i(g_\theta(z))])^2 \\ &= \mathbb{E}[K(x, x')] - 2\mathbb{E}[K(x, g_\theta(z))] + \mathbb{E}[K(g_\theta(z), g_\theta(z'))] \end{aligned}$$

Furthermore, if the kernels are characteristic, then MMD² defines a squared distance (Gretton et al., 2012). Method of moments, on the other hand, is only able to distinguish between probability distributions specified by the model.

Despite robust theory, sample quality has lagged behind adversarial methods, especially if radial basis function (RBF) kernels are used. An explanation perhaps lies in the analysis of MMD² loss as spectral-domain moment matching.

Proposition 1. *Suppose the kernel function $K(x, y) = K(x - y)$ is real, shift-invariant, Bochner integrable, and without loss of generality $K(0)=1$. Then:*

$$\begin{aligned} &\mathbb{E}_{p(x,x')} [K(x, x')] - 2\mathbb{E}_{p(x,y)} [K(x, y)] + \mathbb{E}_{p(y,y')} [K(y, y')] \\ &= \mathbb{E}_{p(w)} [(\mathbb{E}_{p(x)}[\cos(\omega^\top x)] - \mathbb{E}_{p(y)}[\cos(\omega^\top y)])^2] \\ &+ \mathbb{E}_{p(w)} [(\mathbb{E}_{p(x)}[\sin(\omega^\top x)] - \mathbb{E}_{p(y)}[\sin(\omega^\top y)])^2] \end{aligned} \quad (2)$$

where $p(\omega)$ is a probability measure specified by the kernel.

Proof. See Appendix C.1 □

Crucially, for radial basis function kernels, $p(\omega) \propto \exp(-\frac{1}{2}\sigma^2\|\omega\|^2)$. For high-dimensional data, unless the data lie on a spherical shell of appropriate radius, then one likely needs many samples to accurately approximate MMD² distance.

It may be the poor spectral properties of the RBF kernels that have led to poorer samples. More recent work has focused on other kernels – such as sums of RBF kernels at different bandwidths (Sutherland et al., 2016) and rational quadratic kernels (Bińkowski et al., 2018) – and indeed using those kernels improved sample quality. In fact, the proposed Coulomb GAN (Unterthiner et al., 2017) shares a deeper relationship with MMD. It directly minimizes MMD loss using a version of the rational quadratic kernel known as the Plummer kernel to estimate f , and further introduces a discriminator to model the scaled witness function f^* . The upshot is that the generator loss approximates a high-sample biased estimate of MMD loss; see Appendix C.2 for details.

3.2. Adversarial Training

Recent work (Liu et al., 2017) has shown that in practical settings, many GAN objectives are better expressed as gen-

eralized moment matching as discriminators have limited capacity. Given this viewpoint, can asymptotic theory tell us anything about the statistical efficiency of adversarial networks?

Unfortunately, since there are an “infinite” number of moments, we cannot directly apply the above asymptotic theory. We do note that the inner maximization step indirectly “maximizes” G . For clarity, consider the inner maximization step of the Wasserstein GAN:

$$\mathcal{L}(\phi) = \min_{\theta} \max_{\phi} \mathbb{E}_{p(x)}[f_{\phi}(x)] - \mathbb{E}_{p(z)}[f_{\phi}(g_{\theta}(z))]$$

Under the assumption that the generator at θ^* is sufficiently expressive to model the data distribution, then $p(x) \stackrel{d}{=} p(g_{\theta^*}(z))$, $\mathbb{E}_{p(z)}[f_{\phi}(g_{\theta^*}(z))] = \mathbb{E}_{p(x)}[f_{\phi}(x)]$ and:

$$\mathbb{E}[f_{\phi}(g_{\theta}(z))] - \mathbb{E}[f_{\phi}(x)] \approx \mathbb{E}[\nabla_{\theta} f_{\phi}(g_{\theta^*}(z))^{\top} (\theta - \theta^*)]$$

While MoLM optimizes a directional Jacobian of many moments, adversarial training optimizes a directional derivative for a single moment.

Similarly, one can think of discriminator penalties – such as gradient penalties in Wasserstein GANs (Gulrajani et al., 2017), the DRAGAN penalty (Kodali et al., 2017), and average second moment penalty Fisher GANs (Mroueh and Sercu, 2017) – as terms to control Σ . Of the three, Fisher GAN most directly controls the second moment by placing a penalty on f_{ϕ}^2 , but with respect to a mixture distribution of data and samples. That penalties such as the gradient penalty improve other adversarial objectives suggests a deeper connection to controlling Σ .

Also implied by adversarial training is that it matches at most one moment per generator step; due to space constraints, we defer discussion of asymptotics to Appendix B.3.

It is perhaps the connection to statistical efficiency, rather than the choice of a particular distance, that explains the success of adversarial training. It may also explain why direct minimization of Wasserstein distance in the primal yields much poorer samples than minimization in the dual. Developing this hypothesis may better explain why adversarial training works so well.

3.2.1. MOMENT MATCHING IN ADVERSARIAL NETWORKS

Moment matching has also found its way into adversarial training. Salimans et al. (2016) introduced feature matching of activations to stabilize GAN training. Mroueh et al. (2017) match mean and covariance embeddings of a neural network. Mroueh and Sercu (2017) also matches mean embeddings, but constrains the singular value of the covariance matrix of a mixture distribution between data and samples

Table 1. Scores for different metrics for generators trained with random moments/MoLM. For Inception, higher scores are better; for FID, lower scores are better; and for MS-SSIM, scores closer to .379 are better.

Metric/Dataset	CelebA	CIFAR-10
Inception Score	-	2.10/ 6.99
FID	-	160.3/ 33.8
MS-SSIM	.444/ .378	-

to be less than one. Li et al. (2017) attempts to learn invertible and adversarial feature mappings such that one can use an RBF kernel for MMD. Salimans et al. (2018) and Bellemare et al. (2017) also learn adversarial feature mappings in combination with Wasserstein and Energy distance, respectively. Most of these proposals, however, introduce too few moments to train a generator, and require frequent updates of features. The possible exception is MMD GAN, which tries to learn an invertible feature map that can be used with a kernel distance. Invertibility is only enforced through regularization, and the feature mapping is likely not invertible in practice.

4. Experimental Results

We evaluate our method on four datasets: Color MNIST (Metz et al., 2016), CelebA (Liu et al., 2015), CIFAR-10 (Krizhevsky, 2009), and the daisy portion of ImageNet (Rusakovsky et al., 2015). We complement the visual inspection of samples with numerical measures to compare this method to existing work. For CelebA, we use Multi-Scale Structural Similarity (MS-SSIM) (Wang et al., 2003) to show sample similarity within a single class. Higher scores typically indicate mode collapse, while lower indicate higher diversity and better performance. Numbers lower than the test set may imply underfitting. For CIFAR-10, we include the standard Inception Score (IS) (Salimans et al., 2016) and Fréchet Inception Distance (FID) (Heusel et al., 2017).

We aim to answer three questions: 1) does learning moments improve sample quality, 2) what is the effect of including gradient and hidden unit features, and 3) how does sample quality compare to GAN alternatives?

We use convolutional architectures for both our generator and moment networks. To directly compare this algorithm to other methods, unless otherwise noted, we use a DCGAN generator. Direct comparisons using the same moment architecture as a discriminator makes less sense, however, since the set of moments used for adversarial learning come from an output of a network while our method uses low- and high-level information. Thus, we modify the architecture from a standard discriminator, though we only add size-preserving convolution before each stride-two layer of a DCGAN. For details of the specific architectures and hyperparameters used in all our experiments, please see Appendix A in the supplementary material. The models considered here are all




Figure 2. Top row are images from the the dataset (from left to right: Color MNIST, CIFAR-10, CelebA, and the Daisy portion of ImageNet at 128×128 resolution). The middle are samples trained with random moments. The bottom are sampled trained with MoLM.

unconditional.

To answer the first question – does learning moments improve sample quality –, we refer to Figure 2, which highlights the importance of learning moments. While random moments allow the generator to learn some structure of digits on Color MNIST, those moments are not sufficiently good to learn implicit generative models on other datasets. The sample quality is much higher for learned moments on all datasets. We also provide quantitative evidence for CIFAR-10 and CelebA: Table 1 shows a better Inception Score and Fréchet Inception Distance for learned moments compared to random moments on CIFAR-10, and a better MS-SSIM score on CelebA.

For the second question – what is the effect of including gradient and hidden unit features – we refer to Figure 4. For this experiment, we again focused on CIFAR-10 due to more robust metrics compared to other datasets. We tried four types of Adam hyperparameters, and two architectures. Across the board, we found that merely using activations did not work, which is not surprising as activations constitute roughly one-tenth the number of moments needed to train

the generator.³ Using only gradient features allows the model to learn more realistic samples. Using both, however, substantially improve IS and FID.

Finally, we compare Method of Learned Moments to Generative Adversarial Networks, and find MoLM performs as well as, if not better than, its GAN counterparts. On CelebA, shown in Table 2, the Multi-Scale Structural Similarity is as good as, or better than GAN alternatives. Admittedly, this metric is flawed as it only measures sample diversity and not quality of samples. At worst, however, the sample diversity of MoLM is comparable to GANs and the test set.

We find similar results on CIFAR-10. We try two convolutional architectures: the DCGAN, and one – denoted “Conv.” – recently introduced in Miyato et al. (2018). As shown in Table 3, MoLM significantly outperforms gradient penalty and spectrally-normalized GANs on both Inception Score and Fréchet Inception Distance using the Conv. architecture. It also outperforms MMD alternatives using the DCGAN

³Please refer to Figure 5 in the Supplementary Material for samples for different types of moments.




Figure 3. Left pane is Inception Score vs. and right is Fréchet Inception Distance vs. number of objectives solved for different size moment networks. For random moments, there is only a single objective but uses the same number of generator steps.




Figure 4. Inception Score and FID for different moment architectures and learning rates (right third). FID scores for just hiddens were not included because covariance matrix of the inception pool3 layer encountered rank deficiency.

Table 2. Multi-Scale Structural Similarity Results on CelebA for different methods.

Method	GAN	GAN-GP	DRAGAN	WGAN-GP	MoLM (ours)	Test Set
MS-SSIM	.381	.387	.383	.378	.378	.379

architecture, though the comparison is not entirely fair as the kernel sizes are different. Moreover, these latter results are fairly robust to increasing size of the moment network, shown in Figure 3. Inception Score and Fréchet Inception Distance also does not collapse over time.

To illustrate that the method can scale to higher-resolution images, we generate examples from the daisy portion of ImageNet at 128×128 resolution. GANs currently perform conditional image generation on the full dataset; unfortunately, no such conditional version of MoLM currently exists. This preliminary result, however, demonstrates the promise of the algorithm.

5. Discussion

We introduce a method of moments algorithm for training large-scale implicit generative models. We highlight the importance of learning moments and create a stable learning algorithm that performs better than adversarial alternatives. The current algorithm, however, leaves some room for improvement. For example, the moment architectures, slightly modified from discriminator architectures used in adversarial learning, are likely suboptimal. Moreover, the current

Table 3. Comparison of Inception Scores (IS) and FID using 5,000 generated images (-5K) and 50,000 generated images (-50K) for different convolutional architectures. \star from (Miyato et al., 2018). \circ from (Li et al., 2017). Coulomb GAN \dagger from (Unterthiner et al., 2017). Different methods are grouped by generator architecture.

Arch.	Method	IS	FID-5K/50K
4 \times 4 Conv.	GAN-GP \star	$6.93 \pm .11$	37.7/-
	WGAN-GP \star	$6.68 \pm .06$	40.2/-
	SN-GAN \star	$7.58 \pm .12$	25.5/-
	MoLM-1024	$7.55 \pm .08$	25.0/20.3
5 \times 5 DCGAN	MMD-RBF \circ	$3.47 \pm .03$	-/-
	MMD-GAN \circ	$6.17 \pm .07$	-/-
	Coul. GAN \dagger	-	-/27.3
4 \times 4 DCGAN	MoLM-768	$7.56 \pm .05$	31.4/27.3

learning of moments relies on a binary classification heuristic that can almost certainly be improved.

Finally, a connection between adversarial learning and statistical efficiency seems to exist, and exploring this relationship may help us better understand the quiddity of GANs.

Acknowledgments

We were incredibly lucky to have meaningful input that helped shape this work. We would like to thank David Warde-Farley, Aaron van den Oord, Razvan Pascanu, and the anonymous reviewers for trenchant insights, inspiration, and grammar corrections.

References

- A. Anandkumar, D. P. Foster, D. J. Hsu, S. M. Kakade, and Y.-K. Liu. A spectral algorithm for latent dirichlet allocation. In *Advances in Neural Information Processing Systems*, pages 917–925, 2012a.
- A. Anandkumar, D. J. Hsu, and S. M. Kakade. A method of moments for mixture models and hidden markov models. *CoRR*, abs/1203.0683, 2012b. URL <http://arxiv.org/abs/1203.0683>.
- M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein GAN. *arXiv preprint arXiv:1701.07875*, 2017.
- S. Barratt and R. Sharma. A Note on the Inception Score. *ArXiv e-prints*, Jan. 2018.
- M. G. Bellemare, I. Danihelka, W. Dabney, S. Mohamed, B. Lakshminarayanan, S. Hoyer, and R. Munos. The cramer distance as a solution to biased wasserstein gradients. *arXiv preprint arXiv:1705.10743*, 2017.
- M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- D. Duffie and K. J. Singleton. Simulated moments estimation of markov models of asset prices. *Econometrica*, 61(4):929–952, 1993. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2951768>.
- G. K. Dziugaite, D. M. Roy, and Z. Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. *arXiv preprint arXiv:1505.03906*, 2015.
- V. P. Godambe. An optimum property of regular maximum likelihood estimation. *The Annals of Mathematical Statistics*, 31(4):1208–1211, 1960.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 2672–2680, 2014.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pages 5769–5779, 2017.
- A. R. Hall. *Generalized method of moments*. Oxford University Press, 2005.
- C. Han and R. De Jong. Closest moment estimation under general conditions. *Annales d'économie et de statistique*, pages 1–13, 2004.
- L. P. Hansen. Large sample properties of generalized method of moments estimators. *Econometrica: Journal of the Econometric Society*, pages 1029–1054, 1982.
- M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, G. Klambauer, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a nash equilibrium. *arXiv preprint arXiv:1706.08500*, 2017.
- T. Jaakkola and D. Haussler. Exploiting generative models in discriminative classifiers. In *Advances in neural information processing systems*, pages 487–493, 1999.
- N. Kodali, J. Abernethy, J. Hays, and Z. Kira. On convergence and stability of gans. *arXiv preprint arXiv:1705.07215*, 2017.
- A. Krizhevsky. Learning multiple layers of features from tiny images. 2009.
- C.-L. Li, W.-C. Chang, Y. Cheng, Y. Yang, and B. Póczos. Mmd gan: Towards deeper understanding of moment matching network. *arXiv preprint arXiv:1705.08584*, 2017.
- Y. Li, K. Swersky, and R. Zemel. Generative moment matching networks. In *International Conference on Machine Learning*, pages 1718–1727, 2015.
- S. Liu, O. Bousquet, and K. Chaudhuri. Approximation and convergence properties of generative adversarial learning. In *Advances in Neural Information Processing Systems*, pages 5551–5559, 2017.
- Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley. Least squares generative adversarial networks. *arXiv preprint ArXiv:1611.04076*, 2016.
- L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.

- T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BlQRgziT->.
- S. Mohamed and B. Lakshminarayanan. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Y. Mroueh and T. Sercu. Fisher GAN. *CoRR*, abs/1705.09675, 2017. URL <http://arxiv.org/abs/1705.09675>.
- Y. Mroueh, T. Sercu, and V. Goel. Mcgan: Mean and covariance feature matching gan. *arXiv preprint arXiv:1702.08398*, 2017.
- S. Nowozin, B. Cseke, and R. Tomioka. f -GAN: Training generative neural samplers using variational divergence minimization. *arXiv preprint arXiv:1606.00709*, 2016.
- B. A. Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1):147–160, 1994.
- K. Pearson. Asymmetrical frequency curves. *Nature*, 48 (1252):615, 1893.
- K. Pearson. Method of moments and method of maximum likelihood. *Biometrika*, 28(1/2):34–59, 1936.
- A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- M. Rosca, B. Lakshminarayanan, D. Warde-Farley, and S. Mohamed. Variational approaches for auto-encoding generative adversarial networks. *CoRR*, abs/1706.04987, 2017. URL <http://arxiv.org/abs/1706.04987>.
- O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. *arXiv preprint arXiv:1606.03498*, 2016.
- T. Salimans, H. Zhang, A. Radford, and D. Metaxas. Improving GANs using optimal transport. *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rkQkBnJAb>.
- N. Schraudolph. Fast curvature matrix-vector products for second-order gradient descent. *Neural computation*, 14 (7):1723–1738, 2002.
- D. J. Sutherland, H.-Y. Tung, H. Strathmann, S. De, A. Ramdas, A. Smola, and A. Gretton. Generative models and model criticism via optimized maximum mean discrepancy. *arXiv preprint arXiv:1611.04488*, 2016.
- K. Tsuda, M. Kawanabe, G. Rätsch, S. Sonnenburg, and K.-R. Müller. A new discriminative kernel from probabilistic models. In *Advances in Neural Information Processing Systems 14*, pages 977–984, Cambridge, MA, USA, Sept. 2002. Max-Planck-Gesellschaft, MIT Press.
- T. Unterthiner, B. Nessler, G. Klambauer, M. Heusel, H. Ramsauer, and S. Hochreiter. Coulomb gans: Provably optimal nash equilibria via potential fields. *arXiv preprint arXiv:1708.08819*, 2017.
- A. W. Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 1998.
- Z. Wang, E. P. Simoncelli, and A. C. Bovik. Multiscale structural similarity for image quality assessment. In *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Seventh Asilomar Conference on*, volume 2, pages 1398–1402. IEEE, 2003.
- P. G. Wright. *The Tariff on Animal and Vegetable Oils*. The Macmillan Company, 1928.
- J. Zhao and D. Meng. Fastmmd: Ensemble of circular discrepancy for efficient two-sample test. *Neural computation*, 27(6):1345–1372, 2015.