
Metropolis-Hastings Generative Adversarial Networks

Ryan Turner¹ Jane Hung¹ Eric Frank¹ Yunus Saatci¹ Jason Yosinski¹

Abstract

We introduce the Metropolis-Hastings generative adversarial network (MH-GAN), which combines aspects of Markov chain Monte Carlo and GANs. The MH-GAN draws samples from the distribution implicitly defined by a GAN’s discriminator-generator pair, as opposed to standard GANs which draw samples from the distribution defined only by the generator. It uses the discriminator from GAN training to build a wrapper around the generator for improved sampling. With a perfect discriminator, this wrapped generator samples from the true distribution on the data exactly even when the generator is imperfect. We demonstrate the benefits of the improved generator on multiple benchmark datasets, including CIFAR-10 and CelebA, using the DCGAN, WGAN, and progressive GAN.

1. Introduction

Traditionally, density estimation is done with a model that can compute the data likelihood. Generative adversarial networks (GANs) (Goodfellow et al., 2014) present a radically new way to do density estimation: They implicitly represent the density of the data via a classifier that distinguishes real from generated data.

GANs iterate between updating a discriminator D and a generator G , where G generates new (synthetic) samples of data, and D attempts to distinguish samples of G from the real data. In the typical setup, D is thrown away at the end of training, and only G is kept for generating new synthetic data points. In this work, we propose the Metropolis-Hastings GAN (MH-GAN), a GAN that constructs a new generator G' that “wraps” G using the information contained in D . This principle is illustrated in Figure 1.¹

¹Uber AI Labs. Correspondence to: Ryan Turner <ryan.turner@uber.com>.

Proceedings of the 36th International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019. Copyright 2019 by the author(s).

¹Code found at:
github.com/uber-research/metropolis-hastings-gans

The MH-GAN uses Markov chain Monte Carlo (MCMC) methods to sample from the distribution implicitly defined by the discriminator D learned for the generator G . This is built upon the notion that the discriminator classifies between the generator G and a data distribution:

$$D(\mathbf{x}) = \frac{p_D(\mathbf{x})}{p_D(\mathbf{x}) + p_G(\mathbf{x})}, \quad (1)$$

where p_G is the (intractable) density of samples from the generator G , and p_D is the data density implied by the discriminator D with respect to G . If GAN training reaches its global optimum, then this discriminator distribution p_D is equal to the data distribution and the generator distribution ($p_D = p_{\text{data}} = p_G$) (Goodfellow et al., 2014). Furthermore, if the discriminator D is optimal for a fixed imperfect generator, G then the implied distribution still equals the data distribution ($p_D = p_{\text{data}} \neq p_G$).

We use an MCMC *independence sampler* (Tierney, 1994) to sample from p_D by taking multiple samples from G . Amazingly, using our algorithm, one can show that given a perfect discriminator D and a decent (but imperfect) generator G , one can obtain exact samples from the true data distribution p_{data} . Standard MCMC implementations require (unnormalized) densities for the target p_D and the proposal p_G , which are both unavailable for GANs. However, the Metropolis-Hastings (MH) algorithm requires only the ratio:

$$\frac{p_D(\mathbf{x})}{p_G(\mathbf{x})} = \frac{D(\mathbf{x})}{1 - D(\mathbf{x})}, \quad (2)$$

which we can obtain using only evaluation of $D(\mathbf{x})$.

Sampling from an MH-GAN is more computationally expensive than a standard GAN, but the bigger and more relevant training compute cost remains unchanged. Thus, the MH-GAN is best suited for applications where sample quality is more important than compute speed at test time.

The outline of this paper is as follows: Section 2 reviews diverse areas of relevant prior work. In Sections 3.1 and 3.2 we explain the necessary background on MCMC methods and GANs. We explain our methodology of combining these two seemingly disparate areas in Section 4 where we derive the wrapped generator G' . Results on real data (CIFAR-10 and CelebA) and extending common GAN models (DCGAN, WGAN, and progressive GAN) are shown in Section 5. Section 6 discusses implications and conclusions.

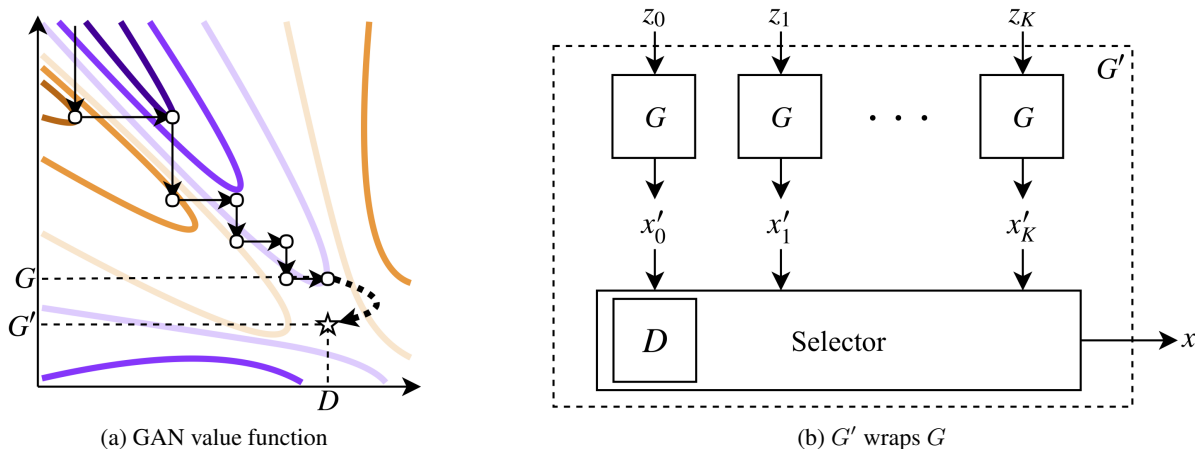


Figure 1. (a) We diagram how training of D and G in GANs performs coordinate descent on the joint minimax value function, shown in the solid black arrow. If GAN training produces a perfect D for an imperfect G , the MH-GAN wraps G to produce a perfect generator G' , as shown in the final dashed arrow. The generator G moves vertically towards the orange region while the discriminator D moves horizontally towards the purple. (b) We illustrate how the MH-GAN is essentially a selector from multiple draws of G . In the MH-GAN, the selector is built using a Metropolis-Hastings (MH) acceptance rule from the discriminator scores D .

2. Related Work

A few other works combine GANs and MCMC in some way. Song et al. (2017) use a GAN-like procedure to train a RealNVP (Dinh et al., 2016) MCMC proposal for sampling an externally provided target p^* . Whereas Song et al. (2017) use GANs to accelerate MCMC, we use MCMC to enhance the samples from a GAN. Similar to Song et al. (2017), Kempinska & Shawe-Taylor (2017) improve proposals in particle filters rather than MCMC. Song et al. (2017) was recently generalized by Neklyudov et al. (2018).

2.1. Discriminator Rejection Sampling

A concurrent work with similar aims from Azadi et al. (2018) proposes discriminator rejection sampling (DRS) for GANs, which performs rejection sampling on the outputs of G by using the probabilities given by D . While conceptually appealing at first, DRS suffers from two major shortcomings in practice. First, it is necessary to find an upper-bound on D over all possible samples in order to obtain a valid proposal distribution for rejection sampling. Because this is not possible, one must instead rely on estimating this bound by drawing many pilot samples. Secondly, even if one were to find a good bound, the acceptance rate would become very low due to the high-dimensionality of the sampling space. This leads Azadi et al. (2018) to use an extra γ heuristic to shift the logit D scores, making the model sample from a distribution different from p_{data} even when D is perfect. We use MCMC instead, which was invented precisely as a replacement for rejection sampling in higher dimensions. We further improve the robustness of MCMC via use of a *calibrator* on the discriminator to get more accurate probabilities for computing acceptance.

3. Background and Notation

In this section, we briefly review the notation and equations with MCMC and GANs.

3.1. MCMC Methods

MCMC methods attempt to draw a chain of samples $\mathbf{x}_{1:K} \in \mathcal{X}^K$ that marginally come from a target distribution p^* . We refer to the initial distribution as p_0 and the proposal for the independence sampler as $\mathbf{x}' \sim q(\mathbf{x}'|\mathbf{x}_k) = q(\mathbf{x}')$. The proposal $\mathbf{x}' \in \mathcal{X}$ is accepted with probability

$$\alpha(\mathbf{x}', \mathbf{x}_k) = \min \left(1, \frac{p^*(\mathbf{x}')q(\mathbf{x}_k)}{p^*(\mathbf{x}_k)q(\mathbf{x}')} \right) \in [0, 1]. \quad (3)$$

If \mathbf{x}' is accepted, $\mathbf{x}_{k+1} = \mathbf{x}'$, otherwise $\mathbf{x}_{k+1} = \mathbf{x}_k$. Note that when estimating the distribution p^* , one must include the duplicates that are a result of rejections in \mathbf{x}' .

Independent samples Many evaluation metrics assume perfectly iid samples. Although MCMC methods are typically used to produce correlated samples, we can produce iid samples by using one chain per sample: Each chain samples $\mathbf{x}_0 \sim p_0$ and then does K MH iterations to get \mathbf{x}_K as the output of the chain, which is the output of G' . Using multiple chains is also better for GPU parallelization.

Detailed balance The detailed balance condition implies that if $\mathbf{x}_k \sim p^*$ exactly then $\mathbf{x}_{k+1} \sim p^*$ exactly as well. Even if \mathbf{x}_k is not exactly distributed according to p^* , the Kullback-Leibler (KL) divergence between the implied density it is drawn from and p^* always decreases as k increases (Murray & Salakhutdinov, 2008). We use detailed balance to motivate our approach to MH-GAN initialization.

3.2. GANs

GANs implicitly model the data \mathbf{x} via a synthetic data generator $G \in \mathbb{R}^d \rightarrow \mathcal{X}$:

$$\mathbf{x} = G(\mathbf{z}), \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d). \quad (4)$$

This implies a (intractable) distribution on the data $\mathbf{x} \sim p_G$. We refer to the unknown true distribution on the data \mathbf{x} as p_{data} . The discriminator $D \in \mathcal{X} \rightarrow [0, 1]$ is a soft classifier predicting if a data point is real as opposed to being sampled from p_G .

If D converges optimally for a fixed G , then $D = p_{\text{data}}/(p_{\text{data}} + p_G)$, and if both D and G converge then $p_G = p_{\text{data}}$ (Goodfellow et al., 2014). GAN training forms a game between D and G . In practice D is often better at estimating the density ratio than G is at generating high-fidelity samples (Shibuya, 2017). This motivates wrapping an imperfect G to obtain an improved G' by using the density ratio information contained in D .

4. Methods

In this section we show how to sample from the distribution p_D implied by the discriminator D . We apply (2) and (3) for a target of $p^* = p_D$ and proposal $q = p_G$:

$$\frac{p_D}{p_G} = \frac{1}{D^{-1} - 1} \quad (5)$$

$$\implies \alpha(\mathbf{x}', \mathbf{x}_k) = \min \left(1, \frac{D(\mathbf{x}_k)^{-1} - 1}{D(\mathbf{x}')^{-1} - 1} \right). \quad (6)$$

The ratio p_D/p_G is computed entirely from the discriminator scores D . If D is perfect, $p_D = p_{\text{data}}$, so the sampler will marginally sample from p_{data} . The use of (6) is further illustrated in Algorithm 1.

A toy one-dimensional example with just such a perfect discriminator is shown in Figure 2. In this example the MH-GAN is able to correctly reconstruct a missing mode in the generating distribution from the tail of a faulty generator.

Calibration The probabilities for D must not merely provide a good AUC score, but must also be well *calibrated*. In other words, if one were to warp the probabilities of the perfect discriminator in (1) it may still suffice for standard GAN training, but it will not work in the MCMC procedure defined in (6), as it will result in erroneous density ratios.

We can demonstrate the miscalibration of D using the statistic of Dawid (1997) on held out samples $\mathbf{x}_{1:N}$ and real/fake labels $y_{1:N} \in \{0, 1\}^N$. If D is well calibrated, i.e., y is indistinguishable from a $y \sim \text{Bern}(D(\mathbf{x}))$, then

$$Z = \frac{\sum_{i=1}^N y_i - D(\mathbf{x}_i)}{\sqrt{\sum_{i=1}^N D(\mathbf{x}_i)(1 - D(\mathbf{x}_i))}} \implies Z \sim \mathcal{N}(0, 1). \quad (7)$$

That is, we expect the Z diagnostic to be a Gaussian in large N for any well-calibrated classifier. This means that for large values of Z , such as when $|Z| > 2$, we reject the hypothesis that D is well-calibrated.

Correcting Calibration While (7) may tell us a classifier is poorly calibrated, we also need to be able to fix it. Furthermore, some GANs (like WGAN) require calibration because their discriminator only outputs a score and not a probability. To correct an uncalibrated classifier, denoted $\tilde{D} \in \mathcal{X} \rightarrow \mathbb{R}$, we use a *held out* calibration set (e.g., 10% of the training data) and either logistic, isotonic, or beta (Kull et al., 2017) regression to warp the output of \tilde{D} . The held out calibration set contains an equal number of positive and negative examples, which in the case of GANs is an even mix of real samples and fake samples from G . After \tilde{D} is learned, we train a probabilistic classifier $C \in \mathbb{R} \rightarrow [0, 1]$ to map $\tilde{D}(\mathbf{x}_i)$ to y_i using the calibration set. The calibrated classifier is built via $D(\mathbf{x}_i) = C(\tilde{D}(\mathbf{x}_i))$.

Initialization We also avoid the burn-in issues that usually plague MCMC methods. Recall that via the detailed balance property (Gilks et al., 1996, Ch. 1), if the marginal distribution of a Markov chain state $\mathbf{x} \in \mathcal{X}$ at time step k matches the target p_D ($\mathbf{x}_k \sim p_D$), then the marginal at time step $k + 1$ will also follow p_D ($\mathbf{x}_{k+1} \sim p_D$). In most MCMC applications it is not possible to get an initial sample from the target distribution ($\mathbf{x}_0 \sim p_D$).

However, for MH-GAN, we have access to real data from the target distribution. By initializing the chain at a sample of real data (the correct distribution), we apply the detailed balance property and avoid burn-in. If no generated sample is accepted by the end of the chain, we restart sampling from a synthetic sample to ensure the initial real sample is never output. To make restarts rare, we set K large (often 640).

Using a restart after an MCMC chain of only rejects has a theoretical potential for bias. However, MCMC in practice often uses chain diagnostics as a stopping criterion, which suffers the same bias potential (Cowles et al., 1999). Alternatively, we could never restart and always report the state after K samples, which will occasionally include the initial real sample. This might be a better approach in certain statistical problems, where we care more about eliminating any potential source of bias, than in image generation.

Perfect Discriminator The assumption of a perfect D may be weakened for two reasons: (A) Because we recalibrate the discriminator, the actual probabilities can be incorrect as long as the decision boundary between real and fake is correct. (B) Because the discriminator is only ever evaluated at samples from G or the initial real sample \mathbf{x}_0 , D only needs to be accurate on the manifold of samples from the generator p_G and the real data p_{data} .




Figure 2. Illustration comparing the MH-GAN setup with the formulation of DRS on a univariate example. This figure uses a p_{data} of four Gaussian mixtures while p_G is missing one of the mixtures. The top row shows the resulting density of samples, while the bottom row shows the typical number of rejects before accepting a sample at that x value. The MH-GAN recovers the true density except in the far right tail where there is an exponentially small chance of getting a sample from the proposal p_G . DRS with $\gamma = 0$ shift should also be able to recover the density exactly, but it has an even larger error in the right tail. These errors arise because DRS must approximate the max D score and use only 10,000 pilot samples to do so, as in Azadi et al. (2018). Additionally, due to the large maximum D , it needs a large number of draws before a single accept. DRS with γ shift is much more sample efficient, but completely misses the right mode as the setup invalidates the rejection sampling equations. The MH-GAN is more adaptive in that it quickly accepts samples for the areas p_G models well; more MCMC rejections occur before accepting a sample in the right poorly modeled mode. In all cases the MH-GAN is more efficient than DRS without γ shift. Presumably, this effect becomes greater in high dimensions.

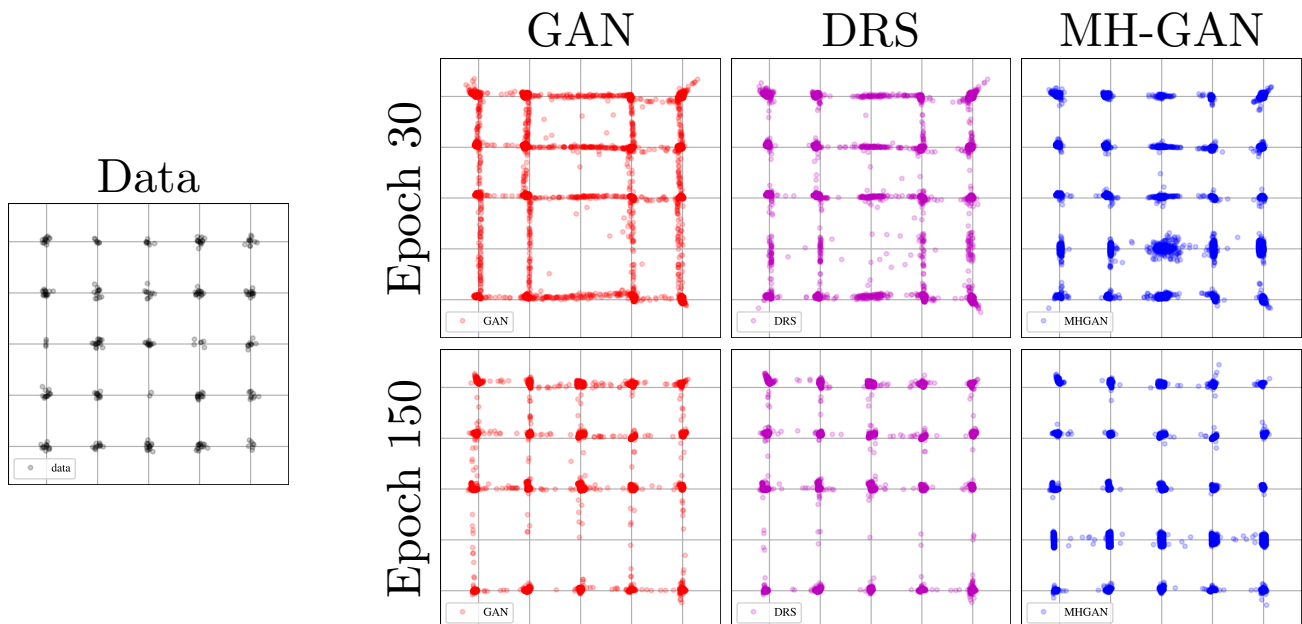


Figure 3. The 25 Gaussians example. We show the state of the generators at epoch 30 (when MH-GAN begins showing large gains) on the top row and epoch 150 (the final epoch) on the bottom row. The MH-GAN corrects areas of mis-assigned mass in the original GAN. DRS appears visually closer to the original GAN than the data, whereas the MH-GAN appears closer to the actual data.

Algorithm 1 MH-GAN

Input: generator G , calibrated disc. D , real samples
Assign random real sample \mathbf{x}_0 to \mathbf{x}
for $k = 1$ **to** K **do**
 Draw \mathbf{x}' from G
 Draw U from Uniform(0, 1)
 if $U \leq (D(\mathbf{x})^{-1} - 1)/(D(\mathbf{x}')^{-1} - 1)$ **then**
 $\mathbf{x} \leftarrow \mathbf{x}'$
 end if
end for
If \mathbf{x} is still real sample \mathbf{x}_0 restart with draw from G as \mathbf{x}_0
Output: sample \mathbf{x} from G'

5. Results

We first show an illustrative synthetic mixture model example followed by real data with images.

5.1. Mixture of 25 Gaussians

We consider the 5×5 grid of two-dimensional Gaussians used in Azadi et al. (2018), which has become a popular toy example in the GAN literature (Dumoulin et al., 2016). The means are arranged on the grid $\mu \in \{-2, -1, 0, 1, 2\}$ and use a standard deviation of $\sigma = 0.05$.

Experimental setup Following Azadi et al. (2018), we use four fully connected layers with ReLU activations for both the generator and discriminator. The final output layer of the discriminator is a sigmoid, and no nonlinearity is applied to the final generator layer. All hidden layers have size 100, with a latent $\mathbf{z} \in \mathbb{R}^2$. We used 64,000 standardized training points and generated 10,000 points in test.

Visual results In Figure 3, we show the original data along with samples generated by the GAN. We also show samples enhanced via the MH-GAN (with calibration) and with DRS. The standard GAN creates spurious links along the grid lines between modes and misses some modes along the bottom row. DRS is able to reduce some of the spurious links but not fill in the missing modes. The MH-GAN further reduces the spurious links and recovers these underestimated modes.

Quantitative results These results are made more quantitative in Figure 4, where we follow some of the metrics for the example from Azadi et al. (2018). We consider the standard deviations within each mode in Figure 4a and the rate of “high quality” samples in Figure 4b. A sample is assigned to a mode if its L_2 distance is within four standard deviations ($\leq 4\sigma = 0.2$) of its mean. Samples within four standard deviations of any mixture component are considered “high quality”. The within standard deviation plot (Figure 4a)

shows a slight improvement for MH-GAN, and the high quality sample rate (Figure 4b) approaches 100% faster for the MH-GAN than the GAN or DRS.

To test the spread of the distribution, we inspect the categorical distribution of the closest mode. Far away (non-high quality) samples are assigned to a 26th unassigned category. This categorical distribution should be uniform over the 25 real modes for a perfect generator. To assess generator quality, we look at the Jensen-Shannon divergence (JSD) between the sample mode distribution and a uniform distribution. This is a much more stringent test of appropriate spread of probability mass than checking if a single sample is produced near a mode (as in Azadi et al. (2018)).

In Figure 4c, we see that the MH-GAN improves the JSD over DRS by $5\times$ on average, meaning it achieves a much more balanced spread across modes. DRS fails to make gains after epoch 30. Using the principled approach of the MH-GAN along with calibrated probabilities ensures a correct spread of probability mass.

5.2. Real Data

For real data experiments we considered the CelebA (Liu et al., 2015) and CIFAR-10 (Torralla et al., 2008) data sets modeled using the DCGAN (Radford et al., 2015) and WGAN (Arjovsky et al., 2017; Gulrajani et al., 2017). To evaluate the generator G' , we plot Inception scores (Salimans et al., 2016) per epoch in Figure 5a after $k = 640$ MCMC iterations. Figure 5b shows Inception score per MCMC iteration: most gains are made in the first $k = 100$ iterations, but gains continue to $k = 400$. This shows that the MH-GAN allows a tunable trade-off between sample quality and computation cost.

In Table 1, we summarize performance (Inception score) across all experiments, running MCMC to $k = 640$ iterations in all cases. Behavior is qualitatively similar to that in Figure 5a. While DRS improves on a direct GAN, MH-GAN improves Inception score more in every case. Calibration helps in every case; and we found a slight advantage for isotonic regression over other calibration methods. Results are computed at epoch 60, and as in Figure 5a, error bars and p-values are computed using a paired t-test across Inception score batches. All results are significantly better than the baseline GAN at $p < 0.05$.

Score distribution In Figure 5c, we show what G' does to the distribution on discriminator scores. MCMC shifts the distribution of the fakes to match the distribution on true images. We also observed that the MH acceptance rate is primarily determined by the overlap of the distributions on D scores between real and fake samples. If the AUC of D is less than 0.90 we see acceptance rates over 20%; but when the AUC of D is 0.95, acceptance rates drop to 10%.

Metropolis-Hastings Generative Adversarial Networks

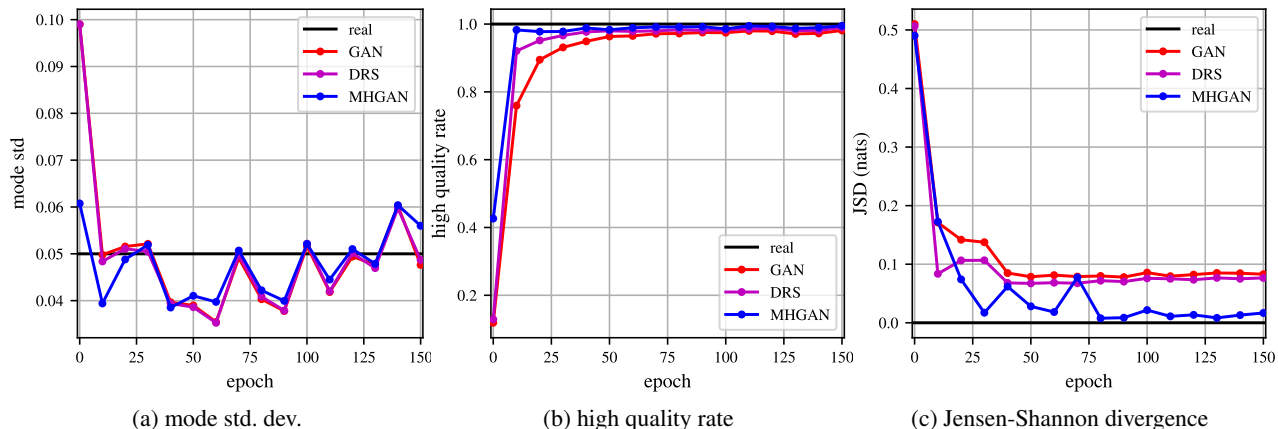


Figure 4. Results of the MH-GAN experiments on the mixture of 25 Gaussians example. (a) On the left, we show the standard deviation of samples within a single mode. The black lines represent values for the true distribution. (b) In the center, we show the high quality rate (samples near a real mode) across different GAN setups. (c) On the right, we show the Jensen-Shannon divergence (JSD) between the distribution on the nearest mode vs a uniform, which is the generating distribution on mixture components. The MH-GAN shows, on average, a $5\times$ improvement in JSD over DRS. We considered adding error bars to these plots via a bootstrap analysis, but the error bars are too small to be visible.




Figure 5. Results of the MH-GAN experiments on CIFAR-10 using the DCGAN. (a) On the left, we show the Inception score vs. training epoch of the DCGAN with $k = 640$ MH iterations. MH-GAN denotes using the raw discriminator scores and MH-GAN (cal) for the calibrated scores. The error bars on MH-GAN performance (in gray) are computed using a t-test on the variation per batch across 80 splits of the Inception score. (b) In the center we show the Inception score vs. number of MCMC iterations k for the GAN at epoch 15. (c) On the right, we show the scores at epoch 13 where there is some overlap between the scores of fake and real images. When there is overlap, the MH-GAN corrects the p_G distribution to have scores looking similar to the real data. DRS fails to fully shift the distribution because 1) it does not use calibration and 2) its “ γ shift” setup violates the validity of rejection sampling.

Table 1. Results showing Inception score improvements from MH-GAN on DCGAN and WGAN at epoch 60. Like Figure 5a, the error bars and p-values are computed using a paired t-test across Inception score batches (higher is better). All results except for DCGAN on CelebA are significant at $p < 10^{-4}$. WGAN does not learn a typical GAN discriminator that outputs a probability, so calibration is actually required in this case.

	DCGAN				WGAN			
	CIFAR-10	p	CelebA	p	CIFAR-10	p	CelebA	p
GAN	2.8789	–	2.3317	–	3.0734	–	2.7876	–
DRS	2.977(77)	0.0131	2.511(50)	<0.0001				
DRS (cal)	3.073(80)	<0.0001	2.869(67)	<0.0001	3.137(64)	0.0497	2.861(66)	0.0277
MH-GAN	3.113(69)	<0.0001	2.682(50)	<0.0001				
MH-GAN (cal)	3.379(66)	<0.0001	3.106(64)	<0.0001	3.305(83)	<0.0001	2.889(89)	0.0266

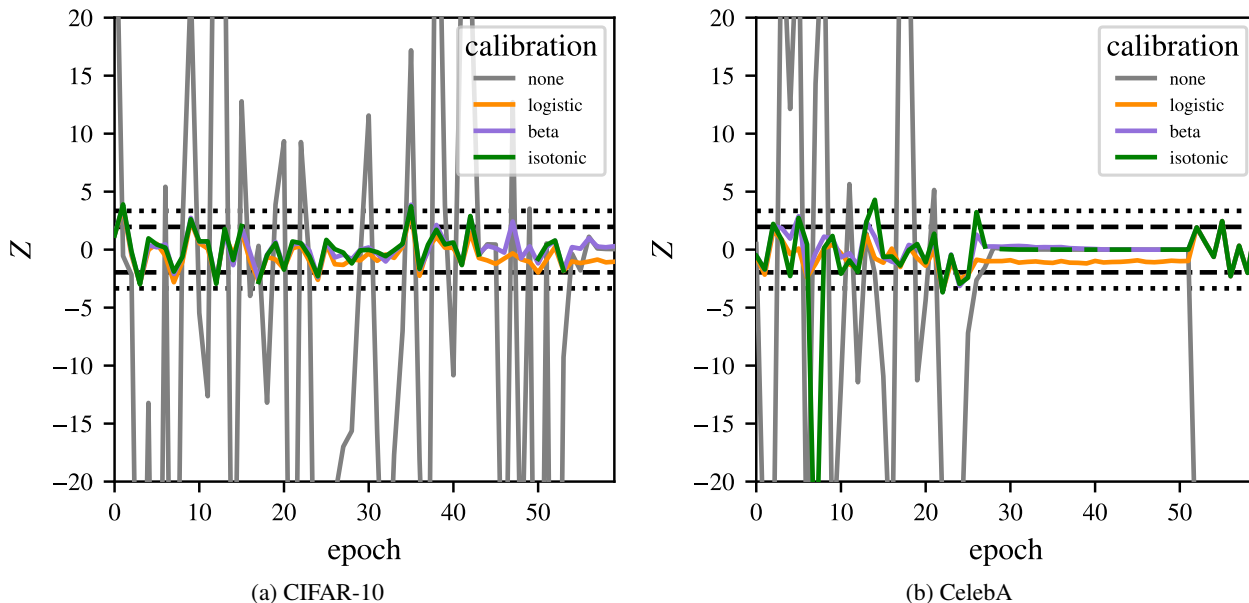


Figure 6. We show the calibration statistic Z (7) for the discriminator on held out data for the DCGAN. The results for CIFAR-10 are shown on the left (a), and CelebA on the right (b). The raw discriminator is clearly miscalibrated being far outside the region expected by chance (dashed black), and after multiple comparison correction (dotted black). All the calibration methods give roughly equivalent results. CelebA has a period of training instability during epochs 30–50 which gives trivially calibrated classifiers.

Calibration results Figure 6 shows the results per epoch for both CIFAR-10 and CelebA. It shows that the raw discriminator is highly miscalibrated, but can be fixed with any of the calibration methods. The Z statistic for the raw discriminator D (DCGAN on CIFAR-10) varies from -77.57 to 48.98 in the first 60 epochs; even after Bonferroni correction at $N=60$, we expect $|Z| < 3.35$ with 95% confidence for a calibrated classifier. The calibrated discriminator varies from -2.91 to 3.60 , showing almost perfect calibration. Accordingly, it is unsurprising that the calibrated discriminator significantly boosts performance in the MH-GAN.

Visual results We also show example images from the CIFAR-10 and CelebA setups in the supplementary material. The selectors (such as MH-GAN) result in a wider spread of probability mass across background colors. For CIFAR-10, it enhances modes with animal-like outlines and vehicles.

5.3. Progressive GAN

To further illustrate the power of the MH-GAN approach we consider the progressive GAN (PGAN) (Karras et al., 2017), which recently produced shockingly realistic images. We applied the MH-GAN to a PGAN using the same setup as with DCGAN, at $k = 800$. We used the pre-trained network of Karras et al. (2017) on CelebA-HQ (1024×1024). Large batches of samples are in the supplementary material.

In Table 2, we use the PGAN as our base GAN and generate random samples from the base, as well as from the addition

of DRS and MH-GAN selectors. The different selectors (DRS and MH-GAN) are run on the same batches of images, so the same images may appear for both generators. Although the PGAN sometimes produces near photorealistic images, it also produces many flawed nightmare like images. To assess image quality, five human labelers manually labeled images as warped or acceptable. Table 2 shows that MH-GAN selects significantly fewer warped images.

Both DRS and MH-GAN show an ability to select just the realistic images. The MH-GAN samples are nearly perfect, while DRS still has many flawed samples.


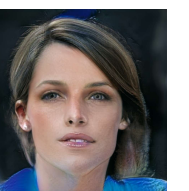





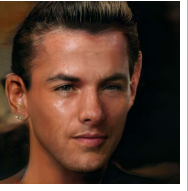
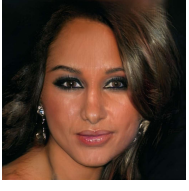
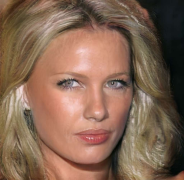

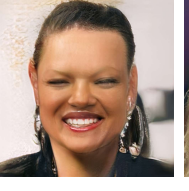

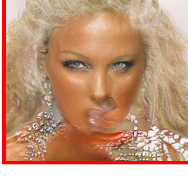
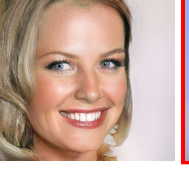
6. Conclusions

We have shown how to incorporate the knowledge in the discriminator D into an improved generator G' . Our method is based on the premise that D is better at density ratio estimation than G is at sampling data, which may be a harder task. The principled MCMC setup selects among samples from G to correct biases in G . This is the only method in the literature which has the property that given a perfect D one can recover G such that $p_G = p_{\text{data}}$.

We have shown the raw discriminators in GANs and DRS are poorly calibrated. To our knowledge, this is the first work to evaluate the discriminator in this way and to rigorously show the poor calibration of the discriminator. Because the MH-GAN algorithm may be used to wrap any other GAN, there are countless possible use cases.

Metropolis-Hastings Generative Adversarial Networks

Table 2. We show 16 *random samples* from the PGAN, calibrated (improved) DRS, and MH-GAN, from the same sequence of G samples. There are 5 cases where the PGAN produces bad warpings (red) while the MH-GAN does not, and 0 cases where the MH-GAN does and the PGAN does not; for DRS, there are 7 where only DRS is warped, and 1 where only MH-GAN is warped. Even with 16 samples, the MH-GAN is better under a one-sided pairwise trinomial test (Coakley & Heise, 1996) at $p = 0.017$ vs DRS and $p = 0.013$ vs PGAN.

PGAN (base)		PGAN with DRS (cal)		PGAN with MH-GAN (cal)	
					
					
					
					
					
					
					
					

Acknowledgements

We thank Rosanne Liu and Zoubin Ghahramani for useful discussions and comments.

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, volume 70, pp. 214–223, 2017.
- Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*, 2018.
- Coakley, C. W. and Heise, M. A. Versions of the sign test in the presence of ties. *Biometrics*, 52(4):1242–1251, 1996.
- Cowles, M. K., Roberts, G. O., and Rosenthal, J. S. Possible biases induced by MCMC convergence diagnostics. *Journal of Statistical Computation and Simulation*, 64(1): 87–104, 1999.
- Dawid, A. P. Prequential analysis. *Encyclopedia of Statistical Sciences*, 1:464–470, 1997.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real NVP. *arXiv preprint arXiv:1605.08803*, 2016.
- Dumoulin, V., Belghazi, I., Poole, B., Mastropietro, O., Lamb, A., Arjovsky, M., and Courville, A. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. Introducing Markov chain Monte Carlo. *Markov Chain Monte Carlo in Practice*, 1:19, 1996.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Proceedings of Advances in Neural Information Processing Systems*, pp. 2672–2680. Curran Associates, Inc., 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of Wasserstein GANs. *arXiv preprint arXiv:1704.00028v3*, 2017.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Kempinska, K. and Shawe-Taylor, J. Adversarial sequential Monte Carlo. In *Bayesian Deep Learning (NIPS Workshop)*, 2017.
- Kull, M., Filho, T. S., and Flach, P. Beta calibration: A well-founded and easily implemented improvement on logistic calibration for binary classifiers. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, volume 54, pp. 623–631, 2017.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 3730–3738, 2015.
- Murray, I. and Salakhutdinov, R. Notes on the KL-divergence between a Markov chain and its equilibrium distribution. 2008.
- Neklyudov, K., Shvechikov, P., and Vetrov, D. Metropolis-Hastings view on variational inference and adversarial training. *arXiv preprint arXiv:1810.07151*, 2018.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X., and Chen, X. Improved techniques for training GANs. In *Proceedings of Advances in Neural Information Processing Systems*, pp. 2234–2242. Curran Associates, Inc., 2016.
- Shibuya, N. Understanding generative adversarial networks. <https://towardsdatascience.com/understanding-generative-adversarial-networks-4dafc963f2ef>, 2017.
- Song, J., Zhao, S., and Ermon, S. A-NICE-MC: Adversarial training for MCMC. In *Proceedings of Advances in Neural Information Processing Systems*, pp. 5140–5150. Curran Associates, Inc., 2017.
- Sugiyama, M., Suzuki, T., and Kanamori, T. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012.
- Tierney, L. Markov chains for exploring posterior distributions. *The Annals of Statistics*, 22(4):1701–1728, 1994.
- Torralba, A., Fergus, R., and Freeman, W. T. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008.