

# Reparameterized Sampling for Generative Adversarial Networks

No Author Given

No Institute Given

## A Appendix

### A.1 Assumptions and Implications

Note that our method needs a few assumptions on the models for our analysis to hold. Here we state them explicitly and discuss their applicability and potential impacts.

**Assumption 1.** *The generator mapping  $G: \mathbb{R}^n \rightarrow \mathbb{R}^m (n < m)$  is injective, and its Jacobian matrix  $\left[\frac{\partial G(\mathbf{z})}{\partial \mathbf{z}}\right]$  of size  $m \times n$ , has full column rank for all  $\mathbf{z} \in \mathbb{R}^n$ .*

For the change of variables in Eqn. 11 and 13 to hold, according to [1], we need the mapping to be injective and its Jaobian should have full column rank. A mild sufficient condition for injectivity is that the generator only contains (non-degenerate) affine layers and injective non-linearities, like LeakyReLU. It is not hard to show that such a condition also implies the full rankness of the Jacobian. In fact, this architecture has already been found to benefit GANs and achieved state-of-the-art results [3]. The affine layers here are also likely to be non-degenerate because their weights are randomly initialized and typically will not degenerate in practice during the training of GANs.

**Assumption 2.** *The discriminator  $D$  offers a perfect estimate the density ratio between the generative distribution  $p_g(\mathbf{x})$  and the data distribution  $p_d(\mathbf{x})$  as in Eqn. 3.*

This is a common, critical, but less practical assumption among the existing sampling methods of GANs. It is unlikely to hold exactly in practice, because during the alternative training of GANs, the generator is also changing all the time, and the a few updates of the discriminator cannot fully learn the corresponding density ratio. Nevertheless, we think it can capture a certain extent information of density ratio which explains why the sampling methods can consistently improve over the baseline at each epoch.

From our understanding, the estimated density ratio is enough to push the generator better but not able to bring it up to the data distribution. This could be the reason why the Inception scores obtained by the sampling methods, can improve over the baselines but cannot reach up to that of real data and fully close the gap, even with very long run of the Markov chains.

Table 5: Fréchet Inception Distance (FID) of different MCMC sampling methods on CIFAR-10 and CelebA.

Method	CIFAR-10		CelebA	
	DCGAN	WGAN	DCGAN	WGAN
GAN	100.363	153.683	227.892	207.545
MH-GAN [4]	100.167	143.426	<b>227.233</b>	207.143
DDLs [2]	145.981	193.558	269.840	232.522
REP-GAN (ours)	<b>99.798</b>	<b>143.322</b>	230.748	<b>207.053</b>




Fig. 6: Visualization of the Markov chains of MH-GAN (top), DDLs (middle), and REP-GAN (bottom) on CelebA with WGAN backbone.

Hence, there is still much room for improvement. To list a few, one can develop mechanisms that bring more accurate density ratio estimate, or relax the assumptions for the method to hold, or establishing estimation error bounds. Overall, we believe GANs offer an interesting alternative scenario for the development of sampling methods.

## A.2 Additional Empirical Results

Here we list some additional empirical results of our methods.

**Fréchet Inception Distance** We additionally report the comparison of Fréchet Inception Distance (FID) in Table 5. We can see the ranks are consistent with the Inception scores in Table 2 and our method is superior in most cases.

**Markov Chain Visualization on CelebA** We demonstrate two Markov chains on CelebA with different MCMC sampling methods of WGAN in Figure 6. We can see that on CelebA, the acceptance ratio of MH-GAN becomes much

higher than that on CIFAR-10. Nevertheless, the sample quality is still relatively low. In comparison, the gradient-based method can gradually refine the samples with Langevin steps, and our REP-GAN can alleviate image artifacts with MH correction steps.

## References

1. Ben-Israel, A.: The change-of-variables formula using matrix volume. *SIAM Journal on Matrix Analysis and Applications* **21**(1), 300–312 (1999)
2. Che, T., Zhang, R., Sohl-Dickstein, J., Larochelle, H., Paull, L., Cao, Y., Bengio, Y.: Your GAN is secretly an energy-based model and you should use discriminator driven latent sampling. *ICML* (2020)
3. Tang, S.: Lessons learned from the training of gans on artificial datasets. *IEEE Access* **8**, 165044–165055 (2020)
4. Turner, R., Hung, J., Saatci, Y., Yosinski, J.: Metropolis-Hastings generative adversarial networks. *ICML* (2019)