Tips and Tricks for Training GANs with Physics Constraints

Luke de Oliveira
Vai Technologies & Lawrence Berkeley National Laboratory
lukedeo@vaitech.io

Michela Paganini
Yale University & Lawrence Berkeley National Laboratory
michela.paganini@yale.edu

Benjamin Nachman
Lawrence Berkeley National Laboratory
bnachman@cern.ch

Abstract

Generative Adversarial Networks (GANs) have seen immense interest and success in recent years. However, most tasks and successes have existed solely within the domain of natural images. In contrast, we provide an overview of modifications and tricks necessary to make GANs work on data from high energy particle physics. We provide select examples of domain-specific thought processes with respect to improving GAN training procedures, aiming to be a resource for researchers in any physical or applied science wishing to apply Generative Adversarial Networks to a choice problem.

1 Introduction

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) provide a game theoretic basis for learning a target distribution $p_{\mathsf{data}}(x)$ that has gained popularity in recent years. Subsequently, the usage of generative modeling in the natural sciences has been gaining popularity, with applications in Cosomology (Mustafa et al., 2017), High Energy Particle Physics (de Oliveira et al., 2017b; Paganini et al., 2017; de Oliveira et al., 2017a), Geology (Chan and Elsheikh, 2017), and Astronomy (Schawinski et al., 2017), among others. A common thread among scientific applications is the notion of acceleration - that is, the utilization of a GAN as a fast sampling method or surrogate to circumvent traditional scientific simulators when the highest fidelity is not required. Although there has been significant preliminary success, these methods still suffer from traditional failure modes that plague GANs (Goodfellow, 2014; Nowozin et al., 2016; Salimans et al., 2016). In fact, many issues that arise in the physical sciences, such as sparsity(Paganini et al., 2017), location dependence (de Oliveira et al., 2017b), and multi-scale behavior across orders of magnitude (Mustafa et al., 2017), can directly exacerbate pre-existing failure modes if not handled carefully.

In this overview, we document select observations, tips, and tricks from the literature that other researchers may find useful while training GANs for scientific data. We pay special attention to implementation details that may get traditionally glossed over, as these insights and tricks are often critical to obtaining usable results. This workshop contribution summarizes empirical techniques that have been extensively tested in related works. We omit the details of the experimentation in favor of a concise collection of tips and tricks for researchers.

Workshop on Deep Learning for Physical Sciences (DLPS 2017), NIPS 2017, Long Beach, CA, USA.

This exposition is organized as follows. Section 1.1 provides a very brief introduction to GANs, Section 2 describes a solution to model sparsity in scientific data, Section 3 poses an attribute conditioning problem in the context of interpolation, and finally, Section 4 mentions some brief observations around training procedures given the observations in Sections 2 and 3.

1.1 Generative Adversarial Networks

Generative Adversarial Networks cast the task of training a deep generative model as a two-player non-cooperative minimax game in which both players are parameterized as deep neural networks. In particular, a *generator* network is tasked with mapping a latent prior $z \sim p_z(z), z \in \mathcal{Z}$ to a reputable synthetic sample that a *discriminator* network is unable to distinguish as being a real (from $p_{\mathsf{data}}(x)$) or fake sample. Though GANs can be shown to minimize a quantity closely related to the Jensen-Shannon divergence, they suffer from a variety of issues related to mode collapse, instability, and lack of reliable metrics.

We now dive in to specific qualities, specifically around image-based representations in the sciences, that are unique for building GANs.

2 Sparsity

In direct opposition to natural images, most data from the sciences is inherently sparse. Even though GANs are not directly designed for this, previous work (de Oliveira et al., 2017b; Paganini et al., 2017; de Oliveira et al., 2017a) has explored this and multiple solutions have been proposed different solutions. A simple approach is to utilize Rectified Linear Units (ReLU) (Maas et al.) to directly induce sparsity in an output layer. However, since ReLU units produce sparse gradients in the generator, care must be given to the training procedure, which is highlighted in in Section 4.

Though utilizing ReLU pointwise nonlinearities induces sparsity, we have no guarantees as to whether or not each generated sample is sparse with an occupancy commensurate with the distribution of sparsity over a full data distribution. A simple solution is to design a quantity closely related to sparsity that the discriminator can utilize to encourage the generator to produce samples at an adequate sparsity. A suggested quantity (de Oliveira et al., 2017a) is defined as follows

$$\mathsf{softsparsity}(X) = \left(\frac{1}{nm}\right) \left\| \frac{|X|^\alpha + |X|^{-\beta}}{|X|^\alpha + |X|^{-\beta} + 1} \right\|_1, \tag{1}$$

where $\alpha, \beta > 0$ and $X \in \mathbb{R}^{m \times n}$. In addition, all powers and $|\cdot|$ operators are assumed to act pointwise on X, and $\|\cdot\|_1$ is the entry-wise 1-norm rather than the induced norm. Examining Eqn. 1, we note that $\mathsf{softsparsity} : \mathbb{R}^{m \times n} \longrightarrow [0, 1]$, and that in the limit as

$$\lim_{\alpha,\beta\to\infty} \mathsf{softsparsity}(X) = \frac{1}{mn} \sum_{i<m} \sum_{j<n} \mathbb{I}[X_{i,j} \neq 0], \tag{2}$$

which is the standard definition of sparsity (or occupancy).

Augmenting the discriminator with this value for generated and data samples allows the discriminator to feed gradient information to the generator. However, we can go one step further and include a minibatch discrimination (Salimans et al., 2016) layer acting just on this value across a minibatch, encouraging the generator to correctly reproduce the distribution of sparsity.

3 Attribute Conditioning

For most scientific applications such as High Energy Physics and Cosmology, we need to not only learn $p_{\mathsf{data}}(x)$, but also approximate $p_{\mathsf{data}}(x|\xi)$, where ξ is a vector of conditioning attributes. This is crucial for most scientific applications, where ξ is of theoretical importance to be able to either condition on or interpolate between.

With respect to parameter interpolation, we assume we have parameters ξ which we would like to interpolate between. Formally, we have a set of values of $\xi_i \in \Xi$, with the key property that Ξ is finite and is sparsely sampled. We would like to, for any $\xi^* \in \mathsf{Conv}(\Xi)$, be able to directly sample from this continuous distribution. This allows dense interpolation in spite of finite samples in a conditioning space.

In this case, categorical conditioning can be taken care of using traditional methods (Mirza and Osindero, 2014; Odena et al., 2016). For continuous values, a simple methodology to condition on continuous characteristics is designed. A separate output/submodel of the discriminator, $h(x)$, is tasked with reconstructing each element of ξ. Define individual distances $d_i(h(x)_i, \xi_i)$ where the index i tracks each feature in conditioning space. We can the require both the discriminator and generator to minimize

$$\mathcal{L}_{\mathsf{conditional}} = \sum_{i=1}^{\mathsf{dim}(\xi)} \lambda_i d_i(h(x)_i, \xi_i), \tag{3}$$

where the λ_i are hyperparameters dictating scaling and relative importance of one vectorial components' loss with respect to the remaining components (de Oliveira et al., 2017a) losses.

4 Training

We now turn to high level observations of training procedures for GANs on scientific data. As noted in Section 2, image-represented data from natural sources tends to be quite sparse. In addition, pixel intensities can span many more orders of magnitude than natural images. This necessitates paying close attention to gradient properties and batch sizes. Sparse images with high dynamic range require a larger batch size in order to smooth out gradients during each update step because most parameters receive no update in a small batch. In addition, the sparsity levels in images can very easily lead to a truth bit being present in generated samples. To solve this, and to discourage this violation of the generator distribution being absolutely continuous with respect to p_{data} (Nowozin et al., 2016), label flipping is used in order create more overlap.

For scientific applications, complete exploration of the data support is essential for any GAN-based system to be useful. Inasmuch as it directly encourages this behavior through batch-level statistics, minibatch discrimination (Salimans et al., 2016) has proven useful to aid in this direction.

Note that many of the observations mentioned in this exposition are irrespective of GAN formulation. Recent work examining alternative divergences for training GANs (Arjovsky et al., 2017; Bellemare et al., 2017; Gulrajani et al., 2017) is completely compatible with the observations elucidated upon here. In particular, although recent work has suggested that most GAN formulations are empirically equivalent (Lucic et al., 2017), we posit that an additional dimension of importance in understanding trade-offs of different GAN formulations should be the transferability and utility in transfer learning or simulator settings, as this provides an application-specific notion of coverage and quality.

5 Conclusion

We have provided a brief exposition into some considerations in applying Generative Adversarial Networks in scientific settings. In particular, we highlighted a subset of ways in which domain understanding, coupled with understanding of GAN dynamics, can lead to improvements to enable progress in fields which use generative modeling as a tool. We hope this exposition will be useful to researchers who work in science as both a reference and inspiration to think creatively about unique problem constraints and how to design good algorithms to either take advantage of or accommodate the nature of domain data.

References

Arjovsky, M., S. Chintala, and L. Bottou
2017. Wasserstein GAN.

Bellemare, M. G., I. Danihelka, W. Dabney, S. Mohamed, B. Lakshminarayanan, S. Hoyer, and R. Munos
2017. The Cramer Distance as a Solution to Biased Wasserstein Gradients.

Chan, S. and A. H. Elsheikh
2017. Parametrization and Generation of Geological Models with Generative Adversarial Networks.

de Oliveira, L., M. Paganini, and B. Nachman
2017a. Controlling Physical Attributes in GAN-Accelerated Simulation of Electromagnetic Calorimeters. In *ACAT (forthcoming)*.

de Oliveira, L., M. Paganini, and B. Nachman
2017b. Learning Particle Physics by Example: Location-Aware Generative Adversarial Networks for Physics Synthesis.

Goodfellow, I., J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio
2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds., Pp. 2672–2680. Curran Associates, Inc.

Goodfellow, I. J.
2014. On distinguishability criteria for estimating generative models.

Gulrajani, I., F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville
2017. Improved Training of Wasserstein GANs.

Lucic, M., K. Kurach, M. Michalski, S. Gelly, and O. Bousquet
2017. Are GANs Created Equal? A Large-Scale Study.

Maas, A. L., A. Y. Hannun, and A. Y. Ng
. Rectifier Nonlinearities Improve Neural Network Acoustic Models.

Mirza, M. and S. Osindero
2014. Conditional Generative Adversarial Nets.

Mustafa, M., D. Bard, W. Bhimji, R. Al-Rfou, and Z. Lukić
2017. Creating Virtual Universes Using Generative Adversarial Networks.

Nowozin, S., B. Cseke, and R. Tomioka
2016. f-GAN: Training Generative Neural Samplers using Variational Divergence Minimization.

Odena, A., C. Olah, and J. Shlens
2016. Conditional Image Synthesis With Auxiliary Classifier GANs.

Paganini, M., L. de Oliveira, and B. Nachman
2017. CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calorimeters with Generative Adversarial Networks.

Salimans, T., I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen
2016. Improved Techniques for Training GANs.

Schawinski, K., C. Zhang, H. Zhang, L. Fowler, and G. K. Santhanam
2017. Generative Adversarial Networks recover features in astrophysical images of galaxies beyond the deconvolution limit.