
Implicit competitive regularization in GANs

Florian Schäfer^{*1} Hongkai Zheng^{*23} Anima Anandkumar¹

Abstract

The success of GANs is usually attributed to properties of the divergence obtained by an optimal discriminator. In this work we show that this approach has a fundamental flaw:

If we do not impose regularity of the discriminator, it can exploit visually imperceptible errors of the generator to always achieve the maximal generator loss. In practice, gradient penalties are used to regularize the discriminator. However, this needs a metric on the space of images that captures visual similarity. Such a metric is not known, which explains the limited success of gradient penalties in stabilizing GANs.

Instead, we argue that the implicit competitive regularization (ICR) arising from the simultaneous optimization of generator and discriminator enables GANs performance. We show that opponent-aware modelling of generator and discriminator, as present in competitive gradient descent (CGD), can significantly strengthen ICR and thus stabilize GAN training without explicit regularization. In our experiments, we use an existing implementation of WGAN-GP and show that by training it with CGD without any explicit regularization, we can improve the inception score (IS) on CIFAR10, without any hyperparameter tuning.

1. Introduction

Generative adversarial networks (GANs): (Goodfellow et al., 2014) are a class of generative models based on a competitive game between a generator that tries to generate realistic new data, and a discriminator that tries to distinguish generated from real data. In practice, both players are parameterized by neural networks that are trained simultaneously by a variant of stochastic gradient descent.

^{*}Equal contribution ¹Caltech ²Shanghai Jiao Tong university

³This work was produced while HZ was a visiting undergraduate researcher at Caltech. Correspondence to: Florian Schaefer <schaefer@caltech.edu>, Hongkai Zheng <devzhk@sjtu.edu.cn>.

The minimax interpretation: Presently, the success of GANs is mostly attributed to properties of the divergence or metric obtained under an optimal discriminator. For instance, an optimal discriminator in the original GAN leads to a generator loss equal to the Jensen-Shannon divergence between real and generated distribution. Optimization over the generator is then seen as approximately minimizing this divergence. We refer to this point of view as the minimax interpretation. The minimax interpretation has led to the development of numerous GAN variants that aim to use divergences or metrics with better theoretical properties.

The GAN-dilemma: However, every attempt to explain GAN performance with the minimax interpretation faces one of the two following problems:

1. Without regularity constraints, the discriminator can always be perfect. This is because it can selectively assign a high score to the finite amount of real data points while assigning a low score on the remaining support of the generator distribution, as illustrated in Figure 1. Therefore, the Jensen-Shannon divergence between a continuous and a discrete distribution always achieves its maximal value, a property that is shared by all divergences that do not impose regularity constraints on the discriminator. Thus, these divergences can not meaningfully compare the quality of different generators.

2. Imposing regularity constraints needs a measure of similarity of images. Imposing regularity on the discriminator amounts to forcing it to map similar images to similar results. To do so, we would require a notion of similarity between images that is congruent with human perception. This is a longstanding unsolved problem in computer vision. Commonly used gradient penalties use the Euclidean norm which is known to poorly capture visual similarity, as illustrated in Figure 2.

We believe that the different divergences underlying the various GAN formulations have little to do with their ability to produce realistic images. This is supported by the large scale studies of Lucic et al. (2017) that did not find systematic differences in the performance of GANs associated with different divergence measures. Understanding of GAN performance is crucial in order to improve training stability and reduce the required amount of hyperparameter tuning.



Figure 1. **The discriminator can always improve:** We want the discriminator confidence to reflect the relative abundance of true and fake data (left). But by picking out individual data points, the discriminator can almost always achieve arbitrarily low loss on any finite data set (right). Even in the limit of infinite data, the slightest misalignment of the supports of generated and real data can be exploited in a similar way.

A way out?: Due to the GAN-dilemma, every attempt at explaining the performance of GANs needs to go beyond the minimax interpretation and consider the dynamics of the training process. In this work, we argue that an implicit regularization due to the simultaneous¹ training of generator and discriminator allows GANs to use the inductive biases of neural networks for the generation of realistic images.

Implicit competitive regularization: We define implicit competitive regularization (ICR) as the introduction of additional stable points or regions due to the simultaneous training of generator and discriminator that do not exist when only training the generator (or discriminator) with gradient descent while keeping the discriminator (or generator) fixed.

It has been previously observed that performing simultaneous gradient descent (SimGD) on both players leads to stable points that are not present when performing gradient descent with respect to either player, while keeping the other player fixed (Mazumdar & Ratliff, 2018). These stable points are not local Nash equilibria, meaning that they are not locally optimal for both players. This phenomenon is commonly seen as a shortcoming of SimGD and modifications that promote convergence only to local Nash equilibria which have been proposed by, for instance, (Balduzzi et al., 2018; Mazumdar et al., 2019). In contrast to this view we believe that ICR is crucial to overcoming the GAN-dilemma and hence to explaining GAN performance in practice by allowing the inductive biases of the discriminator network to inform the generative model.

Summary of Contributions

In this work, we point out that a fundamental dilemma prevents the common minimax interpretation of GANs from explaining their successes. We then show that implicit competitive regularization (ICR), which so far was believed to be a flaw of SimGD, is key to overcoming this dilemma.

¹Here and in the following, when talking about simultaneous training, we include variants such as alternating gradient descent.

Based on simple examples and numerical experiments on real GANs we illustrate how it allows to use the inductive biases of neural networks for generative modelling, resulting in the spectacular performance of GANs.

We then use this understanding to improve GAN performance in practice. Interpreting ICR from a game-theoretic perspective, we reason that strategic behavior and opponent-awareness of generator and discriminator during the training procedure can strengthen ICR. These elements are present in competitive gradient descent (CGD) (Schaefer & Anandkumar, 2019) which is based on the two players solving for a local Nash-equilibrium at each step of training. Accordingly, we observe that CGD greatly strengthens the effects of ICR. In comprehensive experiments on CIFAR 10, competitive gradient descent stabilizes previously unstable GAN formulations and achieves higher inception score compared to a wide range of explicit regularizers, using both WGAN loss and the original saturating GAN loss of Goodfellow et al. (2014). In particular, taking an existing WGAN-GP implementation, dropping the gradient penalty, and training with CGD leads to the highest inception score in our experiments. We interpret this as additional evidence that ICR, as opposed to explicit regularization, is the key mechanism behind GAN performance.

2. The GAN-dilemma

In this section, we study in more detail the fundamental dilemma that prevents the common minimax interpretation from explaining the successes of GANs. In particular, we show how the existing GAN variants fall into one or the other side of the GAN-dilemma.

Metric-agnostic GANs: In the original formulation due to Goodfellow et al. (2014), the two players are playing a zero-sum game with the loss function of the generator given by the binary cross entropy



Figure 2. **The Euclidean distance is not perceptual:** We would like to challenge the reader to order the above three pairs of images according to the Euclidean distance of their representation as vectors of pixel-intensities.³

$$\min_{\mathcal{G}} \max_{\mathcal{D}} \frac{1}{2} \mathbb{E}_{x \sim P_{\text{data}}} [\log \mathcal{D}(x)] + \frac{1}{2} \mathbb{E}_{x \sim P_{\mathcal{G}}} [\log (1 - \mathcal{D}(x))]. \quad (1)$$

Here, \mathcal{G} is the probability distribution generated by the generator, \mathcal{D} is the classifier provided by the discriminator, and P_{data} is the target measure, for example the empirical distribution of the training data. A key feature of the original GAN is that it depends on the discriminator only through its output when evaluated on samples. This property is shared, for instance, by the more general class of f -divergence GANs (Nowozin et al., 2016). We call GAN formulations with this property metric-agnostic.

Metric-informed GANs: To address instabilities observed in the original GAN, Arjovsky et al. (2017) introduced WGAN, with loss function given by

$$\min_{\mathcal{G}} \max_{\mathcal{D}} \mathbb{E}_{x \sim P_{\text{data}}} [\mathcal{D}(x)] - \mathbb{E}_{x \sim P_{\mathcal{G}}} [\mathcal{D}(x)] + \mathcal{F}(\nabla \mathcal{D}) \quad (2)$$

where $\mathcal{F}(\nabla \mathcal{D})$ is infinity if $\sup_x \|\nabla \mathcal{D}(x)\| > 1$ and zero, else. (Gulrajani et al., 2017) propose WGAN-GP, where this inequality constraint is relaxed by replacing \mathcal{F} with a penalty, for instance $\mathcal{F}(\nabla \mathcal{D}) = \mathbb{E} \left[(\|\nabla_x \mathcal{D}\| - 1)^2 \right]$. These GAN formulations are fundamentally different from metric-agnostic GANs in that they depend explicitly on the gradient of the discriminator. In particular, they depend on the choice of metric used to measure the size of $\nabla \mathcal{D}$. Subsequent to WGAN(-GP), which uses the Euclidean norm, other variants such as Sobolev-GAN (Mroueh et al., 2017), Banach-GAN (Adler & Lunz, 2018), or Besov-GAN (Uppal et al., 2019) have been proposed that use different metrics to measure

³The pairs of images are ordered from left to right, in increasing order of distance. The first pair is identical, while the third pair differs by a tiny warping.

gradient size. We refer to these types of GAN formulations as metric-informed GANs.

The problem with metric-agnostic GANs: GANs are able to generate highly realistic images, but they suffer from unstable training and mode collapse that often necessitates extensive hyperparameter tuning. Beginning with (Arjovsky & Bottou, 2017) these problems of the original GAN have been explained with the fact that the supports of the generator distribution and the training data are almost never perfectly aligned. For any fixed generator, the discriminator can take advantage of this fact to achieve arbitrarily low loss, as illustrated in Figure 1. In the case of the Formulation 1, this corresponds to the well known fact that the Jensen-Shannon divergence between mutually singular measures is always maximal. This result extends to all metric-agnostic divergences, simply because they have no way of accessing the degree of similarity between data points on disjoint supports.

Arora et al. (2017); Huang et al. (2017) emphasize that the discriminator is restricted to a function class parameterized by a neural network. However, the experiments of Arjovsky & Bottou (2017) as well as our own in Figure 4 clearly show the tendency of the discriminator to diverge as it achieves near-perfect accuracy. This is not surprising since Zhang et al. (2016) observed that modern neural networks are able to fit even random data perfectly. Arjovsky & Bottou (2017) also show that as the discriminator improves its classification loss, the generator achieves less and less useful gradient information. This is again not surprising, since confidence scores of deep neural networks are known to be poorly calibrated (Guo et al., 2017). Therefore, the outputs of a near-perfect discriminator can not be expected to provide a useful assessment of the quality of the generated samples.

Since GAN optimization is highly non-convex it is natural to ask if GANs find locally optimal points in the form of local Nash or Stackelberg equilibria. This local minmax interpretation has been emphasized by Fiez et al. (2019); Jin et al. (2019), but the experiments of Berard et al. (2019) as well as our own in Figure 4 suggest that good GAN solutions for metric-agnostic GANs are typically not locally optimal for both players. It seems plausible that the discriminator, being highly overparameterized, can find a direction of improvement against most generators.

The problem with metric-informed GANs: The above observation has motivated the introduction of metric-informed GANs that restrict the size of the gradient of the discriminator (as a function mapping images to real numbers). This limits the discriminator’s ability to capitalize on small misalignments between \mathcal{D} and P_{data} and thus makes for a meaningful minimax interpretation even if the two measures have fully disjoint support. However, the Achilles heel of this approach is that it needs to choose a metric to quantify the magnitude of the discriminator’s gradients. Most of the early work on metric-informed GANs chose to measure the size of $\nabla\mathcal{D}$ using the Euclidean norm (Arjovsky & Bottou, 2017; Arjovsky et al., 2017; Gulrajani et al., 2017; Roth et al., 2017; Kodali et al., 2017; Miyato et al., 2018). However, since the discriminator maps images to real numbers, this corresponds to quantifying the similarity of images at least locally by the Euclidean distance of vectors containing the intensity values of each pixel. As illustrated in Figure 2, this notion of similarity is poorly aligned with visual similarity even locally. From this point of view it is not surprising that the generative model of (Chen et al., 2019), based on a differentiable optimal transport solver, produced samples of lower visual quality than WGAN-GP, despite achieving better approximation in Wasserstein metric. As noted by Chen et al. (2019), these observations suggest that the performance of WGAN can not be explained by its relationship to the Wasserstein distance. When comparing a variety of GAN formulations with a fixed budget for hyperparameter tuning, Lucic et al. (2017) did not find systematic differences in their performance. This provides additional evidence that the key to GAN performance does not lie in the choice of a particular divergence between probability measures.

The metric-informed divergences considered so far were all based on the Euclidean distance between images. Other researchers have tried using different metrics on image space such as Sobolev or Besov norms (Adler & Lunz, 2018; Mroueh et al., 2017; Uppal et al., 2019), or kernel maximum mean discrepancy distances (Li et al., 2015; 2017; Bińkowski et al., 2018). However, none of these metrics do a good job at capturing perceptual similarity either, which explains why these variants have not been observed to outperform WGAN(-GP) in general. Researchers in computer

vision have proposed more sophisticated domain-specific distance measures (Simard et al., 1998), kernel functions (Haasdonk & Burkhardt, 2007; Song et al., 2014), and features maps (Dalal & Triggs, 2005). Although computationally expensive, methods from differential geometry have been used for image inter- and extrapolation (Trouvé & Younes, 2005; Berkels et al., 2015; Effland et al., 2018). However, none of these classical methods achieve performance comparable to that of neural network based models, making them unlikely solutions for the GAN dilemma.

A way out: Generative modelling means producing new samples that are similar to the training samples, but not too similar to each other. Thus, every generative method needs to choose how to measure similarity between samples, implicitly or explicitly.

When analyzing GANs from the minimax perspective this assessment of image similarity seems to rely exclusively on the classical metrics and divergences used for their formulation. But modeling perceptual similarity is hard and most commonly used GAN formulations are based on measures of similarity that are known to be terrible at this task. Thus, the minimax point of view can not explain why GANs produce images of higher visual quality than any other method. The key to image classification is to map similar images to similar labels. The fact that deep neural networks drastically outperform classical methods in this tasks leads us to believe that they capture perceptual similarity between images far better than any classical model. We believe that the success of GANs is due to their ability to implicitly use the inductive biases of the discriminator network as a notion of similarity. They create images that look real to a neural network, which acts as a proxy for looking real to the human eye. In the next section we propose a new mechanism, implicit competitive regularization, to explain this behavior.

3. Implicit competitive regularization (ICR)

Implicit regularization: Based on the discussion in the last section, any attempt at understanding GANs needs to involve the inductive biases of the discriminator. However, there is ample evidence that the inductive biases of neural networks do not arise from a limited ability to represent certain functions. Indeed, it is known that modern neural networks can fit almost arbitrary functions (Kolmogorov, 1956; Cybenko, 1989; Zhang et al., 2016). Rather, they seem to arise from the dynamics of gradient-based training that tends to converge to classifiers that generalize well, a phenomenon commonly referred to as implicit regularization (Neyshabur, 2017; Gunasekar et al., 2017; Ma et al., 2017; Azizan et al., 2019; Kubo et al., 2019; Arora et al., 2019).

Implicit regularization is not enough for GANs: The implicit regularization induced by gradient descent lets neural networks prefer sets of weights with good generalization




Figure 3. ICR in the quadratic case: When optimizing only y in Equation (3), it diverges rapidly to infinity, for any fixed y . If however we simultaneously optimize x and y with respective step sizes $\eta_x = 0.09$ and $\eta_y = 0.01$, we converge to $(0, 0)$.

performance. However, the outputs of even a well-trained neural network are typically not informative about the confidence of the predicted class (Guo et al., 2017). Thus, a discriminator trained on finite amounts of real data and data generated by a given generator can be expected to distinguish new real data from new data generated by a similar generator, with high accuracy. However, its outputs do not quantify the confidence of its prediction and thus of the visual quality of the generated samples. Therefore, even considering implicit regularization, a fully trained discriminator does not provide useful gradients for training the generator.

Implicit competitive regularization: We think that GAN training relies on implicit competitive regularization (ICR), an additional implicit regularization due to the simultaneous training of generator and discriminator. When training generator and discriminator simultaneously, ICR selectively stabilizes good generators that would not be stable when training one player while keeping the other player fixed.

Consider the game given by

$$\min_x \max_y x^2 + 10xy + y^2. \quad (3)$$

In this problem, for any fixed x , any choice of y will be sub-optimal and gradient ascent on y (with x fixed) will diverge to infinity for almost all initial values.

What about simultaneous gradient descent? As has been observed before (Mazumdar & Ratliff, 2018), simultaneous gradient descent with step sizes $\eta_x = 0.09$ for x and $\eta_y = 0.01$ for y will converge to $(0, 0)$, despite it being a locally worst strategy for the maximizing player. (See Figure 3 for an illustration.) This is a first example of ICR, whereby the simultaneous optimization of the two agents introduces additional attractive points to the dynamics that are not attractive when optimizing one of the players using gradient descent while keeping the other player fixed.

As outlined in Section 2, the key to the performance of GANs has to lie in the simultaneous optimization process.

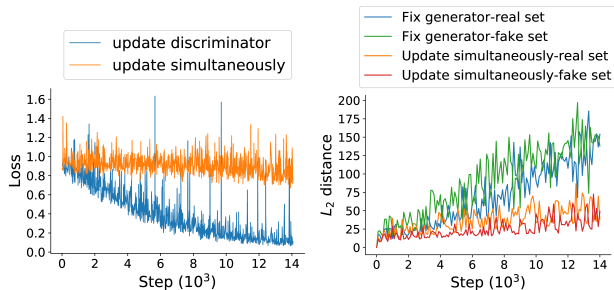


Figure 4. ICR on MNIST: We train a GAN on MNIST until we reach a checkpoint where it produces good images. (First image:) We fix the generator and only train the discriminator, observing that it can reach near-zero loss. When instead training generator and discriminator jointly, the loss stays stable. (Second Image:) When trained individually, the discriminator moves significantly slower when trained jointly with the generator, as measured by its output on a set of thousand reference images.

We now provide evidence that the solutions found by GANs are indeed stabilized by ICR. To this end, we train a GAN on MNIST until it creates good images. We refer to the resulting generator and discriminator as the checkpoint generator and discriminator. We observe that the loss of both generator and discriminator, as well as the image quality, is somewhat stable even though it would diverge after a long time of training. If instead, starting at the checkpoint, we optimize only the discriminator while keeping the generator fixed, we observe that the discriminator loss drops rapidly. For the same number of iterations and using the same learning rate, the discriminator moves away from the checkpoint significantly faster as measured both by the Euclidean norm of the weights and the output on real- and fake images. The observation that the discriminators diverges from the checkpoint faster when trained individually than when trained simultaneously with the generator suggests that the checkpoint, which produced good images, was stabilized by ICR.

4. How ICR lets GANs generate

An (hypo)thesis: In the example in the last section, the checkpoint producing good images was stabilized by ICR. However, we have not yet given a reason why points stabilized by ICR should have better generators, in general. For GANs to produce visually plausible images, there has to be some correspondence between the training of neural networks and human visual perception. Since learning and generalization are poorly understood even for ordinary neural network classifiers, we can not avoid making an assumption on the nature of this relationship. This section relies on the following hypothesis.

Hypothesis How quickly the discriminator can pick up on an imperfection of the generator is correlated with the visual

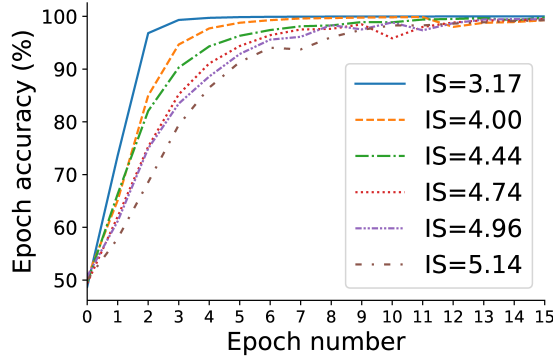


Figure 5. By prematurely stopping the training process, we obtain generators of different image-quality on CIFAR10 (higher inception score (IS) reflects better image quality). We then train a new discriminator against this fixed generator and measure how quickly it increases its classification performance. We use a model trained on the 10-class classification task as starting point for the discriminator to prevent the initial phase of training from polluting the measurements. While all discriminators achieve near-perfect accuracy eventually, the rate of improvement is inversely correlated to inception score of the generator.

prominence of said imperfection.

It is common intuition in training neural network classifiers that more visually obvious patterns are learned in fewer iterations and from less data. It is also in line with the coherent gradient hypothesis of Chatterjee (2020) that explains generalization performance of neural networks with the fact that systematic patterns in the data generate more coherent gradients and are therefore learned faster. While a thorough verification of the hypothesis is beyond the scope of this work we provide some empirical evidence in Figure 5.

This section argues for the following thesis.

Thesis: ICR selectively stabilizes generators for which the discriminator can only improve its loss slowly. By the hypothesis, these generators will produce high quality samples.

An argument in the quadratic case: We begin with the quadratic problem in Equation 3 and model the different speeds of learning of the two agents by changing their step sizes η_x and η_y . In Figure 6 we see that for $(\eta_x, \eta_y) = (0.03, 0.03)$ the two agents slowly diverge to infinity and for $(\eta_x, \eta_y) = (0.01, 0.09)$, divergence occurs rapidly. In general, stable points \bar{x} of an iteration $(x_{k+1} = x_k + F(x_k))$ are characterized by **(1)**: $F(\bar{x}) = 0$ and **(2)** $D_x F(\bar{x})$ having spectral radius smaller than one (Mescheder et al., 2017)[Proposition 3]. For SimGD applied to a zero sum game with the loss of x given by f , these are points with vanishing gradients such that

$$\text{Id} - M := \text{Id} - \begin{pmatrix} \eta_x D_{xx}^2 f & \eta_x D_{xy}^2 f \\ -\eta_y D_{yx}^2 f & -\eta_y D_{yy}^2 f \end{pmatrix}$$

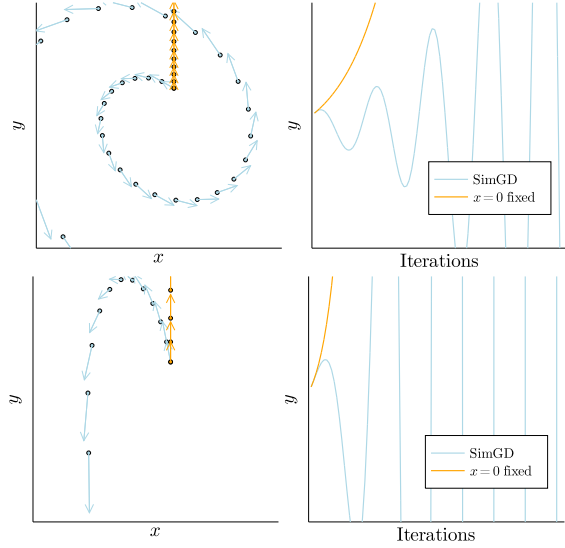


Figure 6. **ICR depends on speed of learning:** When changing the learning rates to $(\eta_x, \eta_y) = (0.03, 0.03)$ (top) or $(\eta_x, \eta_y) = (0.01, 0.09)$ (bottom), SimGD diverges.

has spectral radius smaller than one. For univariate x , y we can set $a := D_{xx}^2 f$, $b := D_{xy}^2 f$, and $c := D_{yy}^2 f$ and compute the characteristic polynomial of M as

$$p(\lambda) = \lambda^2 - (\eta_x a - \eta_y c)\lambda + (-\eta_x \eta_y a c + \eta_x \eta_y b^2).$$

For $\eta_x a > \eta_y c$ and $\eta_x \eta_y b^2 > \eta_x \eta_y a c$ the solutions of this equation have positive real part and therefore the eigenvalues of M have positive real part. By multiplying η_x and η_y by a small enough factor we can obtain a spectral radius smaller than one (c. f. Mazumdar & Ratliff (2018)). Thus, a small enough η_y and large enough mixed derivative b can ensure convergence even for positive c .

If we think of the maximizing player as the discriminator, slow learning (modelled by small η_y) is correlated to good images produced by the generator. Thus, in this interpretation, a good generator leads to ICR stabilizing the point $(0, 0)$ more strongly.

Adversarial training as projection: Surprisingly, ICR allows us to compute a projection with respect to the perceptual distance of a neural network, without quantifying this distance explicitly. Let us consider the following example. We construct a generator \mathcal{G} that maps its 28 weights to a bivariate output. This nonlinear map is modelled as a tiny neural network with two hidden layers, with the final layer restricting the output to the set $\mathcal{S} := \{(e^{s+t}, e^{s-t}) \mid s \in [-\frac{1}{2}, \frac{1}{2}], t \in \mathbb{R}\} \subset \mathbb{R}^2$. We think of this as mapping a set of weights to a generative model that is characterized by only two parameters. In this parameterization, we assume that the target distribution is represented by the point $P_{\text{data}} = (2, 2)$. Importantly, as shown in Figure 7, there is no set of weights that allow the generator

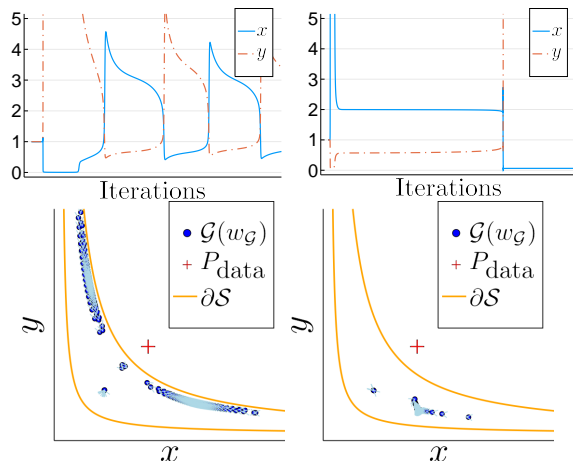


Figure 7. Approximate projection via adversarial training: On the left column, discriminator picks up on errors in the x - and y -direction equally quickly. Therefore, the generator tries to satisfy the criteria alternately, leading to a cyclic pattern. In the right column, the discriminator picks up on errors in the x -direction much more quickly. This causes the generator to try to stay accurate in the x -direction.

to output exactly P_{data} . This is to model the fact that in general, the generator will not be able to exactly reproduce the target distribution. We construct a discriminator \mathcal{D} that maps a generative model (a pair of real numbers) and a set of 28 weights to a real number, by a small densely connected neural network.

We want to model the difference in visual prominence of the two components of P_{data} . To this end, we assume that before before being passed to the discriminator, \mathcal{G} and P_{data} are rescaled by a diagonal matrix $\eta \in \mathbb{R}^{2 \times 2}$. Thus, η determines the relative size of the gradients of \mathcal{D} of the first and second components of the input data. This models the hypothesis that a real discriminator will pick up more quickly on visually prominent features. Importantly, we assume η to be unknown, since we do not have access to a metric measuring "visual similarity to a neural network".

We will now show how adversarial training can be used to approximate a projection with respect to η , without knowing η . We use the loss

$$\min_{w_G \in \mathbb{R}^{28}} \max_{w_D \in \mathbb{R}^{28}} \mathcal{D}(\eta P_{\text{data}}, w_D) - \mathcal{D}(\eta \mathcal{G}(w_G), w_D) \quad (4)$$

and train the two networks using simultaneous gradient descent. For η equal to the identity, we see oscillatory training behavior as \mathcal{G} tries to be accurate first in one, then the other direction. If we instead use $\eta = \begin{pmatrix} 1 & 0 \\ 0 & 10^{-2} \end{pmatrix}$, we are modelling the first component as being more visually prominent. Instead of the oscillatory patterns from before, we observe long periods where the value of the first compo-

nent of $\mathcal{G}(w_G)$ is equal to the first component of P_{data} (see Figure 7). Without knowing η , we have approximated the projection of P_{data} onto \mathcal{S} with respect to the metric given by $(x, y) \mapsto \|\eta(x, y)\|$. To do so, we used the fact that this point is subject to the slowest learning discriminator, and thus the strongest ICR.

We believe that GANs use the same mechanism to compute generators that are close to the true data in the perceptual distance of the discriminator, which in turn acts as a proxy for the perceptual distance of humans.

5. Competitive gradient descent amplifies ICR

How to strengthen ICR: We have provided evidence that GANs' ability to generate visually plausible images can be explained by ICR selectively stabilizing good generators. It is well known that GANs often exhibit unstable training behavior, which is mirrored by the observations in Figures 3, 4 and 7 that ICR often only leads to weak, temporary stability. Thus, it would be desirable to find algorithms that induce stronger ICR than SimGD. To this end, we will find a game-theoretic point of view useful.

Cooperation in a zero-sum game? As discussed in the last section, ICR can stabilize solutions that are locally suboptimal for at least one of the players. Since we did not model either of the two players as altruistic, this behavior may seem puzzling. It is likely for this reason that ICR has mostly been seen as a flaw, rather than a feature of SimGD.

Convergence by competition: The quadratic example in Equation (3) shows that the bilinear term xy is crucial for the presence of ICR. Otherwise, SimGD reduces to each player moving independently according to gradient descent. In fact, the strength of ICR decreases rapidly as $|\alpha|$ and $|\beta|$ diverge to infinity.

The mixed term xy models the ability of each player to retaliate against actions of the other player. In the case of $\beta < 0$, as the maximizing player y moves to plus infinity in order to maximize its reward, it becomes a locally optimal strategy for the minimizing player x to move towards negative infinity in order to minimize the dominant term xy . If $|\beta| \ll 1$ it is favorable for the maximizing player to move back towards zero in order to maximize the dominant term xy . The reason for the maximizing player to stay in the sub-optimal point $y = 0$ (the maximizer of its loss, for $x = 0$) is that the minimizing player can use the mixed term xy to punish every move of y with a counterattack. Thus, the need to avoid counterattacks justifies the seemingly sub-optimal decision of the maximizing player to stay in $y = 0$.

The generator strikes back! This phenomenon is also present in the example of Figure 4. Consider the checkpoint generator from Figure 4 and the over-trained discriminator

that achieves near perfect score against the discriminator. As we can see in Figure 8, training the generator while keeping the over-trained discriminator fixed leads to a rapidly increasing discriminator loss. The over-trained discriminator has become vulnerable to counterattack by the generator! If instead the generator is trained against the checkpoint discriminator, the loss increases only slowly. Thus, ICR can be interpreted as the discriminator trying to avoid counterattack by the generator.

Agent modelling for stronger ICR: The update (x, y) of SimGD applied to the loss function f can be interpreted as the two players solving, at each step, the local optimization problem

$$\min_x x^\top \nabla_x f(x_k, y_k) + \frac{\|x\|^2}{2\eta}, \quad \max_y y^\top \nabla_y f(x_k, y_k) - \frac{\|y\|^2}{2\eta}$$

The terms $x^\top \nabla_x f(x_k, y_k)$, $y^\top \nabla_y f(x_k, y_k)$ express the belief about the loss associated to different actions, based on local information. The quadratic regularization terms express their uncertainty about these beliefs, letting them avoid extreme actions (large steps). However, y (x) does not appear in the local optimization problem of x (y). Thus, the two players are not taking the presence of their opponent into account when choosing their actions. Accordingly, ICR arises only because of the players reaction to, rather than anticipation of each others actions. We propose to strengthen ICR by using local optimization problems that model the players' anticipation of each other's action.

Competitive gradient descent: The updates of competitive gradient descent (CGD) (Schaefer & Anandkumar, 2019) are obtained as Nash equilibria of the local game

$$\begin{aligned} \min_x x^\top \nabla_x f(x_k, y_k) + x^\top [D_{xy} f(x_k, y_k)] y + \frac{\|x\|^2}{2\eta}, \\ \max_y y^\top \nabla_y f(x_k, y_k) + y^\top [D_{yx} f(x_k, y_k)] x - \frac{\|y\|^2}{2\eta}. \end{aligned}$$

Under CGD, the players are aware of each other's presence at every step, since the mixed Hessian $x^\top [D_{xy} f(x_k, y_k)] y$ informs each player, how the simultaneous actions of the other player could affect the loss incurred due to their own action. This element of anticipation strengthens ICR, as indicated by the convergence results provided by Schaefer & Anandkumar (2019). Providing additional evidence, we see in Figure 8 that attempting to over-train the discriminator using CGD leads to a discriminator that is even more robust than the checkpoint discriminator. Applying CGD to the example of Figure 7 also increases the stability of the approximate projection of P_{data} onto \mathcal{S} according to the metric implicit in the discriminator. These results suggest to use CGD to strengthen ICR in GAN training, which we will investigate in the next section. We also expect methods such as LOLA (Foerster et al., 2018) or SGA (Balduzzi et al.,




Figure 8. **ICR and opponent-awareness:** When training the generator for just a few iterations against the over-trained discriminator of Figure 4, the discriminator loss increases rapidly. When attempting to over-train with CGD instead of Adam, the resulting discriminator is even more robust. Similarly, CGD is able to significantly increase the duration for which the generator stays accurate in the (more important) x -direction in Figure 7.

2018; Gemp & Mahadevan, 2018) to strengthen ICR, but a detailed comparison is beyond the scope of this work.

6. Empirical study on CIFAR10

Experimental setup: Based on the last section, CGD strengthens the effects of ICR and should therefore improve GAN performance. We will now investigate this question empirically. In order to make for a fair comparison with Adam, we combine CGD with a simple RMSprop-type heuristic to adjust learning rates, obtaining adaptive CGD (ACGD, see supplement for details). As loss functions, we use the original GAN loss (OGAN) of (1) and the Wasserstein GAN loss function (WGAN) given by

$$\min_G \max_D \mathbb{E}_{x \sim P_{\text{data}}} [\mathcal{D}(x)] - \mathbb{E}_{x \sim P_G} [\mathcal{D}(x)].$$

When using Adam on OGAN, we stick to the common practice of replacing the generator loss by $\mathbb{E}_{x \sim P_G} [-\log(\mathcal{D}(x))]$, as this has been found to improve training stability (Goodfellow et al., 2014; 2016). In order to be generous to existing methods, we use an existing architecture intended for the use with WGAN gradient penalty (Gulrajani et al., 2017). As regularizers, we consider no regularization (NOREG), ℓ_2 penalty on the discriminator with different weights (L2), Spectral normalization (Miyato et al., 2018) on the discriminator (SN), or 1-centered gradient penalty on the discriminator, following (Gulrajani et al., 2017) (GP). Following the advice in (Goodfellow et al., 2016) we train generator and discriminator simultaneously, with the exception of WGAN-GP and Adam, for which we follow (Gulrajani et al., 2017) in making five discriminator updates per generator update. We use the Pytorch implementation of inception score (IS) (Salimans et al., 2016) to compare generator quality.⁴

Experimental results: We will now summarize our main

⁴The Pytorch implementation gives slightly different scores than Tensorflow. We report Tensorflow IS in the supplementary material showing that the relative performance is largely the same.

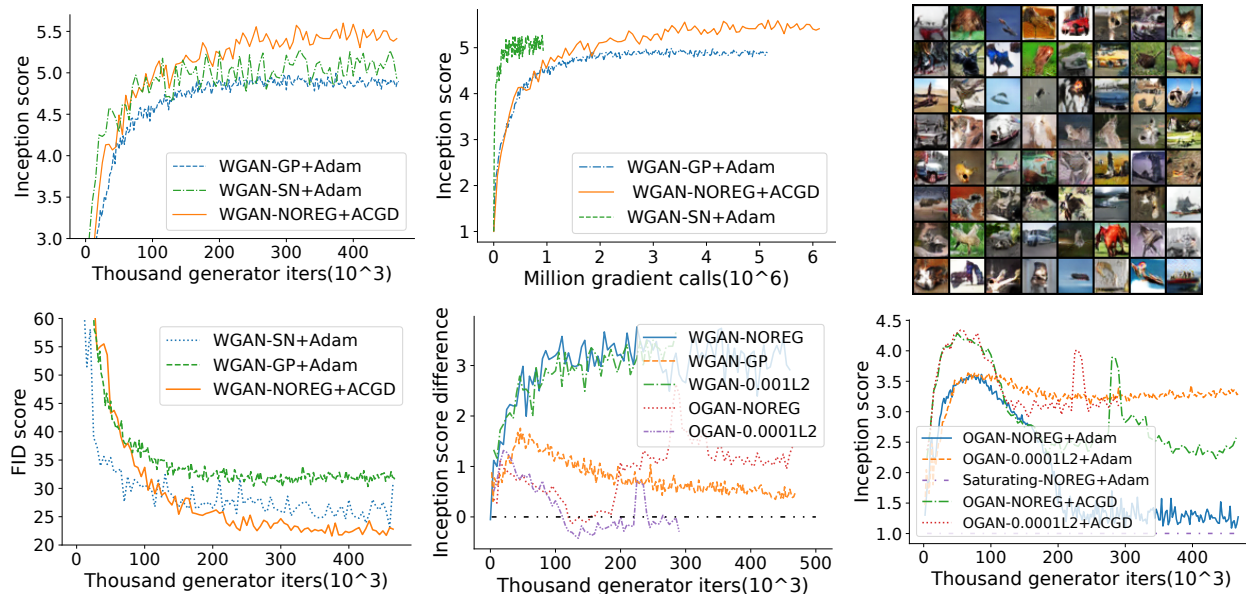


Figure 9. We plot the inception score (IS) against the number of iterations (first panel) and gradient or Hessian-vector product computation (second panel). In the third panel we show final samples of WGAN trained with ACGD and without explicit regularization. In panel four, we compare measure image quality using the Frechet-inception-distance (FID, smaller is better). The results are consistent with those obtained using IS. In panel five, we plot the difference between inception scores between ACGD and Adam (positive values correspond to a larger score for ACGD) over different iterations and models. The only cases where we observe nonconvergence of ACGD are OGAN without regularization or with weight decay of weight 0.0001, as shown in the last panel. The inception score is however still higher than for the same model trained with Adam. When using Adam on the original saturating GAN loss (which we used with ACGD), training breaks down completely.

experimental findings, (see Figure 9). **(1:)** When restricting our attention to the top performing models, we observe that the combination of ACGD with the WGAN loss and without any regularization achieves higher inception score than all other combinations tested. **(2:)** The improvement obtained from training with ACGD persists when measuring image quality according to the Freched-inception-distance (FID) (Heusel et al., 2017). **(3:)** When comparing the number of gradient computations and Hessian-vector products, ACGD is significantly slower than WGAN loss with spectral normalization trained with ADAM, because of the iterative solution of the matrix inverse in ACGD’s update rule. **(4:)** The only instance where we observe erratic behavior with ACGD is when using OGAN without regularization, or with a small ℓ_2 penalty. However, ACGD still outperforms Adam on those cases. In particular training with Adam breaks down completely when using the original saturating loss (as we do for ACGD). **(5:)** When plotting the difference between the inception scores obtained by ACGD and Adam for the same model over the number of iterations, for all models, we observe that ACGD often performs significantly better, and hardly ever significantly worse.

Since CGD strengthens the effects of ICR, the performance improvements obtained with CGD provide further evidence that ICR is a key factor to GAN performance.

7. Conclusion and outlook

In this work, we have pointed out a fundamental flaw present in the static minimax approach to understanding GANs. As an alternative we explain GAN performance with ICR, a mechanism that focuses on the dynamics of simultaneous training. While there is more work left to be done in order to characterize ICR, we provide a number of illustrative experiments on low-dimensional examples and real GANs that supports our conclusions. We also use a game-theoretic interpretation of ICR to identify algorithms such as CGD that can lead to stronger ICR. Indeed, comprehensive experiments on CIFAR10 show systematically improved inception scores and stability when training with CGD, adding further support to our findings.

An important direction for future work is the closer investigation of the generator. Recent work on variational autoencoders (Razavi et al., 2019) and GANs (Karras et al., 2019) suggests that the inductive biases of the generator play an important role, as well. Understanding their interaction with ICR is an important direction of future work. We also hope to better understand the relationship of ICR with local solution concepts such as “proximal equilibria” (Farnia & Ozdaglar, 2020) that emphasize slow improvement of the discriminator.

ACKNOWLEDGMENTS

We would like to thank Houman Owhadi for helpful discussions. A. Anandkumar is supported in part by Bren endowed chair, DARPA PAIHR00111890035, Raytheon, and Microsoft, Google and Adobe faculty fellowships. F. Schäfer gratefully acknowledges support by the Air Force Office of Scientific Research under award number FA9550-18-1-0271 (Games for Computation and Learning) and the Ronald and Maxine Linde Institute of Economic and Management Sciences at Caltech. H. Zheng is supported by Zhiyuan College, Shanghai Jiao Tong University.

References

- Adler, J. and Lunz, S. Banach wasserstein gan. In *Advances in Neural Information Processing Systems*, pp. 6754–6763, 2018.
- Arjovsky, M. and Bottou, L. Towards principled methods for training generative adversarial networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=Hk4_qw5xe.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223, 2017.
- Arora, S., Ge, R., Liang, Y., Ma, T., and Zhang, Y. Generalization and equilibrium in generative adversarial nets (gans). In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 224–232. JMLR. org, 2017.
- Arora, S., Cohen, N., Hu, W., and Luo, Y. Implicit regularization in deep matrix factorization. *arXiv preprint arXiv:1905.13655*, 2019.
- Azizan, N., Lale, S., and Hassibi, B. Stochastic mirror descent on overparameterized nonlinear models: Convergence, implicit regularization, and generalization. *arXiv preprint arXiv:1906.03830*, 2019.
- Balduzzi, D., Racaniere, S., Martens, J., Foerster, J., Tuyls, K., and Graepel, T. The mechanics of n-player differentiable games. *arXiv preprint arXiv:1802.05642*, 2018.
- Berard, H., Gidel, G., Almahairi, A., Vincent, P., and Lacoste-Julien, S. A closer look at the optimization landscapes of generative adversarial networks. *arXiv preprint arXiv:1906.04848*, 2019.
- Berkels, B., Effland, A., and Rumpf, M. Time discrete geodesic paths in the space of images. *SIAM Journal on Imaging Sciences*, 8(3):1457–1488, 2015.
- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- Chatterjee, S. Coherent gradients: An approach to understanding generalization in gradient descent-based optimization. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=ryeFY0EFwS>.
- Chen, Y., Telgarsky, M., Zhang, C., Bailey, B., Hsu, D., and Peng, J. A gradual, semi-discrete approach to generative network training via explicit wasserstein minimization. *arXiv preprint arXiv:1906.03471*, 2019.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- Dalal, N. and Triggs, B. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR’05)*, volume 1, pp. 886–893. IEEE, 2005.
- Effland, A., Rumpf, M., and Schäfer, F. Image extrapolation for the time discrete metamorphosis model: Existence and applications. *SIAM Journal on Imaging Sciences*, 11(1):834–862, 2018.
- Farnia, F. and Ozdaglar, A. Gans may have no nash equilibria. *arXiv preprint arXiv:2002.09124*, 2020.
- Fiez, T., Chasnov, B., and Ratliff, L. J. Convergence of learning dynamics in stackelberg games. *arXiv preprint arXiv:1906.01217*, 2019.
- Foerster, J., Chen, R. Y., Al-Shedivat, M., Whiteson, S., Abbeel, P., and Mordatch, I. Learning with opponent-learning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 122–130. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- Gemp, I. and Mahadevan, S. Global convergence to the equilibrium of gans using variational inequalities. *arXiv preprint arXiv:1808.01531*, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.
- Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. *Deep learning*, volume 1. MIT Press, 2016.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein gans.

- In Advances in neural information processing systems, pp. 5767–5777, 2017.
- Gunasekar, S., Woodworth, B. E., Bhojanapalli, S., Neyshabur, B., and Srebro, N. Implicit regularization in matrix factorization. In Advances in Neural Information Processing Systems, pp. 6151–6159, 2017.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In Proceedings of the 34th International Conference on Machine Learning-Volume 70, pp. 1321–1330. JMLR.org, 2017.
- Haasdonk, B. and Burkhardt, H. Invariant kernel functions for pattern analysis and machine learning. Machine learning, 68(1):35–61, 2007.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In Advances in Neural Information Processing Systems, pp. 6626–6637, 2017.
- Huang, G., Berard, H., Touati, A., Gidel, G., Vincent, P., and Lacoste-Julien, S. Parametric adversarial divergences are good task losses for generative modeling. arXiv preprint arXiv:1708.02511, 2017.
- Jin, C., Netrapalli, P., and Jordan, M. I. Minmax optimization: Stable limit points of gradient descent ascent are locally optimal. arXiv preprint arXiv:1902.00618, 2019.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 4401–4410, 2019.
- Kodali, N., Abernethy, J., Hays, J., and Kira, Z. On convergence and stability of gans. arXiv preprint arXiv:1705.07215, 2017.
- Kolmogorov, A. N. The representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables. Doklady Akademii Nauk SSSR, 108(2):179–182, 1956.
- Kubo, M., Banno, R., Manabe, H., and Minoji, M. Implicit regularization in over-parameterized neural networks. arXiv preprint arXiv:1903.01997, 2019.
- Li, C.-L., Chang, W.-C., Cheng, Y., Yang, Y., and Póczos, B. Mmd gan: Towards deeper understanding of moment matching network. In Advances in Neural Information Processing Systems, pp. 2203–2213, 2017.
- Li, Y., Swersky, K., and Zemel, R. Generative moment matching networks. In International Conference on Machine Learning, pp. 1718–1727, 2015.
- Lucic, M., Kurach, K., Michalski, M., Gelly, S., and Bousquet, O. Are gans created equal? a large-scale study.” arxiv e-prints. arXiv preprint arXiv:1711.10337, 2017.
- Ma, C., Wang, K., Chi, Y., and Chen, Y. Implicit regularization in nonconvex statistical estimation: Gradient descent converges linearly for phase retrieval, matrix completion and blind deconvolution. arXiv preprint arXiv:1711.10467, 2017.
- Mazumdar, E. and Ratliff, L. J. On the convergence of gradient-based learning in continuous games. arXiv preprint arXiv:1804.05464, 2018.
- Mazumdar, E. V., Jordan, M. I., and Sastry, S. S. On finding local nash equilibria (and only local nash equilibria) in zero-sum games. arXiv preprint arXiv:1901.00838, 2019.
- Mescheder, L., Nowozin, S., and Geiger, A. The numerics of gans. In Advances in Neural Information Processing Systems, pp. 1825–1835, 2017.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. arXiv preprint arXiv:1802.05957, 2018.
- Mroueh, Y., Li, C.-L., Sercu, T., Raj, A., and Cheng, Y. Sobolev gan. arXiv preprint arXiv:1711.04894, 2017.
- Neyshabur, B. Implicit regularization in deep learning. arXiv preprint arXiv:1709.01953, 2017.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In Advances in neural information processing systems, pp. 271–279, 2016.
- Razavi, A., van den Oord, A., and Vinyals, O. Generating diverse high-fidelity images with vq-vae-2. In Advances in Neural Information Processing Systems, pp. 14866–14876, 2019.
- Roth, K., Lucchi, A., Nowozin, S., and Hofmann, T. Stabilizing training of generative adversarial networks through regularization. In Advances in neural information processing systems, pp. 2018–2028, 2017.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. In Advances in neural information processing systems, pp. 2234–2242, 2016.
- Schaefer, F. and Anandkumar, A. Competitive gradient descent. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), Advances in Neural Information Processing Systems 32, pp. 7623–7633. Curran Associates, Inc.,

2019. URL <http://papers.nips.cc/paper/8979-competitive-gradient-descent.pdf>.

Simard, P. Y., LeCun, Y. A., Denker, J. S., and Victorri, B. Transformation invariance in pattern recognition—tangent distance and tangent propagation. In Neural networks: tricks of the trade, pp. 239–274. Springer, 1998.

Song, Y., McLoughlin, I. V., and Dai, L.-R. Local coding based matching kernel method for image classification. PloS one, 9(8), 2014.

Trouvé, A. and Younes, L. Metamorphoses through lie group action. Foundations of Computational Mathematics, 5(2):173–198, 2005.

Uppal, A., Singh, S., and Poczos, B. Nonparametric density estimation & convergence rates for gans under besov ipm losses. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), Advances in Neural Information Processing Systems 32, pp. 9086–9097. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9109-nonparametric-density-estimation-convergence-rates-for-gans-under-besov-ipm-losses.pdf>.

Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. arXiv preprint arXiv:1611.03530, 2016.