

Local Convergence of Gradient Descent-Ascent for Training Generative Adversarial Networks

Evan Becker
CS, UCLA
evbecker@ucla.edu

Parthe Pandit
HDSI, UCSD
parthepandit@ucsd.edu

Sundeep Rangan
ECE, NYU
srangan@nyu.edu

Alyson K. Fletcher
Statistics, UCLA
akfletcher@ucla.edu

Abstract—Generative Adversarial Networks (GANs) are a popular formulation to train generative models for complex high dimensional data. The standard method for training GANs involves a gradient descent-ascent (GDA) procedure on a minimax optimization problem. This procedure is hard to analyze in general due to the nonlinear nature of the dynamics. We study the local dynamics of GDA for training a GAN with a kernel-based discriminator. This convergence analysis is based on a linearization of a non-linear dynamical system that describes the GDA iterations, under an *isolated points model* assumption from [2]. Our analysis brings out the effect of the learning rates, regularization, and the bandwidth of the kernel discriminator, on the local convergence rate of GDA. Importantly, we show phase transitions that indicate when the system converges, oscillates, or diverges. We also provide numerical simulations that verify our claims. A full version with complete proofs is available on arXiv [3].

I. INTRODUCTION

Generative Adversarial Networks (GANs) are a class of probabilistic generative models that avoid expensive likelihood computations while still providing good sample quality [4]. To fit complex data distributions, two models (typically deep neural networks) are trained in an alternating manner: a *generator* \mathcal{G} learns a deterministic map from a latent space \mathcal{Z} to the data space \mathcal{X} , while a *discriminator* or *critic* model \mathcal{D} attempts to discern whether a sample belongs to the training dataset or the generated dataset.

The discriminator plays an important, yet poorly understood, role in the training of a GAN. It is well known from [4] that if the discriminator is trained to minimize the cross-entropy between true and generated samples, the generator would minimize the Jensen-Shannon divergence between distributions. Similarly different choices for discriminator loss functions lead to a variety of f -divergences [13] and probability metrics [5].

In practice, however, we apply GDA for training GANs, whereby the discriminator is not allowed to converge, making analysis of the iterative training extremely difficult. Furthermore, training commonly suffers from empirical breakdowns such as mode collapse, in which the entire generated distribution converges to a small portion of the target distribution [14]. The generator may even fail to converge entirely when gradients from the discriminator are too small for the generator to proceed during training. Without an understanding of when and how these phenomena occur, practitioners have to rely on heuristics and extensive hyperparameter tuning based on trial and error procedures [7], [10], [14].

In this work, we characterize the local convergence rates of GAN training when the discriminator is kernel-based. This choice of the discriminator model is motivated by the recently discovered equivalence between wide neural networks and kernel methods via the Neural Tangent Kernel framework [6]. While the discriminators problem is simplified due to the kernel-based discriminator, the overall dynamics of the generated samples remain non-linear and complex, and hence retain many of the properties exhibited by GANs in practice such as mode collapse and divergence [2].

A. Prior Work on Linear GANs and Main Contributions

Stability analysis for GANs under stylized settings goes back to the Dirac-GAN framework from [9], which looked at the local stability of a two-point system using a linear discriminator to demonstrate examples of catastrophic forgetting. Other GAN works use a similar linearization analysis, such as [11], [12]. The isolated points model proposed by [2] allowed for a more complex model while remaining analytically tractable, by letting the generated probability mass differ from the true mass in various isolated regions. We provide new insight into the framework proposed by [2] by going beyond stability analysis and characterizing rates of convergence.

We analyze the local convergence of the non-linear dynamical system that describes the GDA iterates, in settings when the equilibrium is stable. Our analysis is based on a linearization of these non-linear dynamics. We show how changing the kernel width can improve the rate of convergence, while also highlighting a phase transition under which the convergence remains unaffected by changes in the kernel width.

II. MODEL

We investigate the training dynamics of a GAN where the target distribution and the generated distribution are discrete point masses, following the framework of [2]. We highlight key elements of our model below.

A. Target and Generated Distributions

We assume that the target and generated distributions consist of point masses with density functions over $x \in \mathbb{R}^d$ given by

$$\mathbb{P}_r(x) = \sum_{i=1}^{N_r} p_i \delta(x - x_i), \quad \mathbb{P}_g(x) = \sum_{j=1}^{N_g} \tilde{p}_j \delta(x - \tilde{x}_j), \quad (1)$$

where δ is the Dirac delta function, $X = \{x_i\}_{i=1}^{N_r}$ and $\tilde{X} = \{\tilde{x}_j\}_{j=1}^{N_g}$ are the true and generated points, and $\{p_i\}_{i=1}^{N_r}$ and $\{\tilde{p}_j\}_{j=1}^{N_g}$ are their (fixed) probability masses. The problem we consider is learning the locations \tilde{X} so that the generated and true distributions match. Thus the decision variable of the generator model is \tilde{X} . This simplification is justified since we wish to study the role of the discriminator in this work.

B. Kernel Discriminator

The GAN discriminator is a function $f : \mathcal{X} \rightarrow \mathbb{R}$ which predicts whether a sample x is real or fake based on $\text{sign}(f(x))$. In this paper we assume that the discriminator belongs to a Reproducing Kernel Hilbert Space \mathcal{H} corresponding to a bivariate positive definite kernel function $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. The discriminator defines a maximum mean discrepancy (MMD) metric

$$\text{MMD}(\mathbb{P}_r, \mathbb{P}_g) := \max_{\substack{f \in \mathcal{H} \\ \|f\| \leq 1}} \sum_{i=1}^{N_r} p_i f(x_i) - \sum_{i=1}^{N_g} \tilde{p}_i f(\tilde{x}_j) \quad (2)$$

between \mathbb{P}_r and \mathbb{P}_g , which the generator tries to minimize.

C. Minimax Optimization Formulation for Training a GAN

We assume a mini-max loss function similar to [1], [2], [8] of the form:

$$\min_{\tilde{X}} \max_{f \in \mathcal{H}} \mathcal{L}(f, \tilde{X}) \quad (3a)$$

$$\mathcal{L}(f, \tilde{X}) := \sum_{i=1}^{N_r} p_i f(x_i) - \sum_{i=1}^{N_g} \tilde{p}_i f(\tilde{x}_j) - \frac{\lambda}{2} \|f\|_{\mathcal{H}}^2. \quad (3b)$$

The regularization parameter $\lambda > 0$ is some constant, that acts as a Lagrange multiplier for the optimization problem in eq. (2). The loss is a function of the discriminator f and generated samples \tilde{X} .

D. Training Dynamics of Gradient Descent Ascent

We assume the generator performs gradient descent on the above minimax optimization problem with a step size η_g and the discriminator performs gradient ascent with step size η_d . We let (f^t, \tilde{X}^t) denote the discriminator and generated samples in step t .

$$f^{t+1} = f^t + \eta_d \frac{\partial}{\partial f} \mathcal{L}(f^t, \tilde{X}^t) \quad (4a)$$

$$\tilde{X}^{t+1} = \tilde{X}^t - \eta_g \frac{\partial}{\partial \tilde{X}} \mathcal{L}(f^t, \tilde{X}^t) \quad (4b)$$

where the first equation uses the Fréchet derivative with respect to f . This can be simplified since the loss function $\mathcal{L}(f, \tilde{X})$ only consists of linear and quadratic terms of f . Observe that for any u in \mathbb{R}^d , the linear term $f(u) = \langle f, K(u, \cdot) \rangle_{\mathcal{H}}$ due to the reproducing property of the kernel whereby, $\frac{\partial}{\partial f} f(u)$ is the function $v \mapsto K(u, v)$, denoted $K(u, \cdot)$.

Using the loss function in eq. (3), we get the updates

$$f^{t+1} = (1 - \lambda \eta_d) f^t + \eta_d \left(K(\cdot, X) \mathbf{p} - K(\cdot, \tilde{X}_i) \tilde{\mathbf{p}} \right) \quad (5a)$$

$$\tilde{x}_j^{t+1} = \tilde{x}_j^t + \eta_g \tilde{p}_j \nabla f^t(\tilde{x}_j^t), \quad \forall j = 1, 2, \dots, N_g \quad (5b)$$

Notice that (5a) is linear in f , whereby we can simplify these equations further. The following lemma simplifies eq. (5) by eliminating the discriminator f .

Lemma 1 (Training Dynamics). *Assume $f_0 = 0$, the zero function in the RKHS \mathcal{H} . Then the following deterministic dynamical system describes the evolution of the samples generated using eq. (5).*

$$\begin{aligned} \tilde{X}^{t+1} = & \tilde{X}^t + \eta_d \eta_g \sum_{s=0}^t (1 - \lambda \eta_d)^{t-s} \times \\ & \tilde{\mathbf{p}} \odot \left(\nabla_1 K(\tilde{X}^t, X) \mathbf{p} - \nabla_1 K(\tilde{X}^t, \tilde{X}^s) \tilde{\mathbf{p}} \right) \end{aligned} \quad (6)$$

Note that the above dynamical system is nonlinear in \tilde{X} , and is non-Markovian due to dependence of \tilde{X}^{t+1} on $\{\tilde{X}^s\}_{s \leq t}$.

The term $\tilde{\mathbf{p}} \odot \nabla_1 K(\tilde{X}, X) \mathbf{p}$ can be thought of as a *drift*, whereas $\tilde{\mathbf{p}} \odot \nabla_1 K(\tilde{X}, \tilde{X}) \tilde{\mathbf{p}}$ can be thought of as a *diffusion*.

From eq. (6) we can immediately infer a condition for a set of generated points \tilde{X}^* to be in equilibrium

Lemma 2. *A set of points \tilde{X}^* is in equilibrium for the dynamics eq. (6) if and only if*

$$\nabla_1 K(\tilde{X}^*, X) \mathbf{p} = \nabla_1 K(\tilde{X}^*, \tilde{X}^*) \tilde{\mathbf{p}}. \quad (7)$$

Remark 1. The set of equilibrium points depend only on the kernel K and are invariant to the hyperparameters η_d, η_g, λ . However the dynamics and convergence properties of these equilibria depend on η_d, η_g, λ as well as σ .

E. Model and Optimization Hyperparameters

This setting has two model hyperparameters that control the smoothness of the discriminator. The regularization λ controls the \mathcal{H} -norm of the discriminator, which is a measure of global smoothness. In contrast, the kernel bandwidth σ is a measure of the local smoothness. We also have two optimization hyperparameters, which are the learning rate of the generator η_g and the learning rate of the discriminator η_d . In practice often $\eta_g \ll \eta_d$.

III. LOCAL CONVERGENCE AROUND TRUE SAMPLES

A. Assumptions on the kernel function

We assume the the kernel $K(x, x')$ is smooth and, at each true point x_i :

$$\nabla_1 K(x_i, x_i) = \frac{\partial K(x, x_i)}{\partial x} \bigg|_{x=x_i} = \mathbf{0}, \quad (8a)$$

$$-\frac{\partial^2 K(x, x_i)}{\partial x^2} \bigg|_{x=x_i} = \frac{\partial^2 K(x, x')}{\partial x \partial x'} \bigg|_{x=x'=x_i} = \frac{1}{\sigma^2} \mathbf{I}_d \quad (8b)$$

for some $\sigma > 0$ that we call the *kernel width* and represents the curvature of the kernel around $x = x_i$. Note that (8a) and (8b) are satisfied for the standard RBF kernel:

$$K(x, x') = \exp\left(-\frac{1}{2\sigma^2} \|x - x'\|^2\right). \quad (9)$$

Proposition 1. *Under the above assumption, \tilde{X}^* such that $\tilde{x}_j^* = x_i$ for some i , is an equilibrium.*

This follows immediately from the observation in eq. (7). When the assumptions on the kernel eq. (8) are satisfied, both sides of eq. (7) vanish. The results in [2] analyzed the stability of this equilibrium under an isolated points model described below, which localizes the analysis around each true point.

B. Isolated Points Model

We assume the true samples are separated far enough so that there exists a non-empty *isolated neighborhood* V_i around each sample x_i such that,

$$K(x, x') = 0 \text{ for all } x \in V_i \text{ and } x' \in V_j \text{ for all } i \neq j. \quad (10)$$

In other words, the generated points are separated sufficiently far apart such that they are outside the width of the kernel evaluated at another sample. We let \mathcal{N}_i be the set of indices j such that the generated points $\tilde{x}_j^t \in V_i$, for all t .

Thus the dynamics we study can be written as

$$\begin{aligned} \tilde{X}_i^{t+1} &= \tilde{X}_i^t + \eta_d \eta_g \sum_{s=0}^t (1 - \lambda \eta_d)^{t-s} \times \\ &\quad \tilde{p}_i \odot \left(\nabla_1 K(\tilde{X}_i^t, x_i) p_i - \nabla_1 K(\tilde{X}_i^t, \tilde{X}_i^s) \tilde{p}_i \right) \end{aligned} \quad (11)$$

where \tilde{X}_i^t are points generated inside the region V_i , and \tilde{p}_i is the length $|\mathcal{N}_i|$ subvector of \tilde{p} corresponding to these points.

Under this assumption, if $j \in \mathcal{N}_i$ and $k \notin \mathcal{N}_i$, then eq. (11) ignores interaction terms between (\tilde{x}_k, x_i) , and $(\tilde{x}_j, \tilde{x}_k)$, compared to eq. (6). Note eq. (6) tracks $N_r N_g + N_g^2$ interaction terms whereas eq. (11) only tracks $|\mathcal{N}_i| + |\mathcal{N}_i|^2$ terms where $|\mathcal{N}_i|$ is the number of generated points inside V_i .

We will call the updates eq. (11) the *dynamical system in the region V_i* . For the purpose of analysis it is beneficial to write the dynamics involving both the discriminator and the generator as below:

$$f_i^{t+1} = (1 - \lambda \eta_d) f^t + K(\cdot, x_i) p_i - \sum_{i \in \mathcal{N}_i} K(\cdot, \tilde{x}_i) \tilde{p}_i \quad (12a)$$

$$\tilde{x}_j^{t+1} = \tilde{x}_j^t + \eta_g \tilde{p}_j \nabla f^t(\tilde{x}_j) \quad x_j \in V_i \quad (12b)$$

Under the isolated points model, the discriminator satisfies

$$f^t(u) = \sum_{i=1}^{N_r} f_i^t(u), \quad \forall u \quad (13a)$$

$$f^t(x) = f_i^t(x) \quad \forall x \in V_i. \quad (13b)$$

C. Main result

Given a local region V_i , we wish to study the dynamics of the local system given by eq. (12) where \tilde{x}_j are close to x_i for all $j \in \mathcal{N}_i$. That is, all the generated points are close to the true point in that region. To this end, we write the local updates (12) as a mapping

$$(f_i^{t+1}, \tilde{X}_i^{t+1}) = \Phi_i(f_i^t, \tilde{X}_i^t), \quad (14)$$

where $\Phi_i(\cdot)$ represents the update function in (12). Also, let

$$\tilde{X}_i^* = \{\tilde{x}_j^*, j \in \mathcal{N}_i\}, \quad \tilde{x}_j^* = x_i. \quad (15)$$

It is shown in [2] that there exists a parameter vector f_i^* such that (f_i^*, \tilde{X}_i^*) is an *equilibrium point* of $\Phi_i(\cdot)$ in that

$$(f_i^*, \tilde{X}_i^*) = \Phi_i(f_i^*, \tilde{X}_i^*). \quad (16)$$

The condition (16) implies that if $(f_i^t, \tilde{X}_i^t) = (f_i^*, \tilde{X}_i^*)$ for some t , then (f_i^s, \tilde{X}_i^s) will remain at (f_i^*, \tilde{X}_i^*) for all subsequent times $s \geq t$. Let \mathbf{J}_i^* denote the Jacobian of the update mapping $\Phi_i(\cdot)$ at the equilibrium point (f_i^*, \tilde{X}_i^*) and define the spectral radius of the Jacobian

$$\rho_{\max} := \rho_{\max}(\mathbf{J}_i^*) = \max \{ |\rho| \mid \rho \in \text{spec}(\mathbf{J}_i^*) \}, \quad (17)$$

where $\text{spec}(\mathbf{J}_i^*)$ is the spectrum of \mathbf{J}_i^* , i.e., its eigenvalues.

A well-known result of non-linear systems theory [16] is that the equilibrium point (f_i^*, \tilde{X}_i^*) is *locally asymptotically stable* if $\rho_{\max}(\mathbf{J}_i^*) < 1$. Conversely, if $\rho_{\max}(\mathbf{J}_i^*) > 1$, the system can be shown to be *locally unstable* – see [16] for precise definitions. Hence, $\rho_{\max}(\mathbf{J}_i^*)$ can provide necessary and sufficient conditions for local stability. Also, if $\rho_{\max} < 1$ and the system is initialized at (f_i^0, \tilde{X}_i^0) sufficiently close to (f_i^*, \tilde{X}_i^*) then, the components will converge geometrically as

$$\|\tilde{x}_j^t - \tilde{x}_i\| \leq C \rho_{\max}^t \|\tilde{x}_j^0 - \tilde{x}_i\|, \quad (18)$$

for some constant C . Hence, the spectral radius ρ_{\max} also provides a measure of the convergence rate of the system. Our main theorem below applies this result to obtain an exact characterization of the convergence rate of the local dynamics by studying the spectrum of \mathbf{J}^* in terms of the model and optimization hyperparameters.

Recall the model hyperparameters: σ – discriminator kernel width, λ – IPM regularization, η_d – learning rate of discriminator and η_g – learning rate of the generator.

Theorem 1. Consider the isolated neighborhood training dynamics in (12) under the assumptions in Section II in some region V_i . Additionally, assume that the weights of the generated points are equal so that $\tilde{p}_j = \tilde{p}$ for some $\tilde{p} > 0$ and all $j \in \mathcal{N}_i$. Define

$$a := \lambda, \quad b := \frac{\mu \tilde{p} \Delta_i}{\lambda \sigma^2}, \quad c := \frac{\mu \tilde{p} p_i}{\sigma^2}, \quad \mu := \frac{\eta_g}{\eta_d}, \quad (19)$$

and

$$\Delta_i := p_i - \sum_{j \in \mathcal{N}_i} \tilde{p}_j = p_i - |\mathcal{N}_i| \tilde{p}. \quad (20)$$

Then, the eigenvalues of the \mathbf{J}^* are of the form

$$\rho = 1 - \eta_d \nu, \quad (21)$$

where ν is from the set:

$$\nu \in \begin{cases} \{a, b, m \pm \sqrt{m^2 - c}\} & \text{if } |\mathcal{N}_i| > 1 \\ \{a, m \pm \sqrt{m^2 - c}\} & \text{if } |\mathcal{N}_i| = 1. \end{cases} \quad (22)$$

where $m = (a + b)/2$.

The proof of the result is given in Appendix A and builds on the linear analysis in [2]. The theorem above gives an exact characterization of the eigenvalues of the linear system in terms

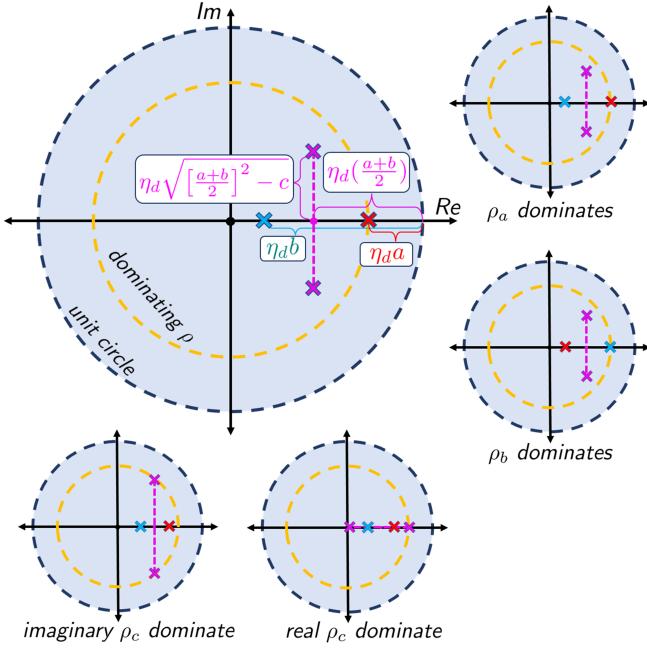


Fig. 1. Eigenvalues of the linearized system in Theorem 1. Here $a := \lambda$ the regularization, $b := \frac{\mu \tilde{p} \Delta_i}{\lambda \sigma^2}$, and $c := \frac{\mu \tilde{p} p_i}{\sigma^2}$, where $\mu := \frac{\eta_g}{\eta_d}$, where σ is the kernel width, and η_d and η_g are learning rates of the discriminator and generator respectively.

of the key parameters including the step sizes and kernel width.

D. Selecting the step size

An immediate consequence of Theorem 1 is that it guides the selection of the step-sizes that ensure local stability. As described above, for local stability, we wish that $|\rho| < 1$ for all ρ in (21). The following provides necessary and sufficient conditions on η_d for this stability condition to occur.

Corollary 1. *Under the conditions in Theorem 1, the spectral radius of the Jacobian, $\rho_{\max}(\text{spec}(\mathbf{J}_i^*)) < 1$, if and only if:*

$$0 < \eta_d < \begin{cases} \min\left\{\frac{2}{a}, \frac{2}{b}, \frac{a+b}{c}\right\} & \text{if } |\mathcal{N}_i| > 1 \\ \min\left\{\frac{2}{a}, \frac{a+b}{c}\right\} & \text{if } |\mathcal{N}_i| = 1. \end{cases} \quad (23)$$

In particular, by choosing η_d small enough, we can always guarantee the system is locally stable when $\lambda > 0$ and $\Delta_i > 0$, meaning that there is at least some regularization and the true point mass exceeds the generated point mass. We can also derive a simple sufficient condition:

Proposition 2 (Sufficient condition for stability). *The training dynamics eq. (12) are stable around equilibrium \tilde{X}_i^* from eq. (15) for all $\Delta_i \in (0, p_i)$ if,*

$$\eta_d < \frac{2}{\lambda}, \quad \text{and} \quad \eta_g < \lambda \sigma^2. \quad (24)$$

The rest of the paper assumes eq. (24) holds and derives convergence rates based on the choice of kernel width σ .

IV. CONVERGENCE RATE AND KERNEL WIDTH

Theorem 1 can also provide insights into the relation of the convergence rate to the system parameters. Specifically, recall from eq. (18) that the spectral radius, $\rho_{\max}(\mathbf{J}_i^*)$ defined in eq. (17), determines the convergence rate of the local dynamics, i.e., ρ_{\max} closer to 1 indicating slower convergence and ρ_{\max} closer to 0 indicates faster convergence. Now, among the values in (21), the ρ that maximizes $|\rho|$ will be one of three values:

$$\rho_a = 1 - \eta_d a, \quad \rho_b = 1 - \eta_d b, \quad \rho_c = 1 - \eta_d (m - \sqrt{m^2 - c}), \quad (25)$$

where $m = (a + b)/2$. The cases when the different values dominate are shown in Figure 1.

It is clear from eq. (25) that controlling the dominant eigenvalue by adjusting the relevant hyperparameters can improve the rate of convergence.

Remark 2 (Saturation with Kernel Width). We now share a phase of the dynamical system where changing σ does not affect the convergence rate. Consider the dynamics (12) with fixed \tilde{p} , and $|\mathcal{N}_i|$. Furthermore, assume η_d is fixed such that the condition from Corollary 1 is satisfied. Then changing the kernel width parameter σ^2 cannot improve the convergence rate in the following settings:

- When all eigenvalues are real and ρ_a dominates. Equivalently $c < \frac{1}{4}(a + b)^2$ and $a < \min\{b, m - \sqrt{m^2 - c}\}$.
- When two eigenvalues are complex and ρ_a dominates. Equivalently $c > \frac{1}{4}(a + b)^2$ and $a < \min\{b, 2m - \eta_d c\}$.

A. Diminishing learning rate

One example regime in which this saturation can clearly be understood is when the learning rate is small, Δ_i is positive, and ρ_c is complex. When the learning rate is small, the magnitude for any eigenvalue of the form $\rho_\nu = 1 - \eta_d \nu$ can be approximated by $|\rho_\nu|^2 \approx 1 - 2\eta_d \text{Re}(\nu) + \mathcal{O}(\eta_d^2)$. This means that we have eigenvalues with approximate magnitudes:

$$|\rho_a|^2 \approx 1 - 2\eta_d \lambda \quad (26a)$$

$$|\rho_b|^2 \approx 1 - 2\eta_d \frac{\mu \tilde{p} \Delta_i}{\lambda \sigma^2} \quad (26b)$$

$$|\rho_c|^2 \approx 1 - \eta_d \left(\lambda + \frac{\mu \tilde{p} \Delta_i}{\lambda \sigma^2} \right) \quad (26c)$$

This yields that the largest eigenvalue is

$$1 - 2\eta_d \cdot \min \left\{ \lambda, \frac{\mu \tilde{p} \Delta_i}{\lambda \sigma^2} \right\} \quad (27)$$

Thus reducing the kernel width σ below $\sqrt{\mu \tilde{p} \Delta_i} / \lambda$, does not lead to changes in the convergence rate $1 - 2\eta_d \lambda$.

V. NUMERICAL RESULTS

In this section, we demonstrate the accuracy of our linearized dynamics by comparing predicted convergence to actual GAN training behavior around local equilibrium.

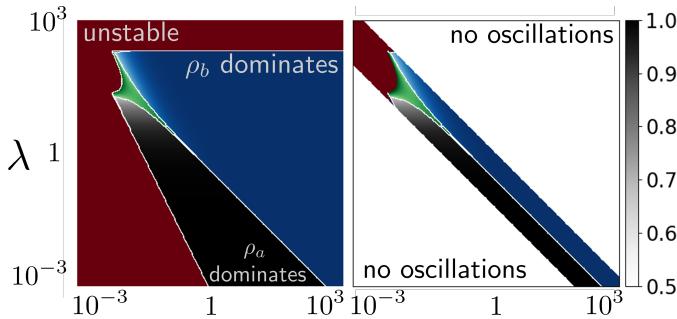


Fig. 2. (Left) Heat map of $|\rho_{\max}|^2$ as a function of (σ, λ) , the kernel width and regularization. The colorbar shows the intensity of ρ_{\max} . Learning rates are $\eta_d = \eta_g = 1e^{-2}$, $p_i = 1$ and $\tilde{p} = 0.8$ whereby $\Delta_i = 0.2$, $\mu = 1$. The region for which $|\rho_{\max}|^2 > 1$ is highlighted in red, where ρ_a dominates is in black, ρ_b dominates is in blue and where ρ_c dominates is in green. In the black region where ρ_a dominates, ρ_{\max} is insensitive to σ . (Right) The region in which ρ_c is complex is selected from the heat map on the left. In this region the system shows oscillatory behavior due to the imaginary part of ρ_c .

a) *Phase transitions*: In Figure 2, we plot the heatmap of dominating eigenvalue magnitude for a range of regularization and kernel width settings. Note that in this figure we use small learning rate ($\eta_d, \eta_g = 1e^{-2}$), meaning firstly that the system is stable for almost all choices of hyperparameters (Figure 2a). In the middle plot (Figure 2b), the majority of fast convergence behaviors occur when ρ_c has imaginary components. To analytically find this region, the condition $m^2 < c$ provides a quadratic inequality in terms of $\gamma = 1/\sigma^2$, from which the roots tell us the exact ranges of kernel widths. When $\Delta = 0$, $\gamma > \frac{\lambda^2}{4\mu\tilde{p}}$, meaning a small enough kernel width will always result in oscillatory behavior. When $\Delta \neq 0$, we have $\gamma \in \frac{\lambda^2}{\Delta^2\tilde{p}\mu} (2p - \Delta \pm 2p\sqrt{1 - \Delta/p})$, meaning extremely small or extremely large kernel widths will have no oscillations. Lastly for the right plot (Figure 2c), we highlight the range of kernel widths that for a given regularization strength do not affect the convergence rate (Remark 2). For positive Δ (more target mass than generated), this region intuitively begins where $\rho_a = \rho_b$: as kernel width shrinks further, the magnitude of both ρ_b and ρ_c shrink, leaving ρ_a fixed and dominating.

In Fig. 3, observe that our approximation matches true training dynamics very precisely when the learning rate is small. Note that in the small learning rate regime, corollary 1 correctly predicts stability for all simulations. In this setting, decreasing the kernel width and increasing regularization can speed up the convergence of the generated point. However, for small regularization, the effect of kernel width is negligible, as predicted by the large saturation region in Fig. 2.

VI. CONCLUSION

In this paper we consider a stylized analysis of GAN training using gradient descent ascent. We assumed that the generator was unconstrained and the discriminator was a kernel model (or equivalently a wide neural network in the kernel regime). The analysis uncovers the role of (i) kernel width (or equivalently network depth), (ii) regularization, and (iii)

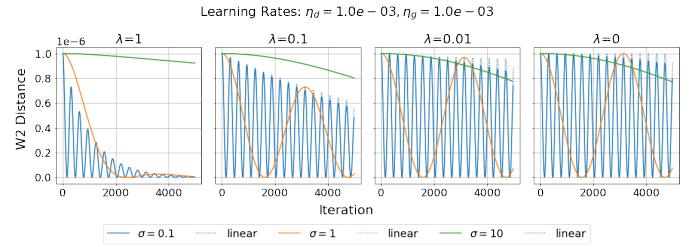


Fig. 3. GAN training behavior with RBF kernel discriminator. One generated point is initialized close to the target, in order to compare predicted local convergence rates (dotted lines) to empirical convergence (solid lines).

learning rates of generator and discriminator, on the stability and local convergence rate of the dynamics.

REFERENCES

- [1] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks, 2017.
- [2] Evan Becker, Parthe Pandit, Sundeep Rangan, and Alyson Fletcher. Instability and local minima in GAN training with kernel discriminators. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022.
- [3] Evan Becker, Parthe Pandit, Sundeep Rangan, and Alyson K Fletcher. Local convergence of gradient descent-ascent for training generative adversarial networks. *arXiv preprint arXiv:2305.08277*, 2023.
- [4] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [5] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012.
- [6] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580, 2018.
- [7] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [8] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. MMD GAN: Towards deeper understanding of moment matching network. *Advances in neural information processing systems*, 30, 2017.
- [9] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge?, 2018.
- [10] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.
- [11] Youssef Mroueh and Truyen Nguyen. On the convergence of gradient descent in gans: Mmd gan as a gradient flow. In *International Conference on Artificial Intelligence and Statistics*, pages 1720–1728. PMLR, 2021.
- [12] Vaishnavh Nagarajan and J. Zico Kolter. Gradient descent gan optimization is locally stable, 2018.
- [13] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 29, 2016.
- [14] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016.
- [15] JR Sylvester. Determinants of block matrices: The mathematical gazette, 84, 460–467, 2000.
- [16] Mathukumalli Vidyasagar. *Nonlinear systems analysis*, volume 42. Siam, 2002.