



HAL
open science

Convergence dynamics of Generative Adversarial Networks: the dual metric flows

Gabriel Turinici

► **To cite this version:**

Gabriel Turinici. Convergence dynamics of Generative Adversarial Networks: the dual metric flows. CADL 2020: Workshop on Computational Aspects of Deep Learning - ICPR 2020, Jan 2021, Milano, Italy. pp.619-634, 10.1007/978-3-030-68763-2_47. hal-03085460

HAL Id: hal-03085460

<https://hal.science/hal-03085460v1>

Submitted on 12 Jul 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Convergence dynamics of Generative Adversarial Networks: the dual metric flows

Gabriel Turinici*

CEREMADE, Université Paris Dauphine - PSL, Paris, France

Jan 10, 2021

Abstract

Fitting neural networks often resorts to stochastic (or similar) gradient descent which is a noise-tolerant (and efficient) resolution of a gradient descent dynamics. It outputs a sequence of networks parameters, which sequence evolves during the training steps. The gradient descent is the limit, when the learning rate is small and the batch size is infinite, of this set of increasingly optimal network parameters obtained during training. In this contribution, we investigate instead the convergence in the Generative Adversarial Networks used in machine learning. We study the limit of small learning rate, and show that, similar to single network training, the GAN learning dynamics tend, for vanishing learning rate to some limit dynamics. This leads us to consider evolution equations in metric spaces (which is the natural framework for evolving probability laws) that we call dual flows. We give formal definitions of solutions and prove the convergence. The theory is then applied to specific instances of GANs and we discuss how this insight helps understand and mitigate the mode collapse.

Keywords: GAN; metric flow; generative network

1 Introduction

Deep generative models are of high interest and used in many applications of deep learning. Among them, the GANs have been one of the most efficient in terms of practical results. The GANs and their convergence are the object of a huge quantity of research papers (4'861 arxiv results mid-October 2020 for "generative adversarial network", 26'110 Google Scholar results). Nevertheless, only very few works concern the behavior of solutions in the general framework of metric spaces or the meaning to be given to the learning trajectory in the limit of a small learning rate. On the other hand, the GANs are known to exhibit unstable convergence behavior (see [18]) and several procedures have been proposed to cure this drawback, among which [5, 9, 12, 19, 30]. In order to contribute to a fundamental understanding of the objects involved, we give in this work a rigorous definition of the concept of solution of the evolution equation associated to a GAN that we call a *dual metric flow*. We identify the hypothesis that guarantee that the discrete solutions converge, when the learning rate τ

*ORCID: 0000-0003-2713-006X, Gabriel.Turinici@dauphine.fr, <http://www.turinici.com>

tends to 0, to a dual metric flow and apply this insight to understand and mitigate the mode collapse phenomena. Finally we give examples that show that the dual flows correspond indeed to procedures used in GAN practice.

1.1 Motivation: W-GANs

The goal of (deep) generative models such as the GANs is to generate new data from some (unknown) distribution given a list of samples drawn from that distribution. To simplify the presentation, we suppose that the distribution to be learned is a set of images. Two objects are important in a GAN: the *Generator* and the *Discriminator*; both are deep networks with fixed, but rich enough, architecture capable of representing a very large class of transformations. For instance, in a Wasserstein-GAN (see [5]), the training has the following form (see figure 1 for an illustration): after initializing (randomly) both the Generator and the Discriminator, the Discriminator is trained first. It takes as input images generated by the Generator (with label "fake") and images from the real database with label "true". It is trained for (one or possibly several) steps in order to achieve good discrimination efficiency between the "fake" and "real" labels. In the next step the discriminator is kept constant and the generator is trained in order to create images which, when run through the (fixed) discriminator obtain as much labels "real" as possible. Then the procedure is repeated till convergence.

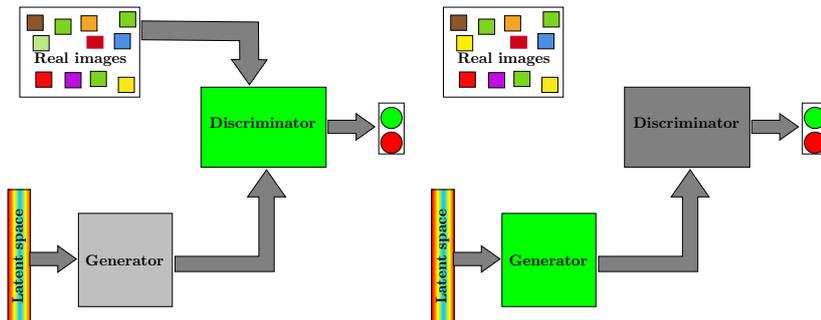


Figure 1: Illustration of the dynamics of a GAN. Color code: in green the part that is active (under training) and in grey the part that is fixed. **Left:** the discriminator is active. **Right:** the generator is active. This is repeated till convergence.

A very natural question is whether such a procedure can converge to a satisfactory solution i.e., if the Generator samples from the right distribution and the Discriminator is able to tell with high precision the quality of any sample. The answer is not always yes, as illustrated by the following simple situation: consider a target distribution which is a Dirac mass centered in some constant x_r . The generator is described by a vector of two real parameters $x \in \mathbb{R}^2$ and the discriminator y has the same format. In the framework of integral probability metrics (see [27]) used in GANs, this simple situation has the following transcription: if the current parameter of the generator is x_n then the next parameter y_{n+1} of the discriminator will be updated to maximize the distance

from x_n to x_r , i.e. $y_{n+1} = y_n + \tau(x_n - x_r)$, where τ is the learning rate. The generator itself will take into account this new discriminator and will move towards the unknown value x_r by taking a step: $x_{n+1} = x_n - \tau y_{n+1}$, where τ is the learning rate. These equations can be also written as

$$\begin{aligned} \frac{y_{n+1} - y_n}{\tau} &= x_n - x_r \\ \frac{x_{n+1} - x_n}{\tau} &= y_{n+1}. \end{aligned} \quad (1)$$

When $\tau \rightarrow 0$ the limit evolution will be

$$\begin{aligned} y'(t) &= x(t) - x_r \\ x'(t) &= y(t). \end{aligned} \quad (2)$$

However, except for very special initial conditions, the system (2) does not have the property that $x(t) \rightarrow x_r$ because $x(t)$ will have a periodic evolution around x_r , see figure 2.

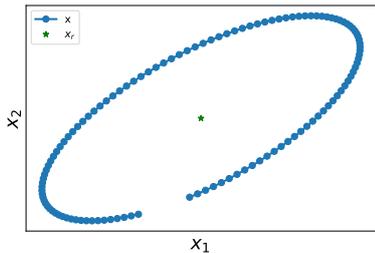


Figure 2: Illustration of the oscillations dynamics of a GAN, see equation (2).

Prompted by this example we aim to analyze in this paper the intrinsic constraints coming from the GAN training in the form of an alternative adversary evolution. More precisely, given the discrete Generator / Discriminator dynamics (similar to equation (1)) we want to write the equivalent limit equation (2).

To do so, we suppose that the architectures of the Generator and Discriminator networks are rich enough so that the Generator can reach with satisfactory precision any target distribution μ_r and the discriminator can realize any mapping that separates an arbitrary pair of distinct distributions in the integral probability metrics sense, i.e. for any distinct distributions the Discriminator can propose a mapping whose averages under the two distributions are different. Of course this is an ideal setting but we take this view in order to better distinguish the effects due to network architectures from those intrinsically included in the GAN convergence protocol.

Let us now introduce the mathematical objects involved in the GAN training. The Generator is a mapping from a given distribution (e.g. the multi-dimensional Gaussian distribution) on some base space (called Latent space) to the space of objects of interest, denoted Ω ; for instance in figure 1 Ω are images). Thus in general the Generator can construct probability distributions on Ω ; denote by $\mathcal{P}_1(\Omega)$ the set of all probability laws on Ω with finite first order

moment. This set can be given the structure of a metric space by introducing a distance; many distances exist and have been used but a popular choice, used by W-GANs is the 1-Wasserstein distance (see [3] for a definition) denoted $d_{W,1}$. Thus formally we work in the metric space $(\mathcal{P}_1(\Omega), d_{W,1})$.

On the other hand the Discriminator constructs a Lipschitz mapping to label the generated samples as fake or real. Mathematically the object is a Lipschitz function $\ell : \Omega \rightarrow \mathbb{R}$. The Discriminator works best when the average $\mathbb{E}_\mu(\ell)$ is as different as possible from the average $\mathbb{E}_{\mu_r}(\ell)$. Again, ℓ lives in a metric space \mathcal{Y} (as opposed to a Hilbert space) with the distance being the Lipschitz distance.

The GAN update of the Generator will move μ_n (n indexes now the learning steps) to some μ_{n+1} along the gradient of the mapping $\mu \mapsto \mathbb{E}_\mu(\ell) - \mathbb{E}_{\mu_r}(\ell)$ (to be minimized, here ℓ being the current Discriminator state). In practice the update is performed with a stochastic descent algorithm (that we take in this work to be the SGD). We will not inquire about the stochastic oscillations but only consider the average state. In full rigor the update rule is difficult to write in the formal, metric space setting: for instance, for the Generator update, the intuitive formula would be $\mu_{n+1} = \mu_n - \tau \nabla_\mu [\mathbb{E}_{\mu_n}(\ell_{n+1}) - \mathbb{E}_{\mu_r}(\ell_{n+1})]$. This formulation has several problems: first the space of probability laws is not a vector space thus the addition and subtraction operator are not well defined. Secondly, the differential structure on the same space is not straightforward to manipulate, i.e., ∇_μ is not easy to work with. For these reasons, we will replace the real dynamics with another, close, one and we model the update by the requirement that μ_{n+1} is the minimizer of

$$\mu \mapsto \frac{d_{\mathcal{X}}(\mu_n, \mu)^2}{2\tau} + [\mathbb{E}_\mu(\ell) - \mathbb{E}_{\mu_r}(\ell)]. \quad (3)$$

The intuitive justification for this formula (very classic in metric space evolution equations [4]) is the following: in a Hilbert space, $d_{\mathcal{X}}(\mu_n, \mu)^2 = \|\mu_n - \mu\|^2$ and taking a general functional $F(\mu)$ in (3) (here $F(\mu) = \mathbb{E}_\mu(\ell) - \mathbb{E}_{\mu_r}(\ell)$) the critical point equations derived from (3) can be written as $\mu_{n+1} = \mu_n - \tau \nabla F(\mu_{n+1})$ which is an implicit gradient descent of step τ starting from μ_n . The advantage of the implicit formulation is that μ_{n+1} defined through (3) can be written for any abstract objects in a metric space (only $d_{\mathcal{X}}(\mu_n, \mu)^2$ and $F(\cdot)$ are required). Same model is used for the Discriminator update steps; all this leads us to consider the main discrete equation that will model the GAN training dynamics ¹

$$\ell_{n+1} = \operatorname{argmin}_{\ell \in \mathcal{Y}} \frac{d(\ell, \ell_n)^2}{2\tau} - [\mathbb{E}_{\mu_n}(\ell) - \mathbb{E}_{\mu_r}(\ell)] \quad (4)$$

$$\mu_{n+1} = \operatorname{argmin}_{\mu \in \mathcal{X}} \frac{d(\mu, \mu_n)^2}{2\tau} + [\mathbb{E}_\mu(\ell_{n+1}) - \mathbb{E}_{\mu_r}(\ell_{n+1})], \quad (5)$$

with $\tau > 0$ the learning rate and \mathcal{X}, \mathcal{Y} two metric spaces. The goal of this paper is to clarify whether, when $\tau \rightarrow 0$, the discrete curves $(\mu_n)_{n \geq 0}, (\ell_n)_{n \geq 0}$, converge to some limit continuous curves; we also want to give a formal definition of the dynamics satisfied by the limit curves.

The GAN protocol described above is not the only one used in the literature. Many other procedures aim to improve convergence, generation quality,

¹see Lemma 1 in appendix for information on the relationship between explicit and implicit numerical schemes.

computing speed, etc. For instance in [12] the authors add a penalty on the gradient in order to make convergence better and the new type of GAN is called WGAN-GP. A full zoology of GAN versions appeared: the Deep Convolutional Generative Adversarial Network (DCGAN, [24]) use deep convolution networks which are better adapted to images, the Coupled GANs proposal (coGAN [21]) uses two generators and two discriminators in order to improve convergence and generation properties, Progressive Growing Generative Adversarial Network (Progressive GAN [16]) improve the generation quality by gradually increasing image resolution (size), Style-Based Generative Adversarial Network (StyleGAN, [17]) better control the Latent space distribution and are able to generate content with given characteristics; other contributions include Cycle-Consistent Generative Adversarial Network (CycleGAN, [31]), Big Generative Adversarial Network (BigGAN, [7]), Pix2Pix [13], and the research is still advancing.

In [10] the authors use even more abstract objects which lead to a loss functional called "Wasserstein of Wasserstein loss"; that is, the most basic objects that are here the images (with our notations elements $\omega \in \Omega$) are not given the usual Euclidian distance, but instead are immersed in a metric space and the Wasserstein distance $d_{W,1}(\omega_1, \omega_2)$ is used to measure the distance (dissimilarity) between two (elementary) images $\omega_1, \omega_2 \in \Omega$. The Wasserstein distance $d_{W,1}(\mu_1, \mu_2)$ is then used a second time in order to discriminate between probability laws $\mu_1, \mu_2 \in \mathcal{P}_1(\Omega)$.

Similar abstraction for the ground distance are to be found in [1] that present Banach space GANs.

All these examples enforce even more the need for an abstract formulation of GAN convergence, that we give below. Finally, see also [14] for implicit procedures relevant to GANs.

1.2 Mathematical setting

We consider \mathcal{X}, \mathcal{Y} two (Polish geodesic) metric spaces (see [8] for an introduction to metric spaces) and $\mathcal{C}(\cdot, \cdot) : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R} \times \mathbb{R}$ a functional (that will stand for the loss functional). Note that $\mathcal{C}(\cdot, \cdot)$ is vector valued, we will denote by $\mathcal{C}^x(\cdot, \cdot)$ and $\mathcal{C}^y(\cdot, \cdot)$ its components. Note that in a GAN we will have

$$\begin{aligned} \mathcal{C}(\ell, \mu) &= (\mathcal{C}^x(\ell, \mu), \mathcal{C}^y(\ell, \mu)) \\ &= (-\mathbb{E}_\mu(\ell) + \mathbb{E}_{\mu_r}(\ell), \mathbb{E}_\mu(\ell) - \mathbb{E}_{\mu_r}(\ell)). \end{aligned} \quad (6)$$

We investigate the equation:

$$\partial_t \begin{pmatrix} x_t \\ y_t \end{pmatrix} + \begin{pmatrix} \nabla_x \mathcal{C}^x(x_t, y_t) \\ \nabla_y \mathcal{C}^y(x_t, y_t) \end{pmatrix} = 0, \quad \begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = \begin{pmatrix} \bar{x} \\ \bar{y} \end{pmatrix}. \quad (7)$$

Such an equation will be called a *dual flow*. The discrete version is defined by the recurrence:

$$x_0^\tau = \bar{x}, \quad x_{k+1}^\tau \in \operatorname{argmin}_{x \in \mathcal{X}} \frac{d(x, x_k^\tau)^2}{2\tau} + \mathcal{C}^x(x, y_k^\tau), \quad k \geq 0. \quad (8)$$

$$y_0^\tau = \bar{y}, \quad y_{k+1}^\tau \in \operatorname{argmin}_{y \in \mathcal{Y}} \frac{d(y, y_k^\tau)^2}{2\tau} + \mathcal{C}^y(x_{k+1}^\tau, y), \quad k \geq 0. \quad (9)$$

These numerical schemes are a distant cousin of some other evolution evolution on metric spaces, namely the evolution flows, see [6, 28]).

From the theoretical point of view, these results are not available with previous techniques from [2, 11, 11, 20, 22, 23, 25]).

2 Theoretical results

2.1 Motivation and literature review

Note that when \mathcal{C}^x is independent of the second argument, i.e.,

$$\mathcal{C}^x(x, y) = E(x), \quad (10)$$

the relation (8) becomes the celebrated implicit Euler-type scheme of Jordan, Kinderlehrer and Otto [15] for the definition of gradient flows in metric spaces

$$\partial_t y_t + \nabla E(y_t) = 0, \quad y_0 = \bar{y}, \quad (11)$$

and received considerable attention (see [4, 26, 29] for instance). However, the situation when E has dependence on other variables has not been treated to the same extent and the related contributions involve gradient flows of time dependent functionals $\mathcal{E}(t, u)$ with a known dependence on time (see [11, 20, 22, 23, 25]). Of course, formally one can set $\mathcal{E}(t, u) = \mathcal{C}(u, y_t)$, and hope to analyze the (x_t, y_t) dynamics in this way. This is not possible for technical reasons (see for instance the discussion in [28]); in particular doing so supposes the knowledge of the dynamics y_t (which is not available) and moreover the dynamics may not be differentiable with respect to time (but remains absolute continuous).

2.2 Basic reminders

The absence of a vector operations in a metric space does no allow to develop fully a differential calculus and requires adaptation of notions of derivative. Accordingly the definition of evolution equations have to use alternative properties.

We recall below the main ideas of such an alternative formulation (see [2]) for the particular case (10)-(11); suppose for a moment that \mathcal{X} is an Euclidian space and E a smooth (C^1 or above) function; then:

$$\begin{aligned} \frac{d}{dt} E(x_t) &= \langle \nabla E(x_t), x'_t \rangle \geq -|\nabla E(x_t)| \cdot |x'_t| \\ &\geq -\frac{1}{2} |x'_t|^2 - \frac{1}{2} |\nabla E|^2(x_t), \end{aligned}$$

or equivalently,

$$\frac{d}{dt} E(x_t) + \frac{1}{2} |x'_t|^2 + \frac{1}{2} |\nabla E|^2(x_t) \geq 0 \quad \forall t, \quad (12)$$

with equality only if x is solution of (11). Therefore asking that

$$\frac{d}{dt} E(x_t) + \frac{1}{2} |x'_t|^2 + \frac{1}{2} |\nabla E|^2(x_t) \leq 0 \quad \forall t, \quad (13)$$

is an equivalent characterization of (11) (more precisely called the EDI formulation). Its integral form is:

$$\forall 0 \leq a \leq b : \\ E(x_b) - E(x_a) + \int_a^b \left(\frac{1}{2} |x'_t|^2 + \frac{1}{2} |\nabla E|^2(x_t) \right) dt \leq 0. \quad (14)$$

The advantage of formulation (14) is that it only uses quantities that can be defined in a metric space (see below for definition of $|x'_t|$ and $|\nabla E|$). The corresponding computation for a bi-variate functional \mathcal{C} is:

$$\forall 0 \leq a \leq b : \int_a^b \left(\frac{d}{dt} \mathcal{C}^x(x_t, \nu) \Big|_{\nu=y_t} \right) dt \\ + \int_a^b \left(\frac{1}{2} |x'_t|^2 + \frac{1}{2} |\nabla_1 \mathcal{C}^x|^2(x_t, y_t) \right) dt \\ + \int_a^b \left(\frac{d}{dt} \mathcal{C}^y(\nu, y_t) \Big|_{\nu=x_t} \right) dt \\ + \int_a^b \left(\frac{1}{2} |y'_t|^2 + \frac{1}{2} |\nabla_2 \mathcal{C}^y|^2(x_t, y_t) \right) dt \leq 0. \quad (15)$$

However this formulation poses specific problems as in general the solution $(x_t, y_t)_{t \geq 0}$ is only absolutely continuous (with respect to time) while, for instance, the manipulation of the term $\frac{d}{dt} \mathcal{C}^x(x_t, \nu) \Big|_{\nu=y_t}$ requires additional assumptions. This will be made precise later.

2.3 Definition of (EDI style) equilibrium flows

Let us recall the following definition:

Definition 1 A curve $x : [0, T] \rightarrow (\mathcal{X}, d)$ is called absolutely continuous if there exists $f \in L^1(0, T)$ such that

$$d(x_{t_1}, x_{t_2}) \leq \int_{t_1}^{t_2} f(t) dt, \quad \forall t_1 < t_2, \quad t_1, t_2 \in [0, T]. \quad (16)$$

For an absolutely continuous curve $(x_t)_{t \in [0, T]}$ the metric derivative of x at r defined by

$$|x'_r| = \lim_{h \rightarrow 0} \frac{d(x_{r+h}, x_r)}{|h|}, \quad (17)$$

exists a.e., belongs to $L^1(0, T)$ and is the smallest L^1 function that verifies (16).

We suppose from now on that \mathcal{C} satisfies the assumption:

(A₁) There exists $C_1 < \infty$ such that $\mathcal{C}^x(y, x), \mathcal{C}^y(y, x) \geq -C_1, \forall x, y \in \mathcal{X} \times \mathcal{Y}$.

For any $\alpha, \beta \in \mathbb{R}, \alpha \leq \beta$, we denote by $\mathcal{S}(\alpha, \beta)$ the set of divisions of the interval $[\alpha, \beta]$. Let $z = (x, y) = (x_t, y_t)_{t \in [0, T]}$ be an absolutely continuous curve in $\mathcal{X} \times \mathcal{Y}$; define for $0 \leq a \leq b \leq T$ and a division $\Delta = \{a = t_0 < t_1 < \dots < t_{N_\Delta} = b\} \in \mathcal{S}(a, b)$:

$$\Upsilon^x(\Delta; z, a, b) = \sum_k \mathcal{C}^x(x_{t_{k+1}}, y_{t_k}) - \mathcal{C}^x(x_{t_k}, y_{t_k}). \quad (18)$$

$$\Upsilon^x(z, a, b) = \liminf_{\Delta \in \mathcal{S}(a, b), |\Delta| \rightarrow 0} \Upsilon^x(\Delta; z, a, b). \quad (19)$$

Similar definitions are introduced for $\Upsilon^y(z, a, b)$ (summing the variations of \mathcal{C}^y along the curve z). Furthermore we denote

$$\Upsilon(\Delta; z, a, b) = \Upsilon^x(\Delta; z, a, b) + \Upsilon^y(\Delta; z, a, b). \quad (20)$$

Remark 1 When \mathcal{X} is e.g., Euclidian and under regularity assumptions on \mathcal{C} it is easy to check that $\Upsilon^x(x, a, b) = \int_a^b \frac{d}{dt} \mathcal{C}^x(x_t, \nu) \Big|_{\nu=y_t} dt$ and the same for Υ^y .

We are now ready to state the formal definition of a solution of (7) in the abstract setting of metric spaces. The particular flavor we use is the so-called "EDI" solution, see [3] for details.

Definition 2 (EDI equilibrium flow) An absolutely continuous curve $z = (x_t, y_t)_{t \in [0, T]}$ is called an EDI-equilibrium flow starting from (\bar{x}, \bar{y}) if $\lim_{t \rightarrow 0} (x_t, y_t) = (\bar{x}, \bar{y})$ and:

$$\begin{aligned} \forall s \geq 0, \quad & \Upsilon(z, 0, s) + \frac{1}{2} \int_0^s |x'_r|^2 + |y'_r|^2 \, dr \\ & + \frac{1}{2} \int_0^s |\nabla_1 \mathcal{C}^x|^2(x_r, y_r) + |\nabla_2 \mathcal{C}^y|^2(x_r, y_r) \, dr \leq 0, \quad (21) \\ \text{a.e. } t > 0, \quad & \forall s \geq t, \quad \Upsilon(x, t, s) + \frac{1}{2} \int_t^s |x'_r|^2 + |y'_r|^2 \, dr \\ & + \frac{1}{2} \int_t^s |\nabla_1 \mathcal{C}^x|^2(x_r, y_r) + |\nabla_2 \mathcal{C}^y|^2(x_r, y_r) \, dr \leq 0, \quad (22) \end{aligned}$$

where the slope $|\nabla_1 \mathcal{C}^x|(x, y)$ of $\mathcal{C}^x(\cdot, \cdot)$ with respect to the first argument evaluated at (x, y) is:

$$|\nabla_1 \mathcal{C}^x|(x, y) = \limsup_{u \rightarrow x} \frac{(\mathcal{C}^x(x, y) - \mathcal{C}^x(u, y))^+}{d(x, u)}, \quad (23)$$

and similarly for $|\nabla_2 \mathcal{C}^y|(x, y)$.

Remark 2 For the particular case of a Hilbert space the definition above coincides with the usual definition of an evolution equation (7).

2.4 Convergence of numerical schemes

Let us denote

$$\mathcal{M}^x(x, y, \tau) = \operatorname{argmin}_{u \in \mathcal{X}} \frac{d(u, x)^2}{2\tau} + \mathcal{C}^x(u, y) \quad (24)$$

$$\mathcal{M}^y(x, y, \tau) = \operatorname{argmin}_{u \in \mathcal{Y}} \frac{d(u, y)^2}{2\tau} + \mathcal{C}^y(x, u). \quad (25)$$

With this definition the numerical scheme in equations (8)-(9) can be written as

$$x_{k+1}^\tau \in \mathcal{M}^x(x_k^\tau, y_k^\tau, \tau), y_{k+1}^\tau \in \mathcal{M}^y(x_{k+1}^\tau, y_k^\tau, \tau). \quad (26)$$

The goal of this contribution is to investigate whether when $\tau \rightarrow 0$ the set $\{(x_k^\tau, y_k^\tau), k \geq 1\}$ converges to a solution of (7) as defined in (21)-(22).

In order to work with meaningful objects, we introduce the following assumption which is the analogue of [2, Assumption 4.8 page 67]:

(A₂) There exists $\bar{\tau} > 0$ such that for any $\tau \leq \bar{\tau}$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$:

$$\mathcal{M}^x(x, y, \tau) \neq \emptyset, \mathcal{M}^y(x, y, \tau) \neq \emptyset. \quad (27)$$

Assuming that assumption **(A₂)** is satisfied, we can define the interpolation *à la de Giorgi* which is a curve $t \in [0, T] \mapsto (x_t^\tau, y_t^\tau)$ such that $(x_0^\tau, y_0^\tau) = (\bar{x}, \bar{y})$ and $\forall t \in]k\tau, (k+1)\tau[$:

$$x_t^\tau \in \mathcal{M}^x(x_{k\tau}^\tau, y_{k\tau}^\tau, t - k\tau), y_t^\tau \in \mathcal{M}^y(x_{(k+1)\tau}^\tau, y_{k\tau}^\tau, t - k\tau). \quad (28)$$

We will need some additional hypothesis:

(A₃) For any $c \in \mathbb{R}$, $r > 0$ and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ the sets $\{u \in \mathcal{X} | \mathcal{C}^x(u, y) \leq c, d(u, x) \leq r\}$ and $\{u \in \mathcal{Y} | \mathcal{C}^y(x, u) \leq c, d(u, y) \leq r\}$ are both compact.

(A₄) The slopes $|\nabla_1 \mathcal{C}^x|$ and $|\nabla_2 \mathcal{C}^y|$ are lower semicontinuous.

(A₅) The function \mathcal{C}^x is Lipschitz with respect to the second argument and \mathcal{C}^y is Lipschitz with respect to the first argument.

(A₆) For any absolutely continuous curve $z = (x_t, y_t)_{t \in [a, b]}$:

$$\Upsilon(x, a, b) \leq \liminf_{\substack{|\Delta_n| \rightarrow 0 \\ z_n = (x_n, y_n) \rightarrow z \\ \sup_n \int_a^b |\dot{x}_n(t)| + |\dot{y}_n(t)| dt < \infty}} \Upsilon(\Delta_n; z_n, a, b), \quad (29)$$

where the convergence of the curves z_n to z is in the uniform (on compacts) norm.

(A₇) There exists $C_L < \infty$ such that for any x, y, u, w :

$$|\mathcal{C}^x(u, v) + \mathcal{C}^x(u, y) - \mathcal{C}^x(w, v) - \mathcal{C}^x(w, y)| \leq C_L d(u, w) d(v, y). \quad (30)$$

and the same for $\mathcal{C}^y(u, v)$.

Remark 3 The assumption **(A₇)** implies **(A₆)** (see [28, Lemma 2]).

With these provisions, the properties of the curves obtained by the numerical scheme (8) are detailed in the Theorem 1.

Theorem 1 Let \mathcal{C} satisfying assumptions **(A₁)**, **(A₂)**, **(A₃)**, **(A₄)**, **(A₅)** and **(A₆)**. Then the set of curves $\{(x_t^\tau, y_t^\tau)_{t \in [0, T]}; \tau \leq \bar{\tau}\}$ defined in (28) is relatively compact in the set of curves in $\mathcal{X} \times \mathcal{Y}$ with local uniform convergence and any limit curve is an EDI equilibrium flow in the sense of Definition 2.

Proof The proof is somehow technical but is a adaptation of the proof of Theorem 1 in [28]: first we show that the map $\tau \mapsto \frac{d(x_\tau, x)^2}{2\tau} + \mathcal{C}^x(x_\tau, y)$ is locally Lipschitz. Then the discrete identity is obtained as in [28, formula (32)] and then estimations similar to [28, formulas (35) and (36)] allow to conclude. ■

Similar results hold for the convex case (see [28, Theorem 2]).

3 Applications

Using Lemma 1 we conclude that if in some circumstances there are ways to use explicit numerical schemes (like for GANs), the convergence is also ensured for the explicit schemes, once the implicit ones converge.

Let us now inquire what are the consequences of the theoretical results for WGAN training. Because the Discriminator is trained first one can consider the variable x to be the Discriminator network parameters that will result in a Lipschitz function $\ell = \ell(x)$ and y to be the Generator parameters that will generate a distribution $\mu = \mu(y)$. Functions \mathcal{C}^x and \mathcal{C}^y are given by (6) thus in particular $\mathcal{C}^x = -\mathcal{C}^y$. With these notations we can apply the Theorem 1 and obtain that, in the limit of a vanishing learning rate, the WGAN training will tend to some evolution curve, both in the space $\mathcal{X} \times \mathcal{Y}$ of parameters, but also in the space of the distributions (where μ belongs) and Lipschitz function (where ℓ belongs).

As a further application, we can investigate the conditions under which the GAN training give rise to a mode collapse. The discussion below is not a mathematical proof but oriented towards a practical understanding. A mode collapse describes, e.g., the situation when a strong Discriminator pushes the Generator to only produce a limited number of samples with a loss in diversity. With our notations, this means that the evolution (x_t, y_t) will be close to a constant (x^∞, y^∞) but the corresponding distribution $\mu(y^\infty)$ is far from μ_r but can be expressed (at least approximately) as a finite sum of Dirac masses $\sum_{a=1}^A p_a \delta_{i_a}$ where $i_a \in \Omega$ are given images. Using the insight from Lemma 2 and under assumption that the generator network is locally injective (i.e., does not generate redundant probability laws) the point $\mu(y^\infty)$ is a critical point of the loss function. But, denoting ℓ^∞ the Lipschitz function corresponding to the Discriminator network, the loss function for the Generator will be $\mathbb{E}_\mu(\ell^\infty) - \mathbb{E}_{\mu_r}(\ell^\infty)$ and the loss of the Discriminator will be $\ell \mapsto \mathbb{E}_{\mu_r}(\ell) - \mathbb{E}_{\mu^\infty}(\ell)$. For such a loss function, the information that μ is a sum of Dirac masses and also a critical point of $\mu \mapsto \mathbb{E}_\mu(\ell^\infty) - \mathbb{E}_{\mu_r}(\ell^\infty)$ implies that moving (in the space of probability laws endowed with the 1-Wasserstein metric) towards any other Dirac mass does not change (decrease) the loss value (to the first order). Therefore the images i_a in the support of the measure μ are necessarily of lowest possible loss value i.e., if the discriminator is good enough, are members of the original "real" image values (in mathematical terms are members of the support of μ_r). On the other hand, if the dynamics of the Discriminator is also blocked in some point ℓ^∞ , this means again that, to the first order, $\mathbb{E}_{\mu^\infty}(\ell) - \mathbb{E}_{\mu_r}(\ell)$ cannot be increased locally when ℓ is slightly perturbed around ℓ^∞ . Or, since μ and μ_r are different we obtain a contradiction. Therefore the mode collapse is **not** a legitimate limit dynamics. We can therefore conclude that if mode collapse happens this is due to a too large time step, to a not strong enough Generator architecture or to numerical traps that can be removed by perturbing slightly the dynamics.

4 Discussion and conclusion

When averaging out the steps of a SGD one obtains the gradient flow of the loss functional. The question that we ask in this paper is what is obtained when one averages out the generator-discriminator dynamics encountered in GANs.

To answer the question we notice that in GANs the ground metric is not always of L^2 type but can be arbitrary (Wasserstein metric as in [10], Banach norm as in [1], etc). Thus we re-formulate the question: when the learning rate becomes smaller and smaller, is there any limit for the curves obtained during the GAN training ? Does this correspond to a dynamical system ? We first give sufficient conditions for this convergence in general metric space when the learning process is composed of implicit steps. On the other side we recall that under mild conditions explicit and implicit steps will be arbitrary close thus converge to the same limit. Therefore the dynamics of GAN training will in general follow the solution of a evolution equation whose details are given explicitly in equation (7). The knowledge of such a fact can help better understand the GAN optimization dynamics and the mode collapse phenomena.

Acknowledgements

A Appendix

A.1 Explicit and implicit numerical schemes in Hilbert spaces

We recall below a standard result on the relationship between explicit and implicit numerical schemes in Hilbert spaces.

Lemma 1 *Let H be a Hilbert space, $f : H \rightarrow \mathbb{R}$ a bounded Lipschitz function with Lipschitz constant L and two numerical schemes defined by the recurrences:*

$$x_{n+1}^E = x_n^E + \tau f(x_n^E), \quad x_0^E = \bar{x} \quad (31)$$

$$x_{n+1}^I = x_n^I + \tau f(x_{n+1}^I), \quad x_0^I = \bar{x} \quad (32)$$

Then for τ small enough:

1. the implicit scheme (32) has a unique solution for any step $n \geq 0$.
2. let $T = N\tau$ for some fixed N , then $\|x_N^I - x_N^E\| \leq C\tau$, with the constant C depending only on f , \bar{x} and T .

Remark 4 *Note that point 2 implies in particular that if, for $\tau \rightarrow 0$, the implicit curves $(x_n^I)_{n \geq 0}$ converge to some limit curve then the explicit curves $(x_n^E)_{n \geq 0}$ converge to the same. However in order to avoid technicalities we will not state precisely what the full curves are and what kind of convergence is obtained.*

Proof Point 1 is obtained by a Picard procedure after observing that the mapping $x \mapsto x_n^E + \tau f(x)$ is a contraction for τ small enough. For the point 2 we make use of the Lipschitz constant of f :

$$\begin{aligned} \|x_{n+1}^I - x_{n+1}^E\| &\leq \|x_n^I - x_n^E\| + \tau L \|x_{n+1}^I - x_n^E\| \\ &\leq \|x_n^I - x_n^E\| + \tau L (\|x_n^I - x_n^E\| + \tau f(x_{n+1}^I)) \end{aligned} \quad (33)$$

Thus, denoting by C_f an upper bound on f :

$$\|x_{n+1}^I - x_{n+1}^E\| \leq (1 + \tau L) \|x_n^I - x_n^E\| + \tau^2 LC_f. \quad (34)$$

If suffices now to use the discrete version of the Gronwall lemma to obtain $\|x_{n+1}^I - x_{n+1}^E\| \leq e^{\tau L}(n+1)\tau^2 LC_f$, and the conclusion follows from:

$$\|x_N^I - x_N^E\| \leq \tau e^{\tau L} TLC_f. \quad (35)$$

A.2 Critical points of gradient flows on intermediary spaces

We investigate in this section a simple situation of a gradient flow of a composed function. Suppose thus an initial space $X_p = \mathbb{R}^n$ (for a GAN the neural network parameter space), an object space $X_o = \mathbb{R}^m$ (for a GAN this will be the space of probability measures where μ belongs and that of Lipschitz functions where ℓ belongs). Consider also two functions $g : X_p \rightarrow X_o$, $f : X_o \rightarrow \mathbb{R}$ and the gradient flow:

$$x'_t = -\nabla_x(f \circ g)(x_t), \quad (36)$$

where for any function we denote by ∇ its differential; for instance $\nabla_o f(o)$ is the gradient of f at the point o , taken as a row vector, $\nabla_x g(x)$ is the $m \times n$ Jacobian matrix of g at x (entry i, j being $\partial g_i / \partial x_j$).

This dynamics in parameter space X_p defines a dynamics $o_t = g(x_t)$ in the object space X_o . We want to investigate the relationship between the dynamics x_t and o_t when the evolution (36) ends up in a stationary point i.e., stalls at some given point $x^\infty \in X_p$ and $o^\infty = g(x^\infty) \in X_o$.

Lemma 2 *Suppose that the functions f and g are of C^1 class (i.e. with continuous derivatives). Then:*

1. denoting $o_t = g(x_t)$ the dynamics in parameter space can also be written

$$x'_t = -(\nabla_o f)(o_t) \cdot (\nabla_x g)(x_t). \quad (37)$$

2. the dynamic in object space X_o is:

$$o'_t = -(\nabla_o f)(o_t) \cdot (\nabla_x g)(x_t) \cdot (\nabla_x g)(x_t)^T. \quad (38)$$

In particular the dynamics in object space is not in general a gradient flow (but will be when $(\nabla_x g) \cdot (\nabla_x g)^T = Id$).

Suppose now that the dynamics (36) is such that for some $t \geq t_1$ we have $x_t = x^\infty$. Then:

3. $o^\infty = g(x^\infty)$ is a critical point of f (i.e., $\nabla_o f(o^\infty) = 0$) as soon as $\nabla_o g^T$ is locally injective around o^∞ (which implies that g is locally injective around o^∞).

Proof Formulas (37) and (38) are derived from (36) using the chain rule. Suppose now x^∞ is a critical stationary point of the dynamics (36). This of course implies that $o_t = g(x_t)$ will also be constant but is not enough to conclude that o^∞ is a critical point of f . But, since the time derivative x'_t in the evolution equation is zero for the constant dynamics $x_t = x^\infty$, we have that $(\nabla_o f)(o^\infty) \cdot (\nabla_x g)(x^\infty) = 0$, which, when $(\nabla_x g)^T(x^\infty)$ is injective will imply that $(\nabla_o f)(o^\infty) = 0$. ■

Note that by the rank theorem the local injectivity of $(\nabla_x g)^T(x^\infty)$ and of g are related.

References

- [1] Adler, J., Lunz, S.: Banach Wasserstein GAN. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 31*, pp. 6754–6763. Curran Associates, Inc. (2018), <http://papers.nips.cc/paper/7909-banach-wasserstein-gan.pdf>
- [2] Ambrosio, L., Gigli, N.: *Modelling and Optimisation of Flows on Networks: Cetraro, Italy 2009*, Editors: Benedetto Piccoli, Michel Rascle, chap. A User’s Guide to Optimal Transport, pp. 1–155. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
- [3] Ambrosio, L., Gigli, N., Savaré, G.: *Gradient flows in metric spaces and in the space of probability measures*. 2nd ed. Basel: Birkhäuser, 2nd ed. edn. (2008)
- [4] Ambrosio, L., Gigli, N., Savaré, G.: *Gradient flows in metric spaces and in the space of probability measures*. 2nd ed. Basel: Birkhäuser, 2nd ed. edn. (2008)
- [5] Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017)
- [6] Blanchet, A., Carlier, G.: Remarks on existence and uniqueness of cournot–nash equilibria in the non-potential case. *Mathematics and Financial Economics* **8**(4), 417–433 (2014). <https://doi.org/10.1007/s11579-014-0127-z>, <http://dx.doi.org/10.1007/s11579-014-0127-z>
- [7] Brock, A., Donahue, J., Simonyan, K.: Large Scale GAN Training for High Fidelity Natural Image Synthesis. arXiv:1809.11096 [cs, stat] (Feb 2019), <http://arxiv.org/abs/1809.11096>, arXiv: 1809.11096
- [8] Burago, D., Burago, Y., Ivanov, S.: *A course in metric geometry*, Graduate Studies in Mathematics, vol. 33. American Mathematical Society, Providence, RI (2001). <https://doi.org/10.1090/gsm/033>, <http://dx.doi.org/10.1090/gsm/033>
- [9] Deshpande, I., Hu, Y.T., Sun, R., Pyrros, A., Siddiqui, N., Koyejo, S., Zhao, Z., Forsyth, D., Schwing, A.: Max-Sliced Wasserstein Distance and its use for GANs. arXiv preprint arXiv:1904.05877 (2019)
- [10] Dukler, Y., Li, W., Lin, A.T., MontÀ^ofar, G.: Wasserstein of Wasserstein loss for learning generative models. In: Chaudhuri, K. (ed.) *Proceedings of the 36th international conference on machine learning*, 9-15 June 2019, Long Beach, California, USA, *Proceedings of machine learning research*, vol. 97, pp. 1716–1725. PMLR, Long Beach, California (2019)
- [11] Ferreira, L.C.F., Valencia-Guevara, J.C.: Gradient flows of time-dependent functionals in metric spaces and applications to pdes. *Monatshefte für Mathematik* pp. 1–38 (2017). <https://doi.org/10.1007/s00605-017-1037-y>, <http://dx.doi.org/10.1007/s00605-017-1037-y>

- [12] Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of Wasserstein GANs. In: Advances in neural information processing systems. pp. 5767–5777 (2017)
- [13] Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-Image Translation with Conditional Adversarial Networks. arXiv:1611.07004 [cs] (Nov 2018), <http://arxiv.org/abs/1611.07004>, arXiv: 1611.07004
- [14] Iwaki, R., Asada, M.: Implicit incremental natural actor critic algorithm. Neural Networks **109**, 103–112 (Jan 2019). <https://doi.org/10.1016/j.neunet.2018.10.007>, <http://www.sciencedirect.com/science/article/pii/S0893608018302922>
- [15] Jordan, R., Kinderlehrer, D., Otto, F.: The variational formulation of the Fokker-Planck equation. SIAM J. Math. Anal. **29**(1), 1–17 (1998). <https://doi.org/10.1137/S0036141096303359>, <http://dx.doi.org/10.1137/S0036141096303359>
- [16] Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive Growing of GANs for Improved Quality, Stability, and Variation. In: International Conference on Learning Representations (2018), <https://openreview.net/forum?id=Hk99zCeAb>
- [17] Karras, T., Laine, S., Aila, T.: A Style-Based Generator Architecture for Generative Adversarial Networks. arXiv:1812.04948 [cs, stat] (Mar 2019), <http://arxiv.org/abs/1812.04948>, arXiv: 1812.04948
- [18] Kodali, N., Abernethy, J., Hays, J., Kira, Z.: On convergence and stability of GANs. arXiv preprint arXiv:1705.07215 (2017)
- [19] Kolouri, S., Pope, P.E., Martin, C.E., Rohde, G.K.: Sliced-Wasserstein autoencoder: an embarrassingly simple generative model. arXiv preprint arXiv:1804.01947 (2018)
- [20] Kopfer, E.: Gradient flow for the Boltzmann entropy and Cheeger’s energy on time-dependent metric measure spaces. ArXiv e-prints (November 2016)
- [21] Liu, M.Y., Tuzel, O.: Coupled Generative Adversarial Networks. arXiv:1606.07536 [cs] (Sep 2016), <http://arxiv.org/abs/1606.07536>, arXiv: 1606.07536
- [22] Mielke, A., Rossi, R., Savaré, G.: Variational convergence of gradient flows and rate-independent evolutions in metric spaces. Milan J. Math. **80**(2), 381–410 (2012). <https://doi.org/10.1007/s00032-012-0190-y>, <http://dx.doi.org/10.1007/s00032-012-0190-y>
- [23] Mielke, A., Rossi, R., Savaré, G.: Nonsmooth analysis of doubly nonlinear evolution equations. Calc. Var. Partial Differential Equations **46**(1-2), 253–310 (2013). <https://doi.org/10.1007/s00526-011-0482-z>, <http://dx.doi.org/10.1007/s00526-011-0482-z>
- [24] Radford, A., Metz, L., Chintala, S.: Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. arXiv:1511.06434 [cs] (Jan 2016), <http://arxiv.org/abs/1511.06434>, arXiv: 1511.06434

- [25] Rossi, R., Mielke, A., Savaré, G.: A metric approach to a class of doubly nonlinear evolution equations and applications. *Ann. Sc. Norm. Super. Pisa Cl. Sci. (5)* **7**(1), 97–169 (2008)
- [26] Santambrogio, F.: Optimal transport for applied mathematicians. *Progress in Nonlinear Differential Equations and their Applications*, 87, Birkhäuser/Springer, Cham (2015). <https://doi.org/10.1007/978-3-319-20828-2>, <http://dx.doi.org/10.1007/978-3-319-20828-2>, calculus of variations, PDEs, and modeling
- [27] Sriperumbudur, B.K., Gretton, A., Fukumizu, K., Schölkopf, B., Lanckriet, G.R.G.: Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research* **11**(Apr), 1517–1561 (2010), <http://www.jmlr.org/papers/v11/sriperumbudur10a.html>
- [28] Turinici, G.: Metric gradient flows with state dependent functionals: The Nash-MFG equilibrium flows and their numerical schemes. *Nonlinear Analysis* **165**, 163–181 (Dec 2017). <https://doi.org/10.1016/j.na.2017.10.002>, <http://www.sciencedirect.com/science/article/pii/S0362546X17302444>
- [29] Villani, C.: Optimal transport. Old and new, *Grundlehren der mathematischen Wissenschaften*, vol. 338. Springer (2009). <https://doi.org/10.1007/978-3-540-71050-9>
- [30] Wu, J., Huang, Z., Acharya, D., Li, W., Thoma, J., Paudel, D.P., Gool, L.V.: Sliced Wasserstein generative models. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3713–3722 (2019)
- [31] Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired Image-To-Image Translation Using Cycle-Consistent Adversarial Networks. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. pp. 2223–2232 (2017)