Convergence dynamics of Generative Adversarial Networks: the dual metric flows

arXiv:2012.10410

Presented as oral paper at the CADL Workshop, Jan. 11 2020, ICPR 2020, Virtual Milano, Italy

Gabriel Turinici

CEREMADE, Université Paris – Dauphine

PSL Research University, Paris, FRANCE

GAN: a dual view

GANs (Generative Adversarial Networks) are used to generate new data (e.g., images) from a database of model data.

It takes the form of a pair of networks, the Generator that creates new data and the Discriminator (Critic) that trains to distringuish the created and original data.

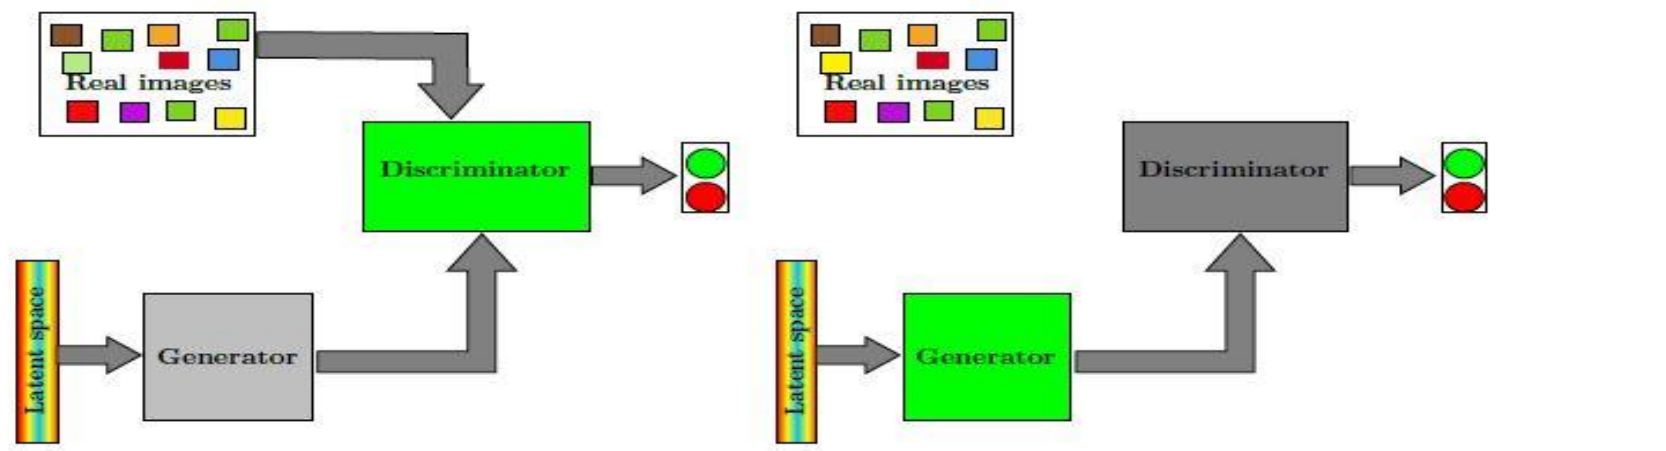

Fig. 1. Illustration of the dynamics of a GAN. Color code: in green the part that is active (under training) and in grey the part that is fixed. **Left:** the discriminator is active. **Right:** the generator is active. This is repeated till convergence.

Congergence of GANs arXiv:2012.10410
CADL / ICPR 2021

GANs: convergence

GANs are computationally intensive (cf. CADL tutorial) and known to have convergence problems, may exhibit mode collapse etc.

Theoretical form is an ellipsoid …

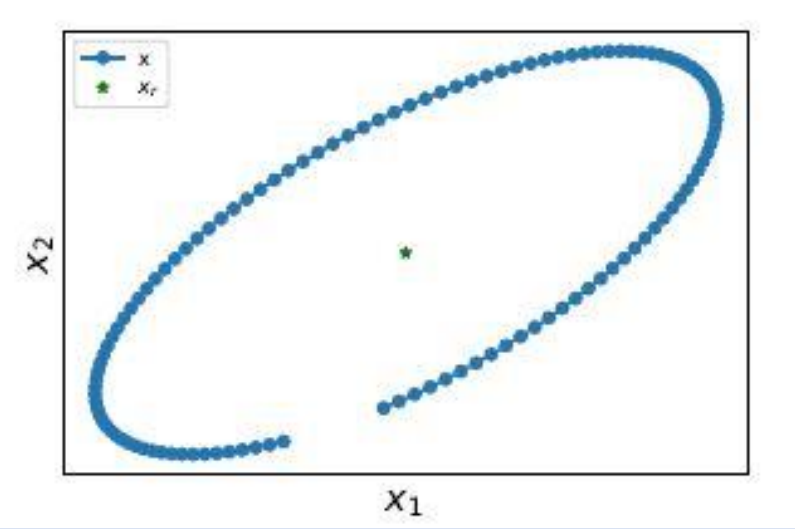

Simple example: (2D GAN … generate a point in RxR)
x = Generator, y= Discriminator

Discrete form equations:
First eq: max discrimination w/r to a reference
Second eq: generate best to mystify the discriminator

$$\frac{y_{n+1} - y_n}{\tau} = x_n - x_r$$
$$\frac{x_{n+1} - x_n}{\tau} = y_{n+1}.$$

Simple example: continuous

$$y'(t) = x(t) - x_r$$
$$x'(t) = y(t).$$

Congergence of GANs arXiv:2012.10410
CADL / ICPR 2021

GANs convergence: motivation (1/2)

Simple example: discrete

$$\frac{y_{n+1} - y_n}{\tau} = x_n - x_r$$
$$\frac{x_{n+1} - x_n}{\tau} = y_{n+1}.$$

continuous

$$y'(t) = x(t) - x_r$$
$$x'(t) = y(t).$$

WGAN update Example

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

'l' = Lipschitz function, 'μ' = probability distribution

QUESTION: what is the continuous equivalent of this discrete form when learning rate τ ->0 ?

GANs convergence: motivation (2/2)

In practice complicated convergence patterns

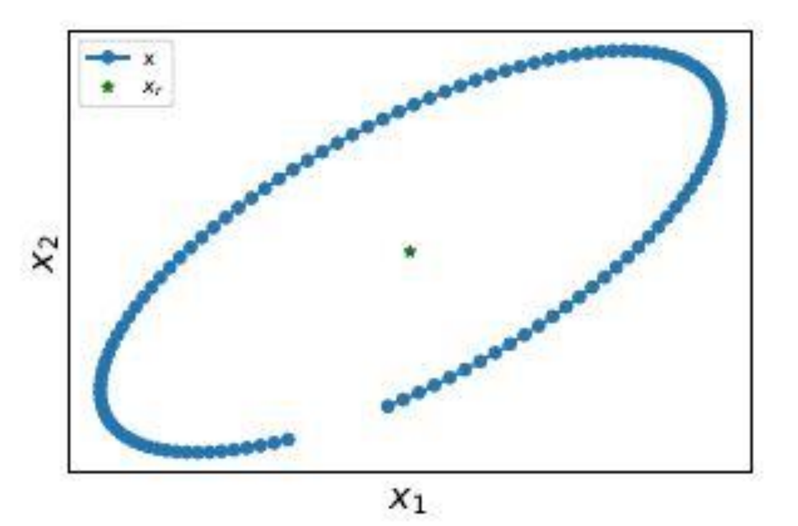

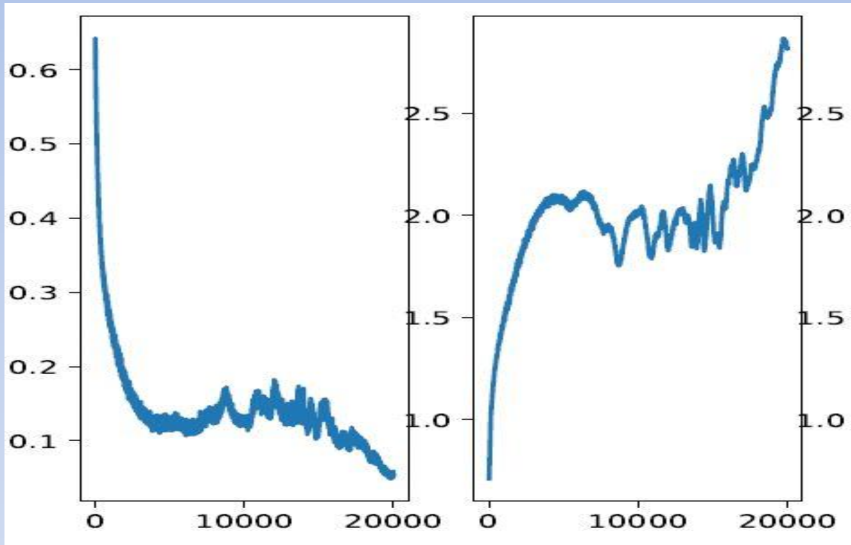

theoretical

empirical

QUESTION: right figure has converged ???

Congergence of GANs arXiv:2012.10410
CADL / ICPR 2021

GANs: convergence technical: metric space

Recall WGAN example

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

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

The intuitive

$$\mu_{n+1} = \mu_n - \tau \nabla_\mu [\mathbb{E}_{\mu_n}(\ell_{n+1}) - \mathbb{E}_{\mu_r}(\ell_{n+1})]$$

generator update … impossible to write for probability distributions (in this form) but OK to write in first form that only uses distances … thus need a metric space equivalent of

$$y'(t) = x(t) - x_r$$
$$x'(t) = y(t).$$

Related notion: gradient flow

GANs: theory

Theorem : the dual (Generator – Discriminator) evolution equation in a suitably chosen metric spaces has a meaning as mathematical object.

The object has intuitive properties, in particular it is the limit of the trajectories for learning rates τ-> 0.

It builds on the notion of gradient flow but need to generalize it because this is NOT a gradient flow.

Related to equilibrium flows (Nonlinear Analysis 165 163-181, 2017)

GAN convergence: applications (1/2)

Convergence criterion : curve has to stabilize in the limit τ-> 0.
E.g.: 2 x more learning steps of ½ size give the same curve ?

Relevant 'x' axis for convergence plots is the cummulative sum of learning rates.

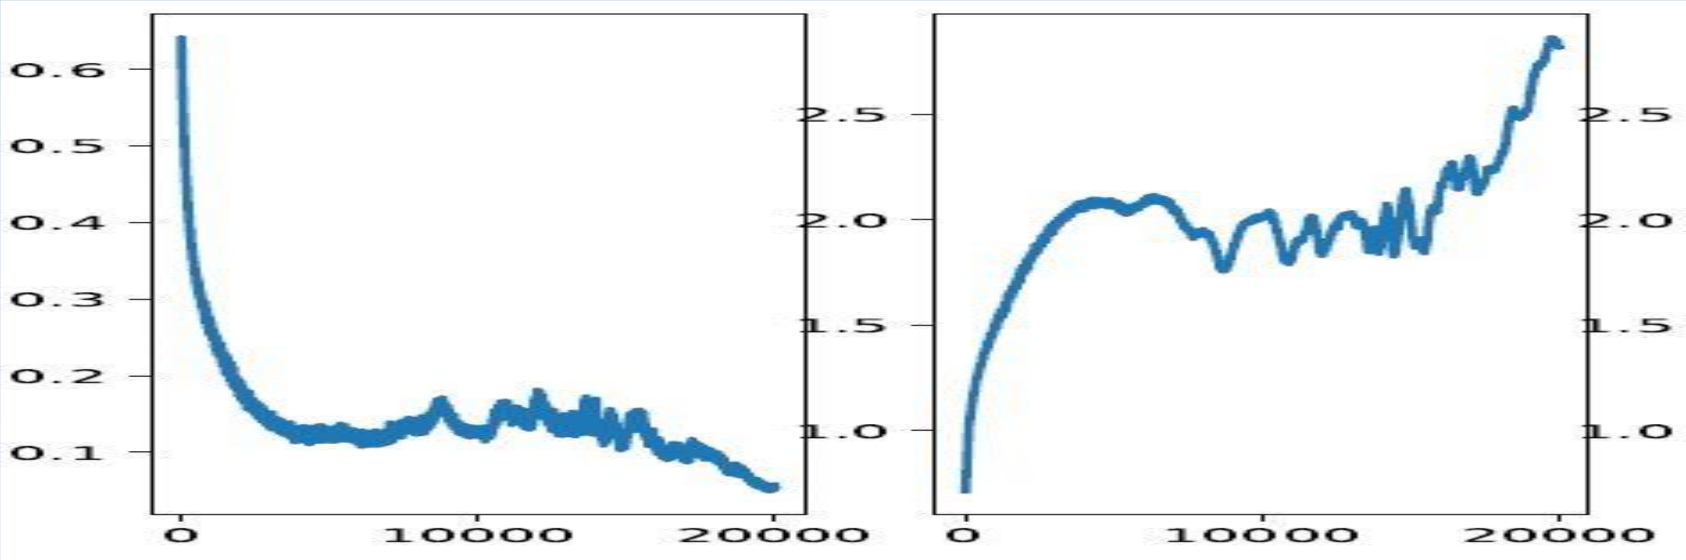

Congergence of GANs arXiv:2012.10410
CADL / ICPR 2021

GAN convergence: applications (2/2): mode collapse

Mode collapse (MNIST example)

generated samples are too similar

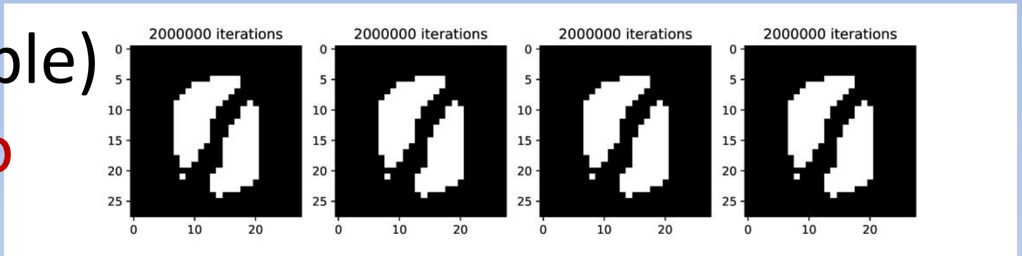

Mode collapse: under hypothesis (…) mode collapse is NOT a limit dynamics (does not satisfy the continuous equation) thus is unstable. One can exit mode collapse by

- changing learning rate
- changing the Generator architecture
- changing the Discriminator architecture
- …

Conclusions and further work

Further applications:

- make automatic use of the convergence information
- Generalize to other settings (VAE, etc.)
- …

Further questions: email

Refs: preprints of the CADL/ICPR Workshop, arXiv:2012.10410