RÉPUBLIQUE FRANÇAISE
Liberté
Égalité
Fraternité

ONERA
THE FRENCH AEROSPACE LAB

Generative Adversarial Networks (GANs) : concept and application to cloudy sky images synthesis

S. Lefebvre (ONERA)
sidonie.lefebvre@onera.fr

Outline

- Intro: Generative Modeling
- Generative Adversarial Networks
  - Principle and interest
  - Some maths
  - Pitfalls & modifications
- Application to clouds synthesis
- Some examples of the use of GANs in astrophysics
- Summary

Generative modeling

Discriminative models vs Generative models

- <u>Discriminative Models</u>

Discriminate between different kinds of data instances, for example classifiers - capture the conditional probability P(Y|X) where Y = labels and X instances

- <u>Generative models</u>

Can generate new data instances - capture the joint probability p(X, Y) in the supervised case, or just p(X) if there are no labels.

GANs are one (clever) kind of generative models

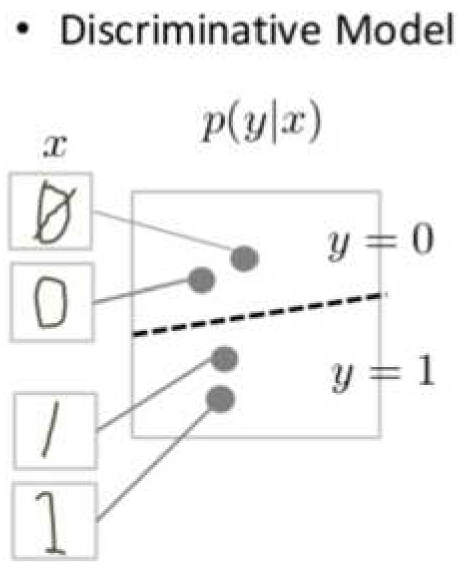

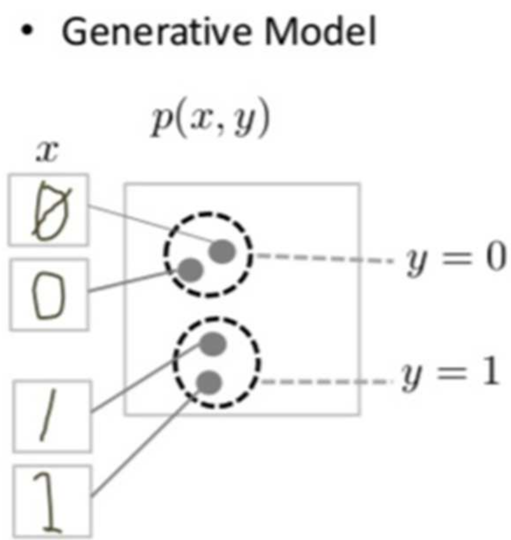

Illustration from https://developers.google.com/machine-learning/gan

Outline

- Intro: Generative Modeling
- Generative Adversarial Networks
  - Principle and interest
  - Some maths
  - Pitfalls & modifications
- Application to clouds synthesis
- Some examples of the use of GANs in astrophysics
- Summary

Principle and interest

Concept proposed in 2014 in the paper **Generative Adversarial Nets** by Ian J. Goodfellow and his colleagues at University of Montreal

"Generative Adversarial networks is the most interesting idea in the last ten years in machine learning" Yann LeCun

GANs are **deep neural network** architectures made of **two different networks**, contesting with each other by playing a zero-sum game

They **learn from mistakes** and try not to make similar errors in the future

Two networks

- **Generator:** learns to generate plausible data from random noise (uniform, Gaussian…)
- **Discriminator**: learns to distinguish the generator's fake data from real data

End: when the generated images are not distinguishable from real images anymore

= Nash equilibrium

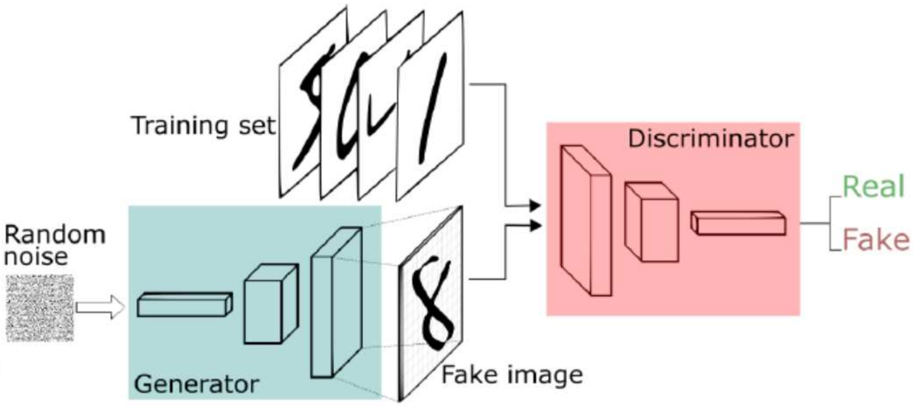

Illustration from https://2018.igem.org/Team:Vilnius-Lithuania-OG/Gan_Introduction

Why are they so popular ?

- Every week, new GAN papers are coming out in lots of topics

A list, not maintained since 2019:
https://github.com/hindupuravinash/the-gan-zoo

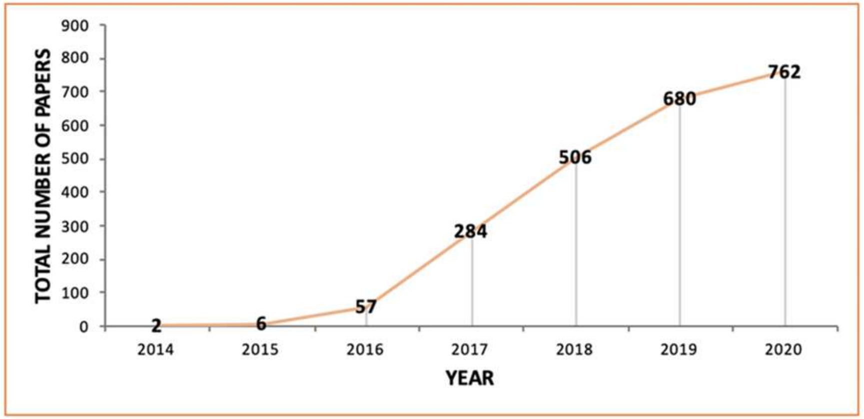

From Z. Farou 10.13140/RG.2.2.24193.48483/2

- Impressive results, improving really fast:

Ian Goodfellow tweet in January 2019
4.5 years of GAN progress on face generation.
https://arxiv.org/abs/1406.2661 https://arxiv.org/abs/1511.06434
https://arxiv.org/abs/1606.07536 https://arxiv.org/abs/1710.10196
https://arxiv.org/abs/1812.04948

Why are they so popular ?

GauGAN, named after post-Impressionist painter Paul Gauguin, creates photorealistic images from segmentation maps, which are labeled sketches that depict the layout of a scene

https://www.nvidia.com/en-us/research/ai-playground/

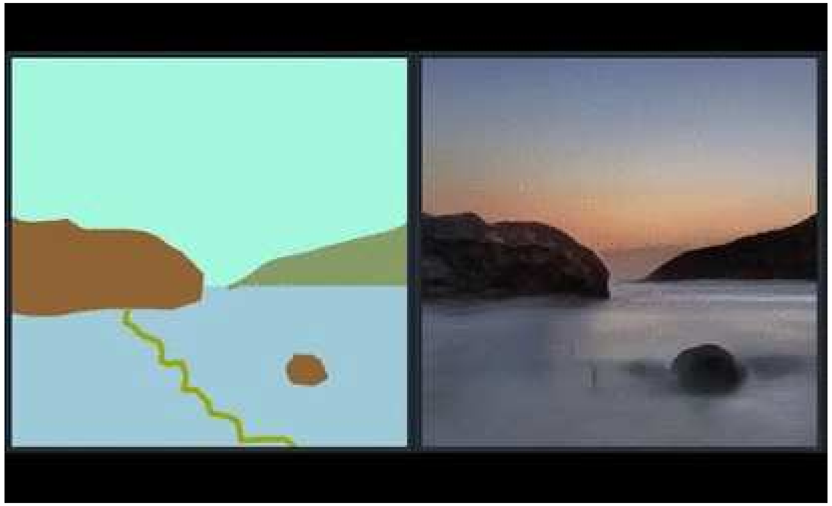

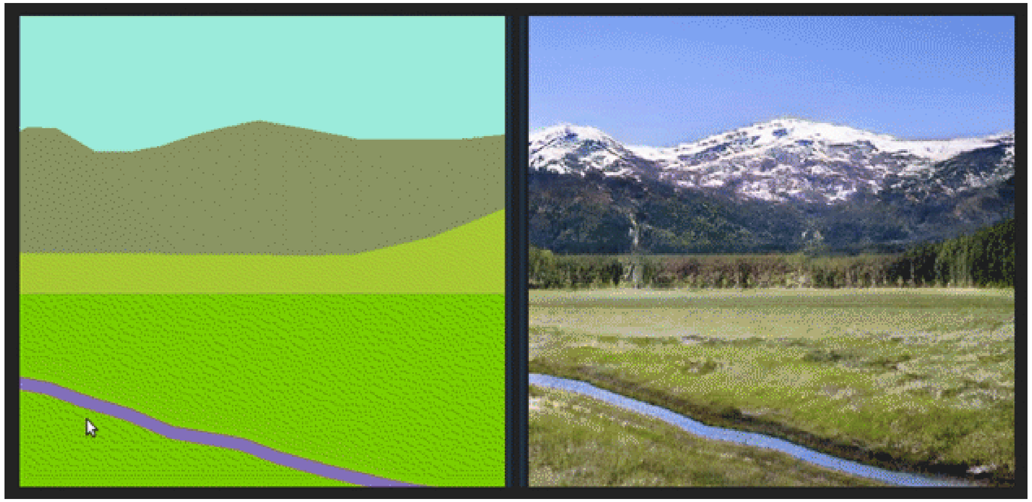

Why are they so popular ?

Large scale GAN training for high fidelity natural image synthesis Brock et al., ICLR'19

Figure 6: Additional samples generated by our model at 512×512 resolution.

Some maths

Generator G:

- maps the noise to the data space
- implicitly defines the distribution $p_{model}(\ .,\theta)$
- G loss penalizes G for producing a sample that D classifies as fake
- G weights update through backpropagation from the discriminator loss through D and G

Backpropagation for D

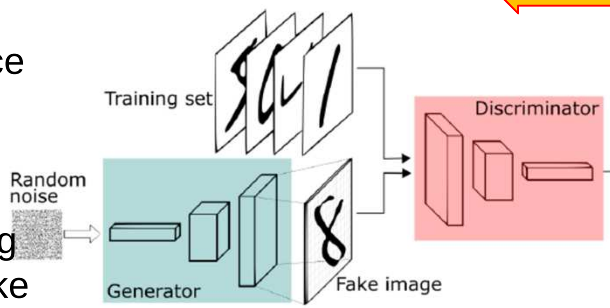

Discriminator D:

- classifies real data + fake data from G
- D loss penalizes missclassifications
- D weights update through backpropagation from the discriminator loss through D

Backpropagation for G

Log loss function to optimize

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbb{E}_{z\sim p_z(z)}[\log(1-D(G(z)))]$$

Discriminator output for real data x

Discriminator output for generated fake data G(Z)

Minimax: Inner maximization by discriminator and outer minimization by generator
=> Alternate discriminator and generator optimization

DCGAN

Unsupervised representation learning with deep convolutional generative adversarial networks, A. Radford et al, ICLR 2016

Simple and efficient network
Mainly composes of convolution layers without max pooling or fully connected layers

This paper gives some guidelines to design a good GAN architecture

A good start (but not the better loss function)

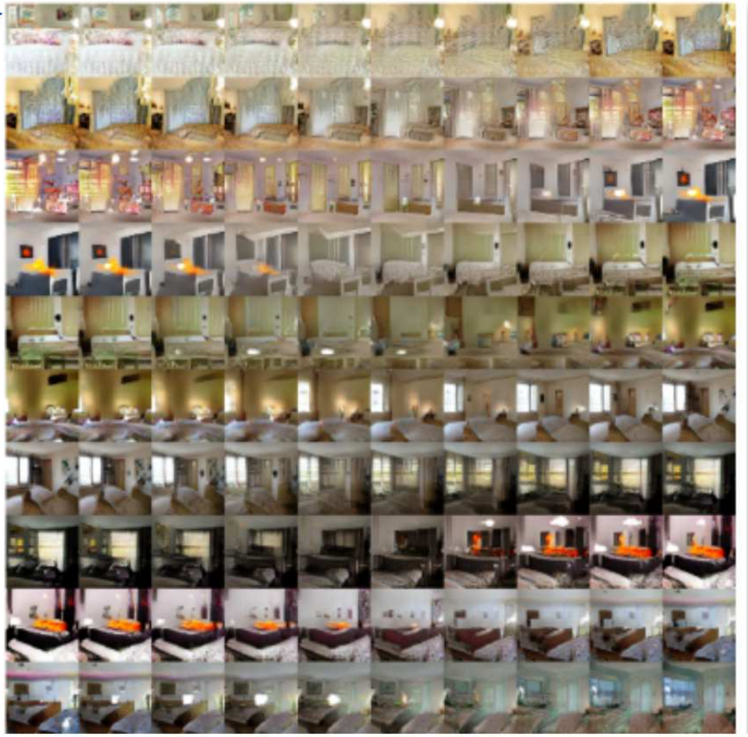

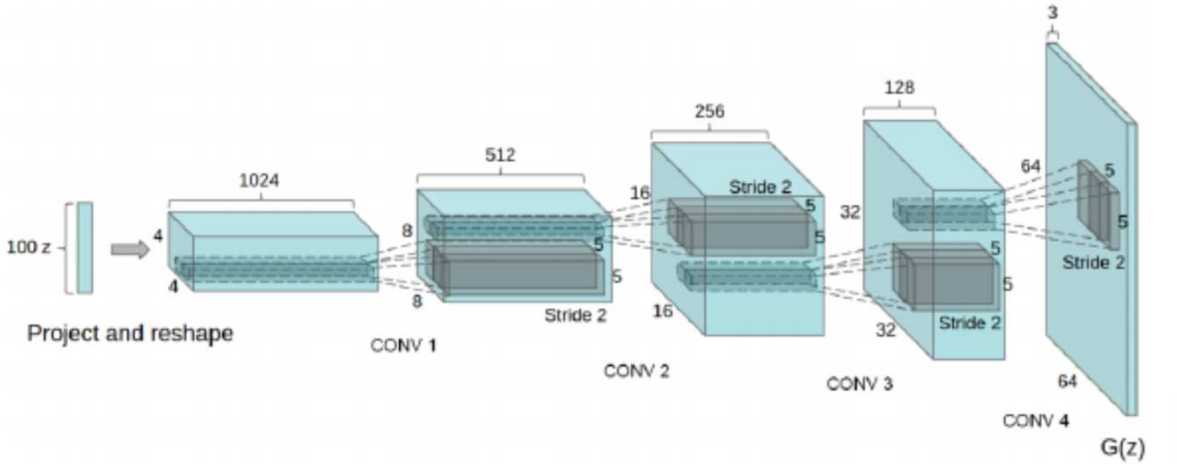

RÉPUBLIQUE FRANÇAISE
ONERA
THE FRENCH AEROSPACE LAB

Pitfalls & modifications

- **Saturation and gradients vanishing**

⇒ **If gradient is too small, it prevents the weights from changing their value**

- **Mode collapse**

During the training, the generator may collapse to a setting where it always produces same outputs

⇒ **It can trick the discriminator but gets stuck in a small space with extremely low variety**

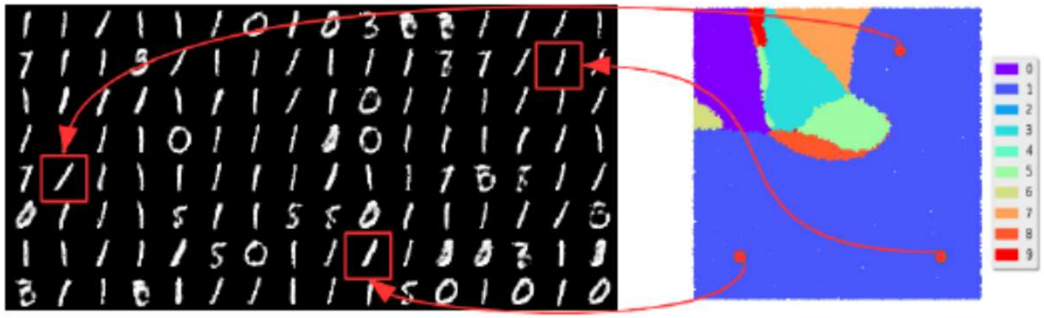

Illustration from Dist-GAN: An Improved GAN using Distance Constraints

Modifications

- **Solutions proposed**

⇒ Change minimax formulation

- Before: minimize $_{\theta_G} \{\mathbb{E}_{z \sim p_z(z)}[\log(1 - D(G(z;\theta_G);\theta_D))]\}$
- After: maximize $_{\theta_G} \{\mathbb{E}_{z \sim p_z(z)}[\log D(G(z;\theta_G);\theta_D)]\}$

⇒ Change loss for **Wasserstein loss**

Wasserstein is a distance function between probability distributions on a given metric space *M*

The Wasserstein loss enables to **train the discriminator to optimality** without vanishing gradients

=> the discriminator doesn't get stuck in local minima

⇒ Regularization : adding noise to discriminator inputs, penalizing discriminator weights

⇒ **Encoder – Decoder architecture**

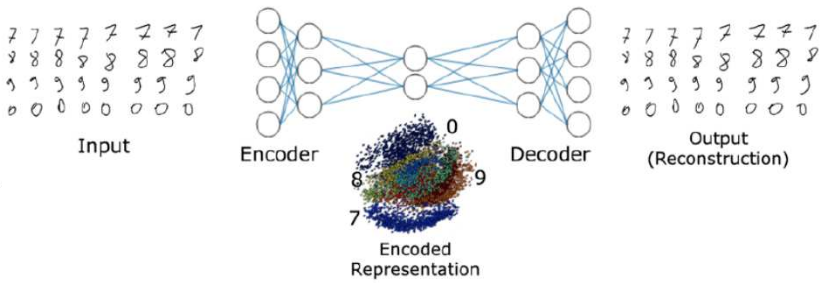

Outline

- Intro: Generative Modeling
- Generative Adversarial Networks
  - Principle and interest
  - Concept
  - Pitfalls & modifications
- Application to clouds synthesis
- Some examples of the use of GANs in astrophysics
- Summary

Goals

- Complete experimental databases to evaluate the performance of optronic sensors

- Keep the **same radiometric levels** and **cloud cover** but with different spatial distribution

- Use ice or liquid water content and thus not many training data

⇒ Input of radiative transfer code
⇒ Images for **different spectral bands**
different viewing angles

- **Keypoints**:

⇒ **Realistic cloud-edges**, as they are the main sources of false alarm for detecting targets on a cloudy sky

⇒ Be able to generate images larger than real input images, without mosaic effect

First results: master internship of P. de Perthuis

We used the GAN proposed by Zhou et al 2018 – Non-Stationary Texture Synthesis by Adversarial Expansion

- Impressive results

- Designed to do zoom

First results: master internship of P. de Perthuis

Learning is done on **1 large image**

The idea is to learn to generate 2k x 2k images from k x k patchs

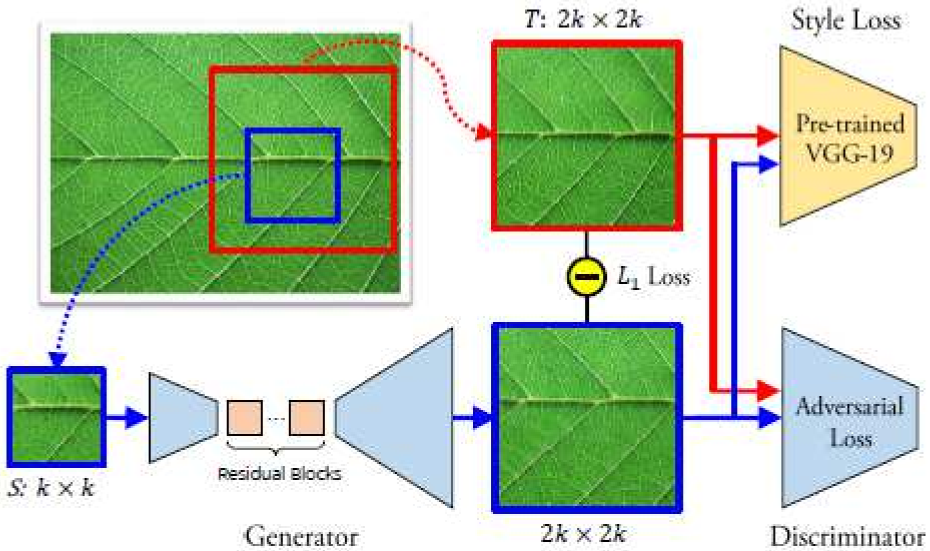

Encoder – decoder architecture

Diversity by using different patchs in one image

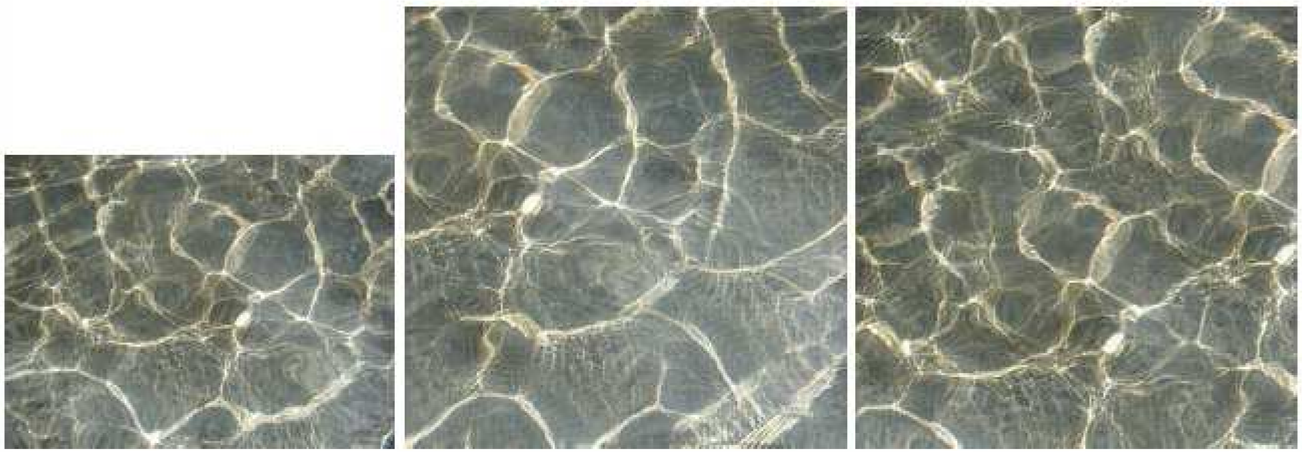

Results

Satellite images of optical thickness

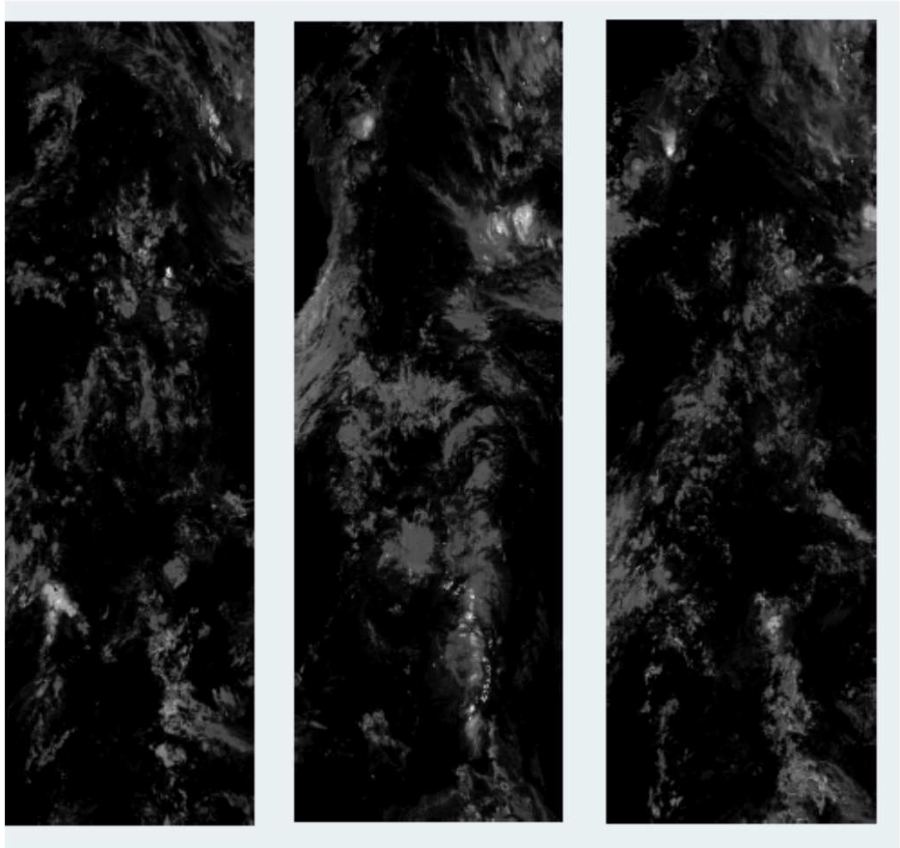

Simulations

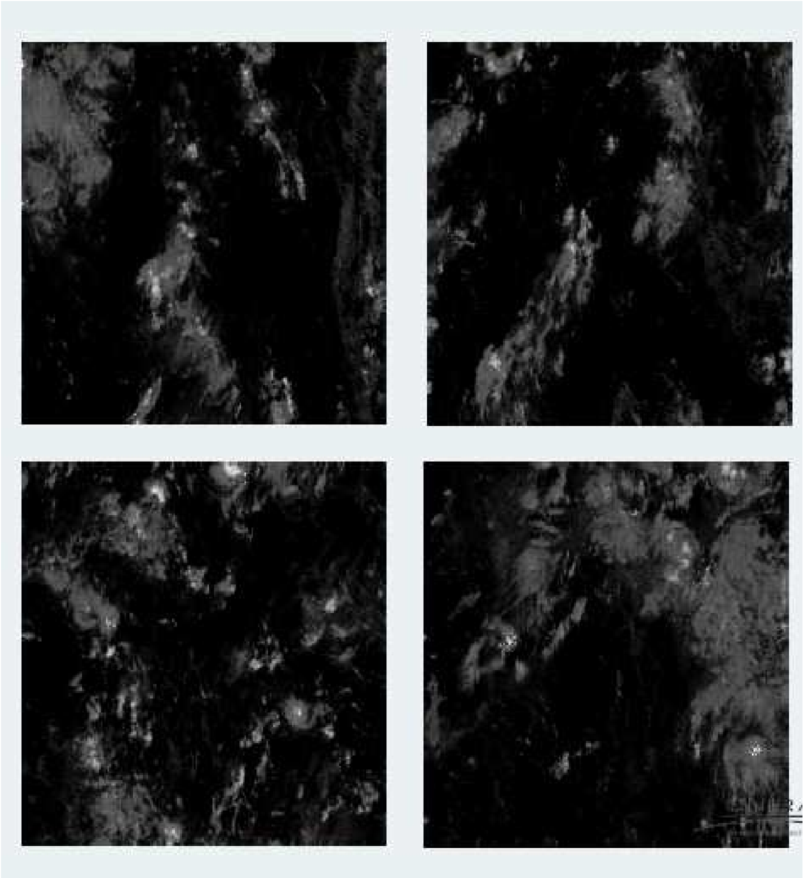

Our input: learning on a database instead of only one image / 2 GANs, the first from random noise

Results

Q : **How can we decide if images simulated by GANs are realistic ?**

- <u>Basic stats</u>: μ, σ, skewness, kurtosis

- <u>Stats for natural images</u>: - **power spectrum** ~1/f^p f spatial frequency

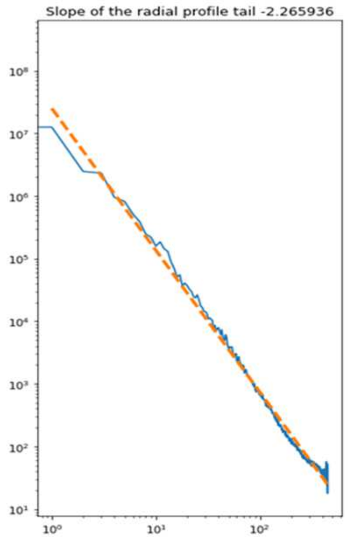

- distribution of **difference between two adjacent pixels**

in rows or columns = generalized Laplace distri $Ce^{-(\frac{x}{\beta})^{\alpha}}$

- <u>Stats for cloud edges</u>: quantiles of distribution of difference between two adjacent pixels in rows or columns at cloud edges

Results

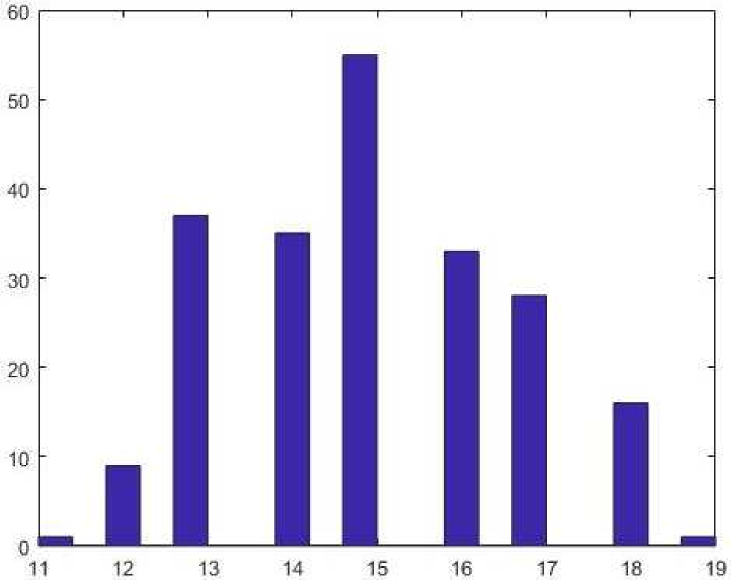

90 % quantile of differences between two adjacent pixels

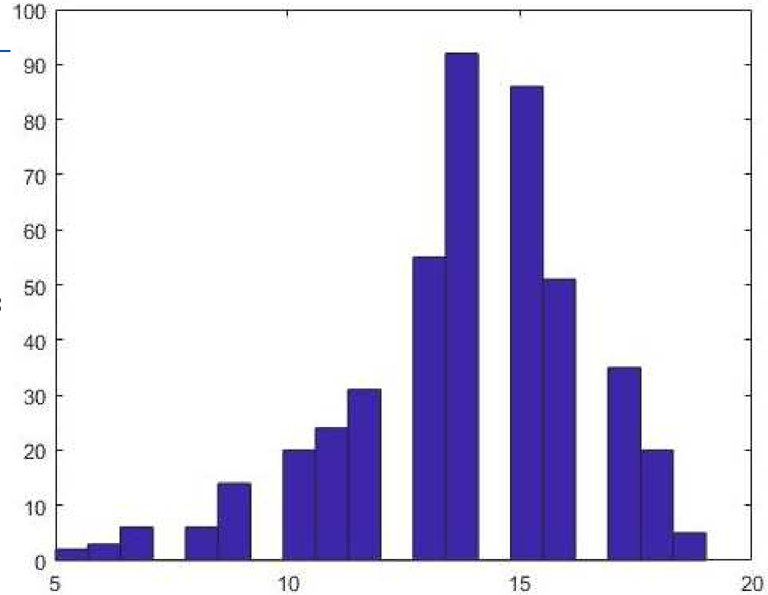

Database

Simulations

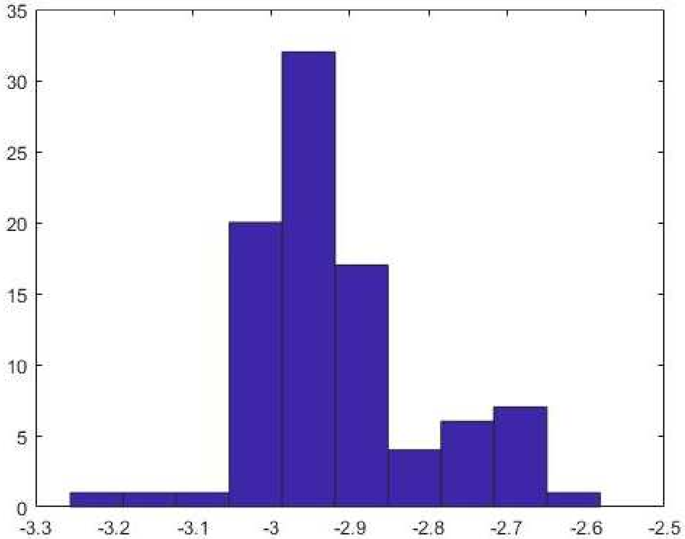

Slope of power spectrum

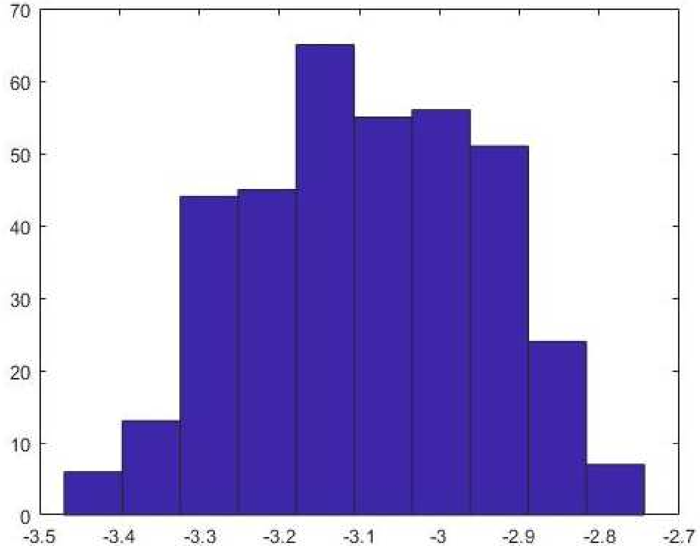

Ongoing work – PhD P. Chatillon (dir. Y. Gousseau Télécom)

Goals:

- Add terms linked to relevant prior physics information in loss function
- Accounting for multiscale effects
- Super-resolution from low resolution images

First results using SinGAN:

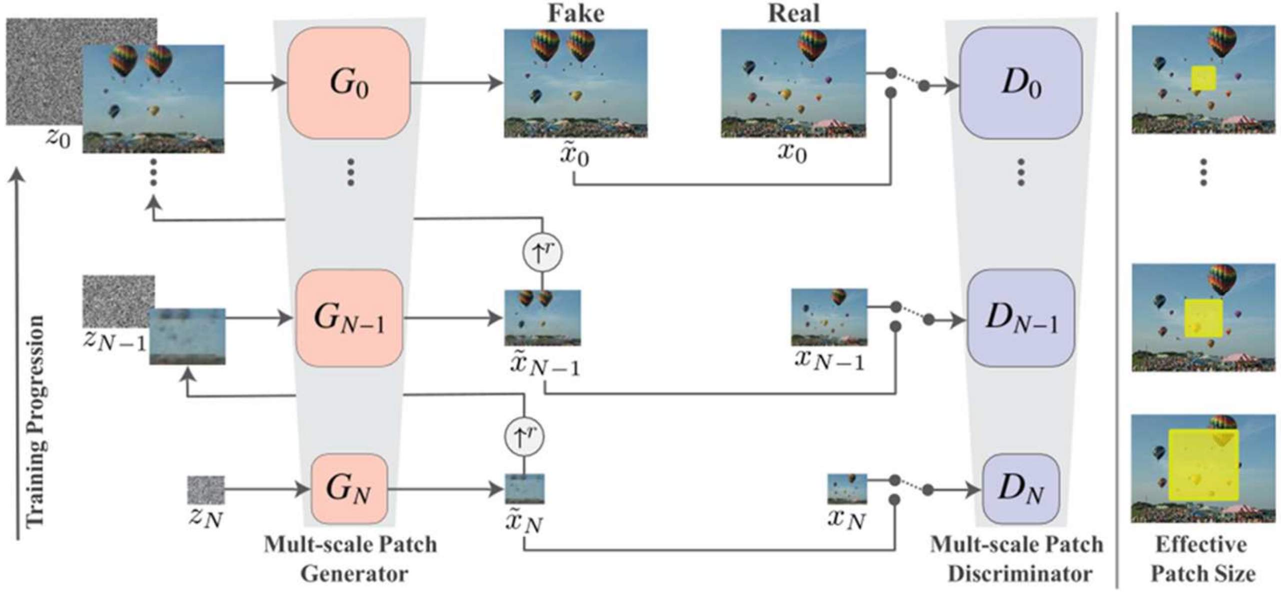

Aim at learning internal stats of patches at different scales within a single image

Trained in a coarse-to-fine fashion

Tamar Rott Shaham, Tali Dekel, Tomer Michaeli: SinGAN: Learning a Generative Model from a Single Natural Image

Ongoing work – PhD P. Chatillon (dir. Y. Gousseau Télécom)

First results:

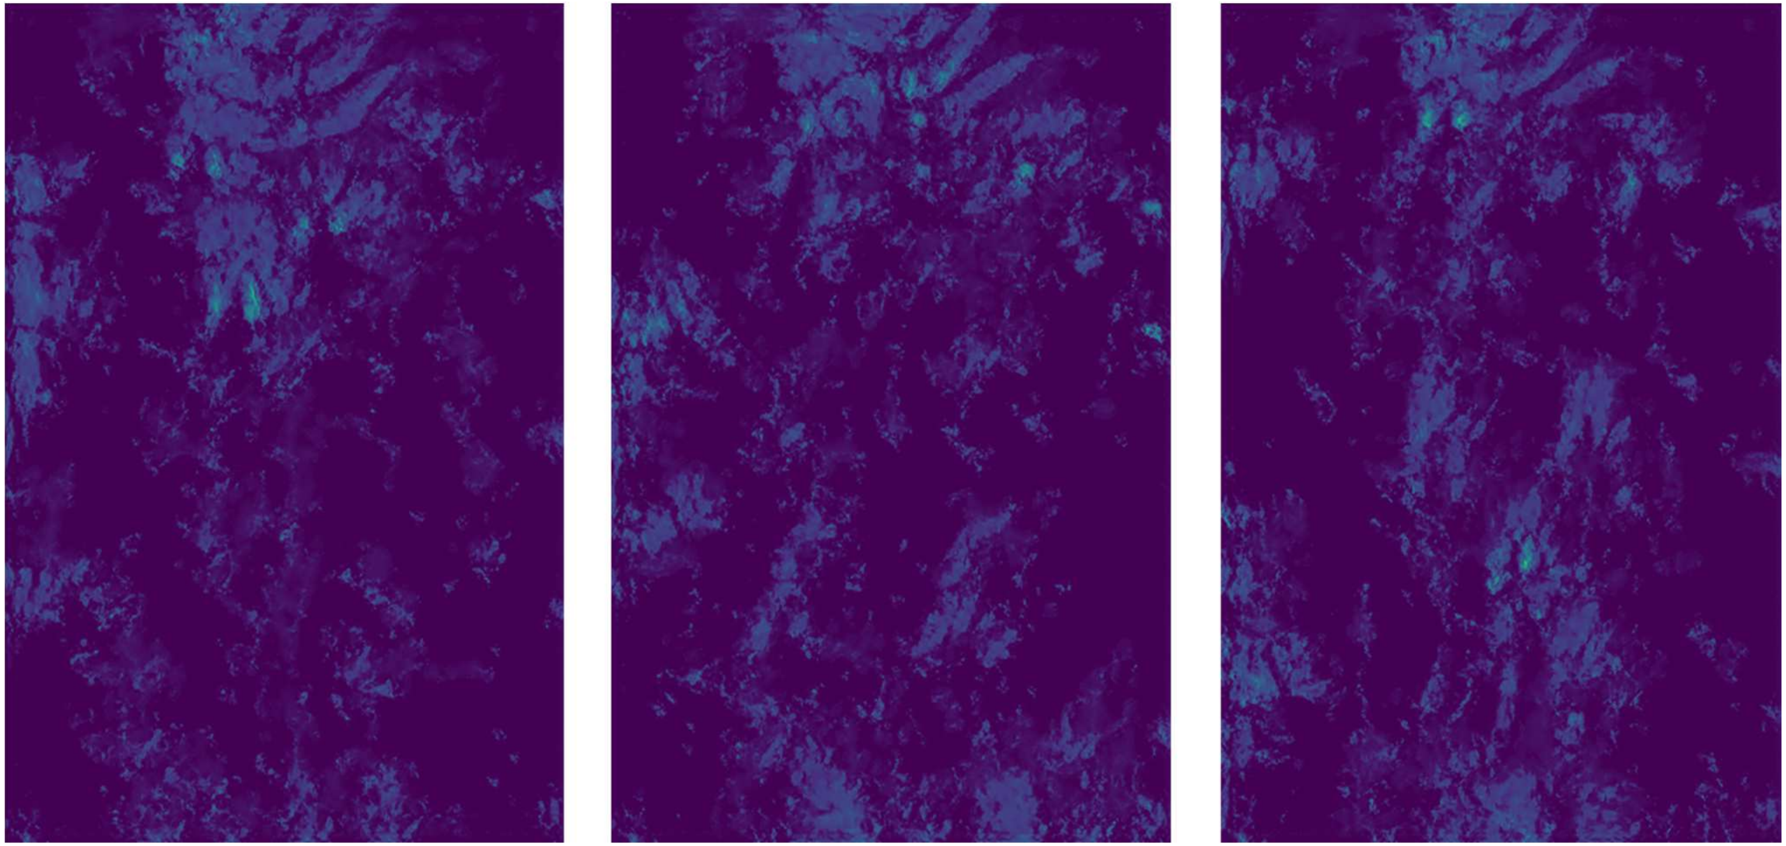

Simulations for satellite images of optical thickness

Pb: lack of diversity - copy of some areas

Outline

- Intro: Generative Modeling
- Generative Adversarial Networks
  - Principle and interest
  - Concept
  - Pitfalls & modifications
- Application to clouds synthesis
- Some examples of the use of GANs in astrophysics
- Summary

Some examples of the use of GANs in astrophysics

- A site listing ~ 300 ML papers in cosmology: https://github.com/georgestein/ml-in-cosmology

Some of the GAN applications cited:

- Fast cosmic web simulations with generative adversarial networks
- Painting halos from 3D dark matter fields using Wasserstein mapping networks
- HIGAN: Cosmic Neutral Hydrogen with Generative Adversarial Networks
- A black box for dark sector physics: Predicting dark matter annihilation feedback with conditional GANs
- Super-resolution emulator of cosmological simulations using deep physical models
- Emulation of cosmological mass maps with conditional generative adversarial networks
- Towards Universal Cosmological Emulators with Generative Adversarial Networks
- AI-assisted super-resolution cosmological simulations
- CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks
- Denoising Weak Lensing Mass Maps with Deep Learning
- Decoding Cosmological Information in Weak-Lensing Mass Maps with Generative Adversarial Networks
- CMB-GAN: Fast Simulations of Cosmic Microwave background anisotropy maps using Deep Learning
- Inpainting Galactic Foreground Intensity and Polarization maps using Convolutional Neural Network

Some examples of the use of GANs in astrophysics

- cosmoGAN: creates high-fidelity, weak gravitational lensing convergence maps
  DOI: 10.1186/s40668-019-0029-9

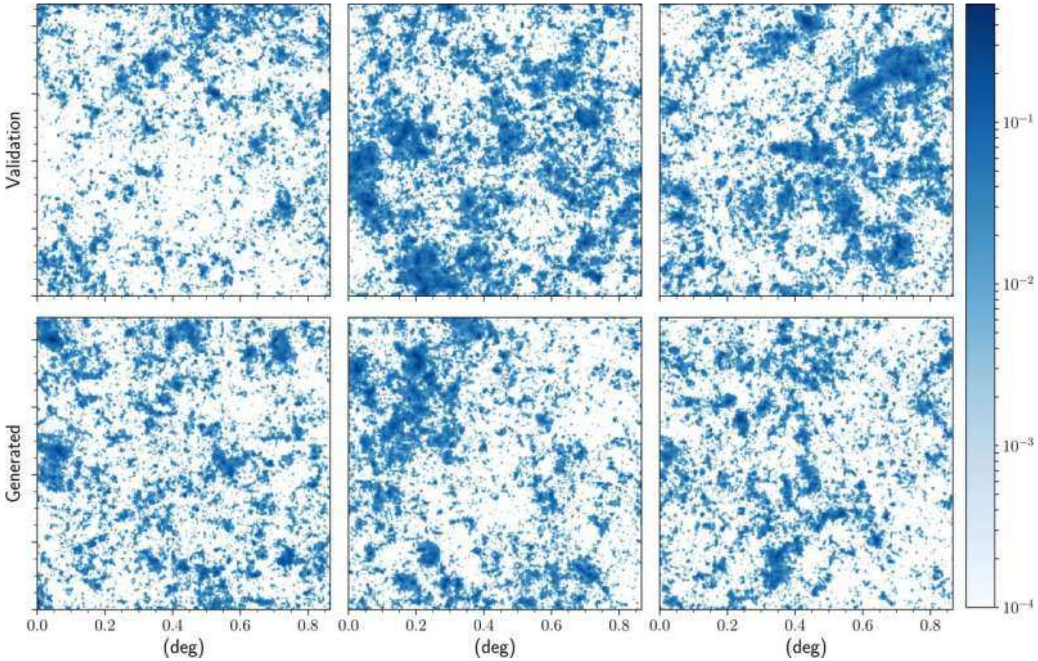

convergence maps are described by the same summary statistics as the fully simulated maps

\+ GANs are accurate and fast

\- They are known to be unstable during training

=> Use of cosmology prior info: typical summary statistics, to evaluate generator

Some examples of the use of GANs in astrophysics

- Generative deep fields: arbitrarily sized, random synthetic astronomical images through deep learning

https://doi.org/10.1093/mnras/stz2886

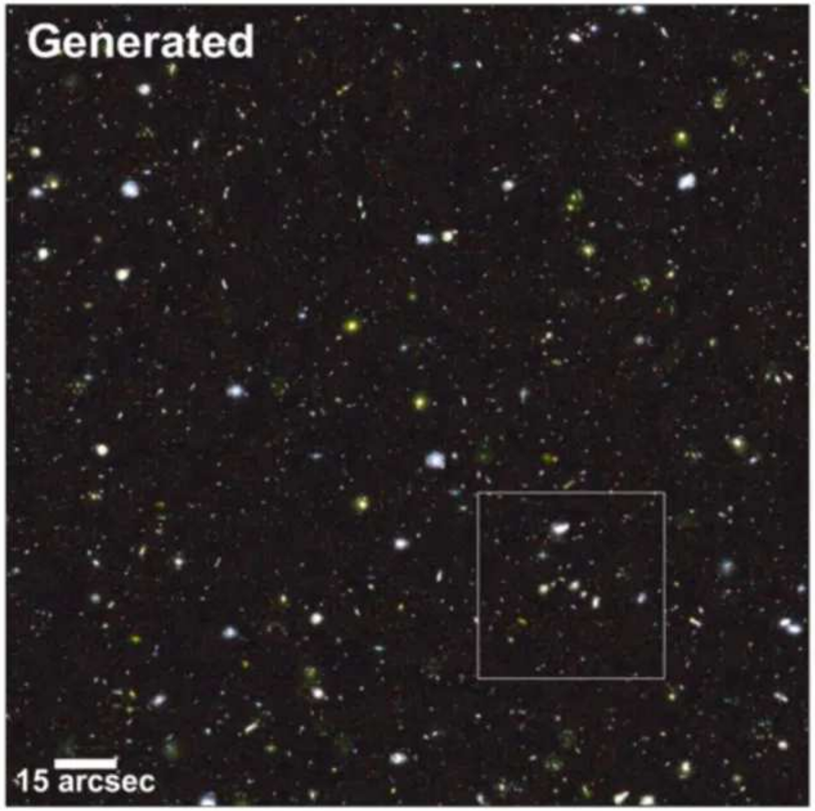

- Galaxy Image Simulation Using Progressive GANs

https://arxiv.org/abs/1909.12160

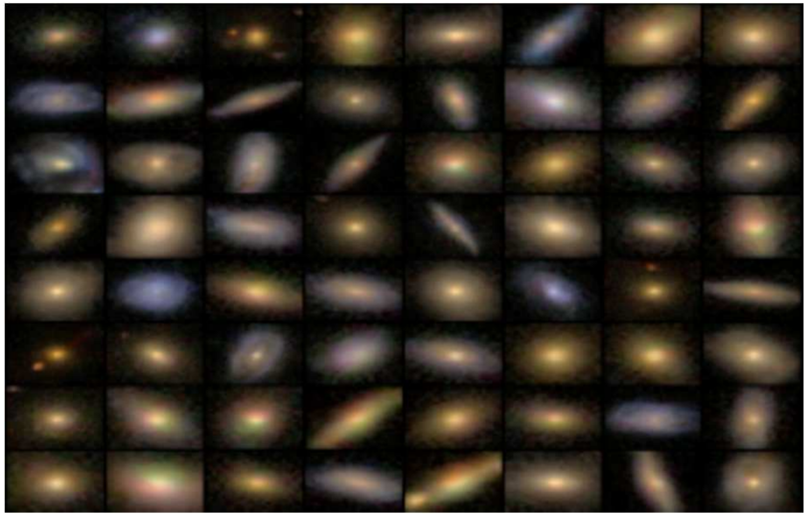

Summary

✓ Generative models

✓ Realistic rendering by optimizing the likelihood

✓ Good results in lots of domains: texture synthesis, style transfer, morphing, super-resolution, denoising, text to image translation…

X Unstable training, mode collapse

=> Wasserstein-GANs, MMD-GANs

X Lack of a proper evaluation metric to inform about training

X Need of multiscale learning => details of textures