Dilated Spatial Generative Adversarial Networks for Ergodic Image Generation

Cyprien Ruffino[1] , Romain Hérault[1] , Eric Laloy[2] and Gilles Gasso[1]

December 6, 2019

[1]Normandie Univ, UNIROUEN, UNIHAVRE, INSA Rouen, LITIS, 76 000 Rouen, France
[2]Belgian Nuclear Research, Institute Environment, Health and Safety, Boeretang 200 - BE-2400 Mol, Belgium

Motivations

- Geostatistical simulation: subsurface modeling
- Models underground water channels
- Need to generate ergodic (texture-like) data
- No global dependencies in the data

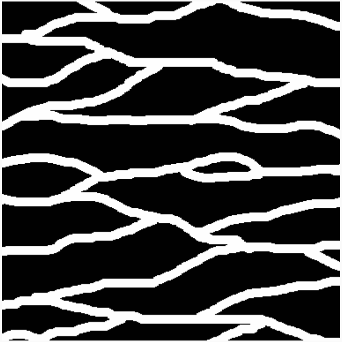

Dataset

Classical geostatistical dataset : $384 \times 384px$ samples from a $2500 \times 2500px$ binary subsurface model

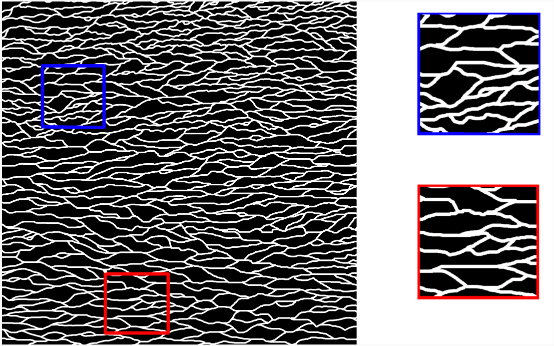

Generative Adversarial Networks[GPAM+14]

Principle

- Two networks : a generator G and a discriminator D
- The **generator** produces synthetic data from a random noise
- The **discriminator** tries to distinguish synthetic data from real ones
- G aims to maximize the error of D

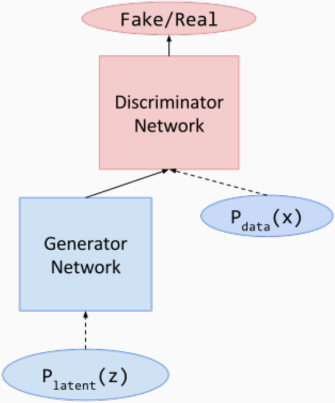

Spatial Generative Adversarial Networks

- Fully-convolutional : generates globally ergodic data
- Used in texture generation tasks
- Changes in the input only affects a local part of the output

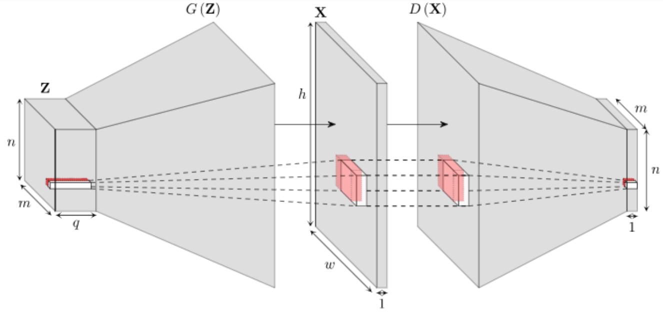

Previous works [LHJL18]

SGAN architecture

- Based on DCGAN [RMC15]
- Generator : 5 deconvolutional layers
- Discriminator : 5 convolutional layers

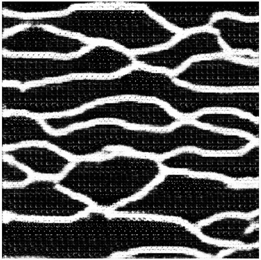

Drawbacks

- Generated data is noisy and blurry
- Ad-hoc solutions : median filtering

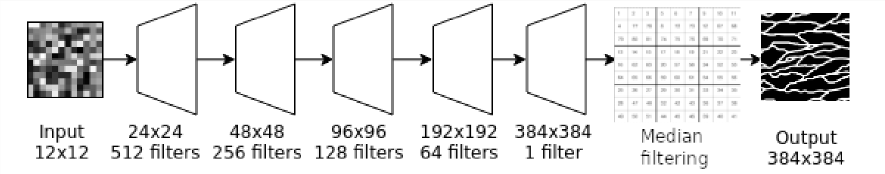

Dilated Convolution [YK15]

Dilated Convolutions

- The filter’s weights are spaced with zeros
- Allows to control the size of the filters without increasing the number of parameters
- Larger receptive field without pooling

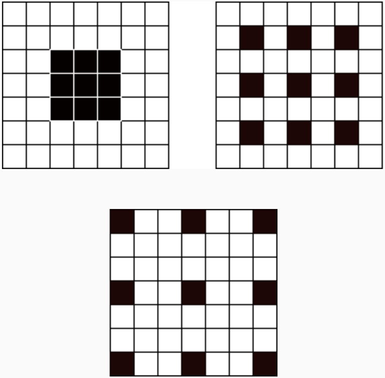

Dilated Spatial GANs

Generator architecture

- Same first five layers as the previous approach
- We add 5 dilated convolutional layers with ReLU activations
- 5×5 filters with dilation from 1 to 5
- This progressively increase the receptive field

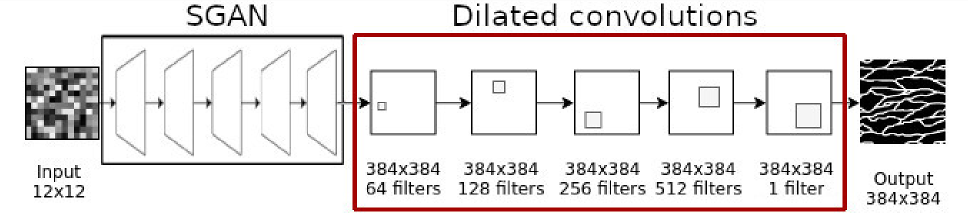

Results: Samples

Previous method

Our method

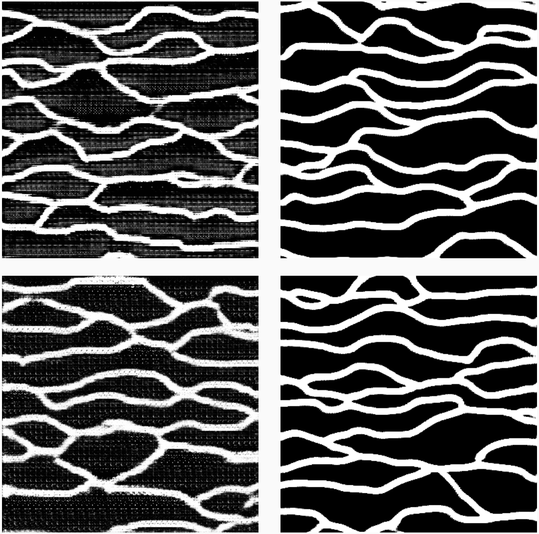

Results: Connectivity functions

Previous method

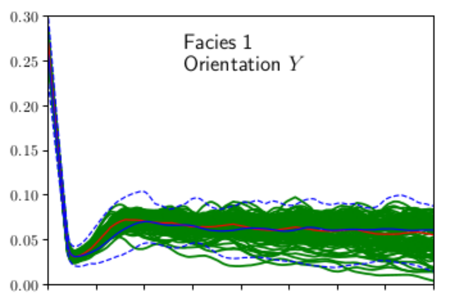

Our method

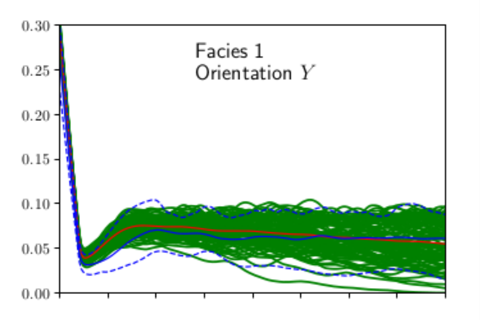

- Represents the geostastistical properties of the data
- Green curves: connectivity functions of generated samples
- Blue curves : max, min and mean values of real samples

Results: Metrics

Total Variation Norm

	Real	SGAN	Our approach
TV_i	$\mathbf{5.37 \times 10^{-2}}$	7.41×10^{-2}	$\mathbf{6.07 \times 10^{-2}}$
TV_a	$\mathbf{5.72 \times 10^{-2}}$	8.02×10^{-2}	$\mathbf{6.53 \times 10^{-2}}$

χ^2 distance between features

	SGAN	Our approach
LBP $R = 1$	10.13	**2.39**
LBP $R = 2$	24.26	**2.33**
HOG	5.79×10^{-4}	$\mathbf{2.37 \times 10^{-4}}$

Conclusion

- Dilated convolutions can enhance the visual quality of the samples
- We obtain significantly better results than the previous method

Extensions

- Applications to similar tasks : natural texture generation

Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio.
Generative Adversarial Nets.
2014.

Nikolay Jetchev, Urs Bergmann, Roland Vollgraf, and Zalando Research.
Texture Synthesis with Spatial Generative Adversarial Networks.
2017.

Eric Laloy, Romain Hérault, Diederik Jacques, and Niklas Linde.
Training-image based geostatistical inversion using a spatial generative adversarial neural network.
Water Resources Research, 54(1):381–406, 2018.

Alec Radford, Luke Metz, and Soumith Chintala.
Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks.
nov 2015.

Fisher Yu and Vladlen Koltun.
Multi-Scale Context Aggregation by Dilated Convolutions.
2015.