Modèles de diffusion pour la synthèse de textures anisotropes

Kinan Abbas[1,2], Patrice Abry[1] and Stephane Roux[1]
[1] *CNRS, ENS de Lyon, LPENSL*, UMR5672, 69342, Lyon cedex 07, France
[2] *XMBAUBLE*, Lyon, France
{firstname.lastname}@ens-lyon.fr

Diffusion Models

- Diffusion models, inspired by thermodynamics [1], generate data by reversing a noise-adding process.
- DDPM [2] uses a Markov chain to add Gaussian noise over T steps, transforming data into noise, then trains a U-Net to denoise it back.

Forward Process: Adds noise iteratively:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1-\beta_t}x_{t-1}, \beta_t I) \quad (1)$$

Reverse Process: Learns to denoise:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (2)$$

Training Loss: Minimizes noise prediction error:

$$\mathcal{L} = \mathbb{E}_{t,x_0,\epsilon}\left[\|\epsilon - \epsilon_\theta(x_t, t)\|^2\right] \quad (3)$$

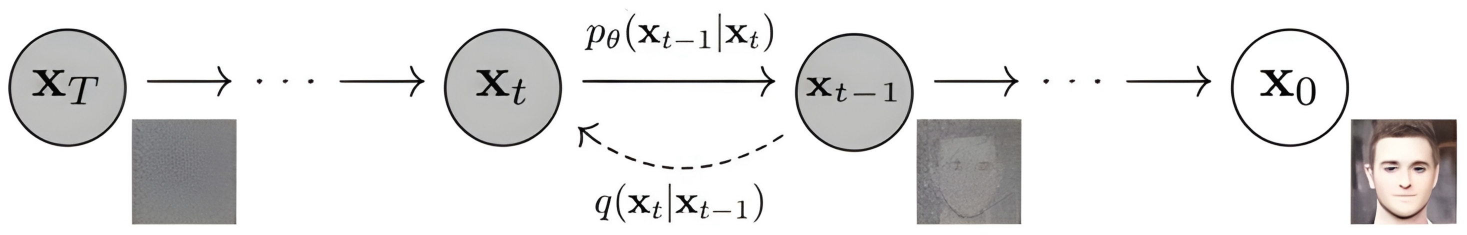

Figure: Forward and reverse diffusion process in DDPM (Source: Ho et al., 2020).

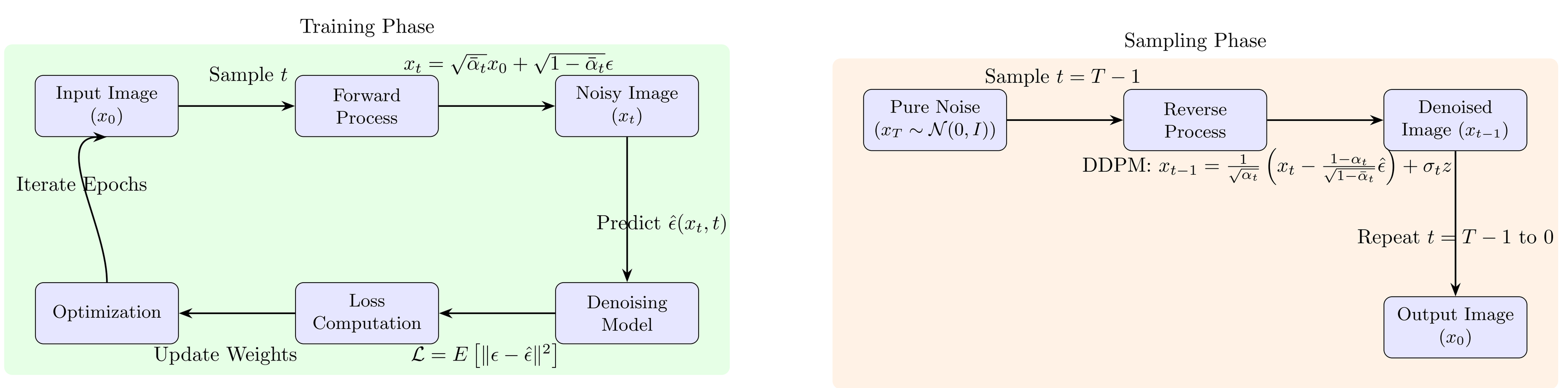

Figure: Block diagram of a diffusion model, showing the training phase (left) where noise is added to input images and the model learns to denoise, and the sampling phase (right) where images are generated from pure noise using the trained model. Equations describe key steps, including the forward process, loss, and reverse process (DDPM).

- The linear noise scheduler progressively increases β_t from a small value (e.g., 0.0001) to a larger value (e.g., 0.02) over T steps [3].
  $\beta_t = \beta_{\text{start}} + \frac{t-1}{T-1}(\beta_{\text{end}} - \beta_{\text{start}})$, where $\beta_{\text{start}} = 0.0001$ and $\beta_{\text{end}} = 0.02$.

Training Configuration

- Based on DDPM [2] with 1000 timesteps, the model trains a U-Net to minimize MSE loss between predicted noise $\epsilon_\theta(\mathbf{x}_t, t)$ and actual noise ϵ, reconstructing 128×128 textures.
- **Training**: 2000 epochs, Adam (lr=0.0002, batch=32) on AD102GL [RTX 6000 Ada] GPU. Tested linear, cosine, quadratic, and sigmoid noise schedulers, trained independently per scheduler.

Architecture:

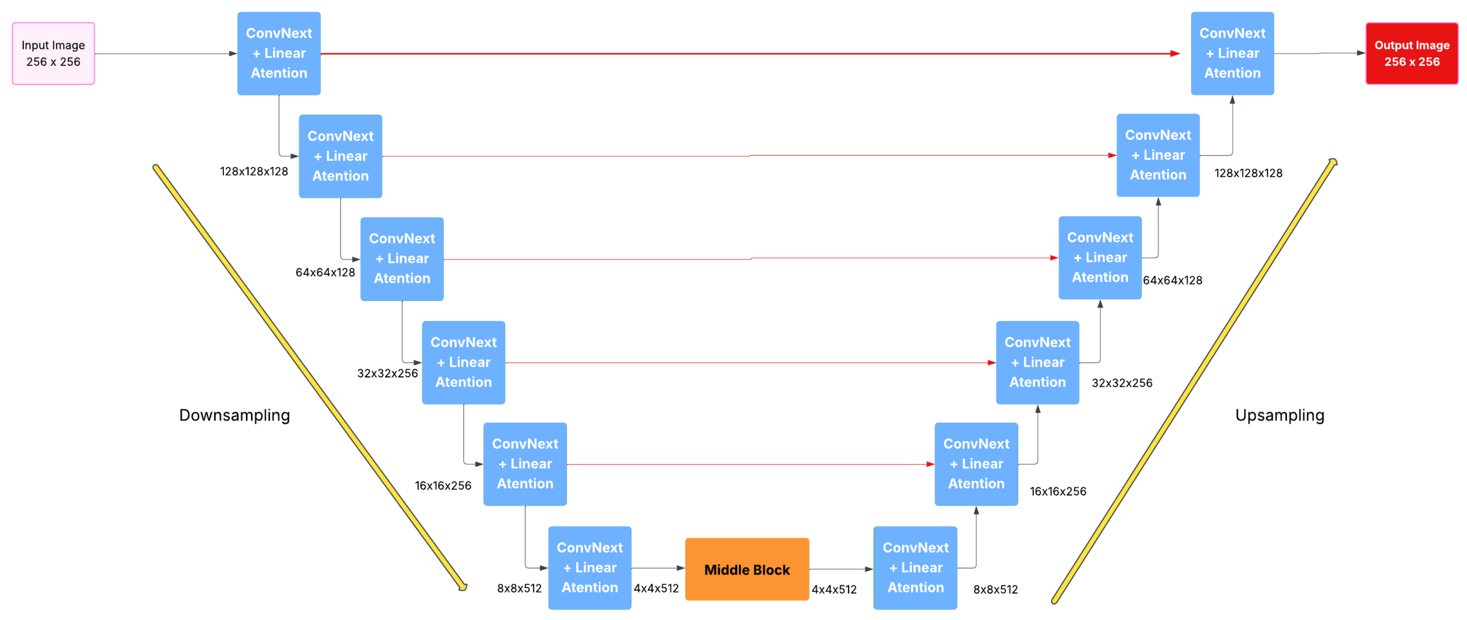

Figure: U-NET with ResNet Blocks

Dataset

- Two datasets, each with 3000 AFBF images generated for different spectral density parametrizations.
- Scenario 1: Anisotropy induced solely by Hurst function h while topothesy uniform.
- Scenario 2: More complex anisotropy induced by directionalities of both h and τ.

Conclusion and Future Work

- This study demonstrates that diffusion models with a linear scheduler effectively synthesize anisotropic textures, as evidenced by the close alignment of generated textures with the training data across all scales and bands.
- The histograms also confirm the model's ability to reproduce not only the ensemble averages but also the variability between realizations of the same ensemble.

Future Work: explore the synthesis of textures with more complex dynamics, for example non-Gaussian, multifractal, multivariate, or heterogeneous.

Anisotropic Self-Similar Textures

- Anisotropic Self-Similar Textures are patterns that look similar at different scales (*self-similar*) and they gave properties that vary with direction (*anisotropic*). They can be modeled by Anisotropic Fractional Brownian Fields (AFBF)
- AFBF are Gaussian processes that model textures with Self-similarity and Anisotropy:

$$X^f(t) = \int_{\mathbb{R}^2} \left(e^{i\langle t,\xi\rangle} - 1\right) \sqrt{f(\xi)}\, \widehat{W}(\xi)\, d\xi \quad (4)$$

where $\widehat{W}(\xi)$ is Gaussian white noise in frequency domain.

- The spectral density $f(\xi)$ controls texture properties:

$$f(\xi) = \tau\left(\frac{\xi}{|\xi|}\right) \|\xi\|^{-(2h(\frac{\xi}{|\xi|})+2)} \quad (5)$$

Where $\frac{\xi}{|\xi|}$ is the Direction of frequency, $\|\xi\|$ is the Magnitude of frequency, $h\left(\frac{\xi}{|\xi|}\right)$ is the Hurst function (roughness) and $\tau\left(\frac{\xi}{|\xi|}\right)$ is the Topothesy (intensity).

- Hurst Function h: Controls roughness in each direction and it ranges: $0 < h < 1$
- Topothesy τ: Adjusts intensity in each direction. Higher τ: Stronger features

Dual-Tree Complex Wavelet Transform (DTCWT)

- Analyze textures by decomposing them into different scales and directions. It is ideal for self-similar and anisotropic textures. It decomposes texture X using mother wavelets $\psi^{(b)}$ for 6 directions ($b = 1, \dots, 6$)

$$c_{j,k}^{(b)} = \langle X, 2^{-j}\psi^{(b)}(2^{-j}(\cdot - k))\rangle \quad (6)$$

- Energy of coefficients [4, 5]:

$$S(j,b) = \sum_k |c_{j,k}^{(b)}|^2 \sim \mathcal{V}(\tau, h, \psi^{(b)})\, 2^{j(2H(h,\psi^{(b)})+1)} \quad (7)$$

- Self-Similarity: Power law indicates consistent scaling. Anisotropy: Variations in $H(h, \psi^{(b)})$ show directional roughness. Variations in $\mathcal{V}(\tau, h, \psi^{(b)})$ show directional intensity.

Results

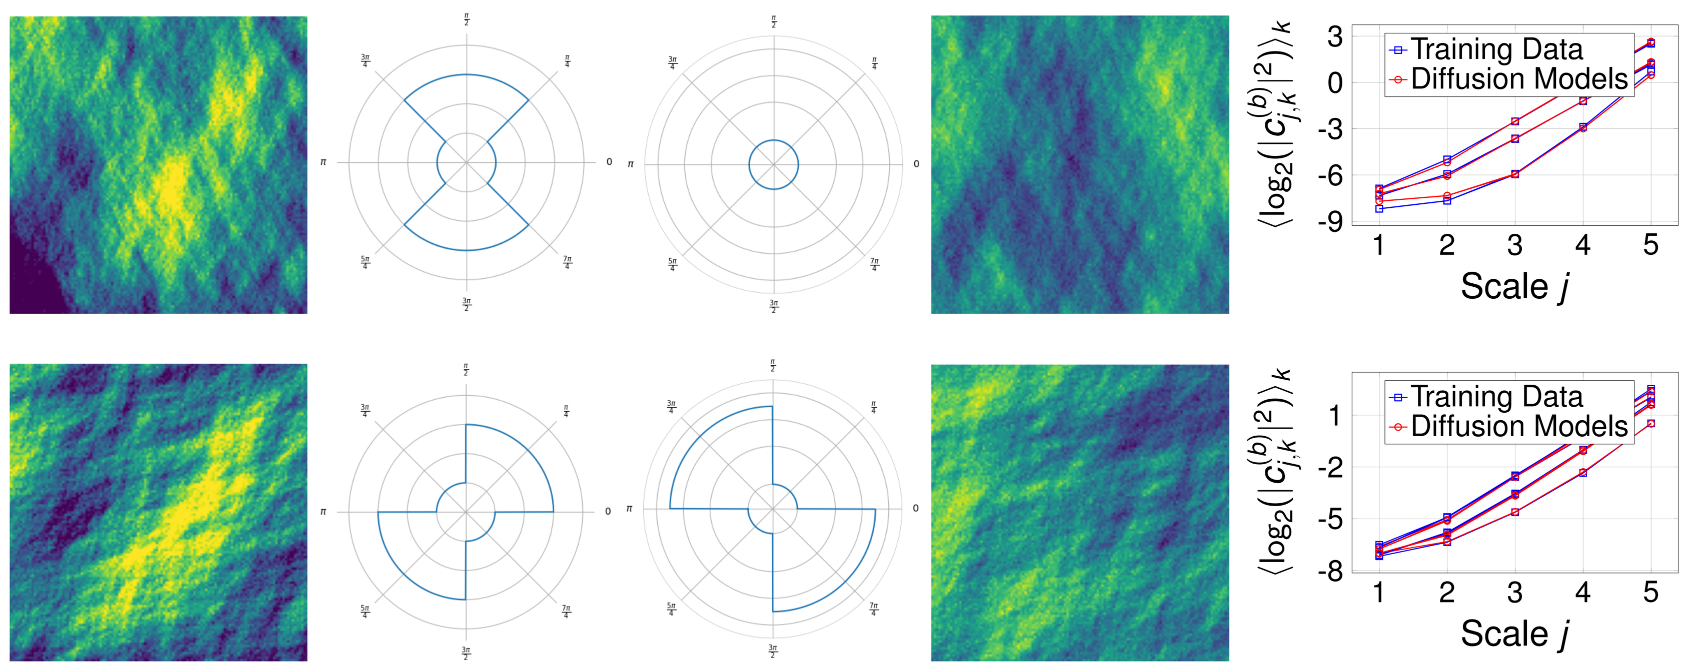

Figure: **Texture Synthesis Performance.** From left to right: (a) Training sample. (b) Hurst function H. (c) Topothesy function τ. (d) Sample generated by the diffusion model. (e) Spatial averages of the base-2 logarithm of the variance of the modulus of the multiband complex wavelet coefficients, for each of the six bands independently, as a function of octaves. The two rows correspond to the scenarios, characterized by different anisotropies.

Table: Statistics of training and generated textures

	Scénario 1		Scénario 2	
Moment	Entraînement	Diffusion	Entraînement	Diffusion
Moyenne	0.00	0.00	0.00	0.00
Variance	0.23	0.20	0.23	0.22
Asymétrie	-0.01	0.00	-0.01	-0.03
Kurtosis	2.38	2.8	2.39	2.70

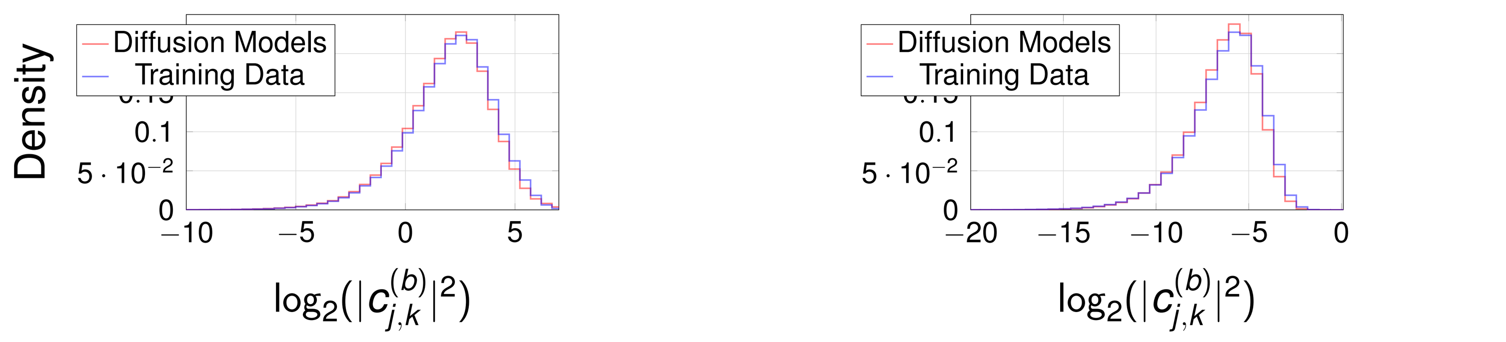

Acknowledgements and References

We sincerely thank Pulsalys for supporting our work. We also acknowledge the support of the Centre Blaise Pascal's IT test platform at ENS de Lyon (Lyon, France) for Machine Learning facilities.

[1] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. "Deep unsupervised learning using nonequilibrium thermodynamics". In: *International conference on machine learning*. pmlr. 2015, pp. 2256–2265.

[2] Jonathan Ho, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models". In: *Advances in neural information processing systems* 33 (2020), pp. 6840–6851.

[3] Kinan Abbas, Patrice Abry, and Stéphane G. Roux. "Diffusion Models for Multifractal Texture Synthesis". In: *The 33rd European Signal Processing Conference (EUSIPCO 2025)*. Palerme (Italie), Italy, Sept. 2025. URL: https://hal.science/hal-05116753.

[4] L. Davy, N. Pustelnik, and P. Abry. "Combining dual-tree wavelet analysis and proximal optimization for anisotropic scale-free texture segmentation". In: *IEEE ICASSP 2023*. Rhodes Island, Greece, 2023.

[5] L. Davy, N. Pustelnik, and P. Abry. "Local wavelet analysis of Gaussian fields: Application to anisotropic scale-free segmentation". In: (2025). Preprint.