Diffusion Models for Multifractal Texture Synthesis

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.
- Denoising diffusion probabilistic models (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:

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

Reverse Process: Learns to denoise:

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

Training Loss: Minimizes noise prediction error:

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

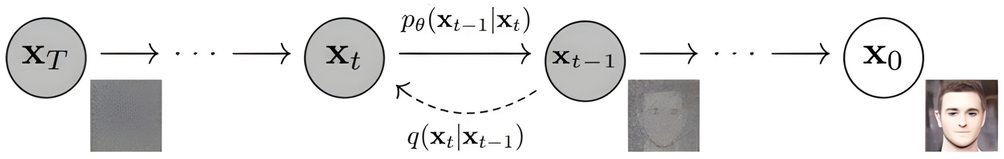

Figure: Forward and reverse diffusion process in DDPM (Source: Ho et al., 2020).

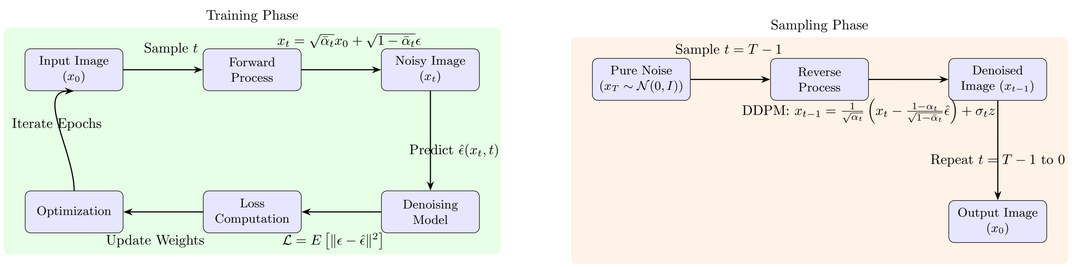

Figure: Block diagram of a diffusion model, showing the training phase (left) where noise is added to input images and the model learns to denoise, and the sampling phase (right) where images are generated from pure noise using the trained model. Equations describe key steps, including the forward process,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

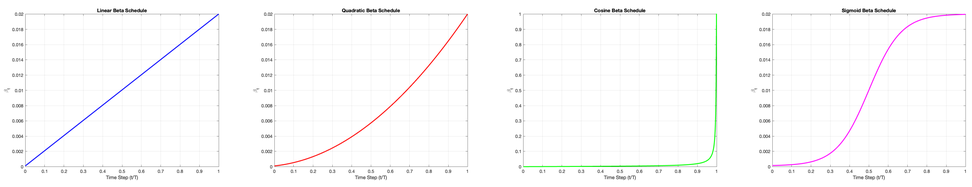

(a) Linear Schedule (b) Quadratic Schedule (c) Cosine Schedule (d) Sigmoid Schedule

Figure: Traditional noise schedulers used in DDPM

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 256×256 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:

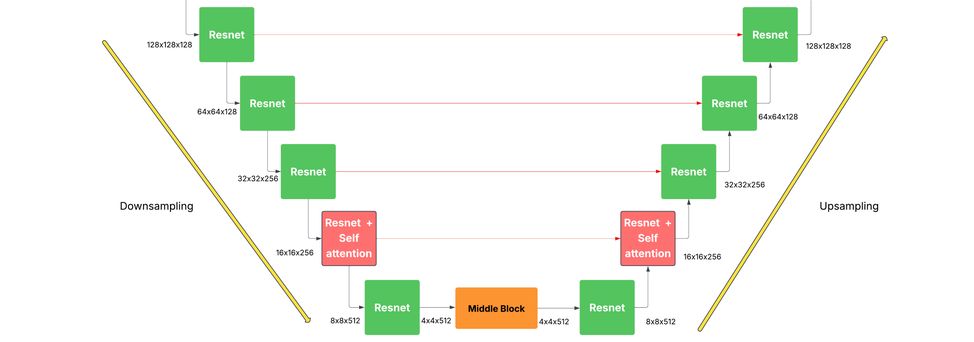

Figure: U-NET with ResNet Blocks

Conclusion and Future Work

- This study explored diffusion neural networks' capability to synthesize multifractal textures using a U-Net trained on 1000 univariate MRW textures with linear, cosine, quadratic, and sigmoid schedulers.
- Texture quality was assessed via wavelet-leader statistics $C_1(j)$ and $C_2(j)$, without embedding these in the loss.
- The linear scheduler slightly outperformed others, effectively reproducing global correlations ($C_1(2^j)$) and multifractality ($C_2(2^j)$).

Future Work: Extend to multivariate multifractal textures or those with anisotropy. Design specialized diffusion architectures to better handle complex statistical structures.

Multifractal Textures

- Multifractal Random Walks (MRW) [3, 4] model real-world textures with scale-free statistics [5]. Defined as:

$$X(\underline{r}) = G_H(\underline{r})e^{\omega_\lambda(\underline{r})}, \tag{4}$$

where $G_H(\underline{r})$ (2D-fGn, Hurst exponent H) and $\omega_\lambda(\underline{r})$ (covariance $C_{MF}\underline{r} = \lambda^2 \log\left(\frac{L}{\|\underline{r}\|+1}\right)$, parameter λ) induce multifractality.

- **Analysis**: Wavelet coefficients $\{d_X(j,\underline{k})\}$ and leaders $L(j,\underline{k})$ yield cumulants:

$$C_1(2^j) = c_1^{(0)} + c_1 \ln 2^j, \quad C_2(2^j) = c_2^{(0)} + c_2 \ln 2^j, \tag{5}$$

with $c_1 = H - \lambda^2/2$, $c_2 = -\lambda^2$, measuring global correlations and multifractality.

Results

- Evaluates texture quality using multifractal analysis, computing cumulants $C_1(j)$ and $C_2(j)$ of wavelet leader logarithms across scales 2^j for 1000 original MRW textures ($H = 0.8$, $\lambda^2 = 0.04$) and generated textures.
- **Qualitative**: Plots averaged $C_1(j)$ and $C_2(j)$ curves to assess scaling behavior alignment between original and generated textures.
- **Quantitative**: Estimates slopes $\hat{c}_1$ and $\hat{c}_2$ from linear regressions of $C_1(j)$ and $C_2(j)$ over scales 2^2 to 2^4. Calculates mean and standard deviation of slopes across 1000 samples to quantify multifractal scaling fidelity.

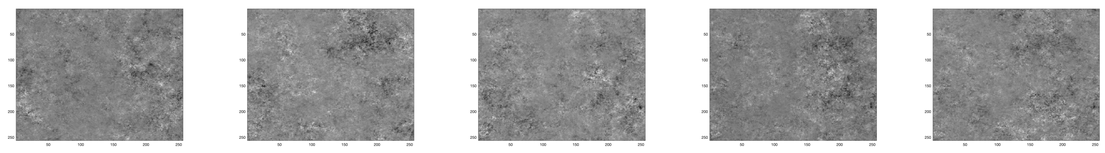

(a) Training Sample (b) Linear (c) Cosine (d) Quadratic (e) Sigmoid

Figure: **Texture Samples**. A sample from the training dataset alongside diffusion model-generated samples using linear, cosine, quadratic, and sigmoid noise schedulers.

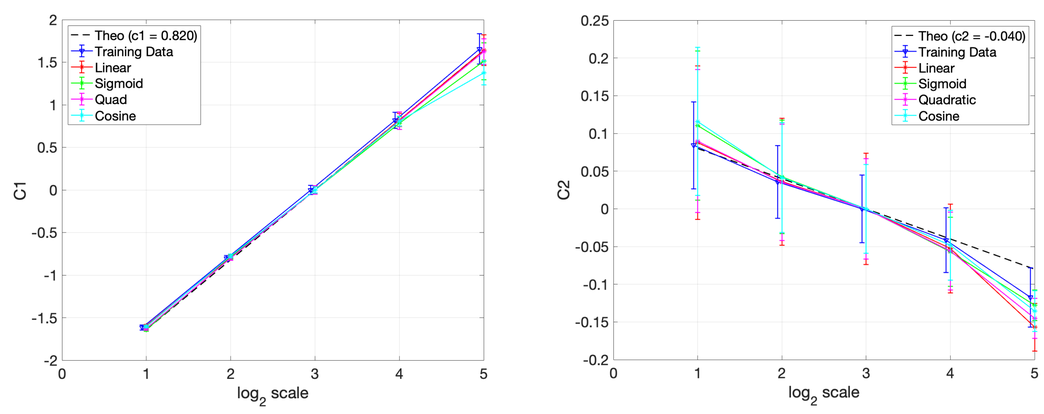

Figure: **Cumulants $C_1(2^j)$ and $C_2(2^j)$ as functions of scales**, averaged across the training set (blue) and diffusion model generated texture sets, obtained from different noise schedulers. Top: $C_1(2^j)$. Bottom: $C_2(2^j)$. The dashed black line materializes the theoretical scaling behavior across scales.

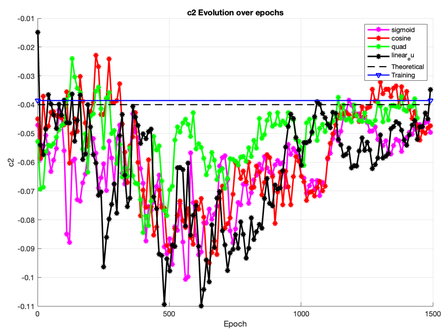

Figure: Cumulant C2 Per Epoch

Table: Means and standard deviations of $\hat{c}_1$ and $\hat{c}_2$, averaged over 1000 realizations for the training and for the diffusion model generated texture sets.

Dataset	$\hat{c}_1$		$\hat{c}_2$	
	Mean	**Std**	**Mean**	**Std**
Training	0.8109	0.0588	-0.0386	0.0242
Sigmoid	0.7924	0.0431	-0.0459	0.0190
Linear	0.7989	0.0438	-0.0417	0.0208
Cosine	0.7901	0.0361	-0.0420	0.0141
Quadratic	0.7851	0.0631	-0.0423	0.0241

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] B. B. Mandelbrot. "Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier". In: *J. Fluid Mech.* 62 (1974), pp. 331–358.

[4] E. Bacry, J. Delour, and J.-F. Muzy. "Multifractal random walk". In: *Phys. Rev. E* 64: 026103 (2001).

[5] S. Jaffard, P. Abry, and H. Wendt. "Irregularities and Scaling in Signal and Image Processing: Multifractal Analysis". In: *Benoit Mandelbrot: A Life in Many Dimensions*. Ed. by Michael Frame and Nathan Cohen. Singapore: World scientific publishing, 2015, pp. 31–116. ISBN: 9789814366069.