

# Beyond Blur: A Fluid Perspective on Generative Diffusion Models

## Supplementary Material

Grzegorz Gruszczyński<sup>1,2</sup>      Jakub Meixner<sup>1,3</sup>      Michal Wlodarczyk<sup>1,4</sup>  
Przemyslaw Musialski<sup>1,5,6</sup>

<sup>1</sup>IDEAS NCBR    <sup>2</sup>Samsung AI Center Warsaw    <sup>3</sup>Polish Academy of Sciences  
<sup>4</sup>Warsaw University of Technology    <sup>5</sup>IDEAS Research Institute    <sup>6</sup>New Jersey Institute of Technology  
g.gruszczyński@samsung.com, {kubameixner, mwlodarcz}@gmail.com, przem@njit.edu

### Appendix

In this appendix, §A contains hyperparameters and experiment setup. §B contains additional samples and interpolations for FFHQ-128, MNIST and LSUN Church datasets. §C contains solver implementation details.

#### A. Hyperparameter settings

We present the complete experimental configuration for our model architecture, including network topology details and optimization parameters. Our implementation utilizes a modified U-Net architecture with residual blocks and multi-head self-attention layers. Table 1 summarizes the dataset-specific configurations. The spatial resolution and batch size were selected to maximize GPU memory utilization while maintaining stable training dynamics.

Table 1. Neural network hyperparameters used during experiments on different datasets.

Dataset	Network param.	Layer multipliers	Base Channels	Learning rate	Resolution	Batch	Attention lvls
FFHQ	210904835	(1, 2, 3, 4, 5)	128	2e-05	128×128	32	(2, 3, 4)
LSUN	261828227	(1, 2, 3, 4, 5)	128	2e-05	128×128	32	(2, 3, 4)
MNIST	42082049	(1, 2, 2)	128	2e-04	28×28	128	(2,)

All experiments were conducted on NVIDIA A100 GPUs using PyTorch 2.6.0. The FFHQ model trained for 1M iterations ( $\approx 146$  hours), MNIST converged within 500k iterations ( $\approx 32$  hours), we additionally train a model for LSUN Church dataset for 500k iterations ( $\approx 80$  hours). We employed random horizontal flipping ( $p=0.5$ ) for FFHQ augmentation, with no augmentation applied to MNIST and LSUN Church.

#### B. Additional samples

We present supplementary experimental results from parameter ablations examining the influence of the Peclet number ( $Pe$ ) on image synthesis quality and sample distribution diversity. Experiments examine Peclet number ( $Pe$ ) impacts on FFHQ-128 ( $Pe \in \{0.0, 0.02, 0.04, 0.06, 0.08, 0.1, 0.12, 0.14\}$ ), MNIST ( $Pe \in \{0.0, 0.02, 0.04, 0.06, 0.08, 0.1\}$ ) and LSUN Church ( $Pe \in \{0.0, 0.02, 0.04, 0.06, 0.08\}$ ). This appendix documents comparative studies conducted on the FFHQ-128, MNIST and LSUN Church datasets, organized as follows: (i) single-initial-state sampling through generation, where ground truth (GT) images are propagated through the forward diffusion process with  $Pe$ -specific dynamics and reconstructed through model inference (Figs. 1, 7, 12); (ii) interpolation trajectories with Peclet ablations, illustrating transition dynamics under varying diffusion constraints (Figs. 2, 8, 13); (iii) multiple initial-state sampling via stochastic generation from diversified initial states, emphasizing  $Pe$ 's role in governing output variability across distinct trajectory initializations (Figs. 3, 4, 9, 14); and (iv) uncured interpolation demonstrating raw model behavior (Figs. 5, 6, 10, 11, 15, 16).

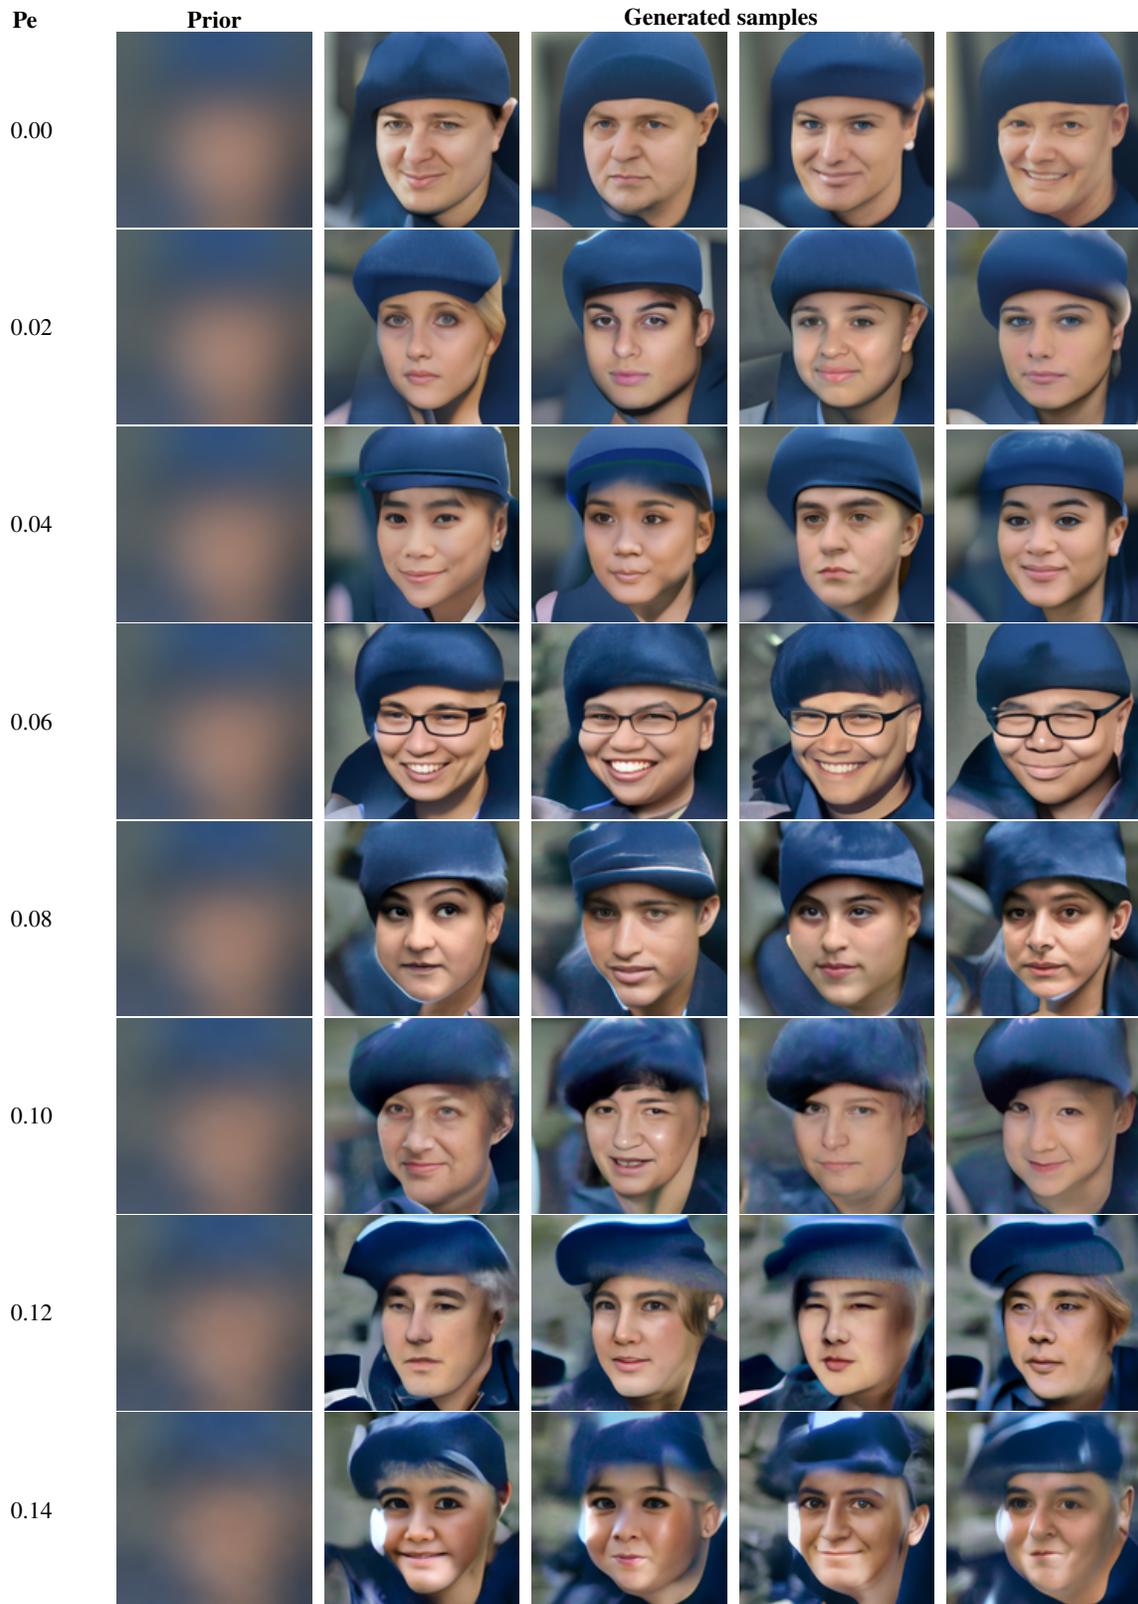


Figure 1. Results for  $\sigma = 20$ , showing inverse processes with varying  $P_e$  numbers. The image prior is consistent across rows for visual comparison, preserving the color palette.



Figure 2. Visual comparison of interpolations between two FFHQ samples. Each undergoes the forward process up to  $\sigma = 20$ , followed by linear interpolation and denoising with SLERP-interpolated generative noise added at each step.

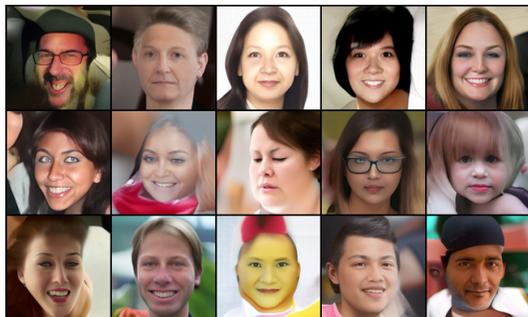


Figure 3. Visual comparison of the results of our method and the IHD method on the FFHQ dataset.

$\sigma = 16$

$\sigma = 20$

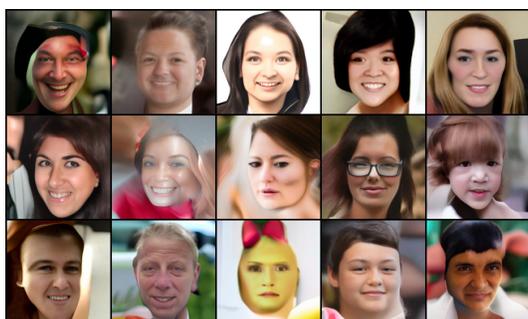
Pe=0.08



Pe=0.10



Pe=0.12



Pe=0.14

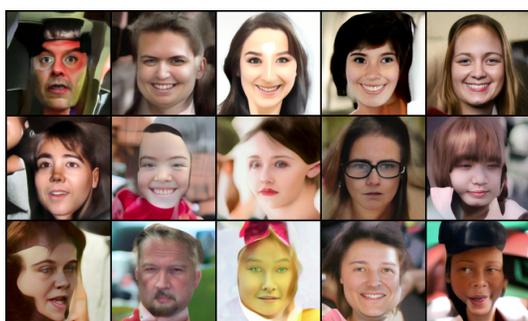


Figure 4. Visual comparison of the results of our method and the IHD method on the FFHQ dataset.



Figure 5. Interpolations between two random images on FFHQ  $128 \times 128$ .  $\sigma = 16$ ,  $Pe=0.8$



Figure 6. Interpolations between two random images on FFHQ  $128 \times 128$ .  $\sigma = 20$ ,  $Pe=0.8$

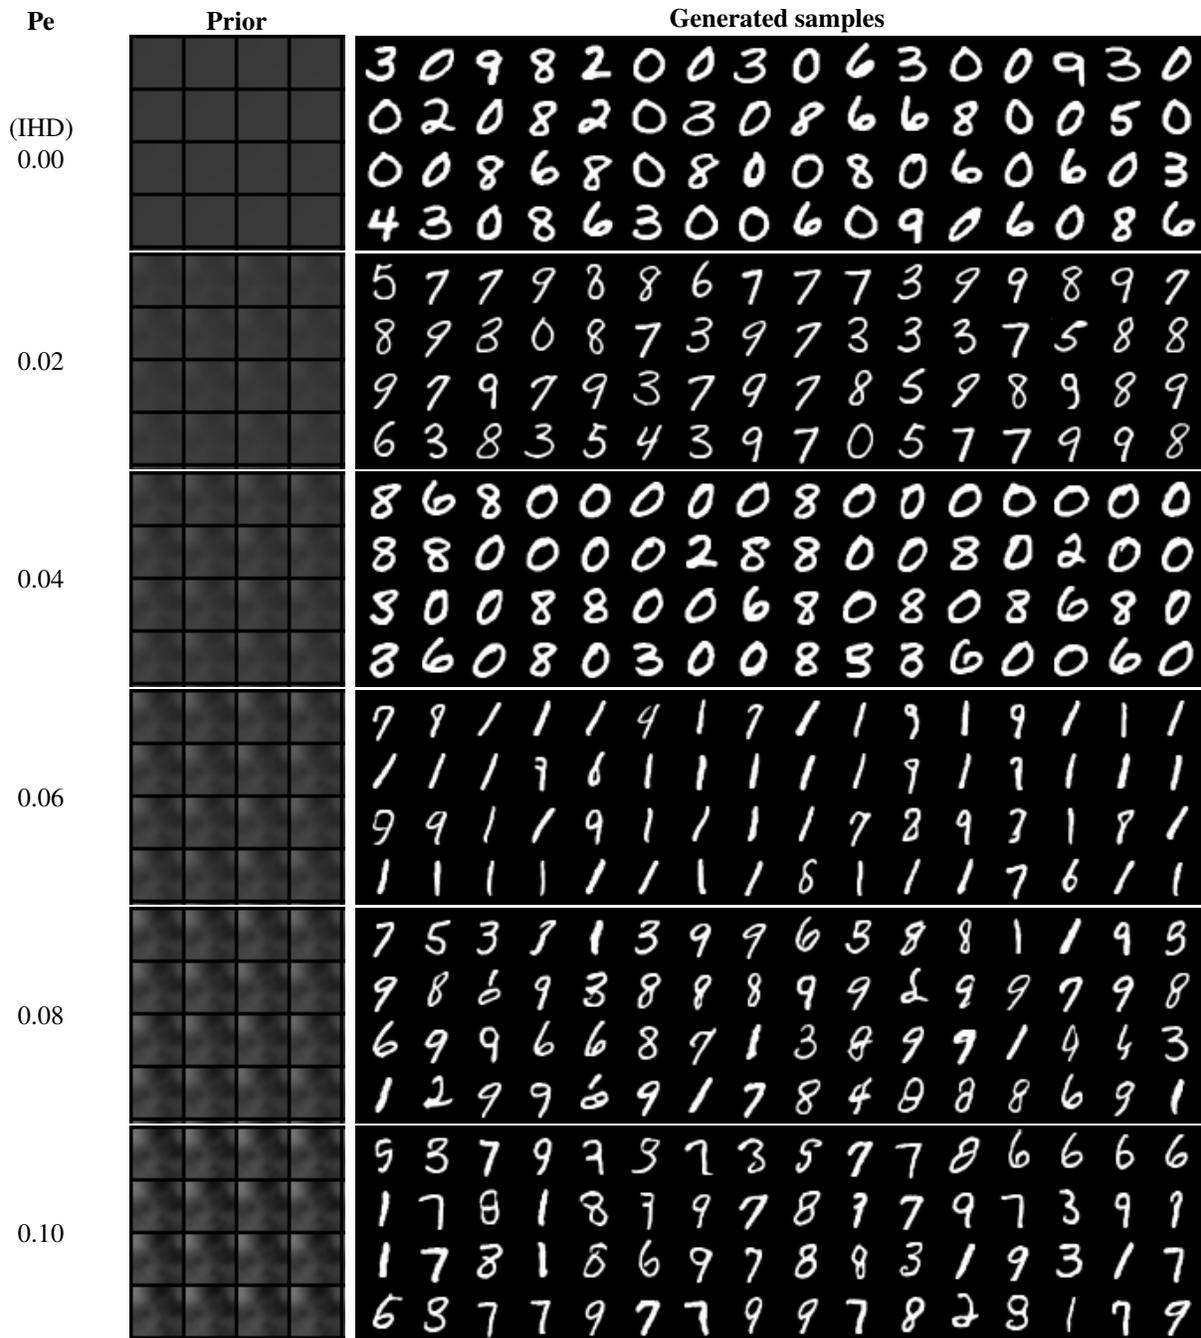


Figure 7. Additional samples with corresponding initial images from MNIST dataset, comparing different Peclet numbers. We can observe the impact of advection in forward process final step, that is the initial state for sampling.

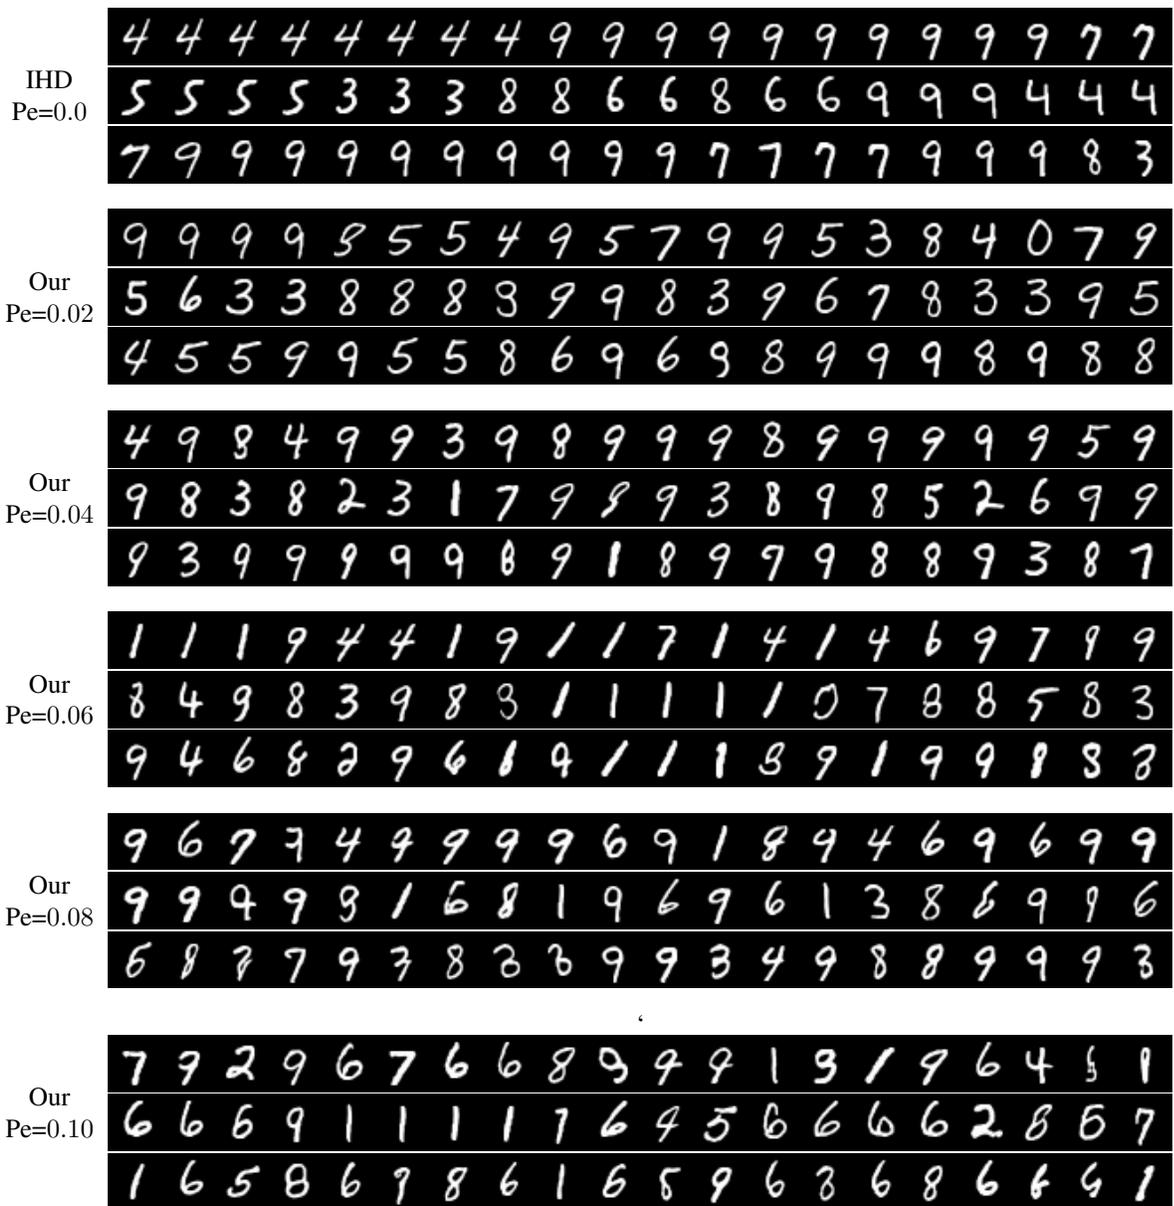


Figure 8. Visual comparison of interpolations between two MNIST samples. Each undergoes the forward process up to  $\sigma = 20$ , followed by linear interpolation and denoising with SLERP-interpolated generative noise added at each step.

5 9 4 1 0 1 9 4 4 7 0 1 7 8 5 8 2 5 1 9 8 1 0 4 9 2 9 7 0 8  
9 4 7 1 9 6 0 1 8 0 6 1 7 0 1 9 1 3 8 9 8 5 0 8 6 3 8 0 8 5  
3 1 8 3 3 7 9 1 9 8 4 8 8 4 9 6 9 5 1 7 1 1 9 8 0 4 0 1 1 4  
9 6 4 0 1 0 5 3 3 8 8 9 9 0 8 3 9 1 3 0 3 3 1 0 2 6 5 9 7 1

(a) (IHD)  $Pe = 0.00$

5 8 4 3 8 1 7 7 7 7 0 1 8 6 8 8 9 9 1 8 6 1 8 9 4 8 3 3 0 7  
8 9 6 1 1 7 8 1 3 8 8 1 9 8 1 7 1 7 7 7 9 3 0 8 9 8 7 3 9 8  
9 9 6 0 9 8 9 7 7 0 9 6 3 9 7 8 9 7 0 1 1 3 8 8 0 9 8 1 7 9  
1 0 7 0 1 0 8 0 6 7 7 5 1 6 4 8 7 1 2 9 3 3 1 3 2 3 9 7 7 1

(b)  $Pe = 0.02$

7 8 5 9 8 1 9 8 1 7 0 9 4 6 7 3 5 7 1 8 1 1 8 1 1 9 8 8 3 7  
9 1 8 1 9 7 8 1 2 6 8 1 9 8 3 3 1 7 8 8 0 6 6 8 9 8 7 9 8 1  
9 1 8 8 9 3 8 1 3 2 0 8 6 9 8 0 8 3 9 1 1 9 8 9 0 9 8 1 1 9  
6 3 9 0 1 8 5 5 9 9 3 6 1 0 1 3 3 1 8 6 8 0 9 0 3 7 8 9 9 1

(c)  $Pe = 0.04$

7 1 3 9 6 1 1 9 9 9 0 3 1 0 1 6 9 7 1 2 1 1 8 1 1 7 3 2 0 7  
1 6 6 1 4 1 8 1 2 2 9 1 9 0 1 1 1 9 7 6 1 6 6 0 1 3 3 2 4 1  
1 1 4 2 9 9 8 1 8 8 1 8 8 1 1 3 2 9 1 1 1 7 1 9 2 9 8 1 1 8  
1 1 1 3 1 8 9 0 3 1 9 4 1 0 1 9 1 1 9 3 9 4 7 0 9 4 6 9 9 1

(d)  $Pe = 0.06$

7 4 7 7 8 1 1 8 3 8 0 6 4 8 4 4 9 9 1 9 9 7 3 1 9 8 7 4 0 8  
6 1 3 1 9 4 9 7 9 8 1 7 4 6 1 8 7 3 8 9 6 9 0 9 7 2 8 0 8 9  
9 7 4 9 9 8 9 9 3 6 5 4 9 9 7 5 8 3 9 1 1 2 5 4 8 4 8 9 1 7  
3 8 3 0 9 3 3 8 3 0 9 9 1 9 1 7 9 9 8 8 9 8 9 0 2 9 9 9 9 1

(e)  $Pe = 0.08$

9 3 6 9 1 4 1 1 9 9 3 1 5 8 1 5 4 5 1 6 5 9 8 6 7 6 0 8 3 1  
6 5 6 1 6 9 3 8 5 6 6 6 3 2 6 8 9 6 6 1 6 6 8 8 9 8 3 2 1 9  
4 1 6 6 7 6 8 1 6 8 6 2 9 9 8 1 0 8 1 3 1 6 7 2 2 9 3 1 6 6  
8 0 6 0 1 6 6 3 6 1 8 6 1 2 9 7 1 1 6 6 8 3 1 0 6 9 6 1 7 4

(f)  $Pe = 0.10$

Figure 9. Visual comparison of the results of our method and the IHD method on the MNIST dataset.

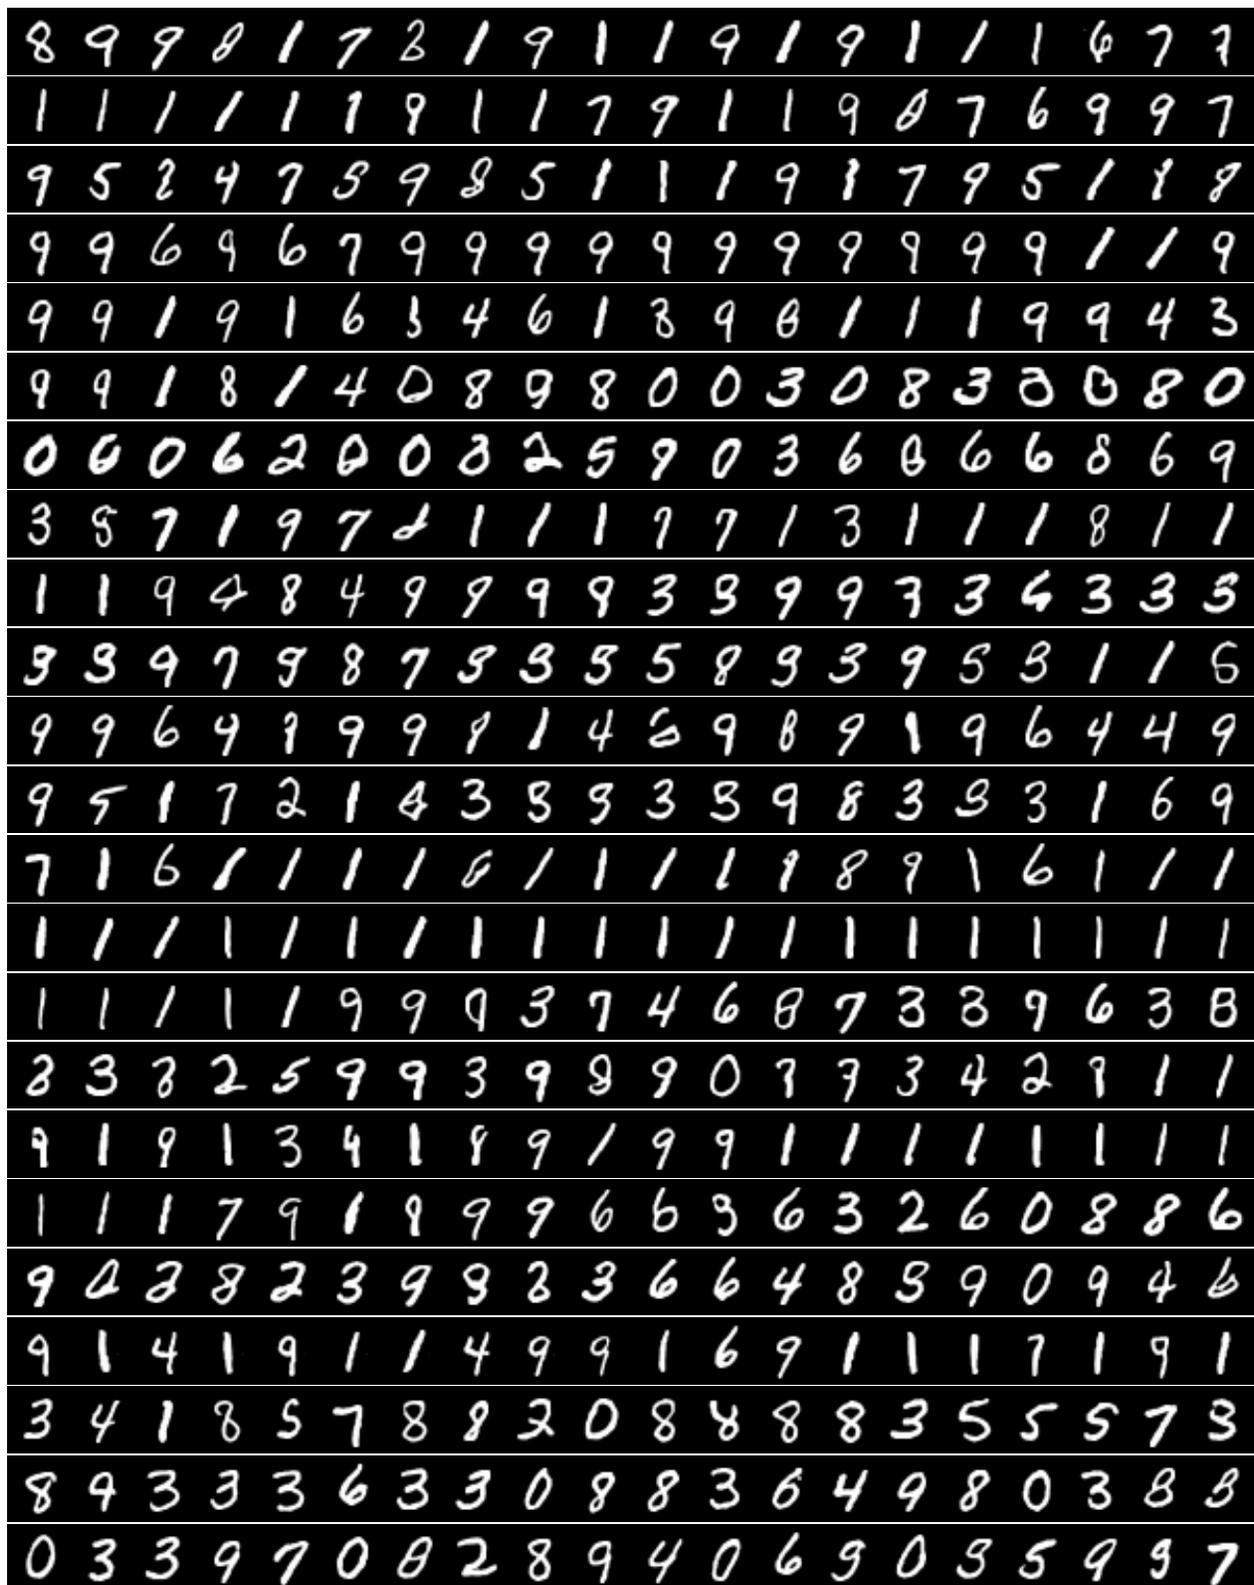


Figure 10. Interpolations between two random images on MNIST,  $\sigma = 20$ ,  $Pe = 0.6$

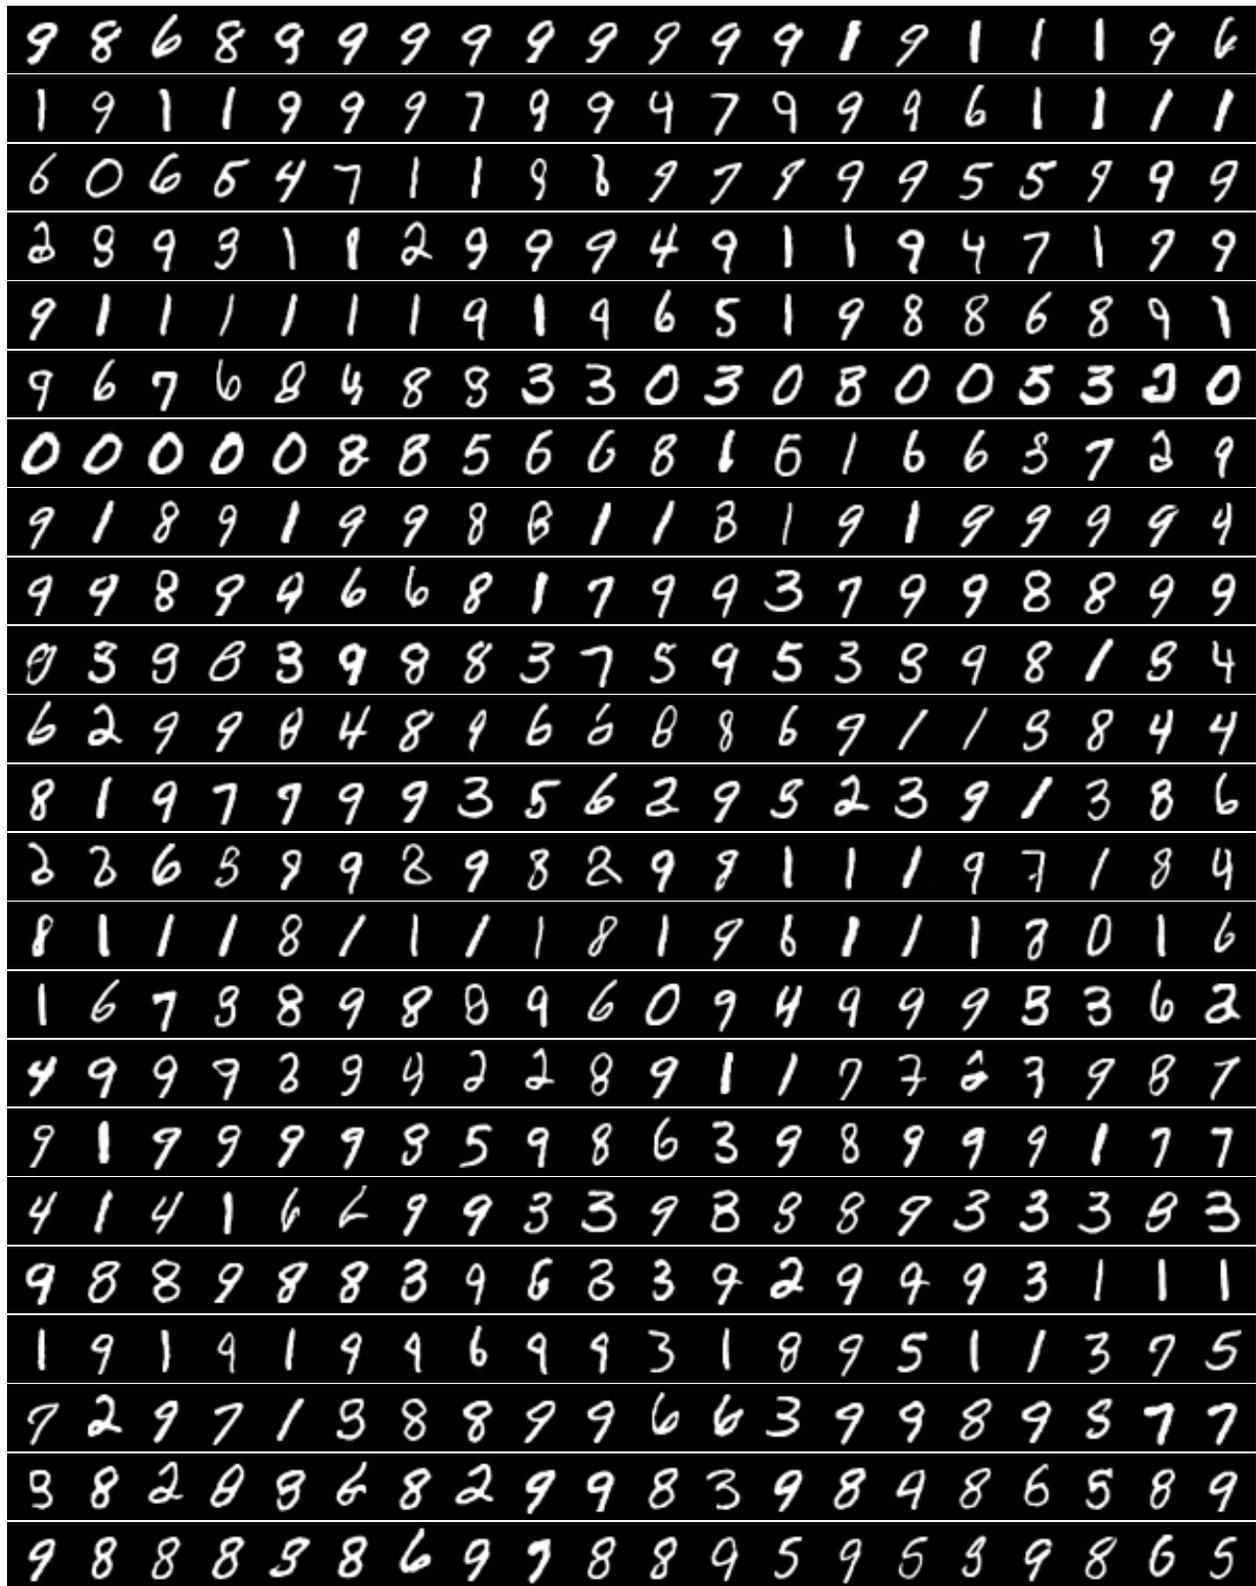


Figure 11. Interpolations between two random images on MNIST,  $\sigma = 20$ ,  $Pe = 0.8$

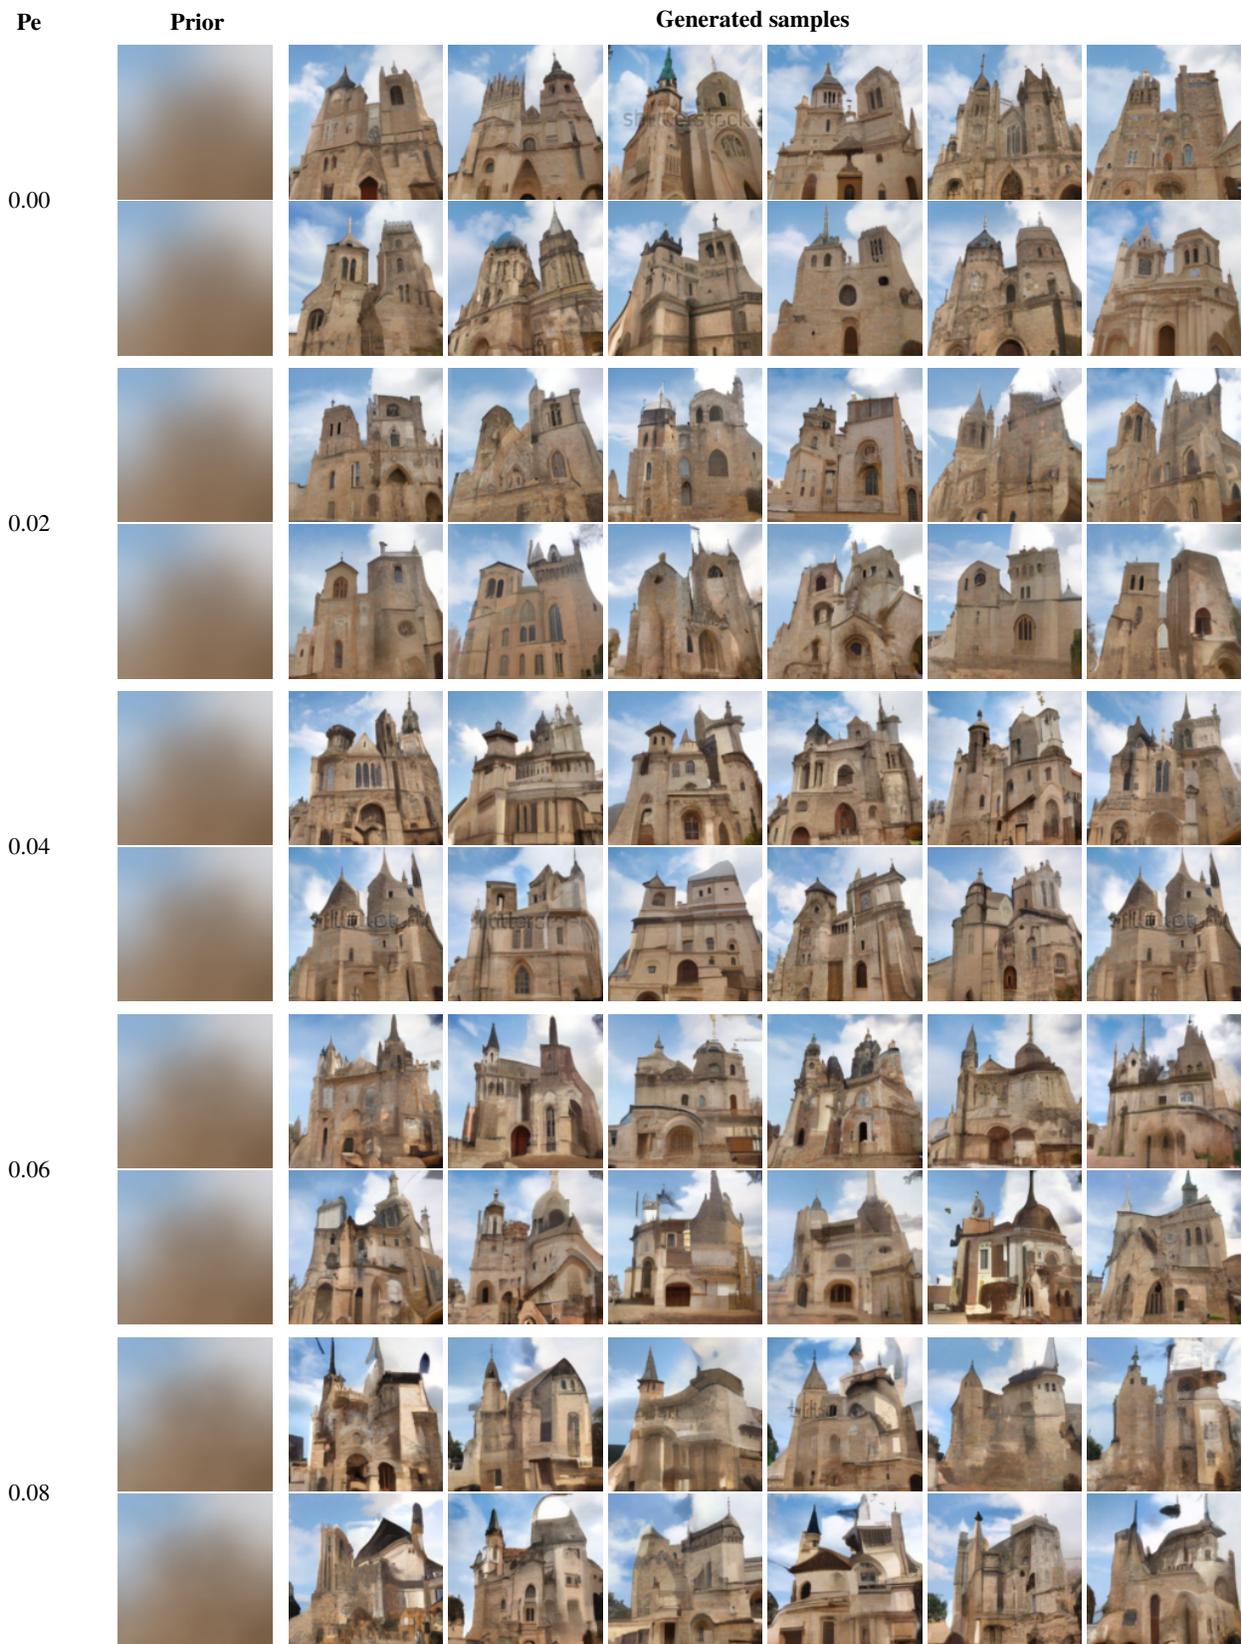


Figure 12. Results for  $\sigma = 20$ , on the LSUN Church dataset, showing inverse processes with varying Pe numbers. The image prior is consistent across rows for visual comparison, preserving the color palette.



Figure 13. Visual comparison of interpolations between two LSUN Church samples. Each undergoes the forward process up to  $\sigma = 20$ , followed by linear interpolation and denoising with SLERP-interpolated generative noise added at each step.

$\sigma = 20$

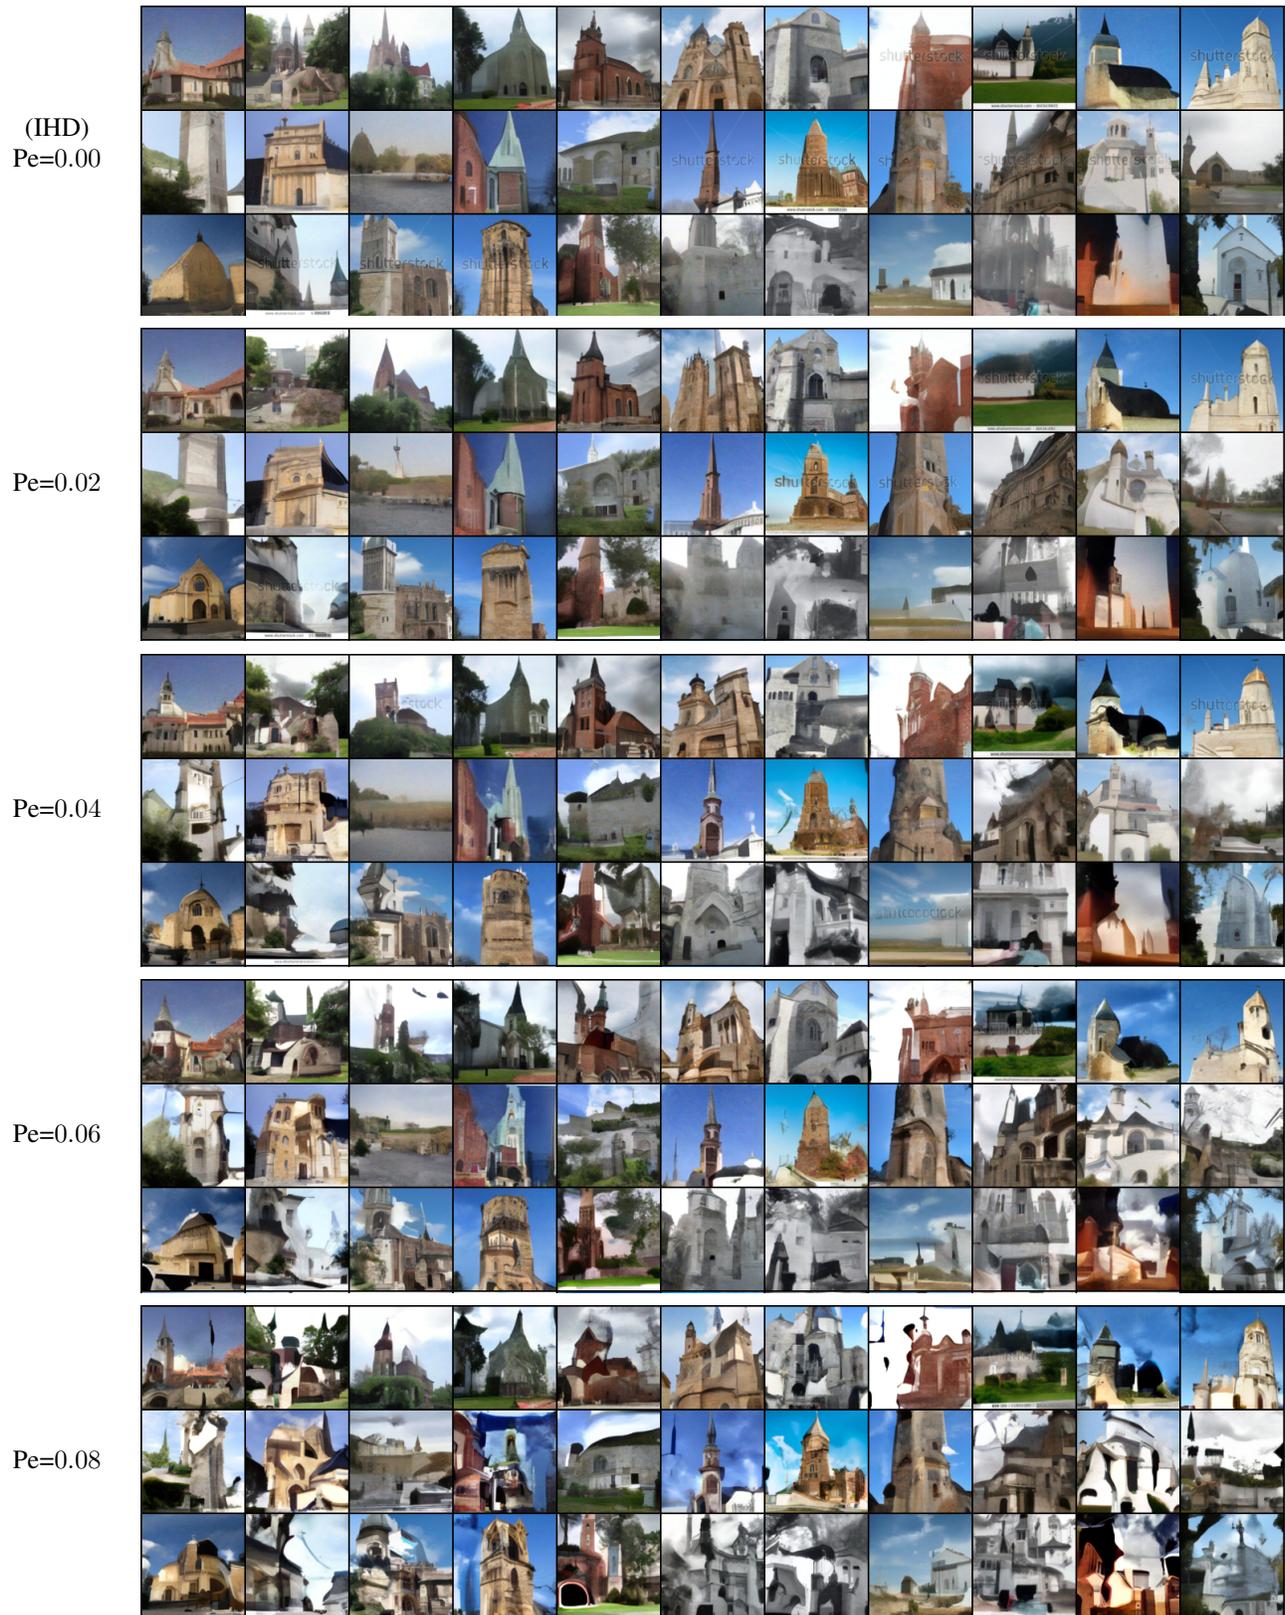


Figure 14. Visual comparison of the results of our method and the IHD method on the LSUN Church dataset.

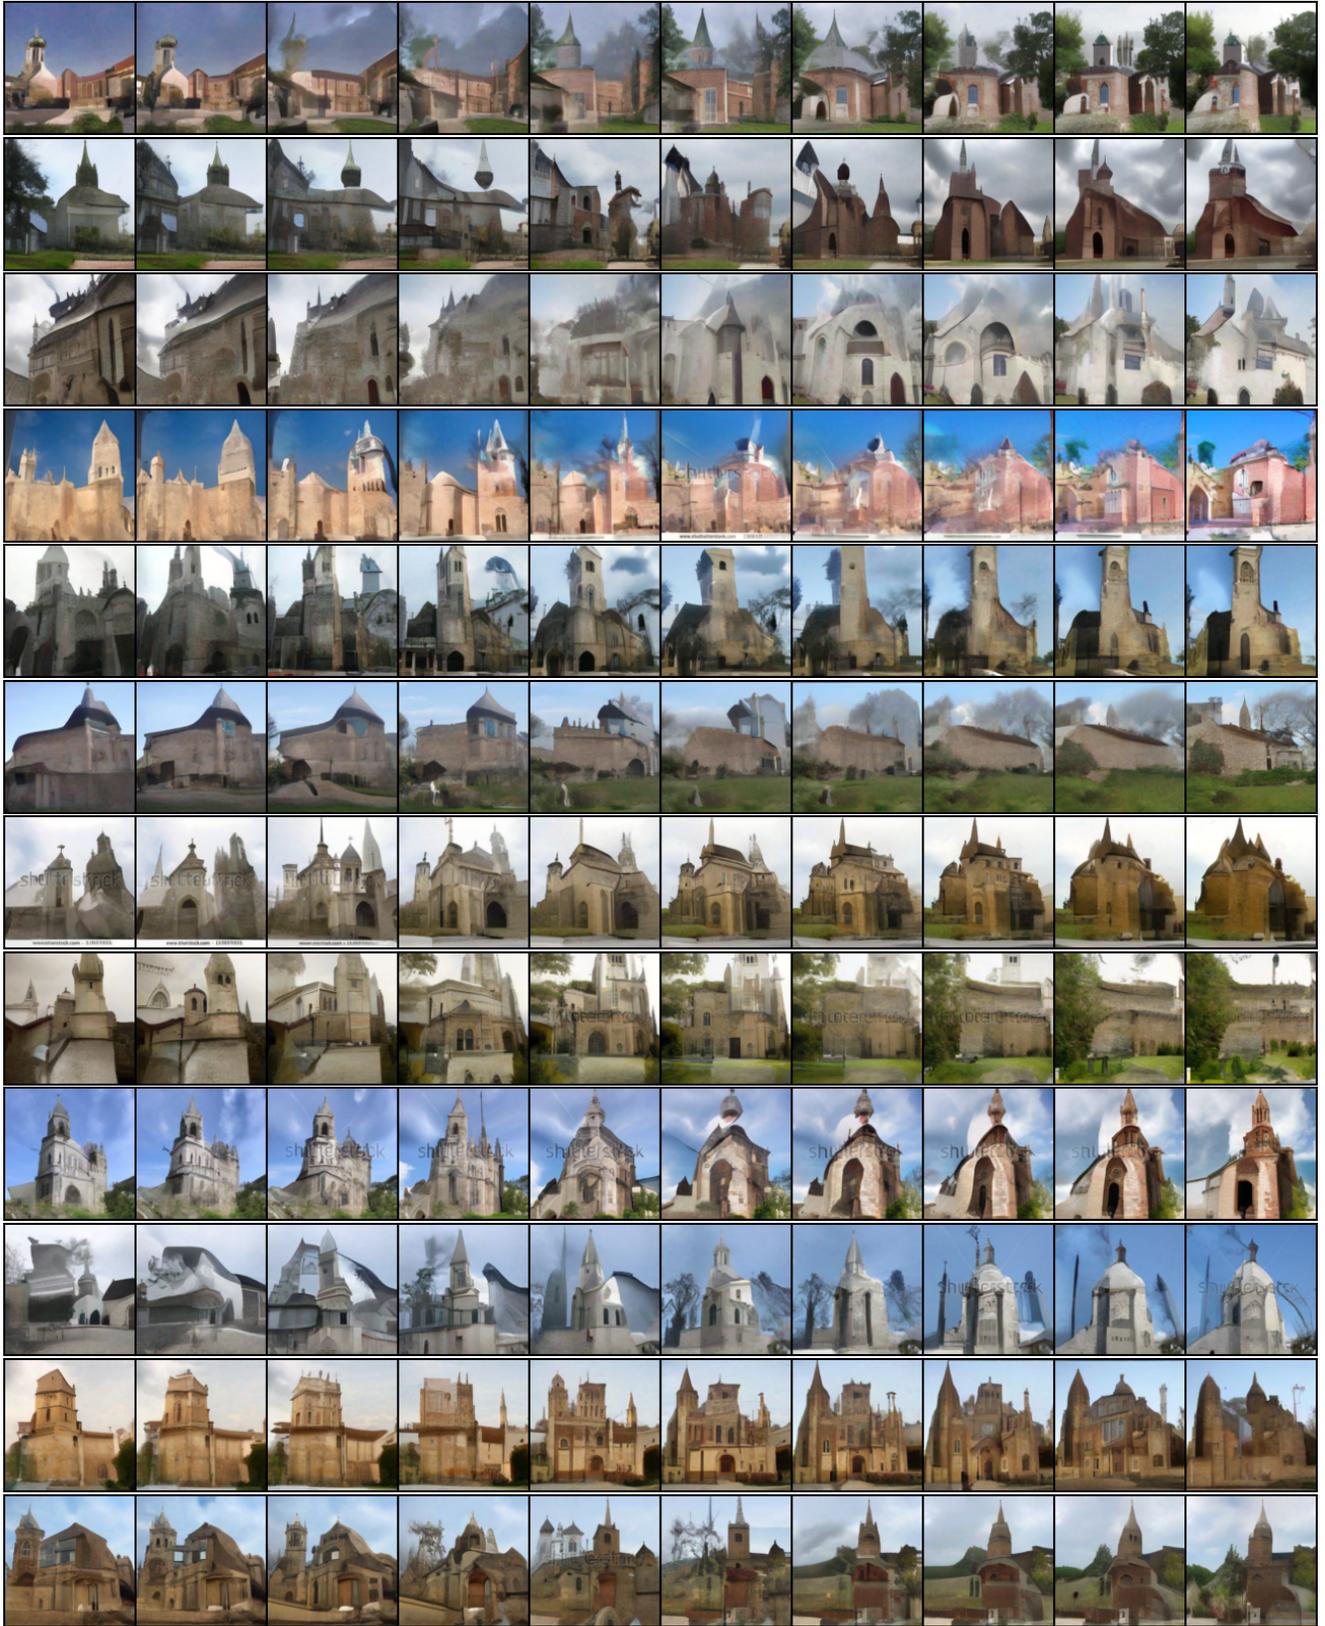


Figure 15. Interpolations between two random images on LSUN Church  $128 \times 128$ .  $\sigma = 20$ ,  $Pe=0.4$



Figure 16. Interpolations between two random images on LSUN Church  $128 \times 128$ .  $\sigma = 20$ ,  $Pe=0.6$

## C. Pseudocode

We insert a pseudocode in Python to illustrate how the Spectral Turbulence Generator and Lattice Boltzmann Method (LBM) work internally.

```
1 import torch as t
2
3 def tanh_limiter(x, min_val, max_val, sharpness=1.0):
4     mid_val = (max_val + min_val) / 2
5     range_val = (max_val - min_val) / 2
6     return mid_val + range_val * t.tanh(sharpness * (x - mid_val) / range_val)
7
8 def limit_velocity_field(u, v, min_val, max_val):
9     velocity_magnitude = t.sqrt(u**2 + v**2)
10    limited_magnitude = tanh_limiter(velocity_magnitude, min_val, max_val)
11
12    small_factor = 1E-9
13    direction_factor = t.where(velocity_magnitude < small_factor, small_factor, limited_magnitude /
14                               velocity_magnitude)
15
16    # Adjust u and v components to match the new limited magnitude while preserving direction
17    upscale = 1.
18    direction_factor *= upscale
19    u_limited = u * direction_factor
20    v_limited = v * direction_factor
21
22    return u_limited, v_limited
23
24 class SpectralTurbulenceGenerator(t.nn.Module):
25     def __init__(self, std_dev_schedule):
26         grid_size = (128, 128)
27         self.std_dev = std_dev_schedule #turbulence intensity scaling factor
28
29         energy_slope = -2.0
30         self.energy_spectrum = lambda k: t.where(t.isinf(k ** (energy_slope)), 0, k ** (energy_slope))
31         self.frequency_range = {'k_min': 2.0 * t.pi / min(grid_size), 'k_max': 2.0 * t.pi /
32                                 (min(self.domain_size) / 1024)}
33
34         # Fourier transform wave numbers
35         kx = (t.fft.fftfreq(grid_size[0], d=1/grid_size[0]) * 2 * t.pi).to('gpu')
36         ky = (t.fft.fftfreq(grid_size[1], d=1/grid_size[1]) * 2 * t.pi).to('gpu')
37         KX, KY = t.meshgrid(kx, ky)
38         self.K = t.sqrt(KX**2 + KY**2).to('gpu')
39
40         # Initialize the phases once and use them in each run
41         self.phase_u = (t.rand(grid_size) * 2 * t.pi).to('gpu')
42         self.phase_v = (t.rand(grid_size) * 2 * t.pi).to('gpu')
43
44         self.amplitude = (t.where(self.K != 0, (self.energy_spectrum(self.K)), 0)).to('gpu')
45         self.amplitude = (t.where((self.K >= self.frequency_range['k_min']) & (self.K <=
46                                 self.frequency_range['k_max']), self.amplitude, 0.0)).to("gpu")
47
48         dt_turb = 1E-4
49         self.omega = dt_turb*self.K
50
51     def generate_turbulence(self, time: int):
52         u_hat = self.amplitude * t.exp(1j * (self.phase_u + self.omega * time))
53         v_hat = self.amplitude * t.exp(1j * (self.phase_v + self.omega * time))
54         u = t.real(t.fft.ifft2(u_hat))
55         v = t.real(t.fft.ifft2(v_hat))
56
57         if self.std_dev[time] < 1E-14:
58             u, v = 0*self.K, 0*self.K #avoid division by 0 in t.std(u)
59         else:
60             u *= self.std_dev[time] / t.std(u)
61             v *= self.std_dev[time] / t.std(v)
```

```

60     u, v = limit_velocity_field(u, v, min_val=-1E-3, max_val=1E-3)
61
62     return u.float(), v.float()

```

```

1  import taichi as ti
2  import taichi.math as tm
3
4  # Fluid solver based on lattice boltzmann method using taichi language
5  # Inspired by: https://github.com/hietwll/LBM-Taichi
6
7  @ti.data_oriented
8  class LBM_ADE_Solver():
9      def __init__(self, config, turbulenceGenerator):
10         self.nx, self.ny = config.domain_size
11         self.turbulenceGenerator = turbulenceGenerator
12
13         self.cs2 = ti.field(ti.f32)(1./3.)
14         self.omega_kin = ti.field(ti.f32, shape=self.max_iter[None])
15         self.omega_kin.from_numpy(1.0 / (3.0* config.kin_visc + 0.5))
16
17         self.rho = ti.field(float, shape=(self.nx, self.ny))
18         self.vel = ti.Vector.field(2, float, shape=(self.nx, self.ny))
19
20         self.f = ti.Vector.field(9, float, shape=(self.nx, self.ny))
21         self.f_new = ti.Vector.field(9, float, shape=(self.nx, self.ny))
22
23         self.Force = ti.Vector.field(2, float, shape=(self.nx, self.ny))
24         self.w = ti.types.vector(9, float)(4, 1, 1, 1, 1, 1 / 4, 1 / 4, 1 / 4, 1 / 4) / 9.0
25         self.e = ti.types.matrix(9, 2, int)(
26             [0, 0], [1, 0], [0, 1], [-1, 0], [0, -1], [1, 1], [-1, 1], [-1, -1], [1, -1])
27
28     def init(self, np_image):
29         self.rho.from_numpy(np_image)
30         self.vel.fill(0)
31
32     def solve(self, iterations):
33         for iteration in range(iterations):
34             self.stream()
35             self.update_macro_var()
36             self.collide_srt()
37             self.vel = self.turbulenceGenerator.generate_turbulence(iteration)
38             self.apply_bounceback_boundary_condition()
39
40     @ti.kernel
41     def stream(self):
42         for i, j in ti.ndrange(self.nx, self.ny):
43             for k in ti.static(range(9)):
44                 ip = i - self.e[k, 0]
45                 jp = j - self.e[k, 1]
46                 self.f[i, j][k] = self.f_new[ip, jp][k]
47
48     @ti.kernel
49     def update_macro_var(self):
50         for i, j in ti.ndrange((1, self.nx-1), (1, self.ny-1)):
51             self.rho[i, j] = 0
52             for k in ti.static(range(9)):
53                 self.rho[i, j] += self.f[i, j][k]
54
55     @ti.kernel
56     def collide_srt(self):
57         omega_kin = self.omega_kin[self.iterations_counter[None]]
58         for i, j in ti.ndrange((1, self.nx - 1), (1, self.ny - 1)):
59             for k in ti.static(range(9)):
60                 feq = self.f_eq(i, j)
61                 self.f_new[i, j][k] = (1. - omega_kin) * self.f[i, j][k] + feq[k] * omega_kin
62
63     @ti.func

```

```
64 def f_eq(self, i, j):
65     eu = self.e @ self.vel[i, j]
66     uv = tm.dot(self.vel[i, j], self.vel[i, j])
67     return self.w * self.rho[i, j] * (1 + 3 * eu + 4.5 * eu * eu - 1.5 * uv)
68
69 @ti.func
70 def apply_bounceback_core(self, i: int, j: int):
71     tmp = ti.f32(0.0)
72     for k in ti.static([1,2,5,6]):
73         tmp = self.f[i, j][k]
74         self.f[i, j][k] = self.f[i, j][k+2]
75         self.f[i, j][k+2] = tmp
76
77     for k in ti.static(range(9)):
78         self.f_new[i, j][k] = self.f[i, j][k]
79
80 @ti.kernel
81 def apply_bounceback_boundary_condition(self):
82     for i in range(0, self.nx):
83         self.apply_bounceback_core(i, 0)
84         self.apply_bounceback_core(i, self.ny-1)
85
86     for j in range(1, self.ny-1):
87         self.apply_bounceback_core(0, j)
88         self.apply_bounceback_core(self.nx-1, j)
```