Everything at Every Scale: Scale-Invariant Diffusion with Continuous Super-Resolution

Zixin Jessie Chen[13*†] **Zhuo Chen**[134*] **Archer Wang**[23] **Jeff Gore**[1] **William T. Freeman**[23]
Congyue Deng[2] **Marin Soljačić**[13]

[1]Department of Physics, Massachusetts Institute of Technology
[2]Department of EECS, Massachusetts Institute of Technology
[3]NSF AI Institute for Artificial Intelligence and Fundamental Interactions
[4]Institute for Data, Systems and Society, Massachusetts Institute of Technology
{jzxchen,chenzhuo,archerdw,congyued,billf,soljacic}@mit.edu

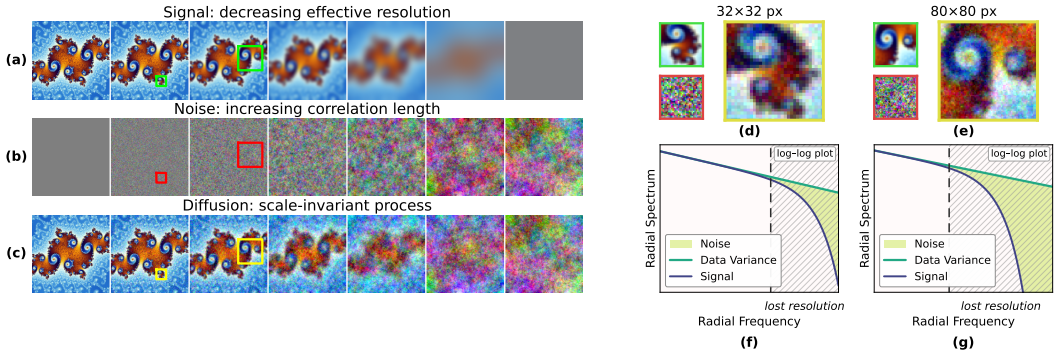

Figure 1: Conceptual illustration of our SKILD on a self-similar fractal image. During the forward process: (a) effective signal resolution decreases; (b) the pixel-space correlation length of the injected noise increases; and (c) for a self-similar field, the process respects the same frequency-space power spectrum across stages. (d-e) A smaller early-time patch is statistically similar to a larger late-time patch. (f-g) Radial power spectrum at corresponding early and later stages. Gray slashed regions indicate modes below the signal-to-noise ratio (SNR) threshold, where resolution is effectively lost.

Abstract

Creating images from noise is image generation; reconstructing fine details from coarse inputs is super-resolution. Despite their practical differences, both can be understood as reversing information loss across scales. We introduce **SKILD**, a **S**cale-invariant **K**-Space **I**mage **L**earning **D**iffusion model that unifies generation and continuous super-resolution within a single unconditional framework. Both natural images and critical physical systems exhibit scale invariance, and we leverage it to design a forward process that attenuates image content from fine to coarse scales while injecting spectrum-matched Gaussian noise, making scale an explicit coordinate of the diffusion dynamics. The same trained reverse process performs generation and continuous super-resolution by varying only the starting timestep: *no task-specific architecture, no conditioning branch, no classifier-free guidance, no retraining per scale factor*. Empirically, SKILD reaches FID 2.61 and Inception Score 9.65 on unconditional CIFAR-10, performs $2\times$–$8\times$ super-resolution on ImageNet from a single unconditional checkpoint while outperforming conditional models across perceptual metrics, and reconstructs critical Ising models whose connected four-point correlations closely track the ground truth.

*Equal contribution
†Corresponding author

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

Scales in images have long been a subject of study in computer vision. Across different scales, images share recurring structure. A zoomed natural image still looks natural, and some natural objects are themselves self-similar, with textures, edges, and structures recurring at different scales. Statistically, this regularity is reflected in natural-image power spectra, which follow approximate power laws over wide frequency ranges [1–4], a signature of approximate scale invariance. The same concept has been studied in parallel in physics, where critical systems display similar scale-invariant behavior, made formal by the renormalization group [5, 6]. This physics perspective also points to a natural way of organizing the transformation between an image and pure noise. Can we take advantage of this scale invariance in diffusion? Rather than corrupting all scales at once, one can erase them in order, one scale at a time. Diffusion, framed in this way, becomes a denoising process respecting self-similarity across scales.

Such a denoising process across scales is, by construction, a progressive super-resolution. At each backward step, finer scales are added back, and running the full reverse process from pure noise produces an image scale by scale. This unifies generation and super-resolution into a single framework. Generation from noise is the extreme case of super-resolution in which the input contains no signal at all; super-resolution is the same reverse process initialized from an intermediate state in which coarser scales have survived. Both are reverse coarse-graining problems, distinguished only by where the reverse process begins.

We realize this idea with **SKILD** (**S**cale-invariant **K**-Space **I**mage **L**earning **D**iffusion), a diffusion model whose forward process corrupts images one scale at a time, from finest to coarsest. Two design choices make this concrete. First, the forward process attenuates high-frequency content before low-frequency content. Second, the noise added at each step carries the spectrum of the dataset itself rather than being white noise, so the model learns to remove noise that statistically resembles the data it learns to generate. Together, these two choices make every intermediate state a coarse-grained, noisy version of the original image in a self-similar manner.

Our contributions are as follows.

- We propose SKILD, a scale-invariant diffusion framework that unifies unconditional generation and continuous super-resolution within a single reverse process. A single, unconditional architecture handles both tasks, replacing what would otherwise be a stack of task-specific architectures, conditioning branches, classifier-free guidance, and per-scale retraining.
- On unconditional CIFAR-10 [7], SKILD is competitive with state-of-the-art diffusion models and achieves the strongest sample quality among frequency-informed diffusion models.
- One trained SKILD checkpoint performs continuous super-resolution at any factor, which we test on ImageNet [8] between $2\times$ and $8\times$. At $4\times$ super-resolution on ImageNet-256, the same model outperforms strong diffusion-based conditional super-resolution baselines on multiple perceptual quality metrics.
- Evaluations on a scientific dataset generated using a critical Ising model show that SKILD reproduces explicit self-similar statistics while a strong diffusion-based conditional super-resolution baseline fails.

2 Related Works

Scale invariance and self-similarity. Scale-space theory analyzes images through continuous smoothing and identifies Gaussian convolution as the canonical linear scale-space operator [9–14]. Natural-image statistics show approximate power-law spectra across scales [1–4, 15], while renormalization group theory describes how distributions transform under coarse-graining and rescaling [5]. These ideas motivate our forward process: attenuation from fine to coarse scales in frequency space, with noise covariance matching the dataset distribution.

Diffusion models across scale and frequency. Diffusion models learn to reverse a noising process [16–18], with various samplers and schedules [19–22]. Several lines of work connect diffusion to multi-scale structure: cascaded and relay models compose resolution-specific conditional stages [23, 24]; other work connects diffusion to renormalization-group flows, optimal transport, or inverse heat dissipation [25–28]. A separate line uses Fourier or wavelet structures to improve controllability,

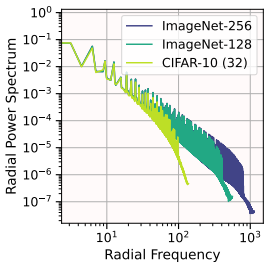

(a) Natural-image variances

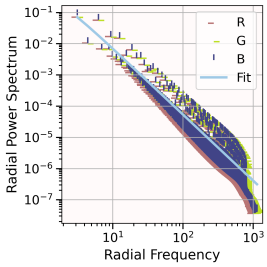

(b) ImageNet-256 variances

Figure 2: Variance power spectra of natural-image datasets. (a) Spectra of CIFAR-10, ImageNet-128, and ImageNet-256 exhibit similar power-law decay over their shared frequency range, indicating approximate scale invariance. (b) Variance of ImageNet-256 computed independently for each color channel (RGB), with a power-law fit recovering the k^{-2} frequency decay.

efficiency, or inductive bias [29–36]. Recent works have also explored image generation as progressive super-resolution in pixel space, replacing additive noise with structured degradations or multi-scale reconstruction processes [37, 38]. Unlike these approaches, SKILD explicitly utilizes self-similarity in frequency space, where the forward process continuously attenuates image statistics from fine to coarse modes. As a result, a single reverse process supports both unconditional generation and continuous super-resolution without conditioning or guidance.

Super-resolution. Beyond classical priors and feed-forward neural methods [39–45], diffusion-based methods often rely on additional conditioning from low-resolution images [46–52]. SKILD requires no extra conditioning: the low-resolution input is an intermediate state of the model's own forward process, and the same reverse process completes the missing fine scales.

3 Preliminaries and Motivations

Standard diffusion. Diffusion models [17, 18] generate samples by reversing a fixed forward noising process. The forward process gradually transforms a data sample $\mathbf{x}_0$ into isotropic Gaussian noise via

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\,\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\,\epsilon, \qquad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \tag{1}$$

where the schedule $\bar{\alpha}_t$ decreases monotonically from 1 at $t = 0$ to nearly 0 at the end of diffusion. Because the marginals are jointly Gaussian, the reverse-time conditional $q(\mathbf{x}_{t-1} \mid \mathbf{x}_t, \mathbf{x}_0)$ is itself Gaussian and analytically tractable. A neural network $\epsilon_\theta(\mathbf{x}_t, t)$ is trained to predict ϵ given $\mathbf{x}_t$ by minimizing $\mathbb{E}\big[\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2\big]$, and substituting this prediction into the reverse posterior gives a tractable sampling step. Iterative denoising starting from pure noise produces samples from the data distribution.

Scale invariance in physics. Critical physical systems, exemplified by the two-dimensional Ising model at its critical temperature, exhibit scale invariance explicitly. Such systems have no characteristic length scale, so configurations look statistically the same after coordinates are coarse-grained and rescaled by any factor $\mathbf{r} \to b\mathbf{r}$. As a consequence, statistical observables follow power laws of the form $O(k) \propto k^{-\alpha}$ with universal exponents α [53, 5, 6], since power laws are the only functions invariant under rescaling up to a multiplicative constant. The renormalization group formalizes this picture. Coarse-graining out fine-scale degrees of freedom and rescaling the result acts as a transformation on probability distributions, and the distribution of a critical system is a fixed point of that transformation.

Power-law spectra of natural images. Natural-image distributions show approximate scale invariance. Their radially averaged power spectra, equivalently the variance per Fourier mode of the dataset $\mathbf{S}_0(\mathbf{k}) = \mathbb{E}\big[|\mathbf{X}_0(\mathbf{k})|^2\big] - \big|\mathbb{E}\big[\mathbf{X}_0(\mathbf{k})\big]\big|^2$, closely follow k^{-2} over a wide frequency range [1–3, 54], on average across a dataset. We confirm this on the datasets used in our experiments. Figure 2 shows the radially averaged variance power spectra for CIFAR-10, ImageNet-128, and ImageNet-256 computed in the discrete cosine transform (DCT) space [55], with the exact transform given in Appendix C. The spectra agree over their shared frequency range and differ mainly near finite-resolution cutoffs. We fit the radial variance with

$$\mathbf{S}_0(\mathbf{k}) = C(\mathbf{k}^2 + \mathbf{k}_0^2)^{-a}, \tag{2}$$

where $\mathbf{k}_0$ regularizes the $\mathbf{k} \to \mathbf{0}$ limit. The fits recover the k^{-2} scaling, and Table C.5 lists the fitted parameters.

Toward scale-invariant diffusion. The observations above raise a natural design question. Given that natural images and critical physical systems share a hierarchy of structure across scales, what would a diffusion forward process look like if it were designed to respect this hierarchy rather than treating all scales on the same footing? We propose such a scale-invariant process in the frequency space in the next section.

4 Scale-Invariant Diffusion in Frequency Space

4.1 Formulation

Forward process. Let $\mathbf{X}_0(\mathbf{k})$ denote the DCT coefficients of an image and let $\mathbf{S}_0(\mathbf{k})$ be the empirical variance spectrum estimated in Section 3. We define the continuous forward marginal

$$\mathbf{X}(\mathbf{k}, t) = \underbrace{e^{-\mathbf{k}^2\lambda(t)/2} \odot \mathbf{X}_0(\mathbf{k})}_{\text{signal}} + \underbrace{\sqrt{1 - e^{-\mathbf{k}^2\lambda(t)}} \odot \sqrt{\mathbf{S}_0(\mathbf{k})}\, \epsilon_t}_{\text{noise}}, \qquad \epsilon_t \sim \mathcal{N}(0, \mathbf{I}), \tag{3}$$

where $\odot$ denotes Hadamard product. The schedule $\lambda(t)$ is a scalar function that monotonically increases in t. As t grows, $e^{-\mathbf{k}^2\lambda(t)/2}$ narrows in frequency space, so high-frequency modes are attenuated before low-frequency ones. The noise prefactor is chosen so that the forward marginal preserves the per-mode covariance $\mathbf{S}_0(\mathbf{k})$ in expectation at every t, and converges to $\mathcal{N}(0, \mathbf{S}_0(\mathbf{k}))$ as the signal term vanishes. In pixel space, the same process convolves the signal with a Gaussian kernel and adds spatially correlated noise whose correlation length grows with t, reflecting the progressive removal of scale structures.

Discretization. All experiments use a DDPM [17] discretization of Eq. (3). For $0 = t_0 < \cdots < t_N = 1$, let $\bar{\boldsymbol{\alpha}}_n(\mathbf{k}) = e^{-\mathbf{k}^2\lambda_n}$, $\boldsymbol{\alpha}_n = \bar{\boldsymbol{\alpha}}_n / \bar{\boldsymbol{\alpha}}_{n-1}$, and $\boldsymbol{\beta}_n = 1 - \boldsymbol{\alpha}_n$. Then

$$\mathbf{X}_n(\mathbf{k}) = \sqrt{\bar{\boldsymbol{\alpha}}_n} \odot \mathbf{X}_0(\mathbf{k}) + \sqrt{1 - \bar{\boldsymbol{\alpha}}_n} \odot \sqrt{\mathbf{S}_0(\mathbf{k})}\, \epsilon_n, \qquad \epsilon_n \sim \mathcal{N}(0, \mathbf{I}). \tag{4}$$

The one-step transition has the same form with $\bar{\boldsymbol{\alpha}}_n$ replaced by $\boldsymbol{\alpha}_n$.

Since all covariances are diagonal in frequency space, the reverse posterior is Gaussian: $q(\mathbf{X}_{n-1} \mid \mathbf{X}_n, \mathbf{X}_0) = \mathcal{N}(\boldsymbol{\mu}_q, \mathbf{S}_0\tilde{\boldsymbol{\beta}}_n)$, with

$$\tilde{\boldsymbol{\beta}}_n = \frac{\boldsymbol{\beta}_n(1 - \bar{\boldsymbol{\alpha}}_{n-1})}{1 - \bar{\boldsymbol{\alpha}}_n}, \qquad \boldsymbol{\mu}_q(n, \mathbf{X}_n) = \frac{1}{\sqrt{\boldsymbol{\alpha}_n}} \left(\mathbf{X}_n - \frac{\boldsymbol{\beta}_n \sqrt{\mathbf{S}_0}}{\sqrt{1 - \bar{\boldsymbol{\alpha}}_n}}\, \epsilon_n \right). \tag{5}$$

Ancestral sampling proceeds by

$$\mathbf{X}_{n-1} = \boldsymbol{\mu}_q(n, \mathbf{X}_n) + \sqrt{\mathbf{S}_0\tilde{\boldsymbol{\beta}}_n} \odot \epsilon_n. \tag{6}$$

The full DDPM and stochastic differential equation (SDE) derivation appear in Appendices A and B.

4.2 Training target and numerical cutoffs

We train an ϵ-prediction network with the loss

$$\mathcal{L} = \mathbb{E}_{n, \mathbf{X}_0, \epsilon}\left[\|\epsilon - \epsilon_\theta(n, \mathbf{X}_n)\|_2^2 \right]. \tag{7}$$

Two implementation details make the finite-resolution process stable. First, very small $\boldsymbol{\alpha}_n(\mathbf{k})$ values can cause large reverse updates for high-frequency modes, so we floor them at 10^{-6} in the ancestral sampler. Second, the zero mode would otherwise have no attenuation or noise because the Gaussian signal filter leaves it untouched. Therefore, we introduce a low-frequency cutoff k_c, and use $\max(\|\mathbf{k}\|, k_c)$ as the schedule for modes with $\|\mathbf{k}\| \le k_c$. This preserves the algebra above while properly handles the low-frequency limit, where scale-invariance is affected by finite size effect.

4.3 Schedules and effective resolution

The schedule $\lambda(t)$ controls how frequencies are attenuated with time. We evaluate two schedules, named by how the damping cutoff in $\mathbf{k}$ moves with time. The log-linear schedule $\lambda(t) = t \cdot 10^{\lambda_i + (\lambda_f - \lambda_i)\, t}$ moves it roughly uniformly on a log scale, and the linear schedule $\lambda(t) = \theta\, t / (\lambda_f (1 -$

$t) + \lambda_i)^2$ moves it roughly uniformly on a linear scale. The multiplicative t ensures $\lambda(0) = 0$. Among the parameters, λ_i primarily sets the high-frequency, early-time behavior, while λ_f, k_c, and θ primarily set the low-frequency, late-time behavior; all four jointly shape the full schedule.

The notion of time-evolving cutoff in $\mathbf{k}$ gives super-resolution a direct interpretation. For a chosen SNR threshold,

$$\mathrm{SNR}_n(\mathbf{k}) = \frac{\bar{\boldsymbol{\alpha}}_n(\mathbf{k})}{1 - \bar{\boldsymbol{\alpha}}_n(\mathbf{k})}, \tag{8}$$

the modes above the threshold define an effective resolution. Starting the reverse process from a timestep where the effective resolution is zero gives image generation; starting from a timestep whose surviving signals correspond to a lower-resolution input gives super-resolution. Because $\lambda(t)$ is continuous before discretization and can be densely sampled in implementation, the effective resolution varies continuously along the schedule, yielding a continuum of super-resolution factors from a single trained model (Figure E.13).

4.4 Connection to scale-space theory and renormalization group

Equation (3) extends the vanilla scale-space operation [10, 11] to frequency space with noise. In pixel space, multiplying DCT modes by $\exp[-\mathbf{k}^2\lambda(t)/2]$ amounts to Gaussian smoothing at scale $\lambda(t)$, the same operator that appears in linear scale-space theory. Crucially, our additional noise term turns the deterministic smoothing into a stochastic coarse-graining process whose final covariance matches the dataset variance spectrum.

The same equation also conceptually resembles a renormalization group (RG) coarse-graining step, where short-distance degrees of freedom are discarded before long-distance structures. We do not claim that our method is an exact RG transformation; rather, we use the RG as analogy: if a dataset exhibits approximate scale invariance, a reverse model trained on the scale-ordered forward process should learn how fine scales are distributed conditioned on coarse scales. The critical-Ising experiment in Section 5.5 tests this idea in a setting with known scale-invariant structure.

5 Experiments

We evaluate whether the same frequency-space diffusion process can serve as an unconditional image generator and a continuous super-resolution model. We test SKILD in three settings: unconditional image generation on CIFAR-10, $2\times$–$8\times$ continuous super-resolution on ImageNet-128 and 256, and a scientific benchmark on the critical two-dimensional Ising model, probing scale invariance directly through four-point correlations. SKILD is competitive with or outperforms strong baselines in each setting.

5.1 Setup

Architecture. All reported models use a score U-Net backbone from the NCSN++ family [18]. Exact channel counts, depths, and attention configurations are in Appendices D and E.

Data. We use CIFAR-10 [7] and ImageNet [8] as released. Critical-Ising configurations are sampled on a 128×128 square lattice using the Wolff cluster algorithm [56]; data-generation details are in Appendix F.

Noise schedule. We test both the log-linear and linear schedules on CIFAR-10 and the linear schedule on ImageNet and Ising experiments. All experiments use $N = 1000$ timesteps; exact schedule parameters are in Appendices D, E, and F.

Training. We train with AdamW and use an exponential moving average of weights at sampling time. Full hyperparameters and compute details are in Appendices D, E, and F.

Evaluation. For CIFAR-10 we report FID [57] and Inception Score (IS) [58] on 50K generated samples. For ImageNet super-resolution we report PSNR, SSIM [59], LPIPS [60], MUSIQ [61], and CLIPIQA [62] from the last checkpoint, on a random 3K-image subset of the ImageNet validation set, following the protocols of [48, 49]. The Ising super-resolution experiment is evaluated on 1K samples from the last checkpoint by a connected four-point correlation, a statistical-physics observable that probes how accurately the model reproduces Ising model structures across scales (Section 5.5).

5.2 Effective resolution protocol

For super-resolution, we choose the reverse starting timestep by the SNR defined in Section 4. We use threshold 0.1 in all ImageNet and Ising super-resolution experiments. This corresponds to applying an effective low-pass filter in the frequency domain, where modes below the effective-resolution cutoff retain signal while higher-frequency modes are attenuated and dominated by noise. At sampling time, the model starts from the exact forward marginal of the paired high-resolution (HR) image at the chosen timestep, then reverses to t_0. This protocol turns super-resolution into a partial reverse diffusion problem rather than a conditional generation problem.

We validate the effective-resolution interpretation by comparing the surviving signal at the chosen timestep with standard bicubic down-up sampling. The MSE and PSNR values in Table E.9 show that SNR of 0.1 produces low-resolution (LR) inputs close to conventional $4\times$ or $8\times$ degradations while preserving consistency with the forward diffusion process.

5.3 Unconditional CIFAR-10 generation

Table 1: Unconditional CIFAR-10 sample quality. Best results within each model group are in **bold**.

Model	**Frequency-agnostic model**				**Frequency/scale-informed model**						
	DDPM [17]	DDPM++ [18]	NCSN++ [18]	EDM [21]	Cold diff.* (deblur) [37]	Cold diff.* (sr) [37]	EqualSNR [32]	WaveDiff [29]	IHDM [26]	DCTdiff [31]	SKILD (ours)
FID ↓	3.17	2.78	2.20	**1.97**	80.08	152.76	13.63	4.01	18.96	5.02	**2.61**
IS ↑	–	9.46	9.64	**9.89**	–	–	–	–	–	7.70	**9.65**

* The [37] results are generated from highly-degraded images which still contains some signal, instead from scratch.

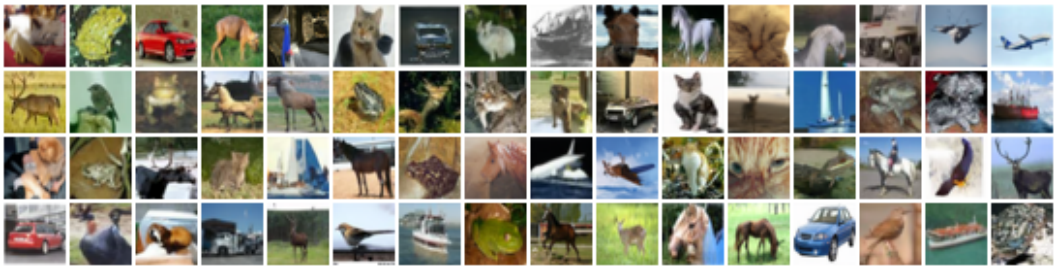

Figure 3: Uncurated samples of generated images on CIFAR-10; more in Appendix G (Figure G.18).

Table 1 compares unconditional CIFAR-10 generation. Figure 3 shows uncurated samples that SKILD generates on CIFAR-10. SKILD is competitive with the state-of-the-art models and achieves the best FID and IS among the frequency or scale-informed models listed, using the linear schedule shown in Figure 4(a).

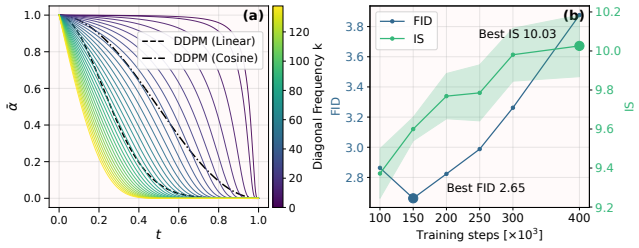

Figure 4: (a) Linear schedule with best FID and IS on CIFAR-10 generation. Lower-frequency modes are attenuated later than higher-frequency ones. (b) Mode collapse with training steps. While IS continues to improve with training, FID converges early on and worsens at later times.

Ablation studies. We conduct several sets of ablation studies and discuss three of them briefly here. Figure 4(b) shows the model over training. A mode collapse appears with training course: FID reaches its best value before the final checkpoint, while IS continues to improve. We interpret this as evidence that low-frequency reconstruction remains the bottleneck for generation on object-centric datasets like CIFAR-10. The limitation section discusses this point directly.

This mode collapse is not inherent to SKILD. It happens at the end of the forward process, where only the lowest frequencies still carry signal (Appendix D.IV.1). Scale invariance breaks down at the lowest frequencies for images of finite size and resolution, much like finite system size and lattice spacing break scale invariance in critical physical systems. This final stage makes up only a small fraction of the diffusion process and needs separate handling. We test two fixes using the same U-Net, optimizer, dataset, training budget, and schedule (Table 2). The first option pretrains a small 19M-parameter VAE [63] on the heavily blurred images from the final diffusion steps, reaching a reconstruction mean squared error of 2.35×10^{-7}. During sampling, the frozen VAE generates this initial low-frequency state and the reverse process runs from there, keeping the whole pipeline unconditional. The second option adds class labels and uses classifier-free guidance [64] with a scale of $s = 1 + w = 1.2$, where w is the guidance weight. Both methods keep training stable through 400K steps, and both FID and IS improve between 200K and 400K steps. Figure 5 shows samples from the class-conditional model. Because the guided model does not use the VAE, the two fixes are independent and can be used together to further improve the performance of model.

Table 2: Two solutions for low-frequency mode collapse on CIFAR-10. Both runs share the same training setup and model configurations. The network used has half the channel count compared to the one in Table 1 and Figure 4(b).

	FID ↓		IS ↑	
Add-ons	200K	400K	200K	400K
Low-frequency VAE	3.76	3.33	9.41	9.63
Class conditioning + CFG ($s = 1.2$)	3.43	3.33	9.48	9.75

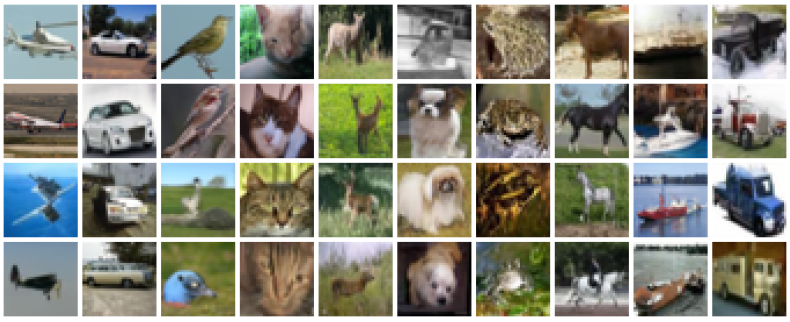

Figure 5: Class-conditional CIFAR-10 samples generated by SKILD with classifier-free guidance ($s = 1.2$). From left to right, each column shows airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck.

In Table D.7, we verify the robustness of SKILD on image generation against a broad range of schedules in the log-linear and linear families. Most schedules reach FID below or near 5 and all reach IS near or above 9 within 400K training steps. This indicates that SKILD is not tuned to a single fragile schedule, albeit the convergence speed is schedule-dependent.

SKILD combines spectrum-matched noise with frequency-ordered attenuation to achieve scale invariance. In Table 3, we remove each component one at a time while keeping the rest of the model configuration and training fixed. First, we replace $\mathbf{S}_0$ with a nearly flat spectrum, which resembles standard white noise rather than the k^{-2} power law of natural images. This change pushes FID at 400K steps up to 20.02. This drastic degradation shows the importance of a spectrum-matched noise, which is necessary to maintain scale invariance in the diffusion process [3]. Second, we keep the true $\mathbf{S}_0$ but replace the frequency-dependent attenuation $\bar{\boldsymbol{\alpha}}_n(\mathbf{k})$ with the standard DDPM schedule [17], which attenuates all frequencies at the same rate. This also hurts both FID and IS significantly. The ablation demonstrates that both components are necessary to ensure scale invariance and achieve high-quality image generation with SKILD.

[3] $\mathbf{S}_0$ is cheap to compute – estimating $\mathbf{S}_0$ on ImageNet takes only a few minutes on a single consumer GPU. It is also similar across natural image datasets.

Table 3: Ablations on the effect of spectrum-matched noise and frequency-dependent attenuation in SKILD. Without either component, FID and IS noticeably degrade.

	FID ↓		IS ↑	
Ablations	200K	400K	200K	400K
flat $\mathbf{S}_0$	18.23	20.02	8.14	7.96
Frequency-agnostic $\bar{\boldsymbol{\alpha}}_n$	6.59	5.35	8.66	8.93

More details of ablations can be found in Appendix D.IV, including FID and IS convergence over training compared to common pixel-space diffusion schedules (Figure D.12), different network predictions, the effectiveness of second-moment sampler, potential of reducing number of diffusion steps, and different numerical cutoffs.

5.4 ImageNet super-resolution

Table 4: $4\times$ super-resolution quality on the ImageNet-Test [48, 49]. Best and second-best results among quality metrics are highlighted in **bold** and underline.

Model	PSNR↑	SSIM↑	LPIPS↓	CLIPIQA↑	MUSIQ↑	# Param. (M)
GAN-based						
BSRGAN [43]	24.42	0.659	0.259	0.581	54.697	16.70
SwinIR [44]	23.99	0.667	0.238	0.564	53.790	16.70
RealESRGAN [45]	24.04	0.665	0.254	0.523	52.538	28.01
Conditional diffusion-based						
LDM-30 [65]	24.49	0.651	0.248	0.572	50.895	113.60
LDM-15 [65]	24.89	0.670	0.269	0.512	46.419	113.60
ResShift [48]	25.01	0.677	0.231	0.592	53.660	118.59
SinSR [49]	24.56	0.657	0.221	0.611	53.357	118.59
IRSDE [51]	24.48	0.602	0.304	0.513	45.382	137.20
DDRM [50]	25.56	0.674	0.471	0.372	24.746	552.80
I2SB [52]	**26.76**	**0.730**	0.206	0.489	53.936	552.80
Unconditional diffusion						
SKILD (Ours)	24.10	0.683	**0.186**	**0.612**	**59.226**	121.12

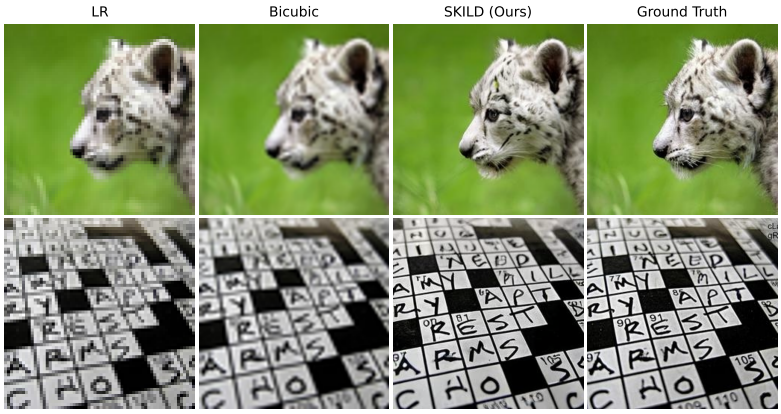

Figure 6: $4\times$ super-resolution samples on ImageNet-256. The model is initialized from a 64×64 low-resolution forward state and reconstructs high-frequency details through the reverse process.

Table 4 reports $4\times$ super-resolution quality on ImageNet. All conditional baselines receive the low-resolution input through an explicit conditioning path, most of which additionally use class labels



Figure 7: Continuous super-resolution on ImageNet-128. (a) Low-resolution inputs at effective resolutions 16×16 through 64×64, corresponding to $8\times$ through $2\times$ super-resolution factors. (b) High-resolution reconstructions produced by the same checkpoint from each effective-resolution starting state. A continuum of super-resolution factors is accessible from a single trained model.

or classifier-free guidance. SKILD uses no conditioning of any kind: it starts from the corresponding forward marginal and runs the same reverse process used for unconditional generation. Despite this simplicity in design, the 256-resolution model achieves the best LPIPS, CLIPIQA, and MUSIQ among the compared methods and the second-best SSIM. PSNR favors methods with higher raw pixel accuracy, while metrics that emphasize human perception favor our method. This gap comes from the perception-distortion tradeoff [66]. Minimizing mean squared error pushes a model to predict the conditional mean, which averages over all the sharp images that could match a given low-resolution input. Models trained this way get higher PSNR by blurring out fine details. Instead of predicting an average, SKILD samples the missing high frequencies through the reverse process, trading raw pixel accuracy for sharper perceptual quality. Two super-resolution samples are shown in Figure 6, with more in Appendix G.

A single trained ImageNet model accommodates a continuum of super-resolution factors by varying only the starting timestep. Figure 7 shows reconstructions at factors from $2\times$ to $8\times$ produced by the same checkpoint, and Figure E.13 plots how the effective resolution decreases continuously with t.

5.5 Scientific benchmark

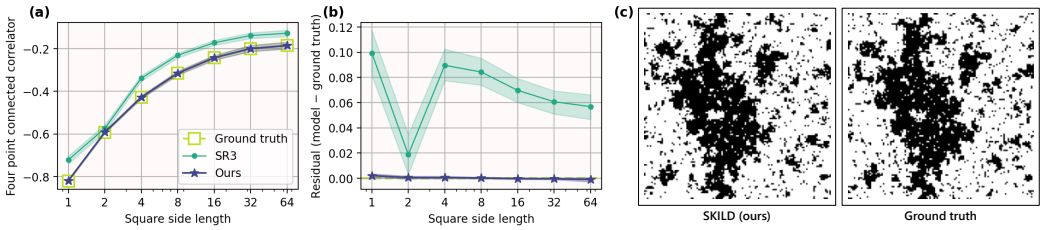

Figure 8: (a-b) Benchmark of four-point correlator accuracy. Our reconstruction closely tracks the ground truth, while SR3 shows a clear deviation. (c) SKILD-reconstructed critical Ising field sample compared to the ground truth.

Natural images are approximately scale-invariant only after averaging over many scenes. Critical physical systems let us ask a stricter question: can a model reconstruct missing fine scales while preserving observables that define the scale-invariant law? We test this on the prototypical two-dimensional Ising model at criticality. Despite its simplicity–placing a spin $s_i \in \{-1, +1\}$ on each lattice site, with nearest neighbors preferring to align–the Ising model serves a foundational role in areas across statistical mechanics, combinatorics, and computational complexity theory. At its critical temperature, the correlation length diverges, the distribution becomes statistically self-similar under RG coarse-graining, and the continuum limit is described by a conformal field theory [53, 5, 6]. Previous works also have applied neural networks to Ising super-resolution and inverse RG [67, 68].

This setting provides a more precise benchmark of scale-invariance than perceptual realism. A visually plausible spin configuration can still have the wrong connected correlations, the wrong response functions, or the wrong universality-class signatures.

Evaluation: connected four-point correlator. We evaluate a connected four-point correlator, equivalently a fourth-order joint cumulant, over the four corner spins of square patches at multiple side lengths. The join cumulant subtracts all pairwise contributions, isolating non-Gaussian dependence that cannot be inferred from the mean, variance, or two-point correlation alone. Higher-order correlations are central observables in critical systems, so matching them across scales is a stronger test than matching visual texture or pixel-level distortion. Data generation, paired evaluation, and the correlator estimation are detailed in Appendix F.

Results. We super-resolve from a 32×32 effective-resolution starting state to a 128×128 critical-Ising field. Figure 8 shows that SKILD's reconstructed four-point correlator closely tracks the ground truth at every patch size, while SR3 [46], a strong diffusion-based conditional super-resolution model, deviates significantly from the ground truth.

6 Conclusions

We introduce SKILD, a scale-invariant frequency-space diffusion model whose forward process is a stochastic coarse-graining operator: fine modes are damped before coarse ones, and the injected noise carries the dataset spectrum. Aligning diffusion with the scale structure of data makes scale an explicit coordinate of the generative process. Unconditional generation and super-resolution then become different starting points of the same reverse trajectory, eliminating task-specific conditioning and guidance. Empirically, SKILD reaches FID 2.61 and IS 9.65 on unconditional CIFAR-10, supports continuous $2\times$–$8\times$ super-resolution from a single ImageNet checkpoint, and reconstructs critical-Ising fields whose connected four-point correlations closes tracks the ground truth.

Discussion. Our method shifts the modeling burden from conditional mappings to the design of a scale-invariant diffusion process. Rather than learning a separate low-to-high-resolution mapping, the model learns a single reverse trajectory over scales, with low-resolution inputs as intermediate forward states. This removes task-specific conditioning or guidance. In the mean time, low-frequency generation becomes a central bottleneck: errors in coarse structure early in the reverse chain propagate and constrain later high-frequency generation. Our results reflect this trade-off: super-resolution benefits from accurate coarse initialization, while unconditional generation remains sensitive to low-frequency modeling. The framework also suggests an evaluation criterion for self-similarity beyond perceptual quality, namely whether fine-scale details remain statistically similar with coarse-scale structure, which is particularly relevant for scientific applications. We provide such an instance through super-resolution experiments on critical-Ising fields.

Limitations and future work. The current model opens a vast range of directions for future work. (i) Currently, sampling requires 1000 ancestral steps; faster samplers or tailored solvers for mode-dependent schedules are a natural next step. (ii) Unconditional generation remains sensitive to low-frequency structure generation, and improvements there should stabilize global structure without sacrificing fine detail. (iii) Our super-resolution protocol uses exact forward marginals as low-resolution inputs; extending to real-world degradations such as unknown camera and compression pipelines is an important direction. (iv) The Ising experiment establishes scale-invariant diffusion on one critical system; extending to additional physical systems and higher-order observables would broaden the scientific benchmark. (v) Our model uses off-the-shelf neural network. New architecture designs tailored to our model could significantly improve its performance while furthering understanding of network designs for diffusion models.

Broader impacts. Unifying generation and super-resolution into a single reverse process can simplify deployment pipelines and reduce reliance on task-specific models. The same capability has dual-use risks: super-resolution can hallucinate plausible but incorrect high-frequency content, with consequences in forensics, medical imaging, and other sensitive settings where outputs require validation. The scientific-data setting illustrates a complementary benefit: scale-invariant models can be evaluated against known physical laws, a style of evaluation we expect to be useful wherever multi-scale structure carries scientific meaning.

Code and reproducibility. The code for reproducing the main results and data of this paper is available at `https://github.com/JazzyCH/SKILD`.

Acknowledgment

The authors acknowledge support from the National Science Foundation under Cooperative Agreement PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions) and the MIT Generative AI Impact Consortium. This work is supported by the Toyota Research Institute University 3.0 Program, and the Department of the Air Force Artificial Intelligence Accelerator under Cooperative Agreement No. FA8750-19-2-1000. ZJC is in part supported by the Kurt Forrest Foundation Fellowship and the Henry Kendall Fellowship. ZC is in part supported by the MathWorks Fellowship and the Henry Kendall Fellowship. AW is in part supported by the National Science Foundation Graduate Research Fellowship. CD is in part supported by the Tayebati Postdoctoral Fellowship. The authors also acknowledge the National Artificial Intelligence Research Resource (NAIRR) Pilot Program, the DeltaAI advanced computing and data resource at the National Center for Supercomputing Applications (supported by NSF Award OAC-2320345 and the State of Illinois), and Lambda Inc. for providing compute resources.

References

[1] Field, D. J. Relations between the statistics of natural images and the response properties of cortical cells. *J. Opt. Soc. Am. A* **4**, 2379–2394 (1987). URL `https://opg.optica.org/josaa/abstract.cfm?URI=josaa-4-12-2379`.

[2] Ruderman, D. L. The statistics of natural images. *Network: computation in neural systems* **5**, 517 (1994).

[3] van der Schaaf, A. & van Hateren, J. Modelling the power spectra of natural images: Statistics and information. *Vision Research* **36**, 2759–2770 (1996). URL `https://www.sciencedirect.com/science/article/pii/0042698996000028`.

[4] Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annual review of neuroscience* **24**, 1193–1216 (2001).

[5] Wilson, K. G. & Kogut, J. The renormalization group and the epsilon expansion. *Physics Reports* **12**, 75–199 (1974).

[6] Belavin, A. A., Polyakov, A. M. & Zamolodchikov, A. B. Infinite conformal symmetry in two-dimensional quantum field theory. *Nuclear Physics B* **241**, 333–380 (1984).

[7] Krizhevsky, A. & Hinton, G. Learning multiple layers of features from tiny images. Tech. Rep., University of Toronto (2009).

[8] Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (2009).

[9] Witkin, A. Scale-space filtering: A new approach to multi-scale description. In *ICASSP '84. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 9, 150–153 (1984).

[10] Koenderink, J. J. The structure of images. *Biological Cybernetics* **50**, 363–370 (1984). URL `https://doi.org/10.1007/BF00336961`.

[11] Lindeberg, T. Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of applied statistics* **21**, 225–270 (1994).

[12] Lindeberg, T. Feature detection with automatic scale selection. *International journal of computer vision* **30**, 79–116 (1998).

[13] Lowe, D. G. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**, 91–110 (2004). URL `https://doi.org/10.1023/B:VISI.0000029664.99615.94`.

[14] Burt, P. J. & Adelson, E. H. The laplacian pyramid as a compact image code. In *Readings in computer vision*, 671–679 (Elsevier, 1987).

[15] Olshausen, B. A. & Field, D. J. Natural image statistics and efficient coding. *Network: Computation in Neural Systems* **7**, 333 (1996). URL `https://doi.org/10.1088/0954-898X/7/2/014`.

[16] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N. & Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265 (pmlr, 2015).

[17] Ho, J., Jain, A. & Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020).

[18] Song, Y. *et al.* Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020).

[19] Nichol, A. Q. & Dhariwal, P. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171 (PMLR, 2021).

[20] Dhariwal, P. & Nichol, A. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021).

[21] Karras, T., Aittala, M., Aila, T. & Laine, S. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems* **35**, 26565–26577 (2022).

[22] Chen, T. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972* (2023).

[23] Ho, J. *et al.* Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research* **23**, 1–33 (2022).

[24] Teng, J. *et al.* Relay diffusion: Unifying diffusion process across resolutions for image synthesis. *arXiv preprint arXiv:2309.03350* (2023).

[25] Cotler, J. & Rezchikov, S. Renormalizing diffusion models. *arXiv preprint arXiv:2308.12355* (2023).

[26] Rissanen, S., Heinonen, M. & Solin, A. Generative modelling with inverse heat dissipation. *arXiv preprint arXiv:2206.13397* (2022).

[27] Masuki, K. & Ashida, Y. Generative diffusion model with inverse renormalization group flows. *arXiv preprint arXiv:2501.09064* (2025).

[28] Sheshmani, A., You, Y.-Z., Buyukates, B., Ziashahabi, A. & Avestimehr, S. Renormalization group flow, optimal transport, and diffusion-based generative model. *Physical Review E* **111**, 015304 (2025).

[29] Guth, F., Coste, S., De Bortoli, V. & Mallat, S. Wavelet score-based generative modeling. *Advances in neural information processing systems* **35**, 478–491 (2022).

[30] Phung, H., Dao, Q. & Tran, A. Wavelet diffusion models are fast and scalable image generators. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10199–10208 (2023).

[31] Ning, M. *et al.* Dctdiff: Intriguing properties of image generative modeling in the dct space. *arXiv preprint arXiv:2412.15032* (2024).

[32] Falck, F. *et al.* A fourier space perspective on diffusion models. *arXiv preprint arXiv:2505.11278* (2025).

[33] Yu, H., Luo, H., Yuan, H., Rong, Y. & Zhao, F. Frequency autoregressive image generation with continuous tokens. *arXiv preprint arXiv:2503.05305* (2025).

[34] Moser, B. B., Frolov, S., Raue, F., Palacio, S. & Dengel, A. Waving goodbye to low-res: A diffusion-wavelet approach for image super-resolution. In *2024 International Joint Conference on Neural Networks (IJCNN)*, 1–8 (IEEE, 2024).

[35] Gao, X., Xu, Z., Zhao, J. & Liu, J. Frequency-controlled diffusion model for versatile text-guided image-to-image translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, 1824–1832 (2024).

[36] Friedrich, P., Wolleb, J., Bieder, F., Durrer, A. & Cattin, P. C. Wdm: 3d wavelet diffusion models for high-resolution medical image synthesis. In *MICCAI workshop on deep generative models*, 11–21 (Springer, 2024).

[37] Bansal, A. *et al.* Cold diffusion: Inverting arbitrary image transforms without noise. *Advances in Neural Information Processing Systems* **36**, 41259–41282 (2023).

[38] Tian, K., Jiang, Y., Yuan, Z., Peng, B. & Wang, L. Visual autoregressive modeling: Scalable image generation via next-scale prediction. *Advances in neural information processing systems* **37**, 84839–84865 (2024).

[39] Freeman, W. & Pasztor, E. Learning low-level vision. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, 1182–1189 vol.2 (1999).

[40] Dong, C., Loy, C. C., He, K. & Tang, X. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* **38**, 295–307 (2015).

[41] Lim, B., Son, S., Kim, H., Nah, S. & Mu Lee, K. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144 (2017).

[42] Wang, X. *et al.* Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0 (2018).

[43] Zhang, K., Liang, J., Van Gool, L. & Timofte, R. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4791–4800 (2021).

[44] Liang, J. *et al.* Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1833–1844 (2021).

[45] Wang, X., Xie, L., Dong, C. & Shan, Y. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, 1905–1914 (2021).

[46] Saharia, C. *et al.* Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence* **45**, 4713–4726 (2022).

[47] Li, H. *et al.* Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* **479**, 47–59 (2022).

[48] Yue, Z., Wang, J. & Loy, C. C. Resshift: Efficient diffusion model for image super-resolution by residual shifting. *Advances in neural information processing systems* **36**, 13294–13307 (2023).

[49] Wang, Y. *et al.* Sinsr: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 25796–25805 (2024).

[50] Kawar, B., Elad, M., Ermon, S. & Song, J. Denoising diffusion restoration models. *Advances in neural information processing systems* **35**, 23593–23606 (2022).

[51] Luo, Z., Gustafsson, F. K., Zhao, Z., Sjölund, J. & Schön, T. B. Image restoration with mean-reverting stochastic differential equations. *arXiv preprint arXiv:2301.11699* (2023).

[52] Liu, G.-H. *et al.* I^2SB: Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872* (2023).

[53] Onsager, L. Crystal statistics. i. a two-dimensional model with an order-disorder transition. *Physical Review* **65**, 117–149 (1944).

[54] Hyvrinen, A., Hurri, J. & Hoyer, P. O. *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision.* (Springer Publishing Company, Incorporated, 2009), 1st edn.

[55] Ahmed, N., Natarajan, T. & Rao, K. R. Discrete cosine transform. *IEEE Transactions on Computers* **C-23**, 90–93 (1974).

[56] Wolff, U. Collective monte carlo updating for spin systems. *Physical Review Letters* **62**, 361–364 (1989).

[57] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B. & Hochreiter, S. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, vol. 30 (2017).

[58] Salimans, T. *et al.* Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, vol. 29 (2016).

[59] Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**, 600–612 (2004).

[60] Zhang, R., Isola, P., Efros, A. A., Shechtman, E. & Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 586–595 (2018).

[61] Ke, J., Wang, Q., Wang, Y., Milanfar, P. & Yang, F. MUSIQ: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5148–5157 (2021).

[62] Wang, J., Chan, K. C. K. & Loy, C. C. Exploring CLIP for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, 2555–2563 (2023).

[63] Kingma, D. P. & Welling, M. Auto-encoding variational Bayes. In *International Conference on Learning Representations* (2014).

[64] Ho, J. & Salimans, T. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022).

[65] Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695 (2022).

[66] Blau, Y. & Michaeli, T. The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6228–6237 (2018).

[67] Efthymiou, S., Beach, M. J. S. & Melko, R. G. Super-resolving the ising model with convolutional neural networks. *Physical Review B* **99**, 075113 (2019).

[68] Shiina, K., Mori, H., Tomita, Y., Lee, H. K. & Okabe, Y. Inverse renormalization group based on image super-resolution using deep convolutional networks. *Scientific Reports* **11**, 9617 (2021).

[69] Jin, Y. *et al.* Pyramidal flow matching for efficient video generative modeling. In *International Conference on Learning Representations* (2025).

[70] Salimans, T. & Ho, J. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations* (2022).

A DDPM formulation

This appendix derives the DDPM posterior used in the main text and lists the training targets supported by our implementation. Figure A.9 shows a forward trajectory on a natural image, illustrating how high-frequency content is removed before low-frequency content under the scale-invariant schedule.

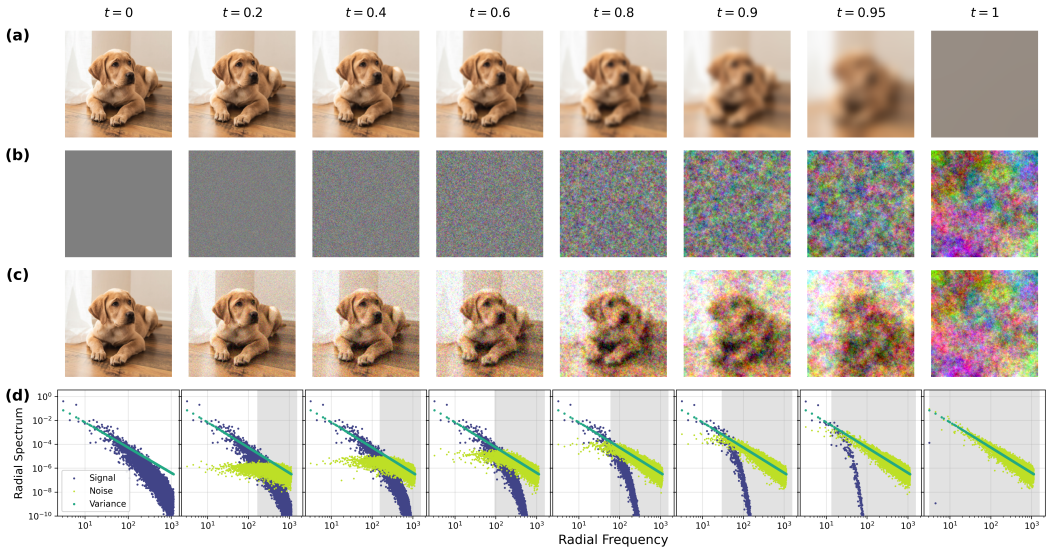

Figure A.9: Forward trajectory of SKILD on a natural image. (a), (b), and (c) show the signal, noise, and full forward process, respectively. (d) shows the radial power spectrum evolution measured from the trajectory. High frequencies are suppressed first, while low-frequency content persists until late timesteps. Gray regions indicate modes below the SNR threshold, where resolution is effectively lost.

A.I Posterior derivation

The forward marginal in main text Eq. (4) can also be written as a one-step Markov transition,

$$\mathbf{X}_n = \sqrt{\boldsymbol{\alpha}_n} \odot \mathbf{X}_{n-1} + \sqrt{\boldsymbol{\beta}_n} \odot \sqrt{\mathbf{S}_0}\, \epsilon_n, \qquad \epsilon_n \sim \mathcal{N}(0, \mathbf{I}), \tag{A.9}$$

with $\bar{\boldsymbol{\alpha}}_n = e^{-\mathbf{k}^2 \lambda_n}$, $\boldsymbol{\alpha}_n = \bar{\boldsymbol{\alpha}}_n / \bar{\boldsymbol{\alpha}}_{n-1}$, $\boldsymbol{\beta}_n = 1 - \boldsymbol{\alpha}_n$, and $\bar{\boldsymbol{\alpha}}_0 = \mathbf{1}$. Note that $\bar{\boldsymbol{\alpha}}_n = \prod_{i=1}^{n} \boldsymbol{\alpha}_i$.

When $\mathbf{X}_0$ is known, the reverse posterior $q(\mathbf{X}_{n-1} \mid \mathbf{X}_n, \mathbf{X}_0)$ is exact. By Bayes' theorem,

$$q(\mathbf{X}_{n-1} \mid \mathbf{X}_n, \mathbf{X}_0) \propto q(\mathbf{X}_n \mid \mathbf{X}_{n-1})\, q(\mathbf{X}_{n-1} \mid \mathbf{X}_0), \tag{A.10}$$

where we used the Markov property and dropped the $\mathbf{X}_{n-1}$-independent factor $q(\mathbf{X}_n \mid \mathbf{X}_0)$. Eq. (A.9) gives

$$q(\mathbf{X}_n \mid \mathbf{X}_{n-1}) = \mathcal{N}(\mathbf{X}_n;\, \sqrt{\boldsymbol{\alpha}_n} \odot \mathbf{X}_{n-1},\, \mathbf{S}_0 \boldsymbol{\beta}_n) \;\propto\; \mathcal{N}(\mathbf{X}_{n-1};\, \mathbf{X}_n / \sqrt{\boldsymbol{\alpha}_n},\, \mathbf{S}_0 \boldsymbol{\beta}_n / \boldsymbol{\alpha}_n)\,, \tag{A.11}$$

where the second form completes the square in $\mathbf{X}_{n-1}$. The forward marginal at $n-1$ is

$$q(\mathbf{X}_{n-1} \mid \mathbf{X}_0) = \mathcal{N}\big(\mathbf{X}_{n-1};\, \sqrt{\bar{\boldsymbol{\alpha}}_{n-1}} \odot \mathbf{X}_0,\, \mathbf{S}_0(1 - \bar{\boldsymbol{\alpha}}_{n-1})\big)\,. \tag{A.12}$$

Multiplying these two Gaussians and using the standard product formula $\tilde{\mathbf{v}} = (\mathbf{v}_1^{-1} + \mathbf{v}_2^{-1})^{-1}$, $\tilde{\mathbf{m}} = \tilde{\mathbf{v}}(\mathbf{m}_1/\mathbf{v}_1 + \mathbf{m}_2/\mathbf{v}_2)$ gives

$$q(\mathbf{X}_{n-1} \mid \mathbf{X}_n, \mathbf{X}_0) = \mathcal{N}\Big(\boldsymbol{\mu}_q,\, \mathbf{S}_0 \tilde{\boldsymbol{\beta}}_n\Big)\,, \tag{A.13a}$$

$$\tilde{\boldsymbol{\beta}}_n = \frac{\boldsymbol{\beta}_n (1 - \bar{\boldsymbol{\alpha}}_{n-1})}{1 - \bar{\boldsymbol{\alpha}}_n}, \tag{A.13b}$$

$$\boldsymbol{\mu}_q = \frac{\sqrt{\bar{\boldsymbol{\alpha}}_{n-1}}\, \boldsymbol{\beta}_n}{1 - \bar{\boldsymbol{\alpha}}_n}\, \mathbf{X}_0 + \frac{\sqrt{\boldsymbol{\alpha}_n}\, (1 - \bar{\boldsymbol{\alpha}}_{n-1})}{1 - \bar{\boldsymbol{\alpha}}_n}\, \mathbf{X}_n. \tag{A.13c}$$

Reverse sampling with the exact posterior reads

$$\mathbf{X}_{n-1} = \boldsymbol{\mu}_q(n, \mathbf{X}_n, \mathbf{X}_0) + \sqrt{\mathbf{S}_0 \tilde{\boldsymbol{\beta}}_n} \odot \epsilon_n. \tag{A.14}$$

A.II Training targets

At sampling time, $\mathbf{X}_0$ is unknown and is replaced by a network estimate $\mathbf{X}_{0,\theta}$ inserted into Eq. (A.13). Our implementation supports four prediction targets: noise ϵ, data $\mathbf{X}_0$, an unwhitened frequency-space velocity $\mathbf{w}$, and a whitened velocity $\mathbf{v}$. The main results use ϵ-prediction; the other targets follow the same posterior algebra and are reported as diagnostic ablations.

ϵ-prediction. Minimizing the MSE against the noise variable,

$$\mathcal{L} = \mathbb{E}_{n,\mathbf{X}_0,\epsilon}\left[\|\epsilon - \epsilon_\theta(\mathbf{X}_n, n)\|_2^2\right], \tag{A.15}$$

inverts main text Eq. (4) to obtain

$$\mathbf{X}_{0,\theta} = \frac{\mathbf{X}_n - \sqrt{1-\bar{\boldsymbol{\alpha}}_n} \odot \sqrt{\mathbf{S}_0}\, \epsilon_{n,\theta}}{\sqrt{\bar{\boldsymbol{\alpha}}_n}}, \tag{A.16}$$

and substituting into Eq. (A.13) gives the ϵ-form of the posterior mean used at sampling,

$$\boldsymbol{\mu}_\theta(n, \mathbf{X}_n) = \frac{1}{\sqrt{\boldsymbol{\alpha}_n}} \left(\mathbf{X}_n - \frac{\boldsymbol{\beta}_n \sqrt{\mathbf{S}_0}}{\sqrt{1-\bar{\boldsymbol{\alpha}}_n}}\, \epsilon_{n,\theta} \right). \tag{A.17}$$

In finite precision we floor very small $\boldsymbol{\alpha}_n(\mathbf{k})$ values during sampling, as described in the main text, to avoid unstable divisions in high-frequency modes.

$\mathbf{X}_0$-prediction. Minimizing the MSE against the data,

$$\mathcal{L} = \mathbb{E}_{n,\mathbf{X}_0,\epsilon}\left[\|\mathbf{X}_0 - \mathbf{X}_{0,\theta}(\mathbf{X}_n, n)\|_2^2\right], \tag{A.18}$$

substitutes $\mathbf{X}_{0,\theta}$ directly into the posterior mean,

$$\boldsymbol{\mu}_\theta(n, \mathbf{X}_n) = \frac{\sqrt{\bar{\boldsymbol{\alpha}}_{n-1}}\, \boldsymbol{\beta}_n}{1-\bar{\boldsymbol{\alpha}}_n}\, \mathbf{X}_{0,\theta} + \frac{\sqrt{\boldsymbol{\alpha}_n}\,(1-\bar{\boldsymbol{\alpha}}_{n-1})}{1-\bar{\boldsymbol{\alpha}}_n}\, \mathbf{X}_n. \tag{A.19}$$

$\mathbf{w}$-prediction. Let $a = \sqrt{\bar{\boldsymbol{\alpha}}_n}$ and $b = \sqrt{1-\bar{\boldsymbol{\alpha}}_n}$. The forward marginal and the velocity target are

$$\mathbf{X}_n = a \odot \mathbf{X}_0 + b \odot \sqrt{\mathbf{S}_0}\, \epsilon_n, \tag{A.20a}$$

$$\mathbf{w}_n = a \odot \sqrt{\mathbf{S}_0}\, \epsilon_n - b \odot \mathbf{X}_0. \tag{A.20b}$$

Given a prediction $\mathbf{w}_{n,\theta}$, the data and colored-noise estimates follow without division (using $a^2 + b^2 = 1$):

$$\mathbf{X}_{0,\theta} = a \odot \mathbf{X}_n - b \odot \mathbf{w}_{n,\theta}, \tag{A.21a}$$

$$\mathbf{z}_{n,\theta} = b \odot \mathbf{X}_n + a \odot \mathbf{w}_{n,\theta}, \tag{A.21b}$$

where $\mathbf{z}_{n,\theta}$ estimates $\sqrt{\mathbf{S}_0} \odot \epsilon_n$. The corresponding loss is

$$\mathcal{L}_w = \mathbb{E}_{n,\mathbf{X}_0,\epsilon}\left[\|\mathbf{w}_n - \mathbf{w}_{n,\theta}(\mathbf{X}_n, n)\|_2^2\right]. \tag{A.22}$$

$\mathbf{v}$-prediction. The whitened velocity divides $\mathbf{w}$ by $\sqrt{\mathbf{S}_0}$,

$$\mathbf{v}_n = a \odot \epsilon_n - b \odot \mathbf{X}_0 / \sqrt{\mathbf{S}_0}. \tag{A.23}$$

For $\mathbf{w}$ and $\mathbf{v}$, the loss can optionally be weighted by the local frequency shell $\bar{\boldsymbol{\alpha}}_{n-1} - \bar{\boldsymbol{\alpha}}_n$. In our ablations, direct prediction of the colored noise $\sqrt{1-\bar{\boldsymbol{\alpha}}_n} \odot \sqrt{\mathbf{S}_0}\, \epsilon_n$ matches the CIFAR-10 FID and IS of ϵ-prediction but converges more slowly, while $\mathbf{X}_0$- and $\mathbf{w}$-prediction converge more slowly still and were not used for the reported metrics.

B SDE formulation

B.I Forward and reverse SDEs

Following the score-matching formulation of [18], we write a forward stochastic differential equation (SDE) and its time-reverse counterpart as

$$\text{forward: } \mathrm{d}\mathbf{X} = \mathbf{f}(\mathbf{X}, t)\, \mathrm{d}t + \mathbf{g}(t) \odot \mathrm{d}\mathbf{w}, \tag{B.24a}$$

$$\text{reverse: } \mathrm{d}\mathbf{X} = \left[\mathbf{f}(\mathbf{X}, t) - \mathbf{g}^2(t) \odot \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\right] \mathrm{d}t + \mathbf{g}(t) \odot \mathrm{d}\bar{\mathbf{w}}. \tag{B.24b}$$

The whole process is computed in frequency space. We use the main text notation: $\mathbf{X}$ for DCT coefficients, $\mathbf{S}_0$ for the empirical spectrum, and $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ for normalized Gaussian noise. We suppress explicit $\mathbf{k}$ arguments where no ambiguity arises.

The continuous-time SDE can be obtained either by differentiating Eq. (3) or, more compactly, by starting from a linear SDE ansatz, integrating it exactly, and matching the resulting mean and variance to the desired forward marginal. We follow the second route, which determines $\mathbf{f}(\mathbf{X}, t)$ and $\mathbf{g}(t)$ mode by mode.

Since $\mathbf{X}$ is evolved under a Gaussian filter in Eq. (3), the SDE ansatz should be linear in $\mathbf{X}$. For each mode $\mathbf{k}$, write

$$\mathrm{d}x = -a(t)\, x\, \mathrm{d}t + b(t)\, \mathrm{d}w. \tag{B.25}$$

Define the integrating factor $m(t) \equiv \exp\Big(\int_0^t a(u)\, \mathrm{d}u\Big)$, so that $\mathrm{d}m = a(t)m(t)\, \mathrm{d}t$. The Itô product rule gives

$$\mathrm{d}(mx) = m\, \mathrm{d}x + x\, \mathrm{d}m + \mathrm{d}[m, x] = m\, b(t)\, \mathrm{d}w, \tag{B.26}$$

since m has finite variation and so $\mathrm{d}[m, x] = 0$. Integrating and dividing by $m(t)$, with $m(0) = 1$,

$$x_t = \underbrace{e^{-\int_0^t a(u)\, \mathrm{d}u}\, x_0}_{\text{signal}} + \underbrace{\int_0^t e^{-\int_s^t a(u)\, \mathrm{d}u}\, b(s)\, \mathrm{d}w_s}_{\text{noise}}. \tag{B.27}$$

Matching $\mathbb{E}[x_t]$ in this expression to the forward marginal Eq. (3) gives

$$e^{-\int_0^t a(u)\, \mathrm{d}u}\, \mathbb{E}[x_0] \overset{!}{=} e^{-k^2 \lambda_t / 2}\, \mathbb{E}[x_0] \quad \Longrightarrow \quad a(t) = \tfrac{1}{2}\, k^2\, \dot{\lambda}(t). \tag{B.28}$$

Matching $\mathrm{Var}[x_t]$ requires $\mathrm{Var}[x_t] = S_0$ at every t when $\mathrm{Var}[x_0] = S_0$, the dataset spectrum:

$$e^{-2\int_0^t a\, \mathrm{d}u}\, S_0 + \int_0^t e^{-2\int_s^t a\, \mathrm{d}u}\, b^2(s)\, \mathrm{d}s \overset{!}{=} S_0. \tag{B.29}$$

Differentiating with respect to t and substituting,

$$-k^2 \dot{\lambda}_t\, S_0\, (1 - e^{-k^2 \lambda_t}) + b^2(t) = S_0\, k^2\, \dot{\lambda}_t\, e^{-k^2 \lambda_t} \quad \Longrightarrow \quad b^2(t) = k^2\, \dot{\lambda}_t\, S_0. \tag{B.30}$$

Restoring all $\mathbf{k}$ modes, we obtain

$$\mathbf{f}(\mathbf{X}, t) = -\tfrac{1}{2}\, \mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{X}_t, \tag{B.31a}$$

$$\mathbf{g}(t) = \sqrt{\mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{S}_0}, \tag{B.31b}$$

and the forward and reverse SDEs read

$$\text{forward: } \mathrm{d}\mathbf{X}_t = -\tfrac{1}{2}\, \mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{X}_t\, \mathrm{d}t + \sqrt{\mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{S}_0} \odot \mathrm{d}\mathbf{w}, \tag{B.32a}$$

$$\text{reverse: } \mathrm{d}\mathbf{X}_t = \Big[-\tfrac{1}{2}\, \mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{X}_t - \mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{S}_0 \odot \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\Big]\, \mathrm{d}t + \sqrt{\mathbf{k}^2\, \dot{\lambda}(t)\, \mathbf{S}_0} \odot \mathrm{d}\bar{\mathbf{w}}. \tag{B.32b}$$

B.II Score function and training objectives

To determine the training target for a score network, we compute the marginal score. Let $\mathbf{A}_t = e^{-\mathbf{k}^2 \lambda(t)/2}$ and $\mathbf{D}_t = \sqrt{\mathbf{S}_0(1 - e^{-\mathbf{k}^2 \lambda(t)})}$ denote the signal and noise prefactors. The forward conditional is $\mathcal{N}(\mathbf{X};\, \mathbf{A}_t \odot \mathbf{X}_0,\, \mathbf{D}_t^2)$, and the marginal is

$$p_t(\mathbf{X}) = \int p_0(\mathbf{X}_0)\, \mathcal{N}(\mathbf{X};\, \mathbf{A}_t \odot \mathbf{X}_0,\, \mathbf{D}_t^2)\, \mathrm{d}\mathbf{X}_0. \tag{B.33}$$

Differentiating under the integral and dividing by $p_t(\mathbf{X})$ gives the score

$$\nabla_{\mathbf{X}} \log p_t(\mathbf{X}) = \mathbb{E}\big[\mathbf{D}_t^{-2} \odot (\mathbf{A}_t \odot \mathbf{X}_0 - \mathbf{X}) \,\big|\, \mathbf{X}\big] = \mathbf{D}_t^{-2} \odot (\mathbf{A}_t \odot \mathbb{E}[\mathbf{X}_0 \mid \mathbf{X}] - \mathbf{X}), \tag{B.34}$$

which gives the Tweedie identity

$$\mathbb{E}[\mathbf{X}_0 \mid \mathbf{X}] = \mathbf{A}_t^{-1} \odot \big(\mathbf{X} + \mathbf{D}_t^2 \odot \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\big). \tag{B.35}$$

The standard score-matching objective is

$$\mathcal{L} = \mathbb{E}_{t,\mathbf{X}_0,\epsilon}\big[\boldsymbol{\omega}(t) \odot \|\mathbf{s}_\theta(\mathbf{X},t) - \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\|_2^2\big] = \mathbb{E}_{t,\mathbf{X}_0,\epsilon}\big[\boldsymbol{\omega}(t) \odot \|\mathbf{s}_\theta(\mathbf{X},t) + \mathbf{D}_t^{-1} \odot \epsilon_t\|_2^2\big], \tag{B.36}$$

where $\mathbf{s}_\theta$ is the score network and $\boldsymbol{\omega}(t)$ is a per-mode weighting function.[4] The regression target $-\mathbf{D}_t^{-1} \odot \epsilon_t$ is the anisotropic analogue of the variance-preserving (VP) diffusion target. Without a low-frequency cutoff, the DC mode $\mathbf{k} = 0$ would be unchanged because neither drift nor noise act on it; in practice we replace $\|\mathbf{k}\|$ by $\max(\|\mathbf{k}\|, k_c)$ for small modes, so the same equations apply with a nonzero low-frequency cutoff.

The choice of $\boldsymbol{\omega}(t)$ determines the prediction target. Three natural choices recover the targets discussed in Appendix A.

ϵ-net weight. Setting $\boldsymbol{\omega}(t) \propto \mathbf{D}_t^2 = \mathbf{S}_0(1 - e^{-\mathbf{k}^2\lambda(t)})$ and writing $\mathbf{s}_\theta = -\mathbf{D}_t^{-1} \odot \epsilon_\theta$, the loss reduces to

$$\mathcal{L} = \mathbb{E}_{t,\mathbf{X}_0,\epsilon}\big[\|\epsilon_\theta(\mathbf{X},t) - \epsilon_t\|_2^2\big]. \tag{B.37}$$

$\mathbf{X}_0$-net weight. Setting $\boldsymbol{\omega}(t) \propto \mathbf{A}_t^{-2} \odot \mathbf{D}_t^4$ and writing $\mathbf{s}_\theta = \mathbf{D}_t^{-2} \odot (\mathbf{A}_t \odot \mathbf{X}_{0,\theta} - \mathbf{X})$ via Tweedie, the loss reduces to

$$\mathcal{L} = \mathbb{E}_{t,\mathbf{X}_0,\epsilon}\big[\|\mathbf{X}_{0,\theta}(\mathbf{X},t) - \mathbf{X}_0\|_2^2\big]. \tag{B.38}$$

KL-minimizing weight. Setting $\boldsymbol{\omega}(t) \propto \mathbf{g}^2(t)$ minimizes the Kullback-Leibler (KL) divergence between the true and learned reverse dynamics over the whole path: a score error $\mathbf{e}(\mathbf{k})$ produces drift error $\mathbf{g}(t) \odot \mathbf{e}(\mathbf{k})$, and for SDEs sharing the same diffusion the path-KL scales as $\frac{1}{2}\int \mathbf{g}^2(t) \odot \mathbb{E}[\mathbf{e}(\mathbf{k})^2]\,\mathrm{d}t$. Modes and times where the reverse drift amplifies score errors then receive proportionally more weight during training.

B.III Samplers

The standard score-matching samplers [18] include the Euler-Maruyama (EM) sampler, the Ordinary Differential Equation (ODE) sampler from probability flows, and the predictor-corrector (PC) sampler. Our experiments use ancestral DDPM sampling (Eq. (A.14)); the SDE samplers below are reference forms.

EM sampler. On a decreasing time grid $1 = t_N > t_{N-1} > \cdots > t_0 = 0$ with $\Delta t_n = t_{n-1} - t_n < 0$, discretizing Eq. (B.32b) gives

$$\mathbf{X}_{n-1} = \mathbf{X}_n + \Big[-\tfrac{1}{2}\,\mathbf{k}^2\,\dot\lambda_n\,\mathbf{X}_n - \mathbf{k}^2\,\dot\lambda_n\,\mathbf{S}_0 \odot \mathbf{s}_\theta(\mathbf{X}_n, t_n)\Big]\,\Delta t_n + \sqrt{\mathbf{k}^2\,\dot\lambda_n\,\mathbf{S}_0} \odot \sqrt{-\Delta t_n}\,\epsilon_n. \tag{B.39}$$

ODE sampler. For any SDE of the form Eq. (B.24a), the probability flow ODE

$$\mathrm{d}\mathbf{X} = \big[\mathbf{f}(\mathbf{X},t) - \mathbf{g}^2(t) \odot \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\big]\,\mathrm{d}t \tag{B.40}$$

matches the SDE marginals $p_t(\mathbf{X})$. Solving in reverse time yields the same target distribution as the reverse SDE, with the Brownian term dropped:

$$\mathbf{X}_{n-1} = \mathbf{X}_n + \Big[-\tfrac{1}{2}\,\mathbf{k}^2\,\dot\lambda_n\,\mathbf{X}_n - \mathbf{k}^2\,\dot\lambda_n\,\mathbf{S}_0 \odot \mathbf{s}_\theta(\mathbf{X}_n, t_n)\Big]\,\Delta t_n. \tag{B.41}$$

PC sampler. The PC sampler refines an EM prediction $\mathbf{X}_{n-1}^p$ via a few iterations of a score-based MCMC corrector at the same time level. A preconditioned Langevin update that leaves $p_t(\mathbf{X})$ invariant is

$$\mathrm{d}\mathbf{X}_\tau = \mathbf{B}(t) \odot \nabla_{\mathbf{X}} \log p_t(\mathbf{X})\,\mathrm{d}\tau + \sqrt{2\,\mathbf{B}(t)} \odot \mathrm{d}\mathbf{w}_\tau, \tag{B.42}$$

which classically uses $\mathbf{B}(t) = \mathbf{I}$ but admits an anisotropic, time-dependent choice given the mode-dependent diffusion. Discretizing with step η_n gives the corrector update

$$\mathbf{X}_{i+1} = \mathbf{X}_i + \mathbf{B}(t_n) \odot \mathbf{s}_\theta(\mathbf{X}_n, t_n)\,\eta_n + \sqrt{2\,\mathbf{B}(t_n)\,\eta_n} \odot \epsilon_n, \tag{B.43}$$

and one full PC step combines the EM predictor with K Langevin corrections initialized at $\mathbf{X}_{i=0} = \mathbf{X}_{n-1}^p$. The corrector reduces the discretization error of the predictor.

[4] Strictly, the loss should reference $\nabla_{\mathbf{X}} \log p_t(\mathbf{X} \mid \mathbf{X}_0)$. The two objectives differ by an $\mathbf{s}_\theta$-independent constant via the Fisher identity $\mathbb{E}[\nabla_{\mathbf{X}} \log p_t(\mathbf{X} \mid \mathbf{X}_0) \mid \mathbf{X}] = \nabla_{\mathbf{X}} \log p_t(\mathbf{X})$, so they share the same minimizer. This is the basis of denoising score matching.

C Scale invariance and power laws in nature and physics

C.I Dictionary between physics and natural images

We provide a brief analogy between field theory and natural images that motivates treating them in the same framework.

A Monte Carlo sample of a scalar field theory can be encoded directly as an image: each pixel at lattice site $\mathbf{r}$ stores the value of the field variable at that site. The full distribution over field configurations is then represented by the dataset of all such images. Conversely, a natural image can be viewed as a single sample from a three-channel scalar field theory, with each pixel storing a field value, the resolution playing the role of the inverse lattice spacing, and the image size playing the role of the physical system size. The infinite-resolution limit corresponds to the continuum limit of the field, and the infinite-size limit corresponds to the thermodynamic limit. The basic dictionary is

- pixel value $\leftrightarrow$ field variable;
- image resolution $\leftrightarrow$ inverse lattice spacing (continuum limit at infinite resolution);
- image size $\leftrightarrow$ system size (thermodynamic limit at infinite size).

A critical field theory is scale-invariant under renormalization-group coarse-graining and rescaling. The corresponding question for images is whether a sufficiently large and high-resolution natural-image dataset is approximately scale-invariant under the analogous operation of low-pass filtering followed by zooming in. The variance spectra in Section 3 indicate that, in expectation across a dataset, the answer is approximately yes.

C.II Power-law benchmark in frequency space

We specify the DCT, IDCT, and frequency normalization used throughout the paper, then summarize the per-dataset variance fits.

DCT and IDCT conventions. The forward transform is the type-II DCT and the inverse is the type-III IDCT. For an image $x \in \mathbb{R}^{H\times W}$,

$$X_{u,v} = \frac{4}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} x_{i,j} \cos\left[\tfrac{\pi}{H}\left(i+\tfrac{1}{2}\right) u\right] \cos\left[\tfrac{\pi}{W}\left(j+\tfrac{1}{2}\right) v\right], \tag{C.44}$$

with inverse

$$x_{i,j} = \sum_{u,v} \gamma_u \gamma_v X_{u,v} \cos\left[\tfrac{\pi}{H}\left(i+\tfrac{1}{2}\right) u\right] \cos\left[\tfrac{\pi}{W}\left(j+\tfrac{1}{2}\right) v\right], \tag{C.45}$$

where $\gamma_0 = 1/2$ and $\gamma_{k>0} = 1$.

Frequency normalization. We define the frequency vector by $\mathbf{k} = (\pi u, \pi v)$, so that mode indices are spaced uniformly by π and the magnitude $\|\mathbf{k}\|$ ranges in $[0, \sqrt{2}\,\pi(H-1)]$ for an $H \times H$ image. Two properties motivate this convention. First, the $4/(HW)$ amplitude factor in the DCT keeps the maximum of the radial power spectrum at the same order of magnitude across resolutions, as seen in Figure C.10 and main text Figure 2. Second, the spacing-fixed convention means that an image of higher resolution extends the spectrum to higher $\mathbf{k}$ rather than rescaling existing modes; in the infinite-resolution limit, the discrete grid fills out the continuum.

Variance fits. Figure C.10 complements the ImageNet-256 fit shown in main text Figure 2(b) with the corresponding fits for CIFAR-10 and ImageNet-128. All three recover an approximate k^{-2} decay across the shared frequency range; deviations at the highest frequencies reflect the finite-resolution cutoff and shrink as resolution grows. Table C.5 summarizes the fitted parameters.

D Supplementary materials for CIFAR-10 experiments

D.I Architecture and training

CIFAR-10 generation uses a score U-Net backbone from the NCSN++ family [18] with discrete DDPM positional embeddings and 8 residual blocks. We scan the base channel count from 128 to

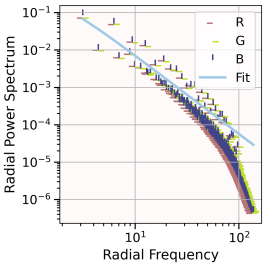

(a) CIFAR-10 variance

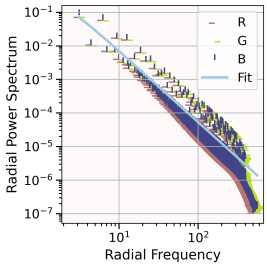

(b) ImageNet-128 variance

Figure C.10: Variance spectra for (a) CIFAR-10 and (b) ImageNet-128, computed independently for each color channel (RGB). Both power-law fits recover the k^{-2} frequency decay reported in previous work. The deviations at the highest frequencies in (a) reflect finite-resolution limitations and shrink in (b) as resolution increases.

Table C.5: Power-law fits of variance power spectra for the natural-image datasets.

Dataset	C	$\mathbf{k}_0^2$	a
CIFAR-10	0.9100 ± 0.0182	1.9406 ± 0.0481	1.0513 ± 0.0021
ImageNet-128	0.9281 ± 0.0016	1.5708 ± 0.0030	1.0590 ± 0.0002
ImageNet-256	0.9322 ± 0.0004	1.5708 ± 0.0008	1.0598 ± 0.0000

256 and report the best configuration. Models are trained with AdamW for 400K steps with batch size 128, learning rate 2×10^{-4}, weight decay 0, and EMA rate 0.999. Checkpoints are saved every 20K steps. Each CIFAR-10 model is trained on a single H100 or GH200 GPU.

D.II Evaluation protocol

Checkpoint selection. We evaluate every saved checkpoint on 50K generated samples. Following the standard convention from [18], we report the checkpoint with the lowest FID during training along with its corresponding IS, rather than averaging across checkpoints. For the run in Table 1, FID reaches its minimum at 160K steps. IS continues to climb at later checkpoints and hits 10.03 at 400K steps, while FID rises to 3.87. Appendix D.IV.1 shows the full training trajectory.

Sampling seeds. Table D.6 evaluates the 160K-step checkpoint from Table 1 across three random seeds, using 50K samples per seed. Seed 42 gives an FID of 2.65 and an IS of 9.63, matching the values noted in Appendix D.IV.1. Table 1 reports seed 44, which achieves the lowest FID. Averaged across the three seeds, the model reaches an FID of 2.632 ± 0.022 and an IS of 9.656 ± 0.033.

Table D.6: FID and IS for the CIFAR-10 checkpoint in Table 1 across three random seeds, using 50K samples per seed. Per-seed IS uncertainty shows the standard deviation across splits, and the bottom row reports the mean and standard deviation across the three seeds.

Seed	FID ↓	IS ↑
42	2.652	9.628 ± 0.117
43	2.637	9.692 ± 0.139
44	2.608	9.647 ± 0.093
Mean	2.632 ± 0.022	9.656 ± 0.033

D.III Noise schedules

We describe the two schedule families tested on CIFAR-10 and motivate their forms. The schedule $\lambda(t)$ controls how the damping cutoff in $\mathbf{k}$ moves with time: a frequency mode is effectively suppressed when $\mathbf{k}^2\lambda(t) \approx \theta$ for some threshold θ, so the damping front is $\mathbf{k}_d(t) = \sqrt{\theta/\lambda(t)}$.

Log-linear schedule. In the continuum limit ($\Delta\mathbf{k} \to 0$), a damping that is uniform on a logarithmic frequency scale respects scale invariance: equal time intervals remove equal logarithmic ranges of

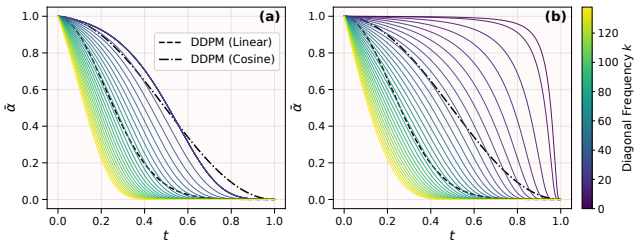

Figure D.11: Best (a) log-linear and (b) linear schedules for CIFAR-10 generation. Different diagonal frequency modes are attenuated independently over time, with higher-frequency modes decaying before lower-frequency ones. The linear and cosine DDPM schedules are shown for reference; they are scalar in frequency space and attenuate all modes uniformly.

modes. The corresponding schedule is

$$\lambda(t) = 10^{\lambda_i + (\lambda_f - \lambda_i)t}, \tag{D.46}$$

which gives $\mathbf{k}_d(t) \propto 10^{-(\lambda_f - \lambda_i)t/2}$, log-linear in t.

Linear schedule. Real datasets are at finite resolution and finite size, neither at the continuum limit nor the thermodynamic limit, so it is also reasonable to damp $\mathbf{k}$ linearly with time:

$$\lambda(t) = \frac{\theta}{(\lambda_f(1-t) + \lambda_i)^2}, \tag{D.47}$$

under which $\mathbf{k}_d(t)$ moves linearly with t.

Boundary fix. Both schedules above have $\lambda(0) > 0$ and $\dot{\lambda}(0) > 0$, so high-frequency modes are damped abruptly at $t = 0$. This discontinuity removes fine details before the model can learn them. We mitigate it by multiplying both schedules by t, ensuring $\lambda(0) = 0$ and a smooth onset. The schedules used in all experiments are

$$\text{log-linear:} \quad \lambda(t) = t \cdot 10^{\lambda_i + (\lambda_f - \lambda_i)t}, \tag{D.48}$$

$$\text{linear:} \quad \lambda(t) = \frac{\theta\, t}{(\lambda_f(1-t) + \lambda_i)^2}. \tag{D.49}$$

Best schedules on CIFAR-10. Figure D.11 shows the best-performing schedule from each family. The exact parameters of the best linear schedule, used to report the FID in the main text, are $\theta = 5.0$, $\lambda_i = 137.7294$, $\lambda_f = 1.57$, $k_c = 3$, and $N = 1000$. The best log-linear schedule uses $\lambda_i = -3.75$, $\lambda_f = -2.0$, $k_c = 31.2$, and $N = 1000$.

D.IV Ablation studies

D.IV.1 Mode collapse

In every CIFAR-10 run without the fixes in Table 2, we observe a clear mode collapse where FID reaches an early minimum and then rises while IS continues to improve. For the linear schedule of Figure D.11(b), FID converges to 2.65 at step 160K with IS 9.63, then degrades to 2.82 at 200K (IS 9.77), 2.98 at 250K (IS 9.79), 3.26 at 300K (IS 9.98), and 3.87 at 400K (IS 10.03). Main text Figure 4(b) plots these trajectories.

We attribute this to instability in the low-frequency portion of the diffusion. Late in the forward process, small changes in the surviving low-frequency modes produce large overall color shifts in the noised image, visible in Figures 1 and A.9. This makes it harder for the model to fix the global color tone and class-level structure of an image, and the consequence is most pronounced on object-centric datasets such as CIFAR-10. The pattern is consistent with image fidelity continuing to improve (IS rising) while the marginal class distribution drifts away from the dataset (FID degrading). The effect is strongest at large k_c, where many low-frequency modes are bundled together at late times.

The low-frequency limit is not strictly scale-invariant, so we expect this regime to need a different treatment. Plausible remedies include chaining the SKILD reverse process with a small VAE or with

Table D.7: **Schedule robustness on CIFAR-10.** Each row reports one frequency-space diffusion schedule. Metrics are shown for 200K and 400K training steps. Across the completed 400K evaluations, most schedules reach FID below or near 5, and all reach Inception Score near or above 9, indicating that SKILD is not tuned to a single fragile schedule. Convergence speed varies by schedule. A dash denotes an unused parameter or an unavailable run. The network used has half the channel count compared to the one in Table 1 and Figure 4(b).

Schedule parameters				FID ↓		IS ↑	
k_c	λ_i	λ_f	θ	200K	400K	200K	400K
Log-linear schedule							
3.0	−4.25	0.0	–	6.77	5.53	8.72	8.98
3.0	−4.25	0.4	–	6.80	5.63	8.63	8.84
3.0	−4.25	0.8	–	7.24	5.66	8.56	8.85
3.0	−3.75	0.0	–	5.65	4.47	8.77	9.07
3.0	−3.75	0.4	–	6.10	4.83	8.74	9.04
3.0	−3.75	0.8	–	6.56	5.17	8.73	8.99
3.0	−3.25	0.0	–	5.72	4.84	8.78	9.07
3.0	−3.25	0.4	–	6.02	4.77	8.77	9.06
3.0	−3.25	0.8	–	6.30	4.77	8.66	8.98
13.4	−4.25	−1.4	–	5.71	4.62	8.81	9.15
13.4	−4.25	−1.0	–	5.97	4.88	8.80	9.06
13.4	−4.25	−0.6	–	6.40	5.35	8.75	8.95
13.4	−3.75	−1.4	–	4.90	4.40	8.95	9.22
13.4	−3.75	−1.0	–	5.08	4.55	8.83	9.10
13.4	−3.75	−0.6	–	5.57	4.62	8.76	8.99
13.4	−3.25	−1.4	–	4.62	5.29	8.97	9.16
13.4	−3.25	−1.0	–	5.21	4.84	8.87	9.13
13.4	−3.25	−0.6	–	5.45	4.71	8.73	9.05
31.2	−4.25	−2.2	–	5.40	4.63	8.86	9.14
31.2	−4.25	−2.0	–	5.76	4.74	8.82	9.05
31.2	−4.25	−1.6	–	6.00	4.93	8.78	8.97
31.2	−3.75	−2.2	–	4.53	4.74	9.15	9.30
31.2	−3.75	−1.6	–	5.39	4.64	8.91	9.05
31.2	−3.25	−2.2	–	4.58	5.74	9.07	9.17
31.2	−3.25	−2.0	–	4.76	5.41	8.97	9.15
31.2	−3.25	−1.6	–	5.30	4.65	8.81	9.08
Linear schedule							
3.0	137.7294	1.57	3.0	5.42	4.41	9.06	9.39
3.0	137.7294	1.57	5.0	4.89	4.27	9.12	9.29
3.0	137.7294	1.57	7.0	6.75	5.33	9.39	9.55
13.4	137.7294	12.0	3.0	5.61	5.14	9.23	9.29
13.4	137.7294	12.0	5.0	4.61	4.99	9.05	9.52
13.4	137.7294	12.0	7.0	4.45	–	9.07	–
31.2	137.7294	29.5	3.0	7.63	7.44	9.18	9.41
31.2	137.7294	29.5	5.0	4.43	5.01	9.12	9.51
31.2	137.7294	29.5	7.0	4.76	5.96	9.02	9.25

a standard pixel-space diffusion model that handles the very-low-frequency modes. The effect should also weaken on higher-resolution and less object-centric natural-image datasets. Table 2 in the main text tests both this VAE approach and a class-conditional model with classifier-free guidance. Both methods keep training stable through 400K steps.

D.IV.2 Schedule robustness

Diffusion models are sensitive to noise schedules [22, 19]. Table D.7 reports a sweep over λ_i, λ_f, k_c, and θ in both the log-linear and linear families. Most schedules reach FID near or below 5 within 400K training steps, and all reach IS at or above 9. Convergence speed depends on the schedule, but the final converged scores are not strongly schedule-dependent, indicating that SKILD does not require a finely tuned schedule for stable training.

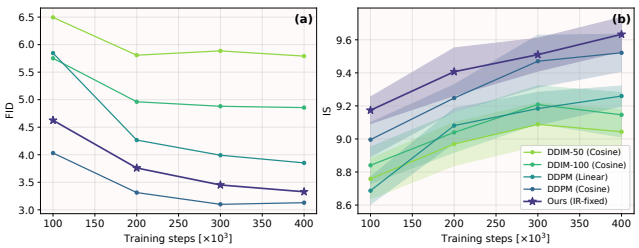

Figure D.12: FID and IS convergence comparison across schedules. (a) The FID of our low-frequency-fixed schedule converges stably over 400K training steps, outperforming the linear DDPM and two DDIM schedules. (b) The IS of our schedule leads the DDPM and DDIM schedules from 100K to 400K training steps.

D.IV.3 Convergence against pixel-space schedules

Figure D.12 compares one of our linear schedules against standard pixel-space DDPM and DDIM schedules. To isolate the effect of the schedule from the mode-collapse instability discussed above, we used a low-frequency-stable variant with $\theta = 2.0$, $\lambda_i = 137.7294$, and $\lambda_f = 22.3$, and used the forward marginal of the ground truth as the sampling input for diagnostics. The frequency-space schedule converges stably over 400K steps, with FID below the linear DDPM and the two DDIM baselines and IS above all pixel-space schedules.

D.IV.4 Additional ablations

Timestep sampler. The second-moment timestep sampler of [20] reduced the training loss faster but did not improve the final FID or IS, consistent with low-frequency recovery being the limiting factor.

Number of diffusion steps. Reducing the number of diffusion steps from 1000 to 500 left FID and IS nearly unchanged. A naive 100-step reduction degraded both metrics. Faster sampling is therefore feasible, but requires solver or distillation work tailored to the mode-dependent schedule.

Numerical cutoffs. We tested numerical cutoffs of $\boldsymbol{\alpha}_n$ in $\{5 \times 10^{-1}, 10^{-4}, 10^{-8}\}$. The smaller cutoffs gave similar convergence with 10^{-6}; 5×10^{-1} degraded sample quality. This is consistent with the low-frequency modes needing room to vary without being over-constrained.

Prediction target. Among the four prediction targets defined in Appendix A, ϵ-prediction (used for the reported metrics) converged fastest. Prediction of the full colored noise $\sqrt{1 - \bar{\boldsymbol{\alpha}}_n} \odot \sqrt{\mathbf{S}_0}\, \epsilon_n$ matched the same FID and IS but required more training steps, while $\mathbf{X}_0$- and $\mathbf{w}$-prediction converged slower still.

D.V Generation and super-resolution from one checkpoint

To test both generation and super-resolution with a single model, we take a 400K-step CIFAR-10 generation checkpoint and run super-resolution from forward states at effective resolutions of 2×2, 4×4, and 6×6. This checkpoint uses the linear schedule from Table D.7 with $k_c = 13.4$, $\lambda_i = 137.7294$, $\lambda_f = 12$, and $\theta = 5$. Table D.8 shows that FID improves steadily as the starting state retains more low-frequency signal to begin with.

Table D.8: FID for unconditional generation and super-resolution to 32×32 using a single CIFAR-10 checkpoint, starting from forward states at different effective resolutions.

	Generation	$2 \to 32$ (16×)	$4 \to 32$ (8×)	$6 \to 32$ (5.3×)
FID ↓	4.99	4.56	4.15	3.70

E ImageNet super-resolution experiment supplements

Architecture and training. ImageNet super-resolution uses a score U-Net backbone from the NCSN++ family [18] with 6 residual blocks per resolution, attention at resolutions 32, 16, and 8,

Table E.9: Agreement between forward-diffusion effective-resolution signals and bicubic down-up images across SNR thresholds. MSE is reported in units of 10^{-4} and PSNR in dB. At SNR $= 0.1$, the MSE for all degradation pipelines is within $O(10^{-4})$ and the PSNR exceeds 30.

	SNR=1		SNR=0.5		SNR=0.1		SNR=0.05		SNR=0.01		SNR=0.005	
Dataset	MSE	PSNR	MSE	PSNR	MSE	PSNR	MSE	PSNR	MSE	PSNR	MSE	PSNR
$4\times$ ImageNet-256	15.9	27.99	10.5	29.79	4.48	33.49	3.84	34.15	4.53	33.44	5.33	32.73
$4\times$ ImageNet-128	21.5	26.67	14.0	28.53	4.87	33.12	3.55	34.50	3.51	34.55	4.23	33.74
$8\times$ ImageNet-128	29.8	25.26	19.7	27.05	7.09	31.49	5.15	32.88	4.87	33.12	5.78	32.38

and channel multipliers $(1, 1, 2, 2, 2, 2)$. We train with AdamW for up to 500K steps with batch size 256, learning rate 10^{-4}, weight decay 10^{-5}, and EMA rate 0.9999, and report metrics from the last checkpoint. The ImageNet-256 model is trained on 8 H100 GPUs; the ImageNet-128 models are trained on a single H100 or GH200 GPU.

Effective-resolution validation. To compare effective-resolution forward states with conventional image resizing, we generate a low-resolution reference by bicubic-downsampling the denormalized image and then upsampling it to the original resolution with bicubic interpolation (`antialias=true`). This bicubic reference validates the SNR-defined effective-resolution interpretation used for super-resolution experiments.

Table E.9 reports the per-dataset MSE and PSNR between the surviving signal at each SNR threshold and the matching bicubic-degraded image. At the threshold SNR $= 0.1$ used in all SR experiments, the MSE is within $O(10^{-4})$ for every degradation pipeline and the PSNR exceeds 30 dB, indicating that the SNR-defined low-resolution input closely matches a conventional $4\times$ or $8\times$ degradation while using the exact forward-process marginals.

Starting from conventional low-resolution images. Our super-resolution protocol starts from the exact forward marginal of the HR image, which carries nearly the same information as a conventional LR image. At the $\mathrm{SNR} = 0.1$ starting point, noise dominates every frequency mode above the effective-resolution cutoff. The noise power exceeds the signal by a factor of 10 at the cutoff and by roughly 10^{32} at the Nyquist frequency, leaving practically no usable high-frequency detail from the HR image. Moreover, the signal attenuation factor $e^{-\mathbf{k}^2\lambda(t)/2}$ acts as a standard Gaussian low-pass filter.

We also verify this directly by constructing a starting state from a standard downsampled LR image. We zero-pad the LR image in DCT space above the cutoff and add the matching schedule noise. Table E.10 compares this reconstructed state against the exact forward state for $4\times$ ImageNet-256 super-resolution ($64 \to 256$), averaged across the ImageNet training set. Across all downsamplers, the two states match within 32 to 36 dB in PSNR and exceed 0.98 in SSIM. Table E.9 provides the complementary comparison between the noise-free forward signal and bicubic down-up images. Pyramidal flow matching [69] applies a similar noise correction to keep different resolutions consistent within a single process.

Finally, we run end-to-end super-resolution from both starting states using a smaller model trained for $4\times$ super-resolution ($32 \to 128$). The resulting metrics in Table E.11 closely track one another, and neither starting state dominates across all metrics.

Super-resolution schedules. The super-resolution experiments use linear schedules with $N = 1000$ steps:

$$
\begin{aligned}
4\times \text{ ImageNet-256:} &\quad \theta = 9.0,\ \lambda_i = 1132.9352,\ \lambda_f = 550.8723,\ k_c = 0;\\
4\times \text{ ImageNet-128:} &\quad \theta = 9.0,\ \lambda_i = 564.2461,\ \lambda_f = 275.4361,\ k_c = 0;\\
8\times \text{ ImageNet-128:} &\quad \theta = 5.0,\ \lambda_i = 564.2461,\ \lambda_f = 102.6489,\ k_c = 0.
\end{aligned}
$$

Figure E.13 shows the corresponding effective-resolution paths.

Sampling cost. Table E.12 compares the step count and per-step wall-clock time of SKILD against diffusion super-resolution baselines, averaged over 20 runs on a single NVIDIA Titan RTX

Table E.10: Comparison between the exact forward starting state and a starting state built from a standard LR image for $4\times$ ImageNet-256 super-resolution ($64 \to 256$) at $\mathrm{SNR} = 0.1$. MSE is measured in DCT space, and all values are averaged over the ImageNet training set.

Downsampler	MSE	PSNR (dB)	SSIM
DCT truncation	6.24×10^{-8}	36.26	0.9902
Lanczos	8.10×10^{-8}	35.17	0.9883
Bicubic (anti-aliased)	9.70×10^{-8}	34.38	0.9865
Area	1.25×10^{-7}	33.24	0.9828
Bilinear (anti-aliased)	1.55×10^{-7}	32.34	0.9803

Table E.11: End-to-end $4\times$ super-resolution ($32 \to 128$) with a smaller model, starting from either the exact forward marginal or an anti-aliased bicubic LR image.

Starting state	MSE ↓	PSNR ↑	SSIM ↑	MS-SSIM ↑	CLIPIQA ↑	MUSIQ ↑
Exact forward marginal	0.006711	22.502	0.6356	0.8824	0.5699	36.338
Bicubic LR image	0.006697	22.419	0.6256	0.8790	0.5932	34.261

GPU. SinSR is distilled into a single step, while the remaining models use multi-step diffusion samplers. Because SR3 has no official code release, we benchmark an unofficial implementation ([https://github.com/novwaul/SR3]). SKILD has the fastest per-step runtime among all compared methods. Its overall sampling time comes from running all 1000 steps of the basic ancestral sampler across our experiments, as we did not optimize sampling speed in this work. Appendix B.III derives the Euler-Maruyama, probability-flow ODE, and predictor-corrector samplers for our process, providing the foundation for fast ODE solvers and step distillation [70]. Since better samplers improve both speed and image quality in standard diffusion [21], we expect similar gains for SKILD.

Table E.12: Sampling step counts and wall-clock time per step, reported as the mean and standard deviation across 20 runs on a single NVIDIA Titan RTX GPU.

	SinSR	ResShift	I2SB	SR3	SKILD
Steps (NFE)	1	15	1000	100	1000
Time per step (ms)	1496.8 ± 1.8	222.2 ± 1.9	252.1 ± 0.8	144.0 ± 0.03	102.8 ± 0.07

F Critical Ising super-resolution details

This appendix details the experiment in Section 5.5. We generate critical Ising configurations and evaluate connected four-point correlations across a range of scales rather than image-quality metrics.

F.I Data

We use the ferromagnetic two-dimensional Ising model on a square lattice,

$$P(s) \propto \exp\left(\beta \sum_{i,j \text{ being neighbors}} s_i s_j\right), \qquad s_i \in \{-1, +1\}, \tag{F.50}$$

at the exact critical inverse temperature

$$\beta_c = \frac{1}{T_c} = \frac{1}{2}\log(1+\sqrt{2}). \tag{F.51}$$

Configurations are sampled with the Wolff cluster algorithm [56], a Markov chain Monte Carlo method for spin systems that avoids the critical slowing down of local spin-flip dynamics. We use 128×128 binary spin fields with values in $\{-1, +1\}$ and periodic boundary conditions.

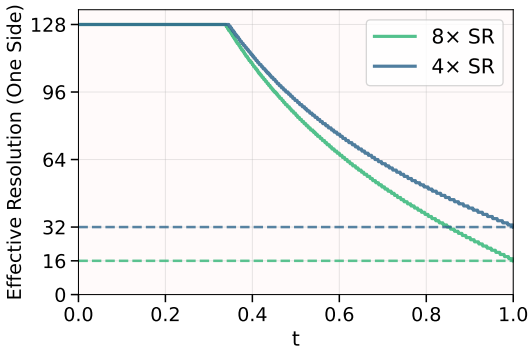

Figure E.13: Effective resolution against time for the ImageNet-128 super-resolution experiments. The $8\times$ experiment ends at an effective resolution of 16, and the $4\times$ experiment ends at 32. For both, the maximum resolution change in a single step is $\Delta R = 1$, so each model supports continuous super-resolution up to its respective final effective resolution.

One transition of the Wolff Markov chain constructs and flips a same-spin cluster. For our ferromagnetic model with coupling $J = 1$, the bond probability at criticality is

$$p = 1 - \exp(-2\beta_c).$$

A transition consists of the following steps:

1. Choose a seed site uniformly at random and let $s_\star$ be its spin. Initialize the cluster and the active frontier to contain only this seed site.
2. Given the current frontier, inspect nearest-neighbor sites with spin $s_\star$ that are not already in the cluster. If a candidate site touches $k \in \{1, 2, 3, 4\}$ frontier sites, add it to the cluster with probability $1 - (1 - p)^k$. This is equivalent to independently activating each bond from a frontier site with probability p.
3. Set the newly added sites as the next frontier and repeat the previous step until the frontier is empty.
4. Flip every spin in the completed cluster: $s_i \leftarrow -s_i$.

We run several independent chains in parallel, each initialized with independent random spins in $\{-1, +1\}^{128\times128}$. Every chain is thermalized for 2,000 Wolff transitions. After thermalization, we save the current configuration from each chain whenever all of them have accumulated at least $2L^2$ flipped spins since the last save. This is equivalent to two lattice sweeps of cluster updates between consecutive saves.

We use 90,000 samples for training and 1,000 held-out samples for testing. We do not use a separate validation set for model selection, since no hyperparameters are tuned on the Ising benchmark. The held-out samples are shared between SKILD and SR3, so all reported correlation differences reflect the models and not the test inputs.

F.II Ground-truth forward initialization

We initialize the reverse process from the exact forward marginal of a high-resolution spin field. Given $\sigma_0(\mathbf{r})$, we take its DCT to obtain $\widehat{\sigma}_0(\mathbf{k})$. For a chosen timestep n_0, the initialization is

$$\widehat{\sigma}_{n_0}(\mathbf{k}) = \sqrt{\bar{\boldsymbol{\alpha}}_{n_0}(\mathbf{k})} \odot \widehat{\sigma}_0(\mathbf{k}) + \sqrt{1 - \bar{\boldsymbol{\alpha}}_{n_0}(\mathbf{k})} \odot \sqrt{\mathbf{S}_0(\mathbf{k})}\, \epsilon(\mathbf{k}), \qquad \epsilon(\mathbf{k}) \sim \mathcal{N}(0, \mathbf{I}), \tag{F.52}$$

which matches the forward marginal of the frequency-space DDPM used during training. The schedule is chosen so that at $n_0 = 1000$ the low-frequency modes corresponding to roughly a 32×32 effective resolution remain largely intact, while higher modes are dominated by the noise term. Reverse sampling from this initialization produces a reconstruction conditioned on the low-frequency content of the paired held-out sample.

The diffusion uses $N = 1000$ steps with a linear schedule, $\theta = 9.0$, and $k_c = 0$. The variance spectrum $\mathbf{S}_0$ is fit by a power law,

$$\mathbf{S}_0(\mathbf{k}) = C(\mathbf{k}^2 + k_0^2)^{-a}, \qquad a = 0.811056, \quad C = 0.26641, \quad k_0^2 = 3.0, \tag{F.53}$$

with a, C, and k_0^2 fit on the Ising training set.

F.III Training hyperparameters

SKILD uses the same NCSN++ backbone as in the ImageNet SR experiments, with one input and output channel, base width 128, channel multipliers $(1, 1, 2, 2, 2, 2)$, six residual blocks per resolution, and attention at resolutions 32, 16, and 8. We train with ϵ-prediction using AdamW with learning rate 10^{-4}, $\beta_1 = 0.9$, weight decay 10^{-5}, 1,000 warmup steps, gradient clipping at 1.0, batch size 256, microbatch size 128, and EMA rate 0.9999. Training uses mixed precision and a uniform timestep sampler. We train for 100,000 optimization steps and use the EMA checkpoint at this final step for all reported results. The Ising model is trained on a single H100 or GH200 GPU.

F.IV Connected four-point correlation

For a fixed side length, each correlation estimate is the empirical mean over all square patches in each image, including all translations and symmetry-equivalent orientations, and is then averaged across images. Let $s_{00}, s_{01}, s_{10}, s_{11}$ denote the four corner spins of a patch. The full four-corner statistic is $G_4 = \mathbb{E}[\, s_{00}\, s_{01}\, s_{10}\, s_{11}\,]$, the edge two-point correlation is $C_a = \mathbb{E}[\, s_{00}\, s_{01}\,]$, and the diagonal two-point correlation is $C_b = \mathbb{E}[\, s_{00}\, s_{11}\,]$. The connected four-point correlation reported in the main text is

$$\kappa_4 = G_4 - 2C_a^2 - C_b^2, \tag{F.54}$$

the fourth-order joint cumulant of the four corner spins. It removes the contribution that pairwise correlations alone explain and tests whether the model reproduces the non-Gaussian critical structure of the field. We compute κ_4 at side lengths $\{1, 2, 4, 8, 16, 32, 64\}$.

F.V Paired evaluation

Each generated sample is paired with the held-out spin field that produced its low-frequency initialization. The left panel of Figure 8 shows paired ground-truth and reconstruction examples; the right panel reports the paired correlation comparison. The baseline is SR3 [46], a conditional diffusion model trained to upsample 32×32 low-resolution Ising fields to their 128×128 ground truth.

For uncertainty estimates we use a paired bootstrap over the 1,000 held-out samples: a single bootstrap index matrix is sampled and applied to all compared methods at every side length, preserving the sample-wise dependence induced by conditioning on the same low-frequency input. We use 1,000 bootstrap resamples and report 99% percentile confidence intervals for the curves in Figure 8.

G Additional image samples

This appendix collects additional super-resolution samples on ImageNet-128 and ImageNet-256 across the super-resolution factors reported in the main text, along with uncurated CIFAR-10 generation samples.

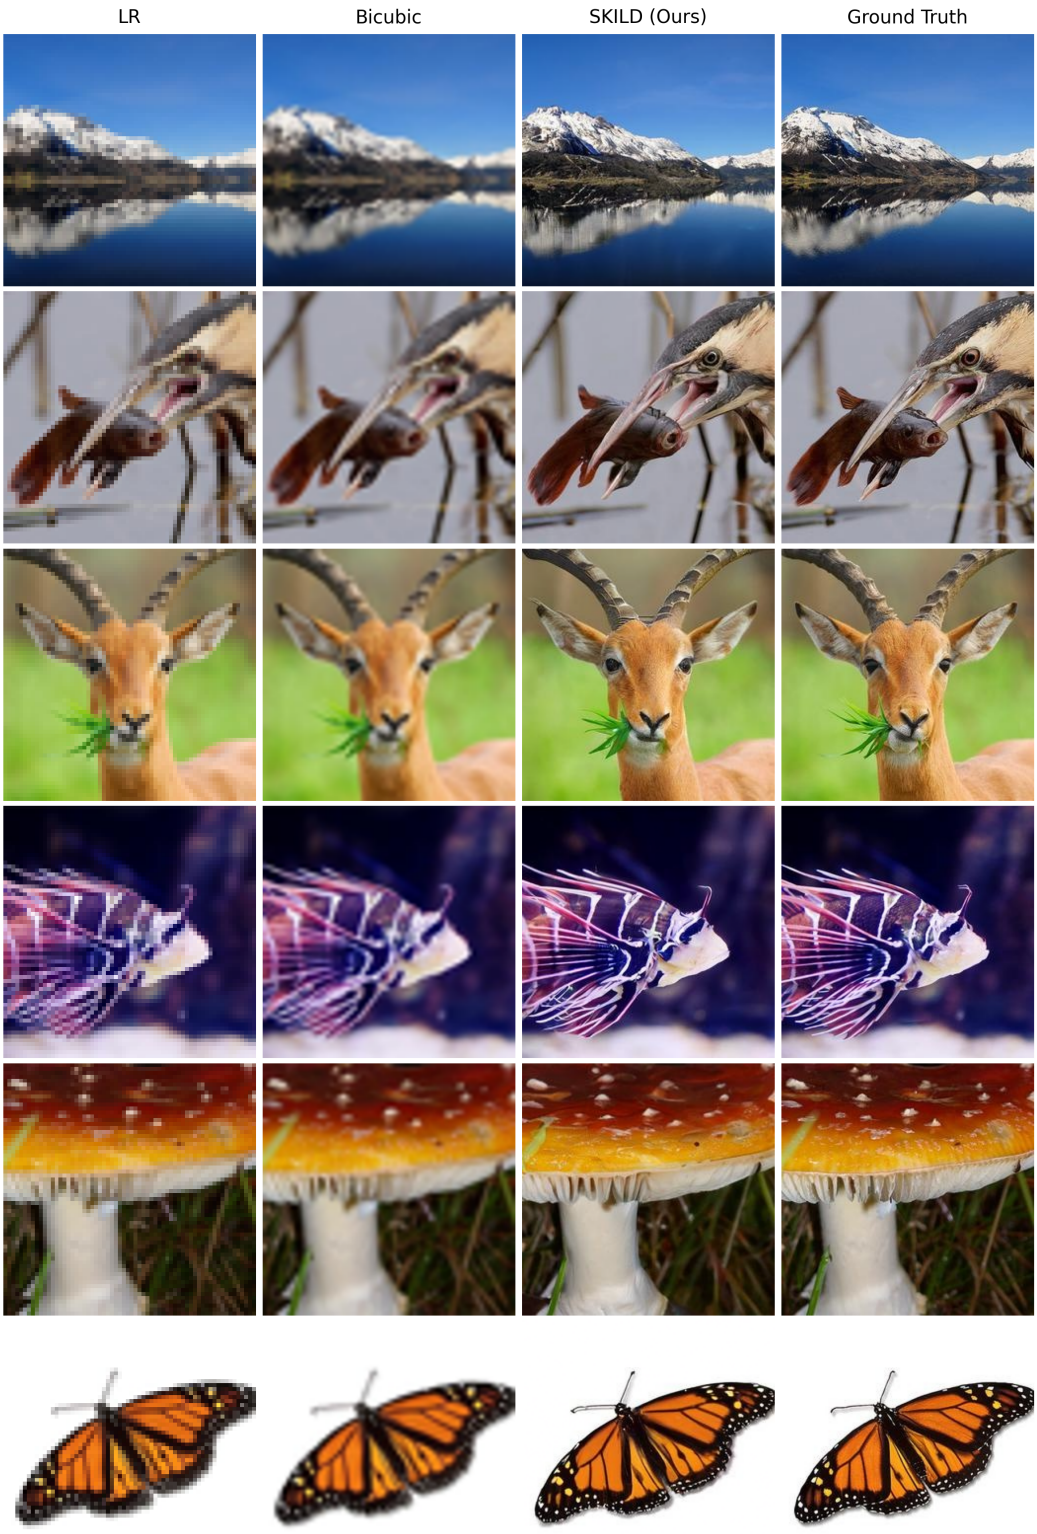

Figure G.14: Additional $4\times$ super-resolution sample comparisons on ImageNet-256.

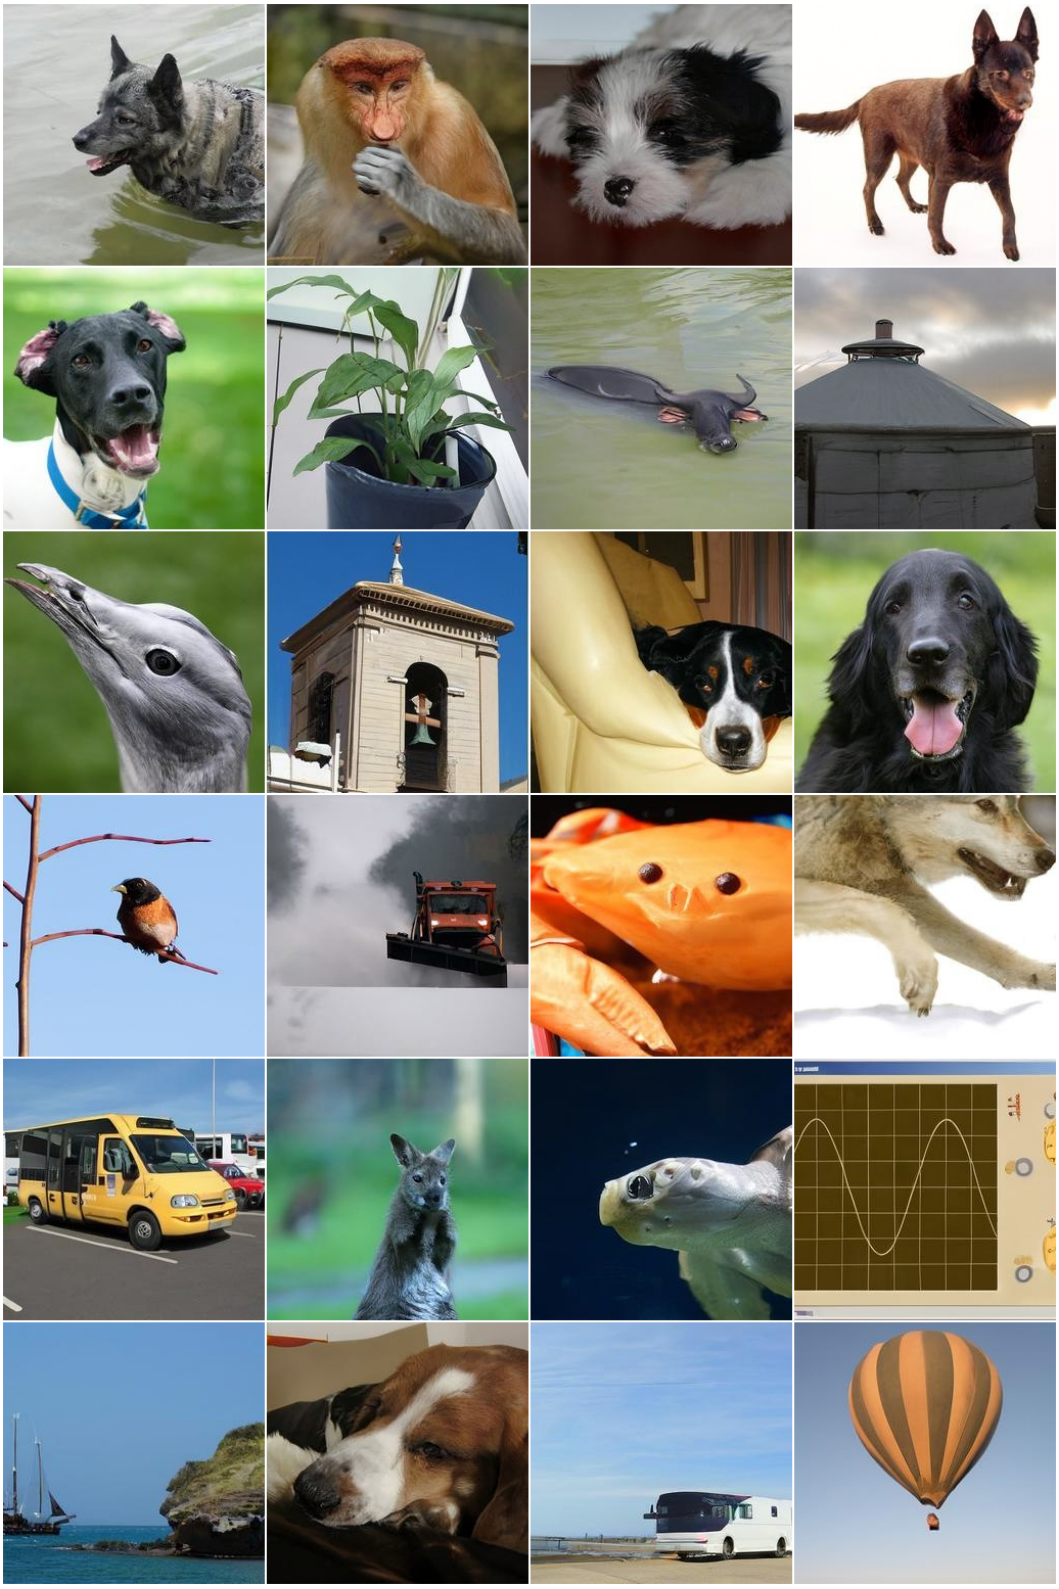

Figure G.15: Additional $4\times$ super-resolution samples on ImageNet-256.

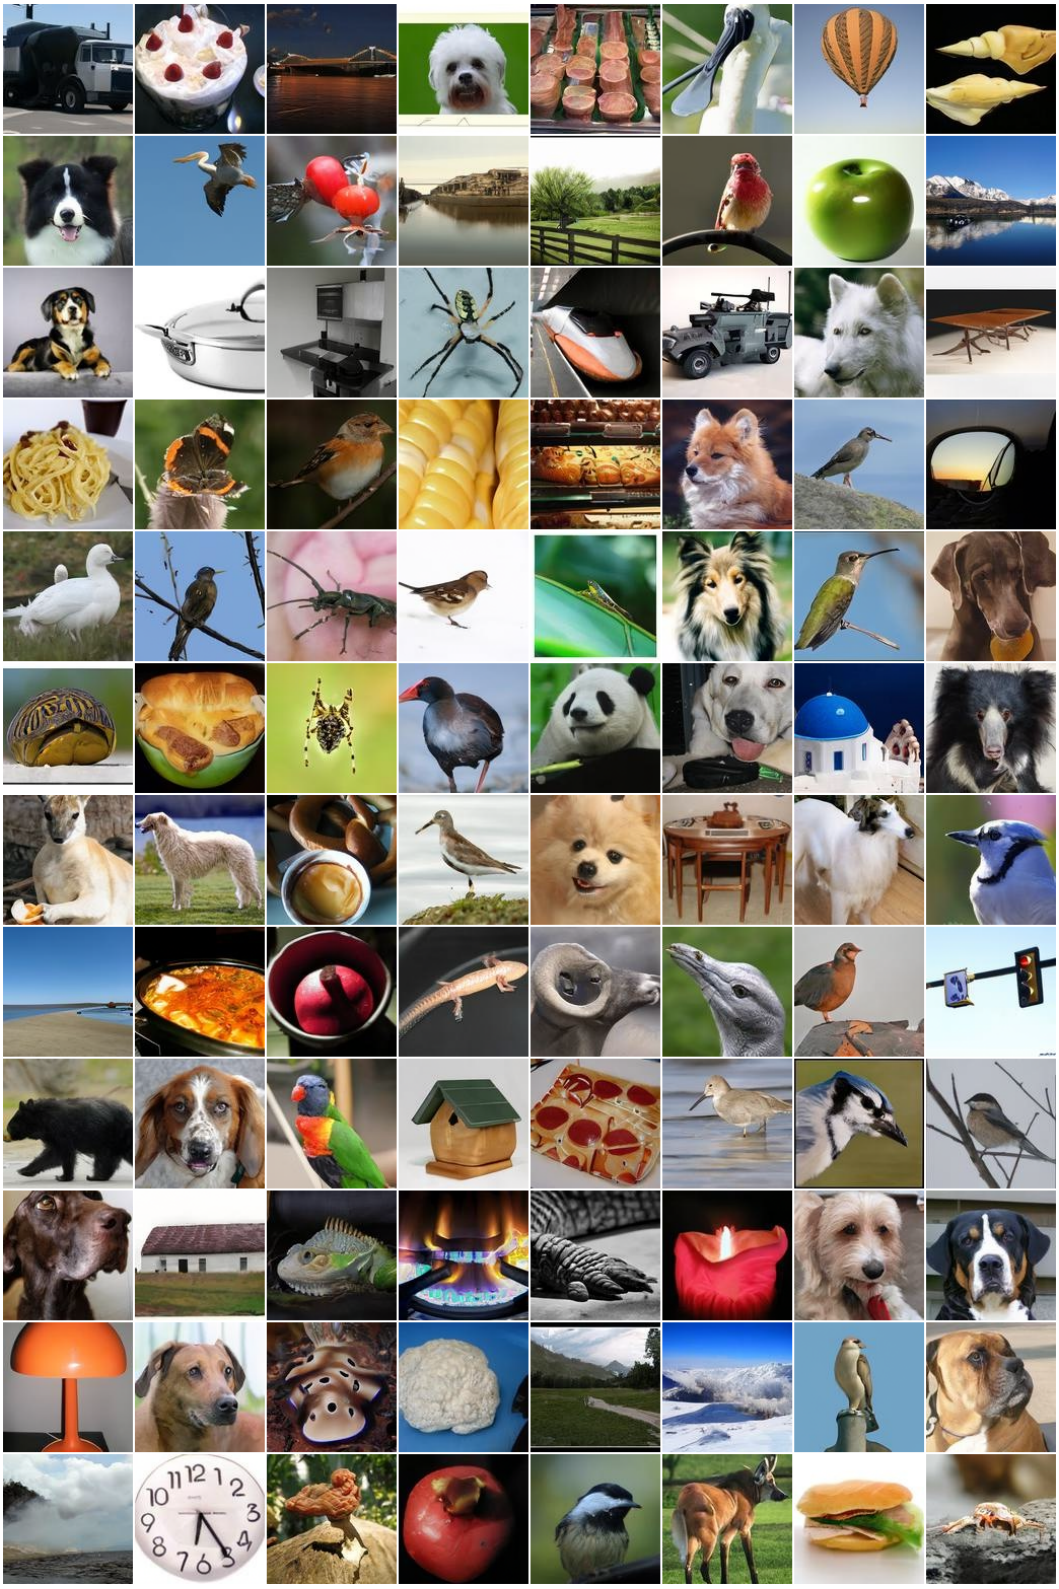

Figure G.16: Additional $4\times$ super-resolution samples on ImageNet-128.

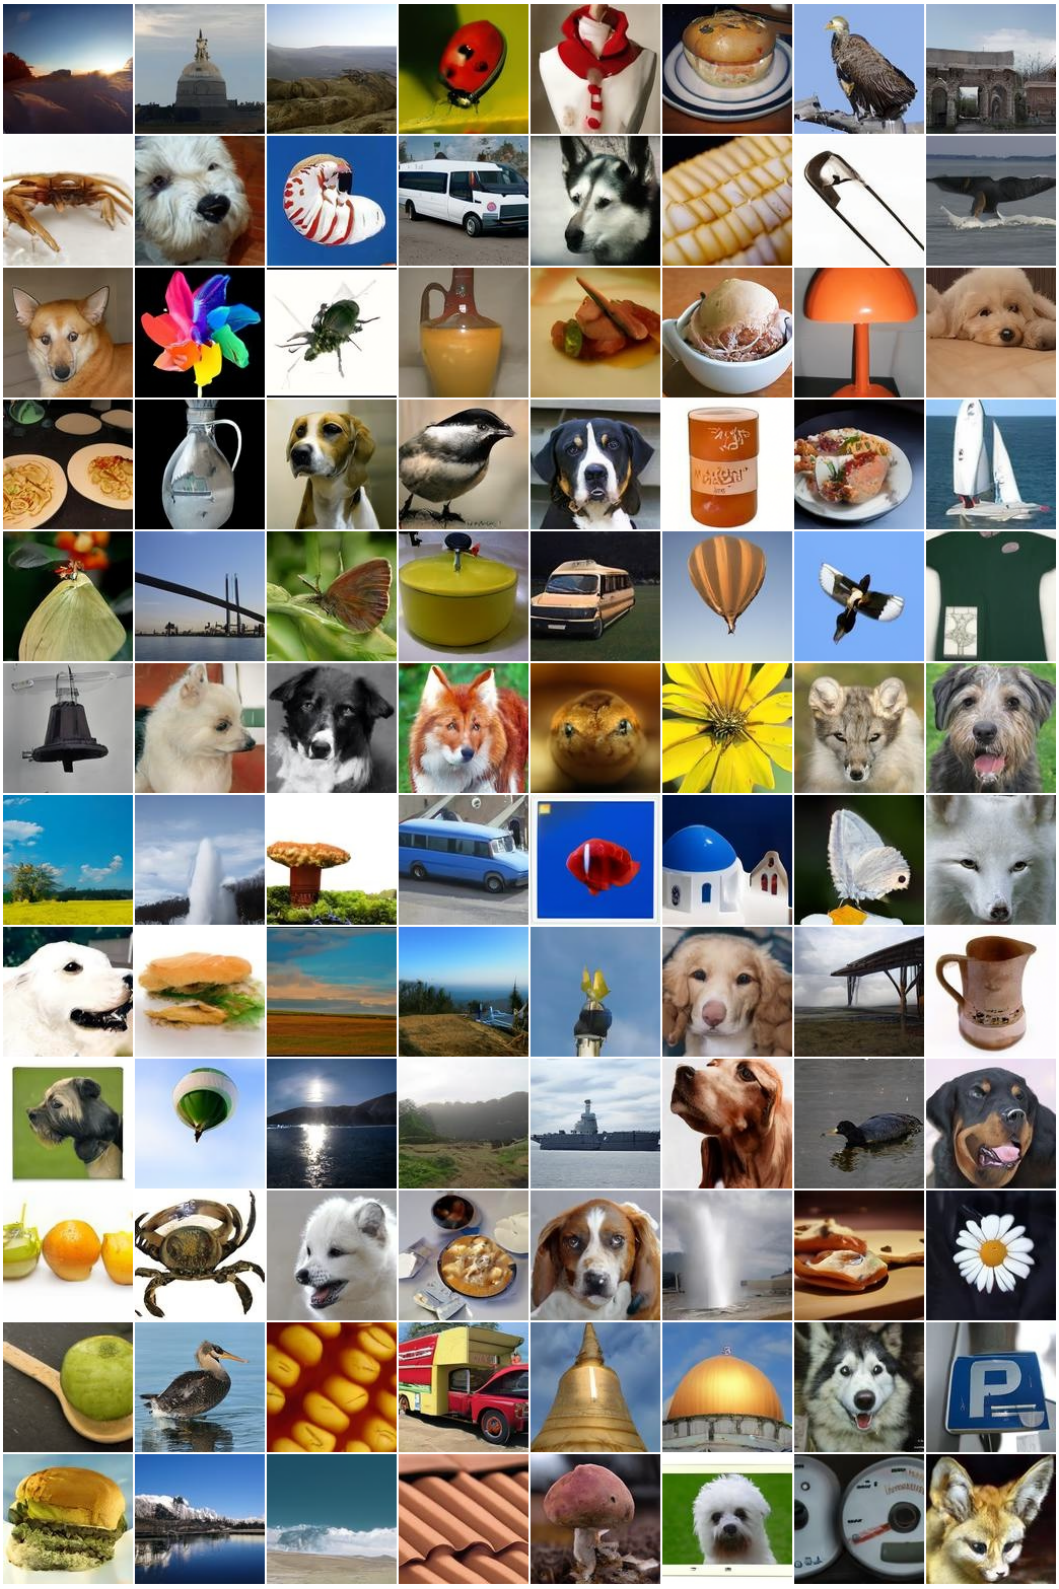

Figure G.17: Additional $8\times$ super-resolution samples on ImageNet-128.

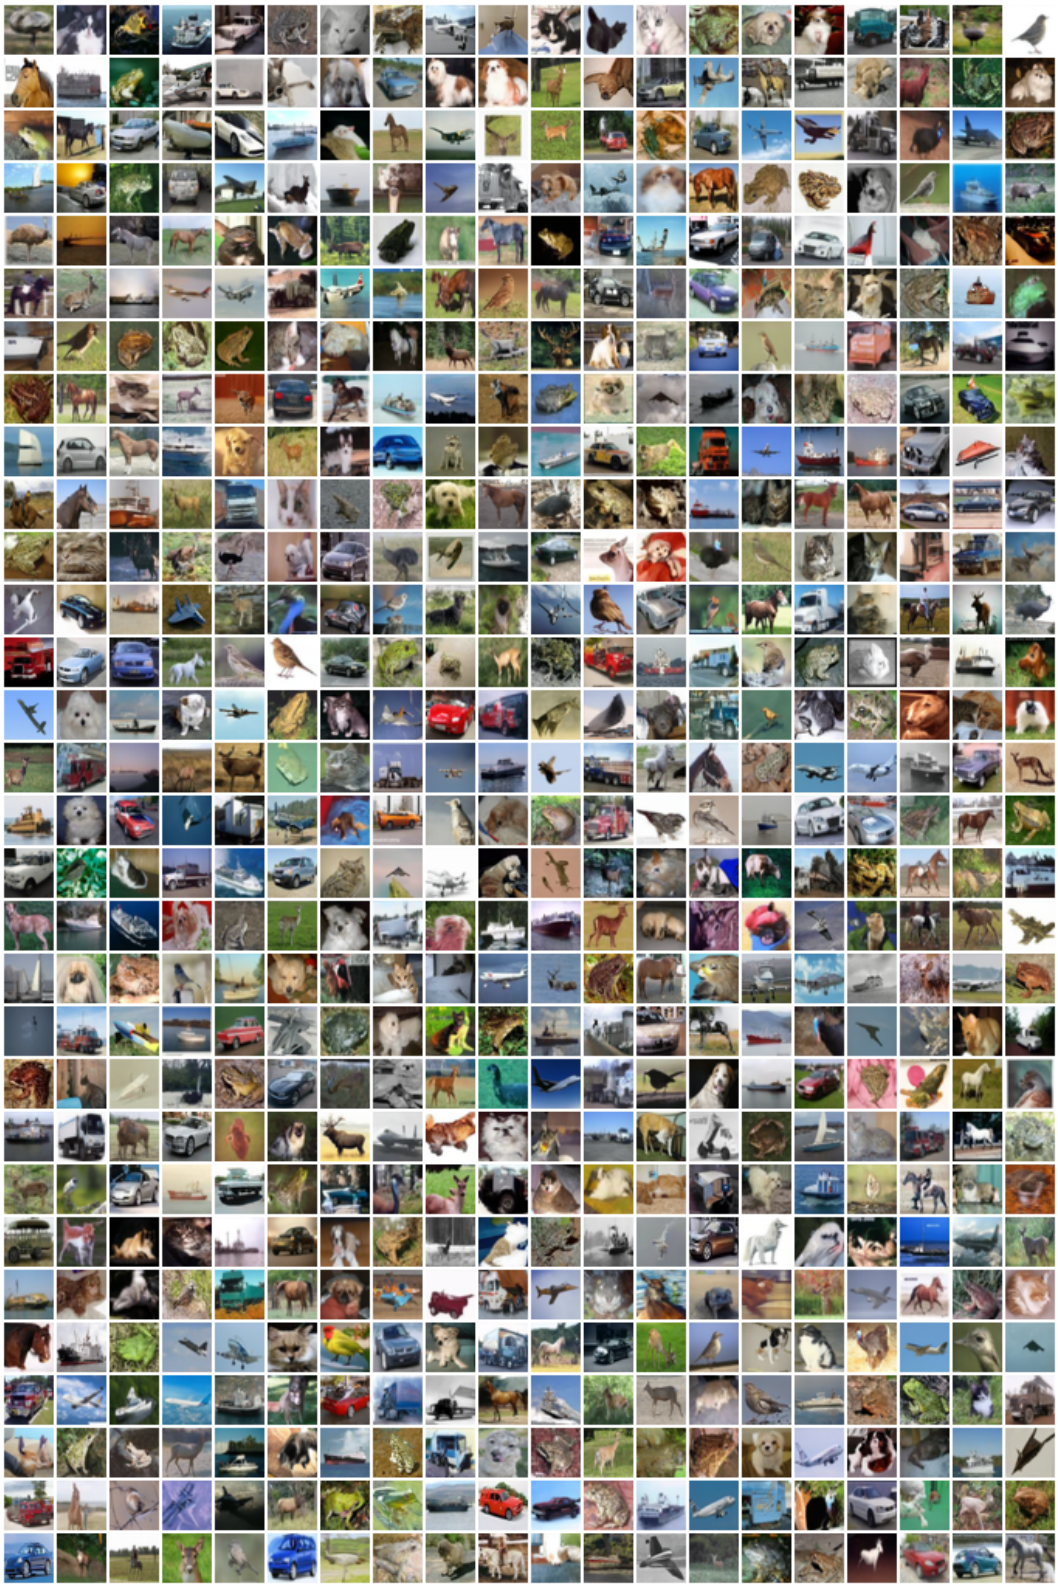

Figure G.18: Uncurated samples of generated images on CIFAR-10.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes].

   Justification: The abstract and introduction state the method, assumptions, benchmark scope, and main empirical claims.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes].

   Justification: The Limitations section discusses sampling cost, low-frequency recovery, and the use of ground-truth forward marginals for super-resolution.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes].

Justification: The main text gives the forward marginal and DDPM equations, and the appendix provides the SDE and posterior derivations.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes].

Justification: The paper reports datasets, model settings, schedules, evaluation protocols, and supplementary details needed to reproduce the main experiments. The code for reproducing the results are included in the supplementary material with the submission of this paper.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes].

   Justification: The code for reproducing the results are included in the supplementary material with the submission of this paper. We will further post the code on GitHub following the publicaiton of the paper.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes].

   Justification: The Experiments section and appendix specify model, optimizer, schedule, dataset, and evaluation details.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes].

   Justification: The scientific-data experiment reports bootstrap confidence intervals for paired correlation estimates; image-generation and ImageNet super-resolution metrics are reported using standard fixed-sample evaluations.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes].

   Justification: The Experiments section and appendices report GPU types and training scales for the main experiments.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes].

   Justification: The work uses standard public image datasets and synthetic scientific data generation, and does not involve human subjects or sensitive personal data.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes].

    Justification: The Limitations and broader-impact discussion cover benefits for multi-scale modeling and risks of plausible but incorrect super-resolution details.

Guidelines:

- The answer [N/A] means that there is no societal impact of the work performed.
- If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A].

    Justification: The paper does not release a high-risk foundation model or scraped dataset; any released code and scientific data will be accompanied by usage notes.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes].

    Justification: The paper cites the external datasets, baselines, and methods used in the experiments.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [Yes].

    Justification: The scientific benchmark is documented in the main text and appendix, including generation and evaluation details.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A].

    Justification: The work does not involve crowdsourcing or research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A].

    Justification: The work does not involve human subjects, so IRB or equivalent approval is not applicable.

    Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A].

    Justification: LLMs are not part of the core methodology, experiments, or scientific claims.

    Guidelines:

    - The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.