

Scale Space Diffusion

Soumik Mukhopadhyay*

soumik@umd.edu

Prateksha Udhayan*

pudhayan@umd.edu

Abhinav Shrivastava

abhinav2@umd.edu

University of Maryland, College Park

Abstract

Diffusion models degrade images through noise, and reversing this process reveals an information hierarchy across timesteps. Scale-space theory exhibits a similar hierarchy via low-pass filtering. We formalize this connection and show that highly noisy diffusion states contain no more information than small, downsampled images - raising the question of why they must be processed at full resolution. To address this, we fuse scale spaces into the diffusion process by formulating a family of diffusion models with generalized linear degradations and practical implementations. Using downsampling as the degradation yields our proposed Scale Space Diffusion. To support Scale Space Diffusion, we introduce Flexi-UNet, a UNet variant that performs resolution-preserving and resolution-increasing denoising using only the necessary parts of the network. We evaluate our framework on CelebA and ImageNet and analyze its scaling behavior across resolutions and network depths. Our website is available at prateksha.github.io/projects/scale-space-diffusion.

1. Introduction

Diffusion models [17, 37] are a class of generative models that achieve image synthesis by reversing an iterative noising process. It has been observed that states at different stages of the diffusion process encode different types of information [29]. As shown along the y-axis of Fig. 1(a), increasing diffusion noise progressively removes fine facial details while retaining only coarse structure. Eventually, with sufficient noising, even this structural information is lost. This illustrates that diffusion timesteps form an intrinsic information hierarchy.

A similar property underlies scale space theory [25], a fundamental subfield of computer vision. Scale spaces also represent image signals in an information-hierarchical manner through successive low-pass filtering. Along the x-axis of Fig. 1(a), we see the loss of details as the resolution de-

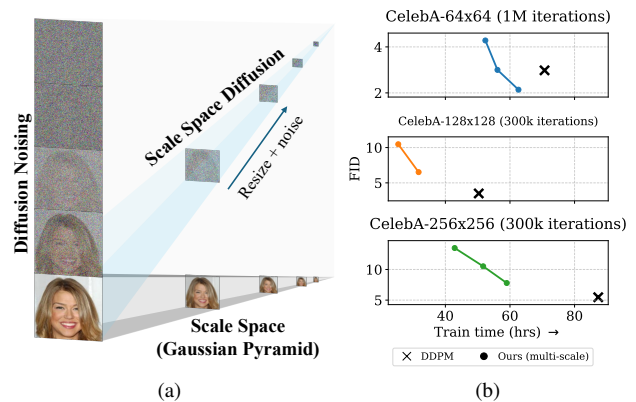


Figure 1. (a) Our proposed Scale Space Diffusion fuses scale spaces into diffusion models. (b) We show trends in image generation performance versus time for our proposed Flexi-UNet for CelebA-64, CelebA-128, and CelebA-256. Multiple point on the same plot represent our models with different number levels (*i.e.*, number of intermediate resolutions). We see immense gains in efficiency with resolution scaling while having reasonable performance.

creases in a Gaussian pyramid, mirroring the information dissipation in the diffusion process. The main distinction lies in the mechanism of information degradation: diffusion uses iterative noising, whereas scale spaces use progressive blurring or downsampling.

We investigate this relationship between diffusion and scale spaces formally through a preliminary mathematical modeling of information in both processes. This reveals striking parallels in their information content, suggesting a fundamental connection between the two. *Intuitively, one may ask why completely noisy images should be processed at high resolution when they contain information equivalent to that of a tiny image.* These parallels indicate that the two axes in Fig. 1(a) correspond to different but compatible ways of information degradation.

In this work, we revisit pixel diffusion to achieve a unification of scale spaces and the diffusion process. Previous attempts at this either operate only at the highest resolution [3, 8, 19], making them computationally inefficient, or rely on simplistic covariance assumptions [1] that may not hold in practice, or perform noisy scale shifting using high-frequency [2] or decorrelation noise [6, 20, 21], which

*Equal contribution.

remain inference-time approximations. Unlike pyramidal flow-matching approaches that approximate scale changes only during inference, our formulation integrates scale transitions directly into the diffusion process. In contrast, we first develop a mathematical theory for diffusion processes under generalized linear degradations, yielding a *family of diffusion processes*. We further illustrate how these can be implemented in modern deep-learning frameworks. Next, using image resizing as the linear degradation, we realize a fusion of scale spaces with diffusion. We term this process Scale Space Diffusion (SSD). Denoising diffusion probabilistic models (DDPM) [17] emerge as a special case of SSD, corresponding to the trivial case of resizing to the same size, *i.e.*, the identity operator. These generalized degradations naturally induce non-isotropic posteriors, which we handle through an implicit sampling procedure.

To realize the general version of Scale Space Diffusion, we require a neural network architecture capable of reversing the downsizing degradation, *i.e.*, it must be able to upsample a noisy state. A naïve approach could use a UNet [34] directly, but this would require even small-scale images to pass through the full network, leading to unnecessary computational cost. To address this, we propose a novel convolutional neural network (ConvNet) architecture that augments the standard UNet to use only the relevant levels of the network. It supports both resolution-preserving diffusion steps and next-resolution upscaling at all stages of a Gaussian pyramid. We denote this architecture as Flexi-UNet.

We analyze our framework and architectures on unconditional image generation using commonly used datasets of CelebA [26] and ImageNet [9]. To study the scaling properties of our method, we conduct experiments at multiple resolutions of CelebA dataset as shown in Fig. 1 (b). We observe that our models are faster during both training and inference while achieving reasonable FID scores. The key contributions of this work are:

1. We uncover and analyze the relationship between the states of diffusion models and the levels of scale spaces.
2. We build the mathematical foundation for a family of generalized linear diffusion processes, and techniques to implement them in modern deep-learning frameworks. With resizing as the choice for the linear degradation, we realize the fusion of diffusion and scale spaces, which we term Scale Space Diffusion.
3. To enable Scale Space Diffusion, we introduce a novel architecture Flexi-UNet capable of handling both resolution-changing as well as resolution-preserving reverse diffusion across multiple resolutions.

2. Related Work

Diffusion Models. Diffusion Models have become the de facto standard for image generation in recent times. Early works such as DDPM [17] achieved high-quality image gen-

eration without adversarial training, but relied on simulating a Markov chain with a large number of steps for sampling. DDIM [36] accelerated the sampling process, while methods such as LDM [33] performed denoising in a compact latent space rather than directly in the pixel-space. Recently, DiTs have become popular, replacing traditional UNet based backbones with transformer architecture [32]. Motivated by the goal of scaling diffusion models for high-resolution image generation while maintaining architectural simplicity and high-frequency image details, we propose an end-to-end Scale Space Diffusion model that performs denoising directly in the pixel domain.

Scale-Space Theory. Scale-space theory [25] is a fundamental concept in computer vision, that provides a framework for multi-scale image representation and analysis. It has been widely used in visual understanding tasks [5, 28]. The underlying idea of representation at multiple scales has been smartly used in the context of generative models to progressively generate images at increasing resolutions. In GAN-based approaches, Progressive GAN [22] has shown excellent results in generating high-resolution images by learning to generate at increasing resolutions during the training process. In some other works such as LAPGAN [10], multiple GANs, one for each scale, are used to upscale the image by producing a residual, similar to a Laplacian pyramid.

Several works in the space of diffusion models have also drawn inspiration from scale-space theory. Cascaded diffusion model [18] consists of a series of diffusion models that generate images of increasing resolutions, where the base model produces a low-resolution image and subsequent super-resolution models refine it using the upsampled version of the low-resolution image as a condition. Matryoshka Diffusion [13] model proposes a diffusion process that denoises inputs at multiple resolutions jointly.

However, none of these approaches directly incorporate scale-space theory in the diffusion process because the noise component of the noisy intermediate state leads to correlated noise pixels at an upsampled state. Some works solve this by adding additional noise at the higher resolution. Relay Diffusion [38] imagines a low-resolution generation as a high-resolution image with block noise and trains a model to denoise it at higher resolution with a weighted combination of block noise and high-resolution noise. Laplacian Diffusion Models [2] train separate models for different resolutions and add a Laplacian residual of high-resolution noise during resolution transitions. However, simply adding high-resolution noise does not fully resolve the distribution mismatch between noisy states at different resolutions. Pyramidal Flow Matching [21] addresses this issue by adding decorrelated noise while also rolling the diffusion process back to a noisier timestep. PixelFlow [6] and Region Adaptive Latent Sampling [20] build on this idea. Bottleneck-Sampling [39] as opposed to increasing scales introduces a bottleneck scale for better generation, while Decomposed

Flow Matching [14] predicts Laplacian residuals of clean images. UDPM [1] tries to add blurring and subsampling into the diffusion process, assumes isotropic posterior covariance to simplify their reverse diffusion derivation, which may not hold, given that the blurring kernels usually overlap in most implementations of resizing. We show through Scale Space Diffusion that end-to-end training of a single diffusion model capable of handling multiple resolutions, with a generalized mathematical formulation for resolution transitions, helps to achieve faster generation, while preserving high-quality.

3. Scale Spaces vis-à-vis Diffusion Timesteps

In this section, we outline the motivation behind our approach, which originates from a simple but compelling intuition. Consider the intermediate states of a diffusion model (Fig. 2(a), bottom) and the scales of a Gaussian pyramid (Fig. 2(b), bottom). If one squints and focuses on the third image from the left along the t -axis, the overall structure of the face begins to emerge, which is remarkably similar to the information present in the images corresponding to smaller spatial scales along the r -axis of the Gaussian pyramid. As we move rightward along either axis (*i.e.*, decreasing t or increasing r), it becomes evident how finer details are added progressively.

This observation suggests a striking correspondence in the information hierarchy between diffusion timesteps and scale-space resolutions (or scales). Our goal is to quantify this correspondence. To do so, we first review the standard diffusion process, and then formalize our intuition by mathematically characterizing the amount of information present across diffusion states.

3.1. Preliminary: Standard Diffusion Process

In standard denoising diffusion probabilistic models (DDPM) [17], the forward diffusion process is modeled as a Markov chain that progressively noises a signal by adding Gaussian noise. For $x_0 \sim q(x_0)$, where $q(x_0)$ is the data distribution, the process is defined as:

$$x_t = \sqrt{\alpha_t}x_{t-1} + \sqrt{1 - \alpha_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (1)$$

where $\{\beta_t\}_{t=1}^T$ is the variance schedule (with $\alpha_t := 1 - \beta_t$). This expression, when applied iteratively over t , leads to an alternative definition that expresses the noisy state as a linear combination of the signal x_0 and the noise ϵ :

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (2)$$

where $\bar{\alpha}_t := \prod_{i=0}^t \alpha_i$. Diffusion models aim to reverse this process by approximating the posterior distribution $q(x_{t-1}|x_0, x_t)$ using a neural network (with parameters θ) that predicts the noise ϵ in Eq. 2. This model, $\epsilon_\theta(x_t, t)$, is trained using a simplified loss function $\mathcal{L}_{\text{simple}} = \mathbb{E}_{x_0, t, \epsilon} [\|\epsilon_\theta(x_t, t) - \epsilon\|_2^2]$. The model can also be parameterized to predict x_0 instead of ϵ .




Figure 2. **Information Analysis.** (a) Amount of information present in a diffusion state as diffusion step t changes. (b) Amount of information present in images at various resolutions (scales).

3.2. Information Degradation in Diffusion and Scale Spaces

Diffusion States. In this section, we formally model the information degradation over the diffusion process. Eq. 2 has two terms – a signal term and a noise term. One way to model the amount of information present in x_t is to compute the percentage of pixels for which the noise term dominates the signal term, *i.e.*, $|\sqrt{1 - \bar{\alpha}_t}\epsilon| > |\sqrt{\bar{\alpha}_t}x_0|$. In other words, we are looking for the probability that $|\epsilon|$ is greater than $s(t)|x_0|$, where $s(t) = \frac{\sqrt{\bar{\alpha}_t}}{\sqrt{1 - \bar{\alpha}_t}}$ is the square root of the signal-to-noise coefficient ratio. We have $P(|\epsilon| > s(t)|x_0|) = (1 - \Phi(s(t)|x_0|)) + \Phi(-s(t)|x_0|) = 2\Phi(-s(t)|x_0|)$, where Φ is the CDF of the standard normal distribution, and the second equality follows from symmetry of the standard normal distribution. Hence, $P(|\epsilon| \leq s(t)|x_0|) = 1 - 2\Phi(-s(t)|x_0|)$. Now to obtain the expected fraction of signal-dominated pixels, a *proxy for information*, we average this probability over the data distribution $q(x_0)$. For simplicity, let us assume $x_0 \sim \mathcal{U}(-1, 1)$. Then the variation of information over timestep t can be written as:

$$\begin{aligned} \text{Info}(t) &= \mathbf{E}_{x_0 \sim \mathcal{U}(-1,1)} [1 - 2\Phi(-s(t)|x_0|)] \\ &= 1 - 2 \int_{-1}^1 p_{\mathcal{U}(-1,1)}(x) \Phi(-s(t)|x|) dx \\ &= 1 - 2 \int_{-1}^1 \frac{1}{2} \Phi(-s(t)|x|) dx \\ &= 1 - \int_{-1}^1 \Phi(-s(t)|x|) dx = 1 - 2 \int_0^1 \Phi(-s(t)x) dx, \end{aligned}$$

where we use the fact that the uniform distribution has density $p_{\mathcal{U}} = \frac{1}{2}$ over $[-1, 1]$, and for the final equality we split the integral about $x = 0$. Using this simplification, $\text{Info}(t)$ can be numerically computed as a function of t , as shown in Fig. 2(a).

Scale Spaces. Similar to the approximation of information across diffusion steps, here we want to approximate the information as a function of image resolution (*i.e.*, scale). A simple way to model this is to assume:

$$\text{Info} \propto \text{Area}.$$

Let us consider a normalized resolution $r \in [0, 1]$, where 0 represents no pixels and 1 represents the highest resolution. Under this assumption, the information can be written as:

$$\text{Info}(r) = r^2.$$

This implies, for example, that if the spatial dimensions of an image are halved, then the information becomes one quarter, which may not be strictly true due to redundancy in pixel space. However, the monotonic trend should still hold. This trend is visualized in Fig. 2(b).

Notice how there is a similarity in the trends of information degradation as t increases versus as r decreases. This analysis quantifies our main intuition regarding the similarity in the information trends across diffusion steps and scale spaces. Given this insight, we aim to leverage this intuition to construct a framework that realizes scale spaces within the current formulation of diffusion models.

In our initial attempts to incorporate scale spaces into diffusion models, we tried to frame this problem as jumping across the same timesteps of independent diffusion processes at varying scales. However, this led to an accumulation of errors during the iterative inference procedure, resulting in suboptimal outputs. Methods such as Pyramidal Flow Matching [6, 20, 21] address this issue by adding decorrelation noise when transitioning across scales and then backtracking in time so that an appropriate noise level is selected. This strategy helps mitigate the error accumulation. Nonetheless, it does not actually resolve the underlying issue – the diffusion process itself is not mathematically modeled to handle scale changes. In this work, we aim to fill this gap.

4. Scale Space Diffusion (SSD)

In this section, we introduce a new family of diffusion processes that use a generalized linear degradation operation for degrading the signal, in addition to the standard additive Gaussian noise. We then show how this formulation can be implemented in deep learning frameworks such as PyTorch [31] for any choice of a linear degradation that is available as a function call. In our case, we choose a downsizing operator as our linear degradation. Next, we present our training and sampling pipelines. Finally, we introduce our architecture that can handle scale-preservation and scale-changing transitions at multiple resolutions.

4.1. Generalized Linear Diffusion Process

4.1.1. Extension to Linear Degradation

We now replace the scalar coefficient of x_{t-1} in Eq. 1, *i.e.*, $\sqrt{\alpha_t}$, with a more generic linear operator M_t . For example, blurring or downsampling can serve as such a linear operator. Let us assume a Gaussian distribution for this updated formulation for the transition distribution $q(x_t|x_{t-1})$ as $x_t = M_t x_{t-1} + \eta_t, \eta_t \sim \mathcal{N}(0, \Sigma_{t|t-1})$. Here, we do not assume $\Sigma_{t|t-1}$ to be isotropic.

Now, repeatedly sampling the next state using the transition distribution, we want to derive an equation analogous to Eq. 2, which provides us x_t given x_0 . It is clear that this will also be a Gaussian distribution $q(x_t|x_0) = \mathcal{N}(\mu_t, \Sigma_t)$. The only constraint we want to enforce is isotropy, *i.e.*, $\Sigma_t = \sigma_t^2 \mathbf{I}$. For the coefficient of x_0 , instead of $\sqrt{\alpha_t} = \sqrt{\alpha_t} \sqrt{\alpha_{t-1}} \dots \sqrt{\alpha_1}$ in Eq. 2, we get $M_{1:t} = M_t M_{t-1} \dots M_1$, *i.e.*, $\mu_t = M_{1:t} x_0$. Hence, $q(x_t|x_0) = \mathcal{N}(M_{1:t} x_0, \sigma_t^2 \mathbf{I})$, which can be expressed as:

$$x_t = M_{1:t} x_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}) \quad (3)$$

Using Theorem 1, similar to blurring diffusion [19], the transition distribution $q(x_t|x_{t-1})$ is given by:

$$x_t = M_t x_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \Sigma_{t|t-1}), \quad (4)$$

$$\text{where } \Sigma_{t|t-1} = \Sigma_t - M_t \Sigma_{t-1} M_t^T.$$

For the isotropic marginals $\Sigma_t = \sigma_t^2 \mathbf{I}$ and $\Sigma_{t-1} = \sigma_{t-1}^2 \mathbf{I}$, we obtain $\Sigma_{t|t-1} = \sigma_t^2 \mathbf{I} - \sigma_{t-1}^2 M_t M_t^T$. For positive semi-definite feasibility we require $\sigma_t^2 \mathbf{I} \succeq \sigma_{t-1}^2 M_t M_t^T$, *i.e.*, $\sigma_t^2 \geq \sigma_{t-1}^2 \lambda_{\max}(M_t M_t^T)$.

As shown in Theorem 2, the reverse diffusion step, *i.e.*, the posterior distribution $q(x_{t-1}|x_t, x_0)$, conditioned additionally on x_0 , is also a normal distribution:

$$\begin{aligned} q(x_{t-1}|x_t, x_0) &= \mathcal{N}(\mu_{t \rightarrow t-1}, \Sigma_{\mu_{t \rightarrow t-1}}), \\ \text{where } \Sigma_{t \rightarrow t-1} &= (\Sigma_{t-1}^{-1} + M_t^T \Sigma_{t|t-1}^{-1} M_t)^{-1}, \text{ and} \\ \mu_{t \rightarrow t-1} &= \Sigma_{t \rightarrow t-1} (\Sigma_{t-1}^{-1} \mu_{t-1} + M_t^T \Sigma_{t|t-1}^{-1} x_t) \end{aligned} \quad (5)$$

Using the Woodbury matrix identity and isotropic covariance assumption, this simplifies to (Theorem 3):

$$\begin{aligned} \Sigma_{t \rightarrow t-1} &= \sigma_{t-1}^2 \mathbf{I} - \frac{\sigma_{t-1}^4}{\sigma_t^2} M_t^T M_t \\ \mu_{t \rightarrow t-1} &= \mu_{t-1} + \frac{\sigma_{t-1}^2}{\sigma_t^2} M_t^T (x_t - M_t \mu_{t-1}) \end{aligned} \quad (6)$$

Please refer to Table 1 for the comparison of our Generalized Linear Diffusion Process framework against DDPM and Blurring Diffusion (BD).

DDPM as a special case of SSD. When $M_t = \sqrt{\alpha_t} \mathbf{I}$ and $\sigma_t = \sqrt{1 - \alpha_t}$, the forward, marginal, and posterior distributions of SSD collapse to those of the DDPM model.

4.1.2. Implementation Details

Choice of M_t . We derived the above framework so that we can introduce scale spaces from Gaussian pyramids into the diffusion process. Although M_t may be any arbitrary linear operator, for our purposes we select it to be a resize operator, which effectively blurs and downsamples the image, and then multiplies it by $a_t = \sqrt{\alpha_t}$, as shown in Algo. 1. Note that this changes the dimensionality of the signal, in contrast

Table 1. Comparison between the formulations of the forward, marginal, and posterior distributions of DDPM and Blurring Diffusion (BD) against our Scale Space Diffusion. For Blurring Diffusion, we use ‘a’ instead of α used in their paper, to not confuse it with the α in DDPM. Note that BD applies a change of variable $u_t = V^T x_t$, where V^T is the Discrete Cosine Transform, before performing diffusion, *i.e.*, diffusion in frequency space. BD and DDPM have equivalent formulations when $a_t = \sqrt{\alpha_t}$ and $\sigma_t = \sqrt{1 - \alpha_t}$. While the formulations share structural similarities, Scale Space Diffusion extends the framework to support general linear degradations (*e.g.*, downscaling), which are not handled by DDPM or BD. We highlight analogous terms with consistent background colors for easier correspondences across different formulations.

Distributions	DDPM [17]	Blurring Diffusion [19]	Scale Space Diffusion
Forward $q(x_t x_{t-1})$	$x_t = \sqrt{\alpha_t} x_{t-1} + \sqrt{1 - \alpha_t} \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$	$u_t = a_{t t-1} u_{t-1} + \sigma_{t t-1} \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$ where $a_{t t-1} = \frac{a_t}{a_{t-1}}$ $\sigma_{t t-1} = \frac{\sigma_t}{\sigma_{t-1}} - a_{t t-1}^2 \sigma_{t-1}^2$	$x_t = \mu_{t t-1} + \eta_{t t-1}, \eta_{t t-1} \sim \mathcal{N}(0, \Sigma_{t t-1})$ where $\mu_{t t-1} = M_t x_{t-1} = M_{1:t} (M_{1:t-1})^{-1} x_{t-1}$, $\Sigma_{t t-1} = \Sigma_t - M_t \Sigma_{t-1} M_t^T$
Marginal $q(x_t x_0)$	$x_t = \sqrt{\alpha_t} x_0 + \sqrt{1 - \alpha_t} \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$	$u_t = a_t u_0 + \sigma_t \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$	$x_t = \mu_t + \eta_t, \eta_t \sim \mathcal{N}(0, \Sigma_t)$ where $\mu_t = M_{1:t} x_0, \Sigma_t = \sigma_t^2 \mathbf{I}$
Posterior $q(x_{t-1} x_t, x_0)$	$x_{t-1} = \tilde{\mu}_{t-1} + \tilde{\beta}_{t-1} \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$ where $\tilde{\beta}_{t-1} = \frac{1 - \alpha_{t-1}}{1 - \alpha_t} \beta_t$ $= \left(\frac{1}{1 - \alpha_{t-1}} + \frac{\alpha_t}{1 - \alpha_t} \right)^{-1}$, $\tilde{\mu}_{t-1} = \frac{\sqrt{\alpha_t(1 - \alpha_{t-1})}}{1 - \alpha_t} x_t + \frac{\alpha_{t-1} \beta_t}{1 - \alpha_t} x_0$ $= \tilde{\beta}_{t-1} \left(\frac{\sqrt{\alpha_t}}{1 - \alpha_t} x_t + \frac{\alpha_{t-1}}{1 - \alpha_{t-1}} x_0 \right)$	$u_{t-1} = \mu_{t \rightarrow t-1} + \sigma_{t \rightarrow t-1} \epsilon, \epsilon \sim \mathcal{N}(0, \mathbf{I})$ where $\sigma_{t \rightarrow t-1}^2 = \left(\frac{1}{\sigma_{t-1}^2} + \frac{\sigma_t^2}{a_{t t-1}^2} \right)^{-1}$, $\mu_{t \rightarrow t-1} = \frac{\sigma_t^2}{\sigma_{t \rightarrow t-1}^2} \left(\frac{a_{t t-1}}{\sigma_{t-1}^2} x_t + \frac{a_{t-1}}{\sigma_{t-1}^2} x_0 \right)$	$x_{t-1} = \mu_{t \rightarrow t-1} + \eta_{t \rightarrow t-1}, \eta_{t \rightarrow t-1} \sim \mathcal{N}(0, \Sigma_{t \rightarrow t-1})$ where $\Sigma_{t \rightarrow t-1} = \left(\Sigma_{t-1}^{-1} + M_t^T \Sigma_t^{-1} M_t \right)^{-1}$ $= \sigma_{t-1}^2 \mathbf{I} - \frac{\sigma_t^2}{\sigma_{t \rightarrow t-1}^2} M_t^T M_t$, $\mu_{t \rightarrow t-1} = \Sigma_{t \rightarrow t-1} \left(M_t^T \Sigma_t^{-1} x_t + \Sigma_{t-1}^{-1} \mu_{t-1} \right)$ $= \mu_{t-1} + \frac{\sigma_t^2}{\sigma_{t \rightarrow t-1}^2} M_t^T (x_t - M_t \mu_{t-1})$

with previous diffusion formulations. However, since we make no assumptions about dimensionality, our framework remains valid regardless. Furthermore, with this choice of M_t , we also define a resolution schedule $r(t)$ that maps diffusion timestep (t) to the corresponding resolution, such that the resolution monotonically decreases as t increases (Fig. 5). Refer to the supplementary Sec. 8.2.1 for another degradation example.

Calculating the Transpose. Since operators like image resizing are implicit, we may not have the matrix form available, making it non-trivial to apply the transpose M_t^T . To address this, we use a vector-Jacobian product of the function call $M_t(\cdot)$, *i.e.*, $M^T v = \text{torch.autograd.grad}(M_t(x), x, \text{grad_outputs}=v) [0]$, which, for linear operators, does not depend on x , as shown in Algo. 2. This computes the derivative of the inner product $\langle v, M_t x \rangle$ with respect to x , *i.e.*, $\nabla_x \langle v, M_t x \rangle = M_t^T v$.

Sampling from a Non-Isotropic Gaussian Distribution. A neat trick to sample from a non-isotropic Gaussian distribution with covariance matrix Σ is to first sample a standard Gaussian noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, and then multiply with the square root of the covariance matrix, so that $\Sigma^{\frac{1}{2}} \epsilon \sim \mathcal{N}(0, \Sigma^{\frac{1}{2}} \mathbf{I} (\Sigma^{\frac{1}{2}})^T) = \mathcal{N}(0, \Sigma)$. In our case, we need to sample noise from $\Sigma_{t \rightarrow t-1}$ from Eq. 6, which depends on implicit operators $M_t(\cdot)$ and $M_t^T(\cdot)$. Thus, we need a way to apply $\Sigma_{t \rightarrow t-1}(\cdot)$ implicitly to a standard Gaussian noise ϵ . For this purpose, we use the Lanczos algorithm [12, 24], which numerically computes $A(x)$ given an implicit symmetric linear operator $A(\cdot)$ and vector x . When the Lanczos algorithm is applied with a square root spectral function over the eigenvalues, we can obtain $A^{\frac{1}{2}} x$. In our case, this gives $\eta_{t \rightarrow t-1} = \Sigma_{t \rightarrow t-1}^{\frac{1}{2}} \epsilon \sim \mathcal{N}(0, \Sigma_{t \rightarrow t-1})$ as shown in Algo. 3.

Algorithm 1 Implicit Linear Operator

```
# M resizes and attenuates signal x
def M(x, a_t, a_t_minus1, size_out):
    return (a_t / a_t_minus1) * F.interpolate(
        x, size=size_out, mode="bilinear",
        align_corners=False, antialias=True)
```

Algorithm 2 Implicit Linear Operator’s Transpose

```
# M_T applies the transpose of M on v
def M_T(M, v, a_t, a_t_minus1, M_input_shape):
    size_out = v.shape()[-2:]
    with torch.enable_grad():
        x = torch.zeros(M_input_shape, requires_grad=True)
        out = M(x, a_t, a_t_minus1, size_out)
        # calculate M^T v = d<v, Mx>/dx
        (g,) = torch.autograd.grad(out, x, grad_outputs=v,
            retain_graph=False)
    return g
```

Algorithm 3 Sampling Non-Isotropic Gaussian Noise

```
# Sample noise from posterior covariance Sigma_{t-->t-1}
def sample_non_isotropic_noise(M, M_T, sigma_t,
    sigma_t_minus1, x):
    rho = (sigma_t_minus1 ** 2) / (sigma_t ** 2)
    # Define matvec operator A
    A = lambda v: v - rho * M_T(M(v))
    # Lanczos approximation of A^(1/2)v
    y = lanczos(A, x, f=lambda l: l.sqrt())
    return sigma_t_minus1 * y
```

4.2. Training and Sampling

To reverse the diffusion process using Eq. 6, our model must predict $\mu_{t-1} = M_{1:t-1} x_0$, which, with our choice of M_t , reduces to a scaled version of an image x_0 at resolution $r(t-1)$. To train such a model, using our Generalized Linear Diffusion Process, we need to first adapt $\mathcal{L}_{\text{simple}}$. When predicting x_0 , the loss becomes $\mathcal{L}_{\text{simple}} = \mathbb{E}_{x_0, t, \epsilon} [s^2(t) \|x_0^\theta(x_t, t) - x_0\|_2^2]$, where $s^2(t)$ is the signal to noise ratio, as shown in [35]. In Min-SNR- γ [15] instead

Algorithm 4 Train

```
def train_iter(x, t, a_t_minus1, model, opt):
    opt.zero_grad()
    t_minus1 = (t-1).clamp(min=0)
    # clean image at res r(t-1) = M_{1:t-1}(x) / a_{t-1}
    x_start_t_minus1 = cumulative_M[t_minus1](x) / a_t_minus1
    # Using Eq.3
    x_t = diffuse(x, t)
    pred_x_start_t_minus1 = model(x_t, t)
    # Using Eq.7
    loss = min_snr_5(t) * ((pred_x_start_t_minus1 -
        x_start_t_minus1) ** 2)
    loss.backward()
    opt.step()
    return loss
```

Algorithm 5 Sampling

```
# get x_{t-1} given x_t
def sample_iter(x_t, t, model):
    pred_x_start_t_minus1 = model(x_t, t)
    mu_t_minus1 = a_t_minus1 * pred_x_start_t_minus1
    # Using Eq.6
    posterior_noise = calculate_posterior_noise(t)
    posterior_mean = calculate_posterior_mean(x_t,
        mu_t_minus1, t)
    x_t_minus1 = posterior_mean + posterior_noise
    return x_t_minus1
```

of the $s^2(t)$ weighting, they use $\min(s^2(t), \gamma)$, with $\gamma = 5$, which improves the performance of x_0 parameterization significantly. Following this, our loss function evaluates to:

$$\mathcal{L} = \mathbb{E}_{x_0, t, \epsilon} \left[\min(s^2(t), \gamma) \left\| x_{0, \theta}^{r(t-1)}(x_t, t) - \frac{1}{a_{t-1}} M_{1:t-1} x_0 \right\|_2^2 \right] \quad (7)$$

where we predict an unscaled μ_{t-1} using a neural network $x_{0, \theta}^{r(t-1)}$ (Algo. 4). Note that the input resolution $r(t)$ of x_t may be smaller than the resolution of the output at $r(t-1)$ as seen in Fig. 3 (left). In standard diffusion training, timesteps are simply sampled uniformly for each batch. However, this is non-trivial in our setting because the $(r(t), r(t-1))$ pairs may not match. To solve this, we first uniformly sample a single t , and if $r(t) = r(t-1)$, then uniformly sample the batch size number of t_i 's that have $r(t_i) = r(t)$. Otherwise, if $r(t) \neq r(t-1)$, then we fill the entire batch with the same t , so there is no size mismatch. Since not all t 's change resolution, many of the M_t 's can be replaced by scalar multiplication with $(a_t/a_{t-1}) = \sqrt{\alpha_t}$.

For sampling (Algo. 5), we start from a random Gaussian noise at the lowest resolution $r(T)$. Our model $x_{0, \theta}^{r(t-1)}$ predicts a clean image at the next resolution $r(t-1)$, using which we can calculate μ_{t-1} and denoise using the posterior distribution (Eq. 6). This also involves sampling from $\Sigma_{t \rightarrow t-1}$, which may not be isotropic, and hence we use Algo. 3 to sample noise from this distribution. Eq. 6 is equivalent to DDPM sampling when $r(t) = r(t-1)$, so the non-isotropic noise sampling can be replaced with normal `torch.randn()` calls for resolution-preserving steps.

4.3. Architecture

We adapt the UNet architecture from Ablated Diffusion Model (ADM) [11] to design our proposed model Flexible-UNet (Flexi-UNet), which supports multi-resolution inputs and outputs to fully realize the scale-space formulation. Because Scale Space Diffusion embeds a resizing operator in the forward diffusion process, the spatial resolution of x_t varies across timesteps, and the reverse model must therefore operate on variable-sized noisy states and sometimes predict a higher-resolution output at the next scale (Fig. 3).

A standard diffusion model, such as ADM [11], is trained to operate at a single fixed resolution throughout all timesteps, and even multiresolution UNet variants only process multiple scales *within* a fixed-resolution diffusion process. In contrast, SSD requires an architecture that natively handles different input resolutions across timesteps. To address this, we explore two architectural designs.

Full UNet (Single Path). The base UNet architecture inherently supports variable-size inputs and outputs, and in principle can operate on any spatial resolution $R \times R$ as long as the kernel sizes, strides, padding, and pooling operations produce valid feature maps at every layer. However, this design has two key limitations for Scale Space Diffusion. First, it requires the input and output resolutions to be equal. In our setting, certain timesteps involve a resolution transition, which would require the model to output at a higher resolution. To handle this, the input must be manually upsampled before entering the UNet whenever such a transition occurs.

Second, the depth of the UNet determines how many distinct spatial scales it can represent. For a UNet with L downsampling blocks, the smallest internal resolution is $\frac{R}{2^{L-1}}$, which fixes the total number of scales to L . This number is typically small and does not grow with the input resolution. For example, the ADM architecture uses 4 feature map resolutions for 64×64 , 5 for 128×128 , and 6 for 256×256 , meaning that across all these models the number of downsampling stages remains fixed at 4. Thus, even at higher resolutions, the network cannot represent more than a handful of scales, limiting the usefulness of a scale-space formulation where many more levels naturally exist.

Flexi-UNet. These limitations motivate our proposed architecture, Flexi-UNet, where different subsets of UNet layers are dynamically activated based on the input resolution. High-resolution inputs traverse the full UNet, while lower-resolution inputs are routed only through the deeper layers, effectively bypassing the early and late blocks. Since each block expects a specific channel dimensionality, we insert 1×1 conv layers to map the input features to the appropriate channel size while preserving spatial resolution.

For denoising steps that do not involve a resolution change, the active pathway through the UNet remains symmetric, using the same number of downsampling and upsampling blocks. When a resolution increase is required, the

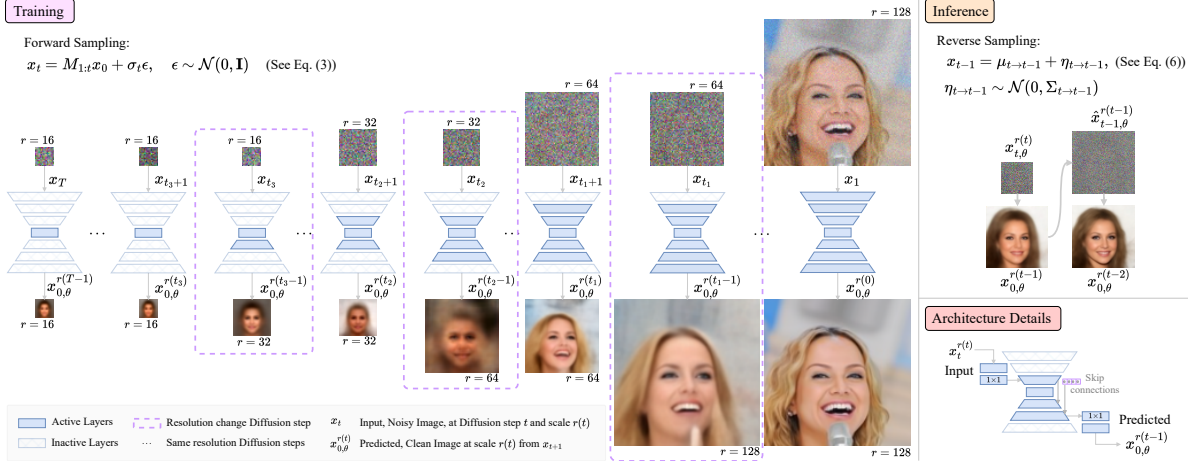


Figure 3. **Overview.** Left: During training x_t 's at resolution $r(t)$ are sampled using Eq. 3, and our model is trained to predict clean image $x_{0,\theta}^{r(t-1)}$ using the loss as in Eq. 7. Our Flexi-UNet is able to process both resolution-preserving and resolution-changing steps at multiple resolution using only parts of the network. Right-top: During sampling, Eq. 6 is used to progressively denoise and upsample to generate images. Right-bottom: Our Flexi-UNet has additional 1×1 Conv layers to take inputs at any UNet encoder block and get outputs from any decoder blocks. For resolution changing, the skip connections are fed with zero-filled tensors.

pathway becomes asymmetric: the model uses one additional upsampling block compared to the number of downsampling blocks encountered. In these cases, the skip connections that would normally come from the bypassed encoder blocks are replaced with zero tensors (Fig. 3). This design allows the model to share parameters across resolutions while supporting valid diffusion dynamics during resolution transitions.

5. Experiments

Datasets. We perform experiments and analyze the performance of Scale Space Diffusion on the CelebA dataset [26] and the ImageNet dataset [9]. The CelebA dataset consists of around 200K training images, while the ImageNet dataset contains around 1.3 million images from 1000 different classes. We use JPEG images for these datasets. We conduct experiments at 64×64 resolution for both CelebA and ImageNet. We additionally show experiments on 128×128 and 256×256 for CelebA dataset. The CelebA experiments helps us understand our method's scalability with increasing resolutions, while ImageNet helps in evaluating the model's ability to learn complex and diverse distributions.

Unconditional Image Generation. We analyze and evaluate Scale Space Diffusion on unconditional image generation, as it allows us to clearly study how scale-space theory integrates with the diffusion process.

Implementation Details. We use the ADM [11] repository as our base codebase and build our baselines (DDPM [17], Blurring Diffusion [19]), as well as Scale Space Diffusion on top of it. For DDPM, we consider two standard parametrizations as our baselines, the ϵ -prediction, and the x_0 -prediction formulation. We train the baseline model with Min-SNR- γ weighting for the x_0 -parametrization to ensure an accurate




Figure 4. **Visual samples.** Top: ImageNet-64 unconditional generation. For the top-most sample we also show model prediction at various scales (8, 16, 32, 64) during SSD. Bottom: CelebA-256 unconditional generation. For the top-most sample we also show model predictions at various scales (8, 16, 32, 64, 128, 256).




Figure 5. **Resolution Schedules.** Mapping diffusion timesteps t to resolution r across 4 scales. Both discrete and continuous variants are shown. The right shows FIDs at 500k iterations (batch size 8). comparison to Scale Space Diffusion. We implement Blurring Diffusion using the pseudo-code provided in their paper.

For the diffusion process, we follow the linear noise schedule proposed in DDPM [17] and use the standard setting of 1000 timesteps. For training, we use AdamW [23, 27] optimizer with a fixed learning rate. We conducted all our

Table 2. **Main Results.** Unconditional image generation results on CelebA dataset over multiple resolutions. Training time is specified in hours. Average GFlops per iteration. Effective batch size is 128 for resolutions 64 and 128, 64 for resolution 256. Here BD refers to Blurring Diffusion [19] and all SSD models use our Flexi-UNet architecture.

Method	CelebA-64 (1M iters)			CelebA-128 (300K iters)			CelebA-256 (300K iters)		
	FID	Time	GFlops	FID	Time	GFlops	FID	Time	GFlops
DDPM- ϵ	2.22	70.30	60.05	4.16	50.50	132.30	5.52	87.31	497.03
DDPM- x_0	2.98	70.71	-	3.50	50.33	-	5.47	87.33	-
BD	2.06	71.79	-	3.67	-	-	4.76	88.08	-
SSD (2L)	2.14	62.63	50.61	-	-	-	-	-	-
SSD (3L)	3.61	56.13	44.27	6.53	31.71	87.38	7.79	59.00	317.36
SSD (4L)	4.28	52.38	38.48	-	-	-	10.52	51.70	272.98
SSD (5L)	-	-	-	10.47	25.41	66.72	-	-	237.70
SSD (6L)	-	-	-	-	-	-	13.50	42.88	209.69

Table 3. **ImageNet-64 Results.** Unconditional image generation results on ImageNet-64 dataset.

Method	FID
DDPM- ϵ	12.82
DDPM- x_0	13.07
Blurring Diffusion	15.34
SSD (2L)	13.08
SSD (4L)	17.89

experiments on NVIDIA H100 and NVIDIA RTX A4000 GPUs. We maintained consistent combinations of learning rate and batch-size across dataset and resolutions. For 64×64 and 128×128 , we used an effective batch size of 128, and trained the models either on a single H100, or on 4 A4000 GPUs, with a per-GPU batch-size of 32. For 256×256 , we used an effective batch size of 64 due to memory constraints and trained them on 2 H100s with a per-GPU batch size of 32. Our learning rate was set to 1×10^{-4} for the 64×64 and 128×128 models, and 5×10^{-5} for 256×256 model, following linear learning rate scaling.

Evaluation. We evaluate our models using the exponential moving average (EMA) weights with a decay rate of 0.9999. We assess the quality of generated images by computing FID [16] scores on 50k samples w.r.t. the training set. We further compare Scale Space Diffusion with the DDPM baseline model in terms of training time, and FLOPs (Floating Point Operations) per forward pass. In addition to FLOPs, we report the sampling latency as the total time to generate a single image. All speed and compute metrics are measured on a single NVIDIA GH200 node.

Main results. Our results are presented in Table 2 and 3. We train the baseline models and Scale Space Diffusion model for 1 million iterations for CelebA-64 and 300k iterations for CelebA-128 and CelebA-256. We report the total training time, average GFLOPs per iteration, and the FID value. We notice that increasing the number of levels significantly decreases training time and GFLOPs. SSD (6L) at 256 resolution takes less than half the time as the baseline DDPM. Table 3 shows that SSD, trained for 1 million iterations,

Figure 6. **Temporal Scaling.** Training time of proposed SSD with our Flexi-UNet across multiple resolutions.




Table 4. **Architecture Ablation.** FID (at 500K iterations on CelebA-64) and Inference time (in secs/generation, 1000 steps, batch size=1, on $1 \times$ A4000) of network architecture variants at 2 and 4 levels.

Method	FID		Inference time	
	res. 64	res. 64	res. 256	res. 256
Full UNet, 2L	2.33	16.19	43.07	
Flexi-UNet, 2L	2.26	15.38	38.99	
Full UNet, 4L	4.90	16.28	34.74	
Flexi-UNet, 4L	4.87	13.43	31.08	

achieves comparable performance to baselines even on the harder ImageNet-64 benchmark. Figure 6 shows that training time for SSD scales well with increasing resolution. Please refer to the supplementary Sec. 8.4 for more comparisons.

Qualitative results. We present qualitative results of our method on ImageNet-64 with SSD (4L) and CelebA-256 with SSD (6L) in Fig. 4. We also show multiple intermediate predictions of the model in SSD.

Individual effectiveness of our mathematical formulation vs architecture. In the supplementary Sec. 8.2, we first show that our Generalized Linear Diffusion Process works, albeit suboptimally, even without Flexi-UNet, and also with an alternate degradation. Approximating anisotropic gaussian noise with isotropic leads to saturation artifacts, showing the need for our anisotropic sampling. Secondly, we show that Flexi-UNet is effective on its own for other formulations of iterative multi-resolution pixel-space generation, albeit suboptimal. We do this by applying it to approximate multi-resolution diffusion as in PyramidalFlow [6, 20, 21].

Resolution Schedule. $r(t)$ specifies the spatial resolution as a function of diffusion timestep. We present 5 different resolution schedules in Figure 5 (left) at 4 levels (64, 32, 16, 8) (refer to the supplementary Sec. 8.3 for details) and note their effects for a CelebA-64 model in Fig. 5 (right). We observe that schedules that spend the least number of timesteps at the higher resolutions train the fastest, but also yield the worst FID (*i.e.* ConvexDecay_2). In contrast, the model trained with ConvexDecay_0.5, which spends the most steps at the highest resolution, achieves the best FID, but requires the longest training time. We use this for all our experiments.

Full UNet vs. Flexi-UNet. In Table 4, we observe that Flexi-UNet has slightly better FID for both 2L and 4L, while being faster than the Full UNet. Hence, we use Flexi-UNet.

Sampling. Use of Lanczos has negligible overhead. Further, SSD does not suffer from performance drop, like DDPM, on sampling steps reduction. (Refer supplementary Sec 8.5.)

Conclusion. We showed that diffusion models and scale spaces share an information hierarchy, and we quantified this connection mathematically. Observing that highly noised diffusion states contain only low-resolution information, we introduced a generalized family of diffusion models that embeds scale-space structure into the forward process, yielding Scale Space Diffusion. To realize this in practice, we proposed the Flexi-UNet architecture and demonstrated its effectiveness on unconditional image generation.

Acknowledgment. This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 140D0423C0076. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

References

- [1] Shady Abu-Hussein and Raja Giryes. Udpn: Upsampling diffusion probabilistic models. *arXiv preprint arXiv:2305.16269*, 2023. 1, 3, 2, 5
- [2] Yuval Atzmon, Maciej Bala, Yogesh Balaji, Tiffany Cai, Yin Cui, Jiaojiao Fan, Yunhao Ge, Siddharth Gururani, Jacob Huffman, Ronald Isaac, et al. Edify image: High-quality image generation with pixel space laplacian diffusion models. *arXiv preprint arXiv:2411.07126*, 2024. 1, 2
- [3] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *NeurIPS*, 2023. 1
- [4] Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22669–22679, 2023. 1
- [5] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 2009. 2
- [6] Shoufa Chen, Chongjian Ge, Shilong Zhang, Peize Sun, and Ping Luo. Pixelflow: Pixel-space generative models with flow. *arXiv preprint arXiv:2504.07963*, 2025. 1, 2, 4, 8, 3, 5
- [7] Katherine Crowson, Stefan Andreas Baumann, Alex Birch, Tanishq Mathew Abraham, Daniel Z Kaplan, and Enrico Shippole. Scalable high-resolution pixel-space image synthesis with hourglass diffusion transformers. In *Forty-first International Conference on Machine Learning*, 2024. 1
- [8] Giannis Daras, Mauricio Delbracio, Hossein Talebi, Alex Dimakis, and Peyman Milanfar. Soft diffusion: Score matching with general corruptions. *Transactions on Machine Learning Research*, 2023. 1
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2, 7
- [10] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *NeurIPS*, 2015. 2
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 6, 7, 1
- [12] Gene H Golub and Gérard Meurant. *Matrices, moments and quadrature with applications*. Princeton University Press, 2009. 5
- [13] Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Joshua M Susskind, and Navdeep Jaitly. Matryoshka diffusion models. In *ICLR*, 2024. 2, 1
- [14] Moayed Haji-Ali, Willi Menapace, Ivan Skorokhodov, Arpit Sahni, Sergey Tulyakov, Vicente Ordonez, and Aliaksandr Siarohin. Improving progressive generation with decomposable flow matching. *arXiv preprint arXiv:2506.19839*, 2025. 3, 5
- [15] Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining Guo. Efficient diffusion training via min-snr weighting strategy. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7441–7451, 2023. 5
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 8
- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 1, 2, 3, 5, 7
- [18] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *JMLR*, 2022. 2
- [19] Emiel Hoogeboom and Tim Salimans. Blurring diffusion models. *ICLR*, 2023. 1, 4, 5, 7, 8
- [20] Wongi Jeong, Kyungryeol Lee, Hoigi Seo, and Se Young Chun. Upsample what matters: Region-adaptive latent sampling for accelerated diffusion transformers. *arXiv preprint arXiv:2507.08422*, 2025. 1, 2, 4, 8, 3
- [21] Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling. *arXiv preprint arXiv:2410.05954*, 2024. 1, 2, 4, 8, 3
- [22] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *ICLR*, 2018. 2, 1
- [23] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [24] Cornelius Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *Journal of research of the National Bureau of Standards*, 45(4):255–282, 1950. 5
- [25] Tony Lindeberg. Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of applied statistics*, 21(1-2):225–270, 1994. 1, 2
- [26] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015. 2, 7
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 7
- [28] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international*

conference on computer vision, pages 1150–1157. Ieee, 1999. 2

- [29] Soumik Mukhopadhyay, Matthew Gwilliam, Yosuke Yamaguchi, Vatsal Agarwal, Namitha Padmanabhan, Archana Swaminathan, Tianyi Zhou, Jun Ohya, and Abhinav Srivastava. Do text-free diffusion models learn discriminative visual representations? In *European Conference on Computer Vision*, pages 253–272. Springer, 2024. 1
- [30] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pages 8162–8171. PMLR, 2021. 6
- [31] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 4
- [32] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205, 2023. 2, 1
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 1, 5
- [34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 2, 1
- [35] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022. 5
- [36] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2
- [37] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *ICLR*, 2021. 1
- [38] Jiayan Teng, Wendi Zheng, Ming Ding, Wenyi Hong, Jianqiao Wangni, Zhuoyi Yang, and Jie Tang. Relay diffusion: Unifying diffusion process across resolutions for image synthesis. *arXiv preprint arXiv:2309.03350*, 2023. 2
- [39] Ye Tian, Xin Xia, Yuxi Ren, Shanchuan Lin, Xing Wang, Xuefeng Xiao, Yunhai Tong, Ling Yang, and Bin Cui. Training-free diffusion acceleration with bottleneck sampling. *arXiv preprint arXiv:2503.18940*, 2025. 2
- [40] XIANG Weilai. FutureXiang/Diffusion, 2025. original-date: 2022-10-18T11:42:46Z. 2