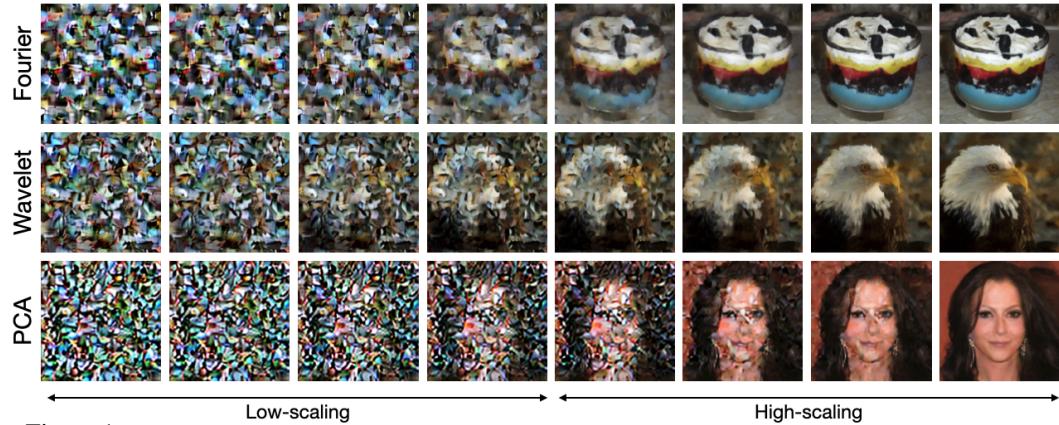


FLOW ALONG THE K -AMPLITUDE FOR GENERATIVE MODELING

005 **Anonymous authors**

006 Paper under double-blind review



023 Figure 1: Generation using K-Flow with three K -amplitude decompositions: Fourier, Wavelet, and PCA.

ABSTRACT

027 In this work, we propose K-Flow, a novel generative learning paradigm that flows
 028 along the K -amplitude domain, where k is a scaling parameter that organizes pro-
 029 jected coefficients (frequency bands), and amplitude refers to the norm of such
 030 coefficients. We instantiate K-Flow with three concrete K -amplitude transfor-
 031 mations: Fourier transformation, Wavelet transformation, and PCA. By incorpo-
 032 rating the K -amplitude transformations, K-Flow enables flow matching across the
 033 scaling parameter as time. We discuss six properties of K-Flow, covering its the-
 034 oretical foundations, energy and temporal dynamics, and practical applications.
 035 Specifically, from the perspective of practical usage, K-Flow allows for steerable
 036 generation by controlling the information at different scales. To demonstrate the
 037 effectiveness of K-Flow, we conduct experiments on both unconditional and con-
 038 ditional image generation tasks, showing that K-Flow achieves competitive perfor-
 039 mance. Furthermore, we perform three ablation studies to illustrate how K-Flow
 040 leverages the scaling parameter for controlled image generation. Additional re-
 041 sults, including scientific applications, are also provided.

1 INTRODUCTION

045 Generative Artificial Intelligence (GenAI) represents a pinnacle achievement in the recent wave of
 046 AI. This field has evolved from foundational methods such as autoregressive models (AR) (Radford,
 047 2018), energy-based models (Hinton, 2002; Carreira-Perpinan & Hinton, 2005; LeCun et al., 2006;
 048 Gutmann & Hyvärinen, 2010; Song & Kingma, 2021), variational auto-encoders (Kingma, 2013),
 049 and generative adversarial networks (Goodfellow et al., 2014), diffusion model (Ho et al., 2020),
 050 to the most cutting-edge flow-matching (FM) framework (Lipman et al., 2022; Liu et al., 2022b;
 051 Albergo & Vanden-Eijnden, 2022).

052 Among these, flow matching (FM) stands out as a density transport method that converts an initial
 053 simple distribution into a complex target distribution through continuous-time flow dynamics. For
 instance, in the context of image generation, FM learns to map a random Gaussian distribution to

054 the pixel-space distribution of images. This process, termed continuous *flow*, is governed by a local-
 055 ized *k-dependent vector field* (or velocity field) and produces a *time-dependent density path*, which
 056 represents the evolution of the probability distribution over time. As a versatile framework, FM
 057 can incorporate a diffusion density path, linking it to established methods such as denoising score
 058 matching (DSM) (Vincent, 2011; Song & Ermon, 2019) and the denoising diffusion probabilistic
 059 model (DDPM) (Ho et al., 2020).

060 **Motivation.** Natural data exhibits an inherent frequency structure, with most of its energy con-
 061 centrated in the low-frequency bands (Abry et al., 1995; Van der Schaaf & van Hateren, 1996).
 062 Reflecting this property, empirical evidence (Dieleman, 2024) shows that DDPMs tend to denoise
 063 from low to high frequencies when transforming white noise (with a uniform frequency spectrum)
 064 into meaningful data, allowing earlier recovery of low-frequency components (Biroli et al., 2024).
 065 Conventional FMs, however, exhibit different path characteristics (Sun et al., 2025), and their fre-
 066 quency progression is not quantitatively established (Figure S2). In parallel, from the reconstruction
 067 perspective, recent research (Kouzelis et al., 2025; Skorokhodov et al., 2025) advocates for explicit
 068 frequency consistency constraints as a regularization strategy for auto-encoders. This supports the
 069 expectation that introducing frequency-aware path into generation can achieve generation quality on
 070 par with or exceeding that of conventional FMs. These observations point to a clear opportunity:
 071 developing generative models that offer fine-grained control in the frequency domain can open new
 072 frontiers in both generation quality and applicability, *e.g.*, frequency editing and restoration.

073 **Key Concepts.** To formalize our approach, we first establish a unified framework that integrates key
 074 frequency-related concepts from the literature, including Fourier frequency analysis and multi-scale
 075 transformations. Central to this framework is the introduction of *K*-amplitude space, parameterized
 076 by a scaling parameter *k*. The *scaling parameter k* is defined as a systematic measure for orga-
 077 nizing frequency bands (or coefficients) across different physical systems and processes (Cardy,
 078 1996; Luijten & Blöte, 1996; Behan et al., 2017; Bighin et al., 2024)¹. Within this framework, we
 079 define *amplitude* as the norm of coefficients obtained by projecting data onto bases corresponding
 080 to different scaling parameters *k*, forming what we term the *K-amplitude space*, or equivalently,
 081 *scaling-amplitude space*.

082 **Our Method.** Such an understanding of scaling parameter and *K*-amplitude space inspires a new
 083 paradigm for generative modeling, which we term **K Flow Matching (K-Flow)**. In essence, K-
 084 Flow performs flow along the *K*-amplitude. There are two main components in K-Flow, and the
 085 first is the *K*-amplitude decomposition. The *K*-amplitude decomposition encompasses a family
 086 of transformations through a linear basis in the *K*-amplitude space, and in this work, we explore
 087 three types: Wavelet, Fourier, and principal component analysis (PCA) decomposition, as illustrated
 088 in Figure 1. Specifically, K-Flow first applies a *K*-amplitude transformation to project data from
 089 the spatial domain into the *K*-amplitude space, where we formulate a novel stochastic interpolant
 090 that naturally accommodates the hierarchical structure. In Appendix A, we provide a comprehensive
 091 analysis of K-Flow through six properties, from theoretical foundations (a & b), energy and temporal
 092 dynamics (c & d) to practical applications (e & f), with a detailed pipeline illustrated in Figure 2.

093 **Our Results.** We demonstrate the effectiveness of K-Flow through extensive experiments on
 094 generation tasks. Qualitatively, our ablation studies reveal the model’s scaling controllability that
 095 aligns with our theoretical motivation, enabling two key capabilities: (1) efficient class-conditional
 096 generation with minimal guidance, where class information is only required during low-*k* inference
 097 stages, which has the potential of reducing computational overhead; and (2) unsupervised frequency
 098 editing through various discretizations of the scaling parameter *k*. Quantitatively, K-Flow achieves
 099 state-of-the-art or comparable performance in both unconditional generation and training-free
 100 image restoration tasks across natural image and scientific datasets.

101 2 BACKGROUND

102 2.1 SCALING PARAMETER *k*, AMPLITUDE, AND *K*-AMPLITUDE DECOMPOSITION

103 Our data generation framework leverages the implicit hierarchical structure of the data manifold. By
 104 ‘implicit’, we refer to the hierarchical characteristics that emerge when a generalized *K*-amplitude

105
 106
 107 ¹We distinguish “scaling parameters” in the context of parameterization from “scale” in general discussions

decomposition is applied, transitioning the representation from the original data space to the K -amplitude space. Illustrations are in Figure 2. More formally, we represent data as a signal $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^m$, or a finite discretization of \mathbb{R}^d and \mathbb{R}^m , where this signal function is equivalent to a vector. For example, In the case of image data, each pixel of one RGB channel can be viewed as a signal mapping from the spatial grid \mathbb{R}^2 to a pixel intensity value in \mathbb{R}^1 . Combining the three channels, they form a vector-valued signal from \mathbb{R}^2 to \mathbb{R}^3 . An alternative approach is to consider data as a high-dimensional vector $\mathbb{R}^{d \times m}$. However, treating data as signal functions better fits the decomposition framework in this work.

Without loss of generality, we take dimension $m = 1$ for illustration. A K -amplitude decomposition involves the decomposition of a function using a complete basis set $\{\mathbf{e}_j\}_{j=1}^n$, where n can be infinite. We introduce a scaling parameter k , which partitions the set $\{\mathbf{e}_i\}_{i=1}^n$ into subsets: $\{\mathbf{e}_i\}_{i=1}^n = \bigcup_k \{\mathbf{e}_k\}$, each with n_k basis. Hence, signal ϕ is expressed as:

$$\phi = \sum_k \phi_k, \quad (1)$$

where $\phi_k := \sum_{j=1}^{n_k} (\phi \cdot \mathbf{e}_{jk}) \mathbf{e}_{jk}$ for $\mathbf{e}_{jk} \in \{\mathbf{e}_k\}$. Inspired by the concept of frequency amplitude, we refer to the norm of ϕ_k as the K -amplitude. The parameter k is termed the scaling parameter as we expect the natural scaling law exists in well-structured data: the amplitude decays as the value of k increases (Field, 1987).

Finally, we define K -amplitude decomposition (or equivalently, K -amplitude transform) \mathcal{F} as the map that sends ϕ to the collection of ϕ_k , and denote the collection of all $\{(\phi \cdot \mathbf{e}_{jk}) \mathbf{e}_{jk}\}_j$ as $\mathcal{F}\{\phi\}(k)$. Then,

$$\mathcal{F}\{\phi\} := \bigcup_k \mathcal{F}\{\phi\}(k). \quad (2)$$

We further assume that this transform has an inverse, denoted by \mathcal{F}^{-1} .

Splitting Probability. Denote the probability of data as p_{data} , then the transformations \mathcal{F} and \mathcal{F}^{-1} induce a probability measure on the associated K -amplitude space. In particular, we denote the induced splitting probability of ϕ_k as $p(k)$ for each scaling parameter k .

In this work, we explore three types of K -amplitude decomposition: Wavelet, Fourier, and principal component analysis (PCA). In Section 2.2, we will provide a classic example using the Fourier frequency decomposition on the three-dimensional space. This example serves to illustrate the construction of the scaling parameter k and K -amplitude.

2.2 EXAMPLE: FOURIER AMPLITUDE DECOMPOSITION

Suppose the data $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$, is drawn from a certain function distribution p_{data} . The challenge of directly fitting the distribution p_{data} is often complex and computationally demanding. Fourier frequency decomposition, however, offers a powerful technique to address this challenge by transforming ϕ into the Fourier space or Fourier domain. In what follows, we will use the terms ‘space’ and ‘domain’ interchangeably.

By applying Fourier frequency decomposition, we express ϕ as a sum of its frequency components. This transformation can potentially unveil the hidden structure within the distribution p_{data} , which is not apparent in the spatial or time domain, and it is thus beneficial for understanding the underlying patterns in the data manifold. To illustrate, the continuous Fourier transform \mathcal{F} of data $\phi(x, y, z) : \mathbb{R}^3 \rightarrow \mathbb{R}$ is expressed as:

$$\mathcal{F}\{\phi\}(k_x, k_y, k_z) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \phi(x, y, z) e^{-2\pi i (k_x x + k_y y + k_z z)} dx dy dz. \quad (3)$$

After this transformation, the spatial variables (x, y, z) are converted into frequency variables (k_x, k_y, k_z) , thereby representing the data in the frequency domain.

Note that the Fourier frequency is characterized by the high-dimensional vector representation (k_x, k_y, k_z) . For our purposes, we aim to distill the notion of frequency into a one-dimensional scaling parameter. Namely, we define the scaling parameter k as the diameter of the expanding ball in Fourier space: $k = \sqrt{k_x^2 + k_y^2 + k_z^2}$. This definition of k provides a simple index that captures the overall scaling parameter of the frequency components in all directions. Moreover, we can

162 decompose the Fourier transform $\mathcal{F}\{\phi\}$ into groups indexed by the scaling parameter k :
 163

$$164 \quad \mathcal{F}\{\phi\}(k) = \bigcup_{\sqrt{k_x^2 + k_y^2 + k_z^2} = k} \mathcal{F}\{\phi\}(k_x, k_y, k_z). \quad (4)$$

$$165$$

166 Intuitively, $\mathcal{F}\{\phi\}(k)$ represents the set of all frequency components that share the same scaling
 167 parameter k . This grouping allows us to examine the contributions of various spatial frequencies of
 168 ϕ when viewed through the lens of frequency k . Furthermore, ϕ_k is just the summation of $\mathcal{F}\{\phi\}(k)$.
 169

170 On the other hand, we can recover ϕ from $\mathcal{F}\{\phi\}$, because the Fourier transform is an invertible
 171 operation: $\phi = \mathcal{F}^{-1}\mathcal{F}\{\phi\}$. Such an invertibility establishes the Fourier transform as a valid example
 172 of K -amplitude decomposition. For discrete data, which inherently possess one highest resolution,
 173 the variables (k_x, k_y, k_z) are situated on a discrete lattice rather than spanning the entire continuous
 174 space. Consequently, the scaling parameter k , derived from these discrete components, is itself
 175 discrete and bounded.

176 2.3 FLOW MATCHING

177 In this work, we primarily focus on the flow matching (FM) generative models and their families
 178 (Lipman et al., 2022; Liu et al., 2022b; Albergo & Vanden-Eijnden, 2022). In FM, the flow
 179 Ψ_t is defined by solutions of an ordinary differential equation (ODE) system with a time-dependent
 180 vector field \mathbf{v} :

$$182 \quad \frac{d}{dt} \Psi_t(x) = \mathbf{v}_t(\Psi_t(x)), \quad (5)$$

183 and we focus on the probability transport aspects of Ψ_t . In particular, the flow provides a means of
 184 interpolating between probability densities within the sample space. Suppose Ψ_t follows an initial
 185 probability p_0 , then for $t > 0$, Ψ_t induces a probability measure p_t : $p_t(B) = p_0(\Psi_t^{-1}(B))$, where
 186 B is a measurable set. Assume that Ψ_t is differentiable, and define a surrogate velocity at time t
 187 as $v_t(x_t, \theta)$ using a deep neural network with parameter θ . Then the vector field matching loss is
 188 defined as:

$$189 \quad \mathcal{L}_{\text{FM}} := \int \int_0^1 dx_0 dt \left\| \frac{d\Psi_t}{dt}(x_t) - v_t(x_t, \theta) \right\|^2. \quad (6)$$

$$190$$

191 By aligning the learned vector field with the true gradient field of the frequency decomposition, this
 192 loss function ensures robust approximation and reconstruction of the data. Additionally, every interpolation
 193 $\pi(x_0, x_1)$ with a time-continuous interpolating function $f_t(x_0, x_1)$ between probabilities
 p_0 and p_1 induces a vector field v_t through the continuity equation:

$$195 \quad \frac{\partial p_t(x_t)}{\partial t} = -\nabla_x (p_t(x_t) v_t(x_t)), \quad (7)$$

$$196$$

197 and v_t is explicitly expressed as: $v_t = \frac{1}{p_t} \mathbb{E}_{\pi(x_0, x_1)} \left[\frac{\partial f_t(x_0, x_1)}{\partial t} \right]$. Although explicit matching of v_t
 198 via the continuity equation is intractable, flow matching permits a conditional version:

$$199 \quad \mathcal{L}_{\text{CFM}} = \mathbb{E}_{\pi(x_0, x_1)} \int_0^t dt \left\| \frac{\partial f_t(x_0, x_1)}{\partial t} - v_t(x_t, \theta) \right\|^2 + \text{constant}. \quad (8)$$

$$200$$

$$201$$

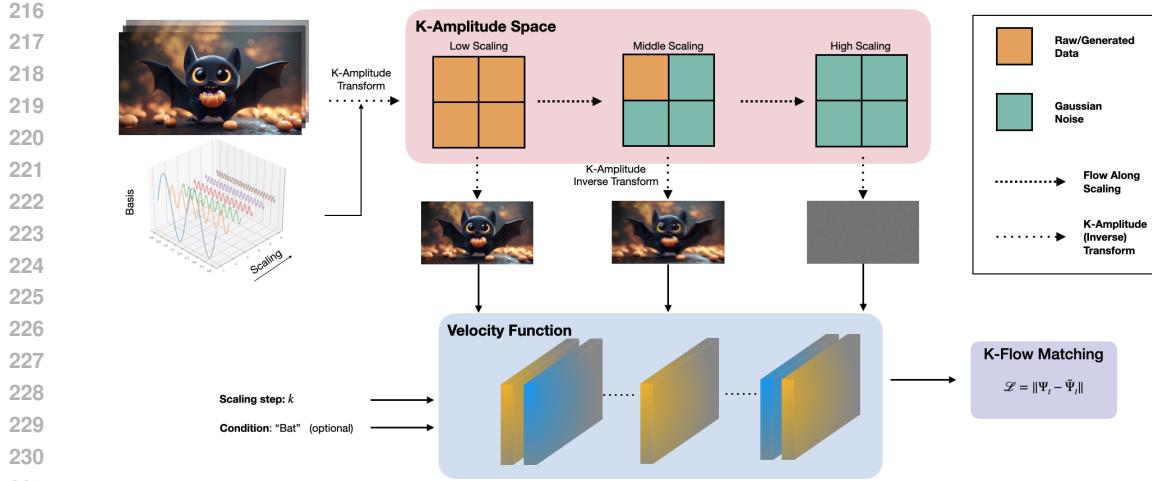
202 As detailed in Section 3, our framework reinterprets the time variable t as the scaling parameter k .
 203 Our goal is to construct a K -amplitude-respected $\pi(x_0, x_1)$ with differentiable functions f_k .

204 3 METHODOLOGY: K-FLOW

205 In this section, we introduce K-Flow. It is constructed from the collection of $\mathcal{F}\{\phi\}(k)$, indexed by
 206 a specific scaling parameter k . As we will demonstrate in Section 3.1, our approach is independent
 207 of the specific construction of the invertible transformation \mathcal{F} and the explicit definition of k . This
 208 flexibility enables us to extend to various K -amplitude decompositions.

209 3.1 K -AMPLITUDE INTERPOLANTS

210 According to the concept of stochastic interpolants (Albergo et al., 2023), all flow models can be
 211 viewed as constructing stochastic paths that interpolate between a known tractable prior distribution
 212 and an unknown target distribution, including flow matching (Lipman et al., 2022), rectified



270 where $\mu(t)$ is a bump function such that $\mu(0) = 1$, $\mu(1) = 0$ and $\mu'(0) = -\mu'(1)$. The antisymmetric property of $\mu'(t)$ ensures that Ψ_k is differentiable from k for all \mathbb{R}^+ , allowing the flow matching loss and other gradient-based optimization techniques. In Equation (10), we have three components:

- 274 1. $\mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} < \lfloor k \rfloor}$ applies to the amplitude components up to the integer part of k .
- 275 2. $\mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \geq \lfloor k \rfloor + 1}$ applies noise padding to components beyond the next integer.
- 276 3. $\mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1)}$ performs linear interpolation of the intermediate amplitude components based on the current t .

278 **Localized Vector Fields.** Instead of directly modeling Ψ_k , we pivot our focus to its conditional gradient field, $\frac{d\Psi_k}{dk}$. By concentrating on the gradient field, we facilitate a dynamic view of how ϕ_k evolves with respect to k . To derive an analytical expression of $\frac{d\Psi_k}{dk}$ conditioned on a given instance pair of data and noise: (ϕ, ϵ) , in what follows, we assume that \mathcal{F} is a linear transform. Then, following Equation (10), we have the conditional vector field:

$$284 \frac{d\Psi_k}{dk}(\phi, \epsilon) = \mathcal{F}^{-1} \left(\mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1)} \cdot \mu'(t)(\epsilon - \mathcal{F}\{\phi\}(k_x, k_y, k_z)) \right), \quad (11)$$

286 for $k \in [\lfloor k \rfloor, \lfloor k \rfloor + 1)$ and $t = k - \lfloor k \rfloor$. Then, following Equation (8), the training objective of 287 K-Flow is to learn the unconditional vector field in Equation (5) by the conditional flow matching:

$$289 \mathcal{L}_{\text{K-Flow}} := \mathbf{E}_{\phi_0} \int_0^K d\phi_0 \, dk \left\| \frac{d\Psi_k}{dk} - v_k(\Psi_k, \theta) \right\|^2. \quad (12)$$

292 By examining Equation (10) closely, we observe that the vector field is naturally localized around 293 a subset of points in the K -amplitude space that satisfy $\sqrt{k_x^2 + k_y^2 + k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1)$. This 294 localization means that the reconstruction at any given k primarily involves K -amplitude components 295 within a narrow frequency band around k . Compared with the flow scheme in the pixel space, the 296 K -amplitude in K-Flow reduces the optimization complexity by restricting the conditional vector 297 field to be within a sub-manifold for each k . This sub-manifold is of low dimensionality, allowing 298 for more focused updates and reducing the optimization space's dimensionality at each step. We will 299 investigate how this localized conditional vector field affects the generation path in Appendix B. The 300 inference computational complexity of our method is discussed in appendix D.

301 We can further generalize the interpolation interval from $(\lfloor k \rfloor, \lfloor k \rfloor + 1)$ to (k_m, k_n) , where k_m 302 and k_n are two integers such that $k_m < k_n$. This adjustment broadens the range for intermediate 303 amplitude components from $\sqrt{k_x^2 + k_y^2 + k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1)$ to $\sqrt{k_x^2 + k_y^2 + k_z^2} \in [k_m, k_n)$. For 304 example, for our experiments, we partition the K -amplitude into two or three parts. See Appendix D 305 for detailed implementations of these partitioning strategies.

308 3.2 EXAMPLES OF K -AMPLITUDE TRANSFORMATION

309 As we can see from Equation (1), all K -amplitude decompositions are achieved through expansion 310 across a complete set of basis functions. However, the behavior of a K -amplitude decomposition 311 (transform) can vary significantly depending on the choice of basis functions. Besides the Fourier 312 transform introduced in Section 2, we provide two representative examples of K -amplitude decom- 313 position: Wavelet transformation, and PCA transformation. More details are in Appendix D.

314 **Wavelet Transform.** Wavelet decomposition (transform) deals with data that are not only scaling- 315 localized but also spatially localized. The scaling parameter of wavelet transform is closely related 316 to the notion of multi-resolution analysis (Mallat, 1989), which provides a systematic way to de- 317 compose a signal into approximations and details at successively finer scales. This hierarchical de- 318 composition is achieved through a set of scaling functions $\omega(x)$, and wavelet functions $\psi(x)$, which 319 together serve as basis functions for the wavelet transformation. More precisely, in the wavelet trans- 320 form, a signal $f(t)$ is expressed as a sum of scaled and translated versions of these basis functions 321 times the corresponding coefficients c and d :

$$323 f(t) = \sum_j c_{k_0, j} \omega_{k_0, j}(t) + \sum_{k \geq k_0} \sum_j d_{k, j} \psi_{k, j}(t), \quad (13)$$

324 where $\omega_{k_0,j}(t)$ and $\psi_{k,j}(t)$ are the scaled and translated scaling and wavelet functions, respectively.
 325 The index j , which originally denotes the translation parameter, groups the basis within each fixed
 326 scaling parameter k naturally. Let $\phi_k := \sum_j d_{k,j} \psi_{k,j}$ for $k > k_0$ and $\phi_k := \sum_j c_{k_0,j} \omega_{k_0,j}$ for
 327 $k = k_0$, then eq. 13 is just one realization of K -amplitude decomposition. Concrete formulas for
 328 different families of wavelet bases, such as Daubechies (db), can be found in Appendix D.
 329

330 In this article, we employ the discrete version of wavelet transform (DWT) as our K -amplitude
 331 transformation \mathcal{F} , which shares the linearity property with the Fourier transform with a bounded
 332 scaling parameter k , providing a structured yet flexible means of decomposing discrete data.
 333

334 **Date-dependent PCA Transform.** Note that Fourier and wavelet decompositions are nonpara-
 335 metric k -amplitude decompositions that are independent of data. While these transformation
 336 methods are powerful, we also aim to find data-dependent decompositions that can capture common
 337 characteristic features specific to a dataset. This motivation leads to principal component analysis
 338 (PCA), a technique widely used for the low-dimensional approximation of the data manifold and
 339 vision features (Izenman, 2012; Chen et al., 2024). Please consult Appendix D for the K -amplitude
 340 realization of PCA transform.
 341

342 4 EXPERIMENTS

343 We conduct a comprehensive experimental evaluation of K-Flow focusing on its technical innova-
 344 tions in k -amplitude adapted generation. Using standard backbone architectures, we perform exten-
 345 sive experiments spanning image and molecular assembly generation tasks and k scaling guidance
 346 editing and restoration. Complete implementation details, experimental configurations, ablation
 347 studies, and scientific generation tasks are provided in Appendix E (Algorithm 1) and Appendix D.
 348

349 4.1 IMAGE UNCONDITIONAL AND CONDITIONAL GENERATION

350 **Image Unconditional Generation** The first
 351 task is to generate random samples after fit-
 352 ting a target data distribution, which is typi-
 353 cally concentrated around a low-dimensional sub-
 354 manifold within the ambient space.
 355

356 **Dataset and Metrics.** We conduct experiments
 357 on the CelebA-HQ (Karras, 2017) dataset with
 358 the resolution of 256×256 . To evaluate the
 359 performance of our proposed method, we em-
 360 ploy two metrics: the Fréchet Inception Dis-
 361 tance (FID) (Heusel et al., 2017), which evaluates
 362 the quality by measuring the statistical similarity
 363 between generated and real images, and Recall
 364 (Kynkänniemi et al., 2019), which measures the diversity of the generated images.
 365

366 **Results.** Table 1 summarizes the comparison between our proposed K-Flow model and other gen-
 367 erative models. For a fair comparison, both the baseline ordinary flow matching (LFM (Dao et al.,
 368 2023)) and our K-Flow flow utilize the same VAE’s latent from (Rombach et al., 2022) and the Dif-
 369 fusion Transformer with the same size (e.g., DIT L/2 (Peebles & Xie, 2023)) as the backbone model.
 370 We can observe that (1) K-Flow achieves the best performance in FID, especially w/ the db6-based
 371 wavelet K-Flow. (2) Although the latent diffusion model (Rombach et al., 2022) gets the highest
 372 score in Recall (diversity), the Fourier and PCA-based K-Flow is comparable with the ordinary la-
 373 tent flow matching. We also evaluated the sensitivity of K-Flow to frequency resolution by varying
 374 the discretization levels of the scaling parameter k on the LSUN Church dataset (Yu et al., 2015),
 375 with comprehensive results provided in Appendix F.
 376

377 **Image Class-conditional Generation** Then we explore how K -amplitude decomposition behaves
 378 when the generation path is conditioned on class labels, where the class label (e.g., dog, cat, fish, etc)
 379 delegates the low-scaling information of each image. This investigation could potentially pave the
 380 way for multi-scaling control, where different scaling components are influenced by specific caption
 381 information. We list the detailed class-conditional generation algorithm in Appendix E.
 382

383 Table 1: Unconditional generation on CelebA-HQ.

Model	FID \downarrow	Recall \uparrow
K-Flow, Fourier-DiT L/2 (Ours)	5.11	0.47
K-Flow, Wave-DiT L/2 (Ours)	4.99	0.46
K-Flow, PCA-DiT L/2 (Ours)	5.19	0.48
LFM, ADM (Dao et al., 2023)	5.82	0.42
LFM, DiT L/2 (Dao et al., 2023)	5.28	0.48
FM (Lipman et al., 2022)	7.34	-
LDM (Rombach et al., 2022)	5.11	0.49
LSGM (Vahdat et al., 2021)	7.22	-
WaveDiff (Phung et al., 2023)	5.94	0.37
DDGAN (Xiao et al., 2021)	7.64	0.36
Score SDE (Song et al., 2020)	7.23	-

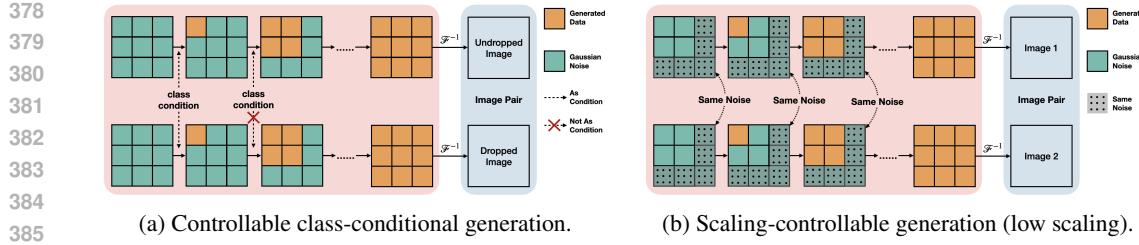


Figure 3: Pipeline of two ablation studies on controllable generation.

Dataset and Metric. We use ImageNet as the middle-size conditional generation dataset (Deng et al., 2009). Beyond evaluating the unconditional FID for the ImageNet dataset, we are also interested in studying how the class control interacts with scaling generation in a quantitative manner.

Results. The results are presented in Table 2. Our primary focus for the FID metric is the classifier-free guidance inference method applied to flow matching models. The data indicates that K-Flow achieves results comparable to LFM. In terms of the recall metric, which assesses the diversity of the generated distribution, our model outperforms the standard LFM. This improvement may be attributed to the fact that the inference path of K-Flow includes a greater number of dimensions during the low-scaling period, as discussed in Appendix B.1.

4.2 IMAGE CONTROLLABLE CLASS-CONDITIONAL GENERATION

The latent flow matching model can implicitly learn low- and high-resolution features (Dao et al., 2023), but the boundary between each resolution is vague, and we cannot explicitly determine which timestep in the inference process corresponds to a specific resolution or frequency. In comparison, our proposed K-Flow enables finer-grained controllable generation. As demonstrated in Figures 3a and 4, K-Flow effectively preserves high-frequency details even when class conditions are omitted during the last 70% of scaling steps, whereas ordinary latent flow exhibits significant blurring. To quantitatively validate this observation, we analyze the conditional discrimination ratio (CDR, formally defined in Appendix F). From Table 2, our model maintains a CDR close to one, indicating robust performance regardless of high-scale condition omission, while conventional LFM shows significantly higher CDR, suggesting performance degradation. These results confirm that K-Flow’s K -amplitude-aware architecture enables more efficient computational resource allocation by exploiting the natural correlation between scaling parameters and the class-label conditional information.

4.3 IMAGE SCALING-CONTROLLABLE GENERATION AND RESTORATION

Our method guarantees that the generation path is disentangled with respect to k (check Appendix B). This allows us to control initial noise at each scaling level (see Algorithm 2), enabling unsupervised editing of different scaling components. We also conduct ablation studies on **Preserving Low Scaling, Modifying High Scaling**. Please check Appendix F for more details and visualizations.

Preserving High Scaling, modifying Low Scaling. This scaling-controllable generation pipeline is illustrated in Figure 3b. It involves sampling multiple images while ensuring that the noise in the high-scaling components remains consistent across all samples. In scaling-controllable image generation, the

Table 2: Class-conditional generation on ImageNet.

Model	FID \downarrow	CDR \downarrow	Recall \uparrow
K-Flow, Wave-DiT L/2 (Ours)	17.8	-	0.56
+ cfg=1.5	4.49	-	0.44
K-Flow, Fourier-DiT L/2 (Ours)	13.0	-	0.57
+ cfg=1.5	2.64	1.49	0.45
LFM, DiT L/2	14.0	-	0.56
+ cfg=1.5	2.77	3.25	0.42
LDM-8 (Rombach et al., 2022)	15.51	-	0.63
LDM-8-G	7.76	-	0.35
DiT-B/2 Peebles & Xie (2023)	43.47	-	-



Figure 5: Results of scaling-controllable generation: Each column presents an image pair that shares high-frequency components while exhibiting distinct low-frequency characteristics.

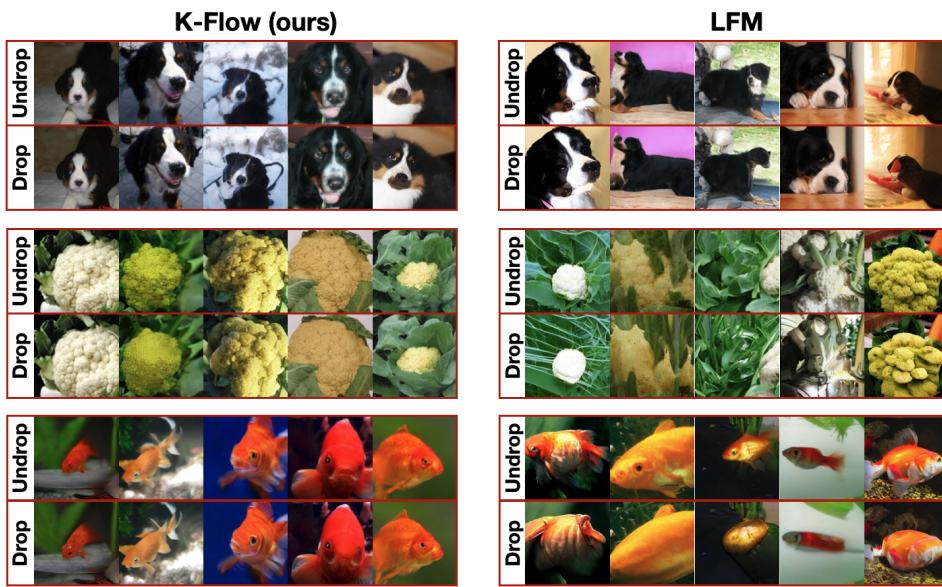


Figure 4: Results of controllable class-conditional generation. ‘Drop’ means we drop the class conditions during the last 70% scaling steps, while ‘undrop’ means we keep the condition all the time.

goal is to maintain consistency in the high-scaling details while allowing variations in the low-scaling context among the generated images, thus this allows K-Flow to achieve unsupervised steerability in a finetuning-free manner. The results on CelebA are presented in Figure 5, where we apply a pretrained Daubechies wavelet (db6-based) K-Flow. It can be observed that facial details, such as eyes, smiles, noses, and eyebrows, remain consistent within each group of images. In contrast, the low-scaling components, including background, gender, age, and hairstyle, vary across the images within the same group. These qualitative results demonstrate how frequency bands naturally correspond to semantic features - facial details persist in high-frequency components while attributes like global background and overall appearance vary in low-frequency components. When applying the same editing protocol to conventional LFM, the results show no such organized frequency-semantic correspondence (Figure S10), highlighting the advantage of K-Flow.

Image Restoration. While traditional semantic metrics cannot directly assess unsupervised frequency-based editing, we quantitatively validate our scaling-aware generation path via image restoration tasks, where frequency-specific changes are objectively measurable. This includes super-resolution, inpainting, and deblurring. From Table S8 in Appendix, K-Flow achieves state-of-the-art performance in terms of PSNR and SSIM metrics on the CelebA dataset, while requiring only half the iterations compared to the baseline PnP-flow (Martin et al., 2024). This demonstrates that our frequency-domain formulation enables more efficient and accurate high-frequency reconstruction. Detailed experimental settings and algorithms are provided in Appendix F.

5 CONCLUSION

In this paper, we introduce K-Flow Matching (K-Flow), an efficient flow-matching model that flows along the K -amplitude for generative modeling. K-Flow naturally generalizes the multi-scales of data (e.g., multi-resolution or frequencies in images) to multi-scales in the K -amplitude space.

Future Directions. As we have verified the effectiveness of K-Flow exclusively on image generation tasks, moving forward, two promising directions are worth exploring. (1) Multimodal generation: This includes tasks such as large-scale data generation guided by dense captions, which could better showcase the steerability of K-Flow by aligning images with natural language inputs. (2) We outline six properties of K-Flow in Section 1, *e.g.*, the amplitude naturally corresponds to energy. While Section 3 briefly discusses how energy is represented in K-Flow, this aspect has not been explored in depth. We believe that such energy term holds potential for integration with the utility of energy-based models in future work.

486
487
ETHICS STATEMENT488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
948
949
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
989
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1097
1098
1099
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1189
1190
1191
1192
1193
1194
1195
1196
1197
1197
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1297
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1389
1390
1391
1392
1393
1394
1395
1396
1397
1397
1398
1399
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1489
1490
1491
1492
1493
1494
1495
1496
1497
1497
1498
1499
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1589
1590
1591
1592
1593
1594
1595
1596
1597
1597
1598
1599
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1689
1690
1691
1692
1693
1694
1695
1696
1697
1697
1698
1699
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1789
1790
1791
1792
1793
1794
1795
1796
1797
1797
1798
1799
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1889
1890
1891
1892
1893
1894
1895
1896
1897
1897
1898
1899
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1989
1990
1991
1992
1993
1994
1995
1996
1997
1997
1998
1999
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2089
2090
2091
2092
2093
2094
2095
2096
2097
2097
2098
2099
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2189
2190
2191
2192
2193
2194
2195
2196
2197
2197
2198
2199
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2289
2290
2291
2292
2293
2294

540 James Alexander Chisholm and Sam Motherwell. Compack: a program for identifying crystal
 541 structure similarity using distances. *Journal of applied crystallography*, 38(1):228–231, 2005.
 542

543 Quan Dao, Hao Phung, Binh Nguyen, and Anh Tran. Flow matching in latent space. *arXiv preprint*
 544 *arXiv:2307.08698*, 2023.

545 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-
 546 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,
 547 pp. 248–255. Ieee, 2009.

548

549 Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a
 550 laplacian pyramid of adversarial networks. *Advances in neural information processing systems*,
 551 28, 2015.

552 S. Dieleman. Diffusion is spectral autoregression. 2024. URL <https://sander.ai/2024/09/02/spectral-autoregression.html>.

553

554 Weitao Du, He Zhang, Tao Yang, and Yuanqi Du. A flexible diffusion model. In *International
 555 Conference on Machine Learning*, pp. 8678–8696. PMLR, 2023.

556

557 Weinan Ee. A proposal on machine learning via dynamical systems. *Communications in Mathematics and Statistics*, 5:1–11, 02 2017. doi: 10.1007/s40304-017-0103-z.

558

559 Carlos Esteves, Mohammed Suhail, and Ameesh Makadia. Spectral image tokenizer. *arXiv preprint*
 560 *arXiv:2412.09607*, 2024.

561

562 David J. Field. Relations between the statistics of natural images and the response properties of cor-
 563 tical cells. *J. Opt. Soc. Am. A*, 4(12):2379–2394, 12 1987. doi: 10.1364/JOSAA.4.002379. URL
 564 <https://opg.optica.org/josaa/abstract.cfm?URI=josaa-4-12-2379>.

565

566 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
 567 Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information
 568 processing systems*, 27, 2014.

569

570 Saulius Grazulis, Daniel Chateigner, Robert T. Downs, Alex F. T. Yokochi, Manuel Quirós, Luca
 571 Lutterotti, Elena Manakova, Justinas Butkus, Peter Moeck, and Armel Le Bail. Crystallography
 572 open database – an open-access collection of crystal structures. *Journal of Applied Crystallography*,
 573 42(4):726–729, 2009. doi: 10.1107/S0021889809016690.

574

575 Jitao Gu, Shuangfei Zhai, Yizhe Zhang, Miguel Angel Bautista, and Josh Susskind. f-dm: A multi-
 576 stage diffusion model via progressive signal transformation. *arXiv preprint arXiv:2210.04955*,
 577 2022.

578

579 Hongyu Guo, Yoshua Bengio, and Shengchao Liu. Assembleflow: Rigid flow matching with in-
 580 erial frames for molecular assembly. In *The Thirteenth International Conference on Learning
 581 Representations*, 2025. URL <https://openreview.net/forum?id=jckKNzYYA6>.

582

583 Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle
 584 for unnormalized statistical models. In *Proceedings of the thirteenth international conference on
 585 artificial intelligence and statistics*, pp. 297–304. JMLR Workshop and Conference Proceedings,
 586 2010.

587

588 Alfred Haar. Zur theorie der orthogonalen funktionensysteme. *Mathematische Annalen*, 71(1):
 589 38–53, 1911.

590

591 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter.
 592 Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in
 593 neural information processing systems*, 30, 2017.

594

595 Geoffrey E Hinton. Training products of experts by minimizing contrastive divergence. *Neural
 596 computation*, 14(8):1771–1800, 2002.

597

598 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
 599 neural information processing systems*, 33:6840–6851, 2020.

594 Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Gradient step denoiser for convergent
 595 plug-and-play. *arXiv preprint arXiv:2110.03220*, 2021.
 596

597 Alan Julian Izenman. Introduction to manifold learning. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(5):439–446, 2012.
 598

599 Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song,
 600 Yadong Mu, and Zhouchen Lin. Pyramidal flow matching for efficient video generative modeling.
 601 *arXiv preprint arXiv:2410.05954*, 2024.
 602

603 Jalal Karam. On the zeros of daubechies orthogonal and biorthogonal wavelets. *Applied Mathematics*, 03:778–787, 01 2012. doi: 10.4236/am.2012.37116.
 604

605 Tero Karras. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
 606

607 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Proc. NeurIPS*, 2022.
 608

609 Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021.
 610

611 Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
 612

613 Theodoros Kouzelis, Ioannis Kakogeorgiou, Spyros Gidaris, and Nikos Komodakis. Eq-vae: Equivariance regularized latent space for improved generative image modeling, 2025. URL <https://arxiv.org/abs/2502.09509>.
 614

615 Tuomas Kynkänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved
 616 precision and recall metric for assessing generative models. *Advances in neural information
 617 processing systems*, 32, 2019.

618 Hugo Lavenant and Filippo Santambrogio. The flow map of the fokker–planck equation does not
 619 provide optimal transport. *Applied Mathematics Letters*, 133:108225, 2022.
 620

621 Yann LeCun, Sumit Chopra, Raia Hadsell, M Ranzato, Fujie Huang, et al. A tutorial on energy-based learning. *Predicting structured data*, 1(0), 2006.
 622

623 Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image
 624 generation using residual quantization. In *Proceedings of the IEEE/CVF Conference on Computer
 625 Vision and Pattern Recognition*, pp. 11523–11532, 2022.
 626

627 Jiarui Lei, Xiaobo Hu, Yue Wang, and Dong Liu. Pyramidflow: High-resolution defect contrastive
 628 localization using pyramid normalizing flow. In *Proceedings of the IEEE/CVF conference on
 629 computer vision and pattern recognition*, pp. 14143–14152, 2023.
 630

631 Zhengqi Li, Richard Tucker, Noah Snavely, and Aleksander Holynski. Generative image dynamics.
 632 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 633 24142–24153, 2024a.
 634

635 Zongjian Li, Bin Lin, Yang Ye, Liuhan Chen, Xinhua Cheng, Shenghai Yuan, and Li Yuan. Wf-vae:
 636 Enhancing video vae by wavelet-driven energy flow for latent video diffusion model. *arXiv preprint arXiv:2411.17459*, 2024b.
 637

638 Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching
 639 for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
 640

641 Shengchao Liu, Hongyu Guo, and Jian Tang. Molecular geometry pretraining with se (3)-invariant
 642 denoising distance matching. *arXiv preprint arXiv:2206.13602*, 2022a.
 643

644 Shengchao Liu, Yanjing Li, Zhuoxinran Li, Zhiling Zheng, Chenru Duan, Zhi-Ming Ma, Omar
 645 Yaghi, Animashree Anandkumar, Christian Borgs, Jennifer Chayes, et al. Symmetry-informed
 646 geometric representation for molecules, proteins, and crystalline materials. *Advances in neural
 647 information processing systems*, 36, 2024a.

648 Shengchao Liu, Divin Yan, Hongyu Guo, and Anima Anandkumar. Equivariant flow matching
 649 framework for learning molecular cluster crystallization. In *ICML 2024 Workshop on*
 650 *Geometry-grounded Representation Learning and Generative Modeling*, 2024b. URL <https://openreview.net/forum?id=lCVqpQvr41>.

651

652 Sifei Liu, Shalini De Mello, and Jan Kautz. Cosae: Learnable fourier series for image restoration.
 653 In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

654

655 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
 656 transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022b.

657

658 Erik Luijten and Henk WJ Blöte. Finite-size scaling and universality above the upper critical di-
 659 mensionality. *Physical review letters*, 76(10):1557, 1996.

660 Stephane G Mallat. Multiresolution approximations and wavelet orthonormal bases of $l^2(r)$. *Trans-
 661 actions of the American Mathematical Society*, 315(1):69–87, 1989.

662

663 Sérgolène Martin, Anne Gagneux, Paul Hagemann, and Gabriele Steidl. Pnp-flow: Plug-and-play
 664 image restoration with flow matching. *arXiv preprint arXiv:2410.02423*, 2024.

665

666 Wael Mattar, Idan Levy, Nir Sharon, and Shai Dekel. Wavelets are all you need for autoregressive
 667 image generation. *arXiv preprint arXiv:2406.19997*, 2024.

668

669 Yves Meyer. *Ondelettes et opérateurs. I: Ondelettes*, 1990.

670

671 Yves Meyer. *Wavelets and operators: volume 1*. Number 37. Cambridge university press, 1992.

672

673 George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lak-
 674 shminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine
 675 Learning Research*, 22(57):1–64, 2021.

676

677 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of
 678 the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.

679

680 Hao Phung, Quan Dao, and Anh Tran. Wavelet diffusion models are fast and scalable image gener-
 681 ators. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 682 pp. 10199–10208, 2023.

683

684 Ashwini Pokle, Matthew J Muckley, Ricky TQ Chen, and Brian Karrer. Training-free linear image
 685 inverses via flows. *arXiv preprint arXiv:2310.04432*, 2023.

686

687 Alec Radford. Improving language understanding by generative pre-training. 2018.

688

689 Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with
 690 vq-vae-2. *Advances in neural information processing systems*, 32, 2019.

691

692 Sucheng Ren, Qihang Yu, Ju He, Xiaohui Shen, Alan Yuille, and Liang-Chieh Chen. Flowar: Scale-
 693 wise autoregressive image generation meets flow matching. *arXiv preprint arXiv:2412.15205*,
 2024.

694

695 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 696 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-
 697 ence on computer vision and pattern recognition*, pp. 10684–10695, 2022.

698

699 Dohoon Ryu and Jong Chul Ye. Pyramidal denoising diffusion probabilistic models. *arXiv preprint
 700 arXiv:2208.01864*, 2022.

701

702 Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad
 703 Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern anal-
 704 ysis and machine intelligence*, 45(4):4713–4726, 2022.

705

706 Ivan Skorokhodov, Sharath Girish, Benran Hu, Willi Menapace, Yanyu Li, Rameen Abdal, Sergey
 707 Tulyakov, and Aliaksandr Siarohin. Improving the diffusability of autoencoders. *arXiv preprint
 708 arXiv:2502.14831*, 2025.

702 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
 703 *Advances in neural information processing systems*, 32, 2019.
 704

705 Yang Song and Diederik P Kingma. How to train your energy-based models. *arXiv preprint*
 706 *arXiv:2101.03288*, 2021.

707 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 708 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
 709 *arXiv:2011.13456*, 2020.

710 Qiao Sun, Zhicheng Jiang, Hanhong Zhao, and Kaiming He. Is noise conditioning necessary for
 711 denoising generative models?, 2025. URL <https://arxiv.org/abs/2502.13129>.

713 Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. Visual autoregressive modeling:
 714 Scalable image generation via next-scale prediction. *arXiv preprint arXiv:2404.02905*, 2024.

715 Hoai-Chau Tran, Duy MH Nguyen, Duy M Nguyen, Trung-Tin Nguyen, Ngan Le, Pengtao Xie,
 716 Daniel Sonntag, James Y Zou, Binh T Nguyen, and Mathias Niepert. Accelerating transformers
 717 with spectrum-preserving token merging. *arXiv preprint arXiv:2405.16148*, 2024.

719 Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In
 720 *Neural Information Processing Systems (NeurIPS)*, 2021.

721 van A Van der Schaaf and JH van van Hateren. Modelling the power spectra of natural images:
 722 statistics and information. *Vision research*, 36(17):2759–2770, 1996.

724 Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*,
 725 23(7):1661–1674, 2011.

727 Yair Weiss and William T Freeman. What makes a good model of natural images? In *2007 IEEE*
 728 *conference on computer vision and pattern recognition*, pp. 1–8. IEEE, 2007.

729 Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with
 730 denoising diffusion gans. *arXiv preprint arXiv:2112.07804*, 2021.

732 Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun:
 733 Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv*
 734 *preprint arXiv:1506.03365*, 2015.

735 Yasi Zhang, Peiyu Yu, Yaxuan Zhu, Yingshan Chang, Feng Gao, Ying Nian Wu, and Oscar Leong.
 736 Flow priors for linear inverse problems via iterative corrupted trajectory matching. *arXiv preprint*
 737 *arXiv:2405.18816*, 2024.

739 Qinqing Zheng, Matt Le, Neta Shaul, Yaron Lipman, Aditya Grover, and Ricky TQ Chen. Guided
 740 flows for generative modeling and decision making. *arXiv preprint arXiv:2311.13443*, 2023.

741 Zhenyu Zhou, Defang Chen, Can Wang, and Chun Chen. Fast ode-based sampling for diffusion
 742 models in around 5 steps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
 743 *Pattern Recognition*, pp. 7777–7786, 2024.

744 Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc
 745 Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the*
 746 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1219–1229, 2023.

747

748

749

750

751

752

753

754

755

756 THE USE OF LARGE LANGUAGE MODELS
757758 We employ an LLM to refine the language and edit the draft of this paper, including:
759760 • Correcting grammatical errors, punctuation, and spelling.
761 • Improving sentence structure to enhance clarity, flow, and readability.
762 • Suggesting alternative phrasing for more precise and professional academic expression.763 All modifications suggested by the LLM were critically reviewed, vetted, and approved by the au-
764 thors. The final text accurately reflects our own ideas, arguments, and research findings.
765766 A PROPERTIES OF K-FLOW
767768 (a) K-Flow provides **a first-principle way to organize the scaling k** . Unlike perception-based
769 computer vision tasks, which often favor certain scaling (frequency) bands, a K -amplitude based
770 generative model strives for an optimal organization of all scalings to ensure that the final generated
771 sample is of high fidelity. By constructing K -amplitude scaling-based vector fields, the integrated
772 flow naturally incorporates all scaling information, and the conditional flow matching training ob-
773 jective provides a perfect trade-off of accuracy-efficiency inside localized scalings. We will also
774 demonstrate how different discretizations of K-Flow with related works, highlighting the connec-
775 tions and integrations with existing methods in the field.776 (b) K-Flow enables **multi-scale modeling in the K -amplitude space**. Compared to the original data
777 space, such as the pixel space in images, the K -amplitude space provides a more natural perspective
778 for defining and analyzing multi-scale information, namely, K -amplitude decomposition empowers
779 K-Flow for effective multi-scale modeling. By decomposing the feature representation into multiple
780 scaling components in the K -amplitude space, K-Flow associates each scaling with an amplitude.
781 Higher values of K -amplitude correspond to higher-frequency information, capturing fine-grained
782 details, while lower values encode lower-frequency information, representing more coarse-grained
783 features. Let us take the image for illustration. Images inherently exhibit a hierarchical structure,
784 with information distributed across various resolution levels. Low-resolution components capture
785 global shapes and background information, while high-resolution components encode fine details
786 like textures, often sparse and localized. By projecting these components into the K -amplitude
787 space, K-Flow captures such hierarchical information effectively and naturally, enabling precise
788 modeling of the interplay between scales.789 (c) K-Flow supports **a well-defined scale along with energy**. The amplitude is also used to reflect
790 the *energy* level at each scale of the data. In physics, it is proportional to the square of the amplitude.
791 In comparison, for the modeling on the original data space, though we can inject application-specific
792 inductive bias, such as multiple pixel resolutions for images, they do not possess a natural energy
793 concept.794 (d) K-Flow interprets **scaling as time**. From elucidating the design space of the traditional flow
795 matching perspective, K-Flow re-defines the artificial time variable (or the signal-to-noise ratio vari-
796 able proposed in (Kingma et al., 2021)) as the ordering index of frequency space. In this context, the
797 artificial time variable effectively controls the traversal through different levels of a general notion
798 of frequency decompositions, scaling each frequency component appropriately. This perspective
799 aligns with the concept of renormalization in physical systems, where behavior across scales is sys-
800 tematically related.801 (e) K-Flow supports the **fusion of intra-scaling and inter-scaling modeling**. K-Flow flows across
802 scaling as time, and namely, K-Flow naturally merges the intra- and inter-scaling during the flow
803 process. Thus the key module turns to the smooth interpolant, as will be introduced in Section 3.
804 This is in comparison with existing works on multi-modal modeling (Burt & Adelson, 1987; Tian
805 et al., 2024; Atzmon et al., 2024), where the special design of the intra-scaling and inter-scaling is
806 required.807 (f) K-Flow supports **explicit steerability**. The flow process across scales enables K-Flow to control
808 the information learned at various hierarchical levels. This, in turn, allows finer-grained control of
809 the generative modeling, facilitating more precise and customizable outputs. By understanding and
leveraging K-Flow’s steerability, its utility can be significantly enhanced across diverse domains,

including Artificial Intelligence-Generated Content (AIGC), AI-driven scientific discovery, and the safe, responsible development of AI technologies.

B DISCUSSION

B.1 FROM CONDITIONAL TO UNCONDITIONAL PATH IN K-FLOW

In Section 3, our frequency-localized path is defined at the conditional level ($\frac{d\Psi_k}{dk}(\phi, \epsilon)$), and it is only related to the unconditional vector field ($v_k(\Psi_k, \theta)$ in eq. (12)) through the equivalence of conditional flow matching and unconditional flow matching at the loss level (Lipman et al., 2022). In this section, we try to study the splitting property of the unconditional K -amplitude vector field.

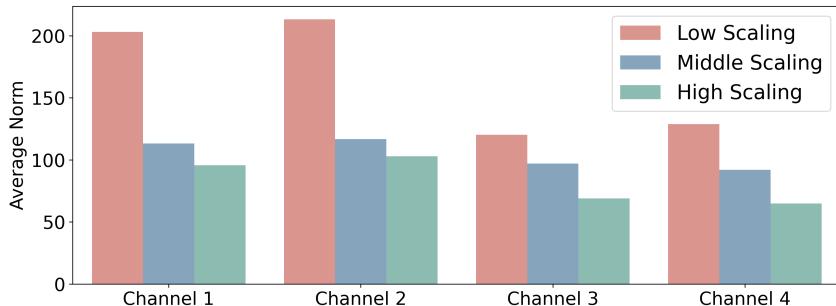
By the K -amplitude decomposition, the transformed data probability p_{data} satisfies the telescoping property:

$$p_{data} = p(k_0)p(k_1|k_0)\dots p(k_{max}|k_{max}-1, \dots, k_0), \quad (14)$$

with k_0 and k_{max} denoting the lowest and highest scaling. Then, according to the definition of our proposed K-Flow Ψ_k , the interpolated probability at scaling step t is also localized:

$$p_t(\cdot) = p(k_0) \cdots p_t(\cdot | \lfloor k \rfloor, \dots, k_0) p_\epsilon(\lfloor k \rfloor + 1) \cdots p_\epsilon(k_{max} | k_{max} - 1, \dots, k_0), \quad (15)$$

where p_ϵ denotes the distribution of the initial noise and $t \in [|k|, |k| + 1]$. Combining Equation (15), the localization property of the bump function, and Lemma 1 of (Zheng et al., 2023), the unconditional vector field has an explicit form: $v_t(\Psi_k) = a_t \cdot \Psi_k + b_t \nabla \log p_t(\Psi_k)$, where a_t and b_t are hyper-parameters determined by the bump function we choose.



Supplementary Figure S1: On the low-scaling hypothesis. The graph illustrates the relative norm distribution for each scaling component as defined by the wavelet decomposition in the latent space. It can be observed that the low-scaling component exhibits a significantly higher norm (energy), nearly twice that of the high-scaling component.

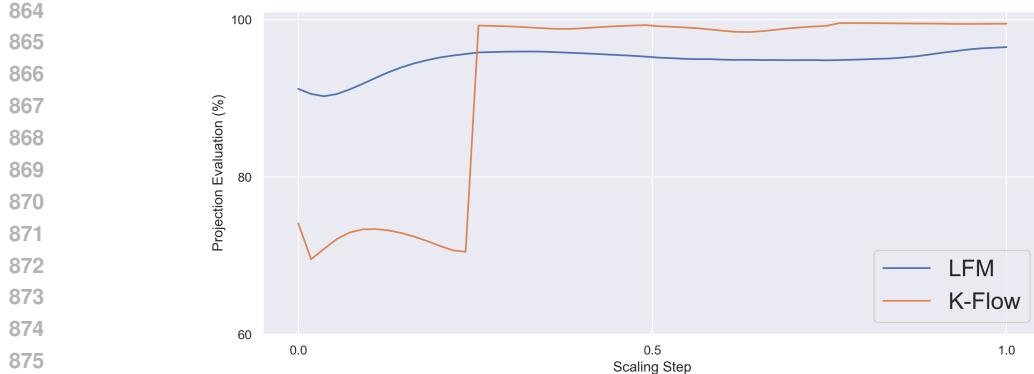
Noise Splitting. A key characteristic of flow models is their deterministic nature after the initial noise sampling. Specifically, once the initial noise is sampled, the flow follows a fixed path to generate the final data sample. According to Equation (15), during scaling step t : (1) the scaling components below $[k]$ remain unchanged; (2) the scaling components above $[k]$ remain unchanged; (3) The distribution of higher scaling components maintains the same characteristics as their initial noise distribution.

By these observations, we now investigate how segmented initial noise in the K-Flow space influences the final output of the K-Flow flow. Suppose we discretize scaling parameter k into two parts: $\mathcal{F}\{\Psi_k\} = \{\phi_{\text{low}}(k), \phi_{\text{high}}(k)\}$. When flowing along the low-scaling component, the vector field v_k can be re-expressed in a conditional form:

$$v_k(\Psi_k) = v_k(\phi_{\text{low}}(k), c) \quad (16)$$

where constant c represents the (static) initial noise for the high-scaling part. This noise-conditioned property in the k -amplitude domain leads us to explore whether fixing the high-scaling noise and altering the low-scaling noise allows for unsupervised editing of relative low-scaling semantics in an image. Indeed, we observed this phenomenon, the qualitative results will be discussed in section 4.3.

From Figure 5, we observe that a targeted common high-scaling initial noise guides our K-Flow flow toward generating human faces with similar detail but varying low-level content. See the experiment section for a more detailed analysis.



Supplementary Figure S2: Projection Error Comparison with Different flow Models. The graph illustrates the PCA projection errors of two trained models throughout the entire inference process, with distinct segments marked by dashed lines. The red and blue lines represent the original latent flow matching (LFM) and the K-Flow with two amplitude components, respectively. The projection error is quantified by the reconstruction error for each generation step from the PCA compression, using the first two principal components. Owing to the scaling-aware nature of our flow, the low-amplitude portion (the initial part of the curve) resides in a relatively high-dimensional space, resulting in higher projection errors for the two-dimensional PCA projection.

B.2 THE EFFECT OF SCALING STEP k FOR IMAGE RECONSTRUCTION

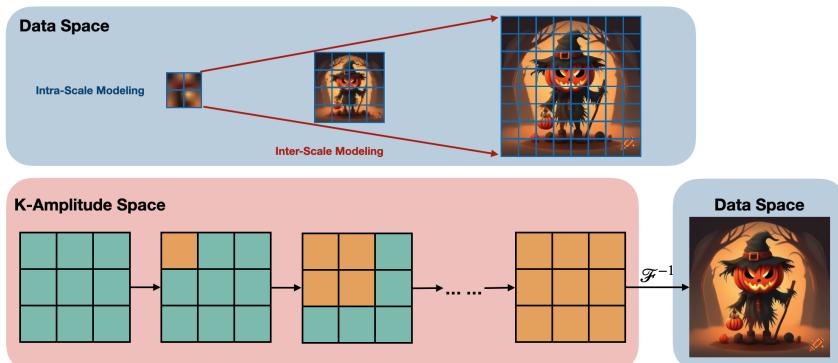
K-Flow's ability to leverage the low-dimensional structure of data is primarily enabled by its K-Flow localization property. This enables a strategic path through low-dimensional spaces, which can be directly compared with the generation path of conventional flow models. In our model, this path incorporates an explicit frequency hierarchy, which hypothesizes that the low-frequency components - concentrated in the earlier stages of the model - may share more dimensions in common, particularly from a semantic perspective, than the high-frequency components positioned later in the generative process. Conversely, an ordinary flow model may exhibit a more uniform distribution of dimensionality across the entire generative path.

Motivated by this hypothesis, we conduct a case study using PCA to approximate the dimension of the generation trajectory $\{\Psi_k\}_{k=k_0}^{k_{max}}$. As illustrated in Figure S2, we measure how closely the dimension of the generation path aligns with a two-dimensional subspace spanned by the first two components of the model's PCA decomposition, denoted by $\{\tilde{\Psi}_k\}_{k=k_0}^{k_{max}}$. Inspired by (Zhou et al., 2024), the reconstruction ratio is defined by $1 - \|\Psi_k - \tilde{\Psi}_k\|_2 / \|\Psi_k\|_2$. In other words, a higher value of the reconstruction ratio indicates that the model's dimension is closer to two. Therefore, the trend of the error curve with respect to the scaling parameter k reveals a distinct separation in the effective dimension between low- and high-scaling components. Evidently, the low-scaling segments display more semantic consistency and thus, occupy a larger dimensional space, whereas the high-scaling segments converge to a more confined or lower-dimensional structure.

It is important to note that this exploration into the dimensionality of generative paths is practically meaningful. Previous study (Zhou et al., 2024) has shown that the effectiveness of distilling a generative model with fewer steps from a pre-trained diffusion model theoretically depends on the model's dimensionality at each step, as informed by the high-dimensional Mean Value Theorem. The observations from Figure S2 provide empirical support for this concept. Specifically, the ability of K-Flow to maintain a lower-dimensional structure in high-scaling components suggests a promising approach for fast sampling distillation methods.

B.3 RELATED WORK DISCUSSION

The field of generative modeling has seen significant advancements in recent years, driven by a variety of frameworks, including adversarial generative networks (GAN) (Goodfellow et al., 2014), variable autoencoders (VAE) (Kingma, 2013), and normalizing flows (Papamakarios et al., 2021). In this work, we focus on continuous normalizing flow generative models (Chen et al., 2018), with



Supplementary Figure S3: Comparison of multi-scale modeling: pixel data space and K-Amplitude space.

particular emphasis on the conditional flow matching training scheme, which originates from the denoising score matching training framework (Vincent, 2011).

Both diffusion models and continuous flow matching models aim to lower the complexity of directly optimizing the log-likelihood of data by introducing an additional stochastic path. However, as proved in (Lavenant & Santambrogio, 2022), the canonical path for diffusion models and rectified flows is not optimal. This realization motivates our introduction of frequency decomposition as a key design element in generative models.

By breaking down the formula of our K-Flow vector field with respect to the scaling parameter k , we can summarize three successful factors as general principles for scaling modeling.

- A good K -amplitude decomposition can leverage the problem's inherent biases towards certain scaling bands. For generative tasks, it is crucial that all K-Flow bands are effectively modeled to ensure the generation of high-quality, controllable outputs. In addition, the computational resources required may vary between different scales, thus necessitating careful consideration of resource allocation.
- Modeling within each scaling component, which is formulated in our K-Flow-localized vector fields.
- Modeling bridges along different scalings, which is achieved through our flow ODE and the (time) K-Flow embedding block for the U-Net or DIT architecture.

This approach to inter- and intra-modeling for K -amplitude is also applicable to scenarios emphasizing certain frequencies or scalings. For instance, (Li et al., 2024a) enhanced oscillatory motion control in video generation by discarding the high-frequency component of the Fourier decomposition. As discussed in Section 3, the scaling parameter of spatially localized wavelet (multi-resolution) decomposition is closely linked to image resolution. Notable contributions in this domain include (Atzmon et al., 2024) and (Lei et al., 2023), which introduced a multi-stage resolution for fine-grained editing, and (Jin et al., 2024), which concentrated on efficient video generation. We provide a systematic review of frequency- or scaling-based generative approaches in Appendix C, highlighting key developments in this direction.

In related research on auto-regressive modeling, (Mattar et al., 2024) presented wavelets as an effective auto-regressive unit, while (Tian et al., 2024) focused on the scale as a key element for image auto-regression. A significant example is (Phung et al., 2023), which transitioned the latent space from pixel to wavelet space for generative models using wavelet diffusion. However, their method employed the same conditional noising schedule for score matching as traditional diffusion models. In contrast to their approach, our proposed K-Flow integrates wavelet decomposition as a multi-channel module within the neural network architecture for training diffusion models. Additionally, our work extends the notion of wavelet space to the more general K -amplitude space.

We also want to highlight another research line that has recently caught the attention: the auto-regressive modeling over the pixel space for image generation. One classic work is VAR Tian et al. (2024). It introduces a hierarchical down and up sample paradigm that models images in a coarse-to-fine manner across multiple resolutions and models the data distribution in an auto-regressive manner. In contrast, our proposed K-Flow integrates the flow paradigm for density estimation and

972 leverages the K -amplitude space as a stronger inductive bias, as illustrated in Figure S3. Another
 973 related work is the auto-regressive flow model proposed in (Ren et al., 2024), that implements con-
 974 ditional flow matching sequentially across scales. Although (Ren et al., 2024) shares some common
 975 terminology with our work (e.g., scales, flow matching), K-Flow provides significant advantages
 976 through its unified flow process, architecture-agnostic design, and theoretically grounded frequency
 977 domain framework. On the practical implementation side, Unlike (Ren et al., 2024), which requires
 978 separate flow matching for each scale and relies on specific architectures (autoregressive trans-
 979 formers that treat scaling as conditional input), our approach implements a single coherent flow that
 980 connects all frequency scales during inference while maintaining architecture flexibility.

981 **Summary.** In summary, K-Flow is a more general framework, with its three key factors potentially
 982 benefiting generation-related tasks like super-resolution and multi-resolution editing. For example,
 983 (Liu et al.) utilized a learnable Fourier transform to construct a harmonic module in the bottleneck
 984 layer of an autoencoder. We provide a comprehensive list of related works in Appendix C.

985 B.4 CONNECTING K-FLOW WITH SSL REPRESENTATION AND GENERATION

986 From the above discussion, we have seen how pretrained vision models leverage the sparsity and
 987 locality of natural data in various K -amplitude domains for perception and generation-based tasks.
 988 In the realm of unsupervised learning, (Liu et al., 2022a; 2024a; Chen et al., 2024) explore whether
 989 generative-based representations, particularly those derived from denoising diffusion models, can
 990 achieve parity with contrastive-based representation learning methods for downstream tasks. A key
 991 observation from their findings Chen et al. (2024), which aligns with our approach of employing
 992 K -amplitude decomposition (the PCA instance), is the revelation that the most powerful repres-
 993 entations are obtained through denoising within a latent space, such as the compressed PCA space.
 994 Another merit of PCA is that denoising along the PCA directions can achieve faster convergence for
 995 denoising, which is revealed in (Du et al., 2023).

996 To transition from unsupervised representation learning to real data generation, incorporating all K -
 997 amplitude scalings is essential. Rather than compressing or amplifying specific scaling bandwidths,
 998 generative tasks require novel organization or ordering of all frequencies. Besides our flow-based
 999 frequency generation approach, (Tian et al., 2024) connects different scales (which can be inter-
 1000 preted as the wavelet K -amplitudes) using residual connections with an auto-regressive training
 1001 objective. Residual connections, as a discretization of ordinary differential equations (ODEs) pro-
 1002 posed in (Ee, 2017), suggest that (Tian et al., 2024)'s approach can be seen as a special discretization
 1003 of our K-Flow with a flexible flow matching training objective.

1004
 1005 Supplementary Table S1: Comparison among PCA, contrastive, and generative SSL.

	Basis Learning	Reconstruction Learning
PCA SSL	Non-parameterized, Determined By Data	Parameterized
Contrastive SSL	Parameterized	N/A
Generative SSL	Parameterized	Parameterized

1026 C RELATED WORK ON FREQUENCY, AND MULTI-SCALING
10271028 There have been multiple research lines on studying generative modeling, especially in terms of
1029 multi-scale modeling. In this work, we would like to summarize them as the following three venues.
10301031 C.1 MULTI-SCALE IN PIXEL RESOLUTION, FLOW AND DIFFUSION
10321033 **Laplacian Pyramid and Laplacian Operator.** In mathematics, the Laplacian operator computes
1034 the second derivative of a function, emphasizing regions with significant intensity changes, such
1035 as edges or high-frequency details. Similarly, the Laplacian Pyramid (Burt & Adelson, 1987) de-
1036 composes an image into multiple scales, extracting the low-frequency components (smooth regions)
1037 through downsampling. The high-frequency details, such as edges and textures, are modeled as the
1038 residuals between adjacent resolution layers. The primary objective of the Laplacian Pyramid is to
1039 represent these residuals across scales in a hierarchical fashion.
10401041 **LAPGAN (Laplacian Generative Adversarial Networks)** (Denton et al., 2015) adopts the Lapla-
1042 cian pyramid idea into the generative adversarial network (GAN) framework (Goodfellow et al.,
1043 2014). By focusing on learning residuals between successive levels of resolution, it effectively gen-
1044 erates high-quality super-resolution images.
10451046 **SR3 (Super-Resolution via Repeated Refinement)** (Saharia et al., 2022) leverages DDPM (De-
1047 noising Diffusion Probabilistic Models) (Ho et al., 2020) and DSM (Denoising Score Match-
1048 ing) (Vincent, 2011; Song & Ermon, 2019) for high-resolution image generation. Specifically, SR3
1049 enhances low-resolution images to high-resolution by utilizing multiple cascaded conditional diffu-
1050 sion models. In this framework, the low-resolution images serve as conditions, and the model’s aim
1051 is to predict the corresponding high-resolution images as outputs.
10521053 **PDDPM (Pyramidal Denoising Diffusion Probabilistic Models)** (Ryu & Ye, 2022) is a follow-up
1054 work of SR3, and it improves the model by only modeling one score network. The key attribute to
1055 enable this is by adding the fractional position of each pixel to the score network, and such fractional
1056 position information can be naturally generalized to different resolutions.
10571058 **f-DM** (Gu et al., 2022) is developed concurrently with PDDPM and shares the approach of utilizing
1059 only one diffusion model. It distinguishes itself by explicitly applying a sequence of transformations
1060 to the data and emphasizing a resolution-agnostic signal-to-noise ratio within its diffusion model
1061 design.
10621063 **Edify Image** (Atzmon et al., 2024) is a state-of-the-art model capable of generating photorealistic,
1064 high-resolution images from textual prompts (Atzmon et al., 2024). It operates as a cascaded pixel-
1065 space diffusion model. To enhance its functionality, Edify Image employs a downsampling process
1066 that extracts low-frequency components and creates three distinct resolution levels, ranging from
1067 low to high frequency, with the original image representing the highest frequency level. Another
1068 key innovation of Edify Image is its meticulously crafted training and sampling strategies at different
1069 resolutions, utilizing attenuated noise schedules.
10701071 C.2 MULTI-SCALE IN PIXEL RESOLUTION, VAE AND AR
10721073 **VQ-VAE2 (Vector Quantized VAE 2)** (Razavi et al., 2019) enforces a two-layer hierarchical struc-
1074 ture, where the top layer captures global features such as object shapes and geometry, while the
1075 bottom layer focuses on local details like texture. It models data density within the variational
1076 autoencoder (VAE) framework (Kingma, 2013) and incorporates an autoregressive (AR) module to
1077 enhance the prior for improved generative performance.
10781079 **RQ-VAE (Residual-Quantized VAE)** (Lee et al., 2022) integrates recursive quantization into the
1080 VAE framework. It constructs a representation by aggregating information across D layers, where
1081 the first layer provides a code embedding closely aligned with the encoded representation, and each
1082 subsequent layer refines this by reducing the quantization error from the previous layer. By stacking
1083 D layers, the accumulated quantization error is minimized, enabling RQ-VAE to offer a coarse-to-
1084 fine-grained approach to modeling. For modeling, the general pipeline follows the VAE framework,
1085 while each latent code is decomposed into D layers and is predicted in an autoregressive manner.
1086

1080
 1081 **VAR (Visual AutoRegressive)** (Tian et al., 2024) introduces a novel paradigm for density estimation
 1082 by decomposing images into multiple resolutions across various scales. This approach is inspired by
 1083 the hierarchical nature of human perception, where images are interpreted progressively from global
 1084 structures to finer details. Leveraging this concept, VAR models the entire image in a coarse-to-fine
 1085 manner, adhering to the principles of multi-scale hierarchical representation.

1086 **C.3 MULTI-SCALE IN FREQUENCY, AR, VAE, AND DIFFUSION**

1087
 1088 **WaveDiff (Wavelet Diffusion)** (Phung et al., 2023) leverages the discrete wavelet transform to shift
 1089 the entire diffusion process into the wavelet spectrum. Its primary objective is to reduce model
 1090 complexity by operating in the transformed spectrum space instead of the pixel domain.

1091
 1092 **PiToMe (Protect Informative Tokens before Merging)** (Tran et al., 2024) is a token merging
 1093 method designed to balance efficiency and information retention. PiToMe identifies large clusters of
 1094 similar tokens as high-energy regions, making them suitable candidates for merging, while smaller,
 1095 more unique, and isolated clusters are treated as low-energy and preserved. By interpreting attention
 1096 over sequences as a fully connected graph of tokens, PiToMe leverages spectral graph theory to
 1097 demonstrate its ability to preserve critical information.

1098
 1099 **WF-VAE (Wavelet Flow VAE)** (Li et al., 2024b) is a parallel work that injects the Wavelet transform
 1100 into the backbone model of the VAE framework for extracting the multi-scale pyramidal features.
 1101 We need to emphasize that WF-VAE introduces frequency decomposition as an inductive bias into
 1102 the backbone model to simulate energy flow. In contrast, our K-Flow retains the backbone archi-
 1103 tecture and instead injects the K -amplitude as the realm for energy flow. In other words, K-Flow
 1104 incorporates the multi-scale concept through the time domain.

1105
 1106 **SIT (Spectral Image Tokenizer)** (Esteves et al., 2024) is a parallel work to ours that processes the
 1107 spectral coefficients of input patches (image tokens) obtained through a discrete wavelet transform.
 1108 Motivated by the spectral properties of natural images, SIT focuses on effectively capturing the high-
 1109 frequency components of images. Furthermore, it introduces a scale-wise attention mechanism,
 1110 referred to as scale-causal self-attention, which is designed to improve the model’s expressiveness
 1111 across multiple scales.

1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

1134
1135

D METHOD DETAILS

1136
1137

D.1 FOURIER TRANSFORM AS A K -AMPLITUDE DECOMPOSITION

1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149

We have shown how to build the K -Amplitude through the Fourier space in Section 2.2. In the discrete setting, the Fourier transform is realized by basis functions of the form $W_N^{kn} = e^{-j\frac{2\pi}{N}kn}$, where N is the length of the sequential data. An effective K -amplitude decomposition exploits this structure by aligning with the inherent hierarchical structure of the data manifold. For example, if most of the energy or amplitudes are concentrated in the low-scaling range, the generative capability of the flow can be enhanced by allocating more steps or resources to these low frequencies (this hypothesis is even true in the latent space, as it's demonstrated in fig. S1). Conversely, fewer steps can be allocated to high frequencies that carry minimal mass or information. For the Fourier transform, this tendency is evident in the analysis of natural images, which often exhibit the celebrated $1/f$ spectrum phenomenon (Weiss & Freeman, 2007). This phenomenon suggests that energy diminishes with increasing scaling parameter, meaning that low-scaling components hold the majority of the signal's information content.

1150
1151

D.2 PCA TRANSFORM AS A K -AMPLITUDE DECOMPOSITION

1152
1153
1154

From the K -amplitude perspective, PCA is an eigen-decomposition obtained by the data covariance matrix. The covariance matrix is given by:

1155
1156

$$\mathbf{C} = \frac{1}{n} \mathbf{X}_{\text{centered}}^\top \mathbf{X}_{\text{centered}},$$

1157
1158
1159
1160
1161
1162

where $\mathbf{X}_{\text{centered}} = \mathbf{X} - \bar{\mathbf{X}}$ is the centered data matrix. In this context, the principal components reveal the relative importance of each transformed direction. To translate PCA into a K -amplitude decomposition, we define the scaling parameter k as the relative order of the principal components. For implementation, we utilize the eigenvalue decomposition of \mathbf{C} for PCA, and the eigenvalues in their descending ordering define the scaling parameter k .

1163

D.3 DWT TRANSFORM AS A K -AMPLITUDE DECOMPOSITION

1164
1165
1166
1167

The Discrete Wavelet Transform (DWT) (Akansu & Haddad, 1992) is utilized to decompose a signal at multiple scales, capturing both time and frequency characteristics. It involves scaling and translating wavelets.

1168
1169
1170
1171

The DWT decomposes the input signal into approximation and detail coefficients:

- Given a discrete signal $x[n]$ (expressed by a finite-dimensional vector), use the scaling function $\phi(t)$ and wavelet function $\psi(t)$ to generate coefficients:

1172
1173

$$c_k[j] = \sum_n x[n] \cdot \phi_{k,j}[n], \quad d_k[j] = \sum_n x[n] \cdot \psi_{k,j}[n].$$

1174
1175
1176
1177

Here, $c_k[j]$ are the approximation coefficients at scale k , and $d_k[j]$ are the detail coefficients at scale k . Comparing with our definition of K-Flow decomposition, k is just a discrete scaling parameter.

The inverse transform then reconstructs the original signal from the coefficients:

1178
1179
1180

$$x[n] = \sum_k c_j[k] \phi_{j,k}[n] + \sum_k d_j[k] \psi_{j,k}[n]$$

1181
1182

Recursive Relationship between different Scales (k) Different levels of decomposition are recursively related:

1183
1184
1185
1186
1187

1. $k = 1$: A single level decomposition results in approximation coefficients c_1 and detail coefficients d_1 ;
2. $k = 2$: A two-level decomposition first produces coefficients c_1 and d_1 . Then, the approximation coefficients c_1 are further decomposed into a second level of approximation coefficients c_2 and detail coefficients d_2 .

1188 For $k = 2$, the decomposition looks like: $x[n] \rightarrow (c_2, d_2, d_1)$, where d_1 represents the high-
 1189 frequency components (level 1 detail coefficients) and c_1 is the low-frequency component (level 1
 1190 approximation coefficients). Further decomposing c_1 yields c_2 (level 2 approximation coefficients)
 1191 and d_2 (level 2 detail coefficients). This recursive relationship illustrates why we can effectively
 1192 take a finite maximum scaling parameter k_{max} and still own an inverse transform.
 1193
 1194

1195 **Practical Design choice.** In this paper’s experiments, especially the Wavelet version of K-Flow
 1196 flow, we take the k_{max} to be one or two. One means decomposing the data into two scales, and two
 1197 means decomposing the data into three scales.
 1198

1199 **Pre-conditioning the data based on scaling** As illustrated in fig. S1, the energy distribution exhibits
 1200 significant heterogeneity across different frequency bands, which consequently leads to non-uniform
 1201 vector field norms in our localized K-Flow. To address this training instability, two approaches can
 1202 be considered: First, following the methodology proposed in conventional diffusion models (Karras
 1203 et al., 2022), we could incorporate input-output preconditioning modules into the neural architecture.
 1204 However, this approach necessitates modifications to the backbone network structure, potentially
 1205 affecting model compatibility and transfer learning capabilities.
 1206

1207 In this paper, we propose a more flexible solution through component-wise normalization of the
 1208 multi-scale decomposed data. Specifically, after performing K -amplitude decomposition, we compute
 1209 the statistical moments (mean and standard deviation) for each discretized scaling component
 1210 independently. This normalized representation is then processed through the flow, followed by an
 1211 inverse normalization step to restore the original scale. This pre-processing approach effectively
 1212 stabilizes the training dynamics while preserving the architectural integrity of the backbone model.
 1213
 1214

1214 D.4 IMPLEMENTATION DETAILS OF K-FLOW VECTOR FIELD

1216 K-Flow is architecture-agnostic in terms of its vector field implementation, making it compatible
 1217 with classical architectures such as U-Net (Song et al., 2020) and Vision Transformers (Peebles
 1218 & Xie, 2023) that are widely adopted in continuous normalizing flows and diffusion models. The
 1219 integration of our method requires only one targeted modification: replacing the conventional time-
 1220 embedding module with a K -amplitude-embedding module, where the temporal input is substituted
 1221 by the scaling parameter k . This modification enables direct incorporation of scaling information
 1222 while preserving the original architectural benefits, though we leave the exploration of specialized
 1223 architectures for K -amplitude flow as future work.
 1224

1225 For practical implementation, we provide several variants of bump functions in this subsection to
 1226 facilitate exploration of the design space, with the complete training algorithm detailed in Algo-
 1227 rithm 1. For additional insights on the K -amplitude localization property and its implications for
 1228 model efficiency, we refer readers to Appendix D.5.

1229 **Remarks.** Despite this model-agnostic nature, the unique K -amplitude localization property of
 1230 Equation (11) offers an opportunity to design more efficient models. For instance, consider points
 1231 that lie outside the support of function $\mathbb{I}_{\sqrt{k_x^2 + k_y^2 + k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1]}$. In these regions, their derivative re-
 1232 mains zero, indicating that they do not contribute to the optimization process for the corresponding
 1233 scaling band. This selective activation allows us to focus computational efforts solely on the val-
 1234 ues within the support of the indicator function, $\mathbb{I}_{\sqrt{k_x^2 + k_y^2 + k_z^2} \in [\lfloor k \rfloor, \lfloor k \rfloor + 1]}$. By doing so, the values
 1235 outside this region can be treated as static conditions, providing a fixed context.
 1236

1237 **Scaling Discretization.** In the main text, we assume, by default, that the scaling parameter k takes
 1238 integer values: $k \in \{0, 1, 2, \dots, k_{max}\}$. Thus, the differentiable vector field v_k for continuous k is
 1239 defined by interpolating between $\lfloor k \rfloor$ and $\lfloor k \rfloor + 1$.
 1240

1241 We now extend this concept to a more general setting where k may take a limited set of integer
 1242 values within the range from 0 to k_{max} . Suppose k_m and k_n represent two specific integer values
 1243 for k . We demonstrate how to extend k continuously within the connected interval $[k_m, k_n]$. Let
 1244 $t := k - k_m$. The differentiable version of ϕ_k is then expressed as:

$$\begin{aligned}
\Psi_{k_m+t} = & \mathcal{F}^{-1} \left(\mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} < k_m} \cdot \mathcal{F}\{\phi\}(k_x, k_y, k_z) + \left(1 - \mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \geq k_n} \right) \cdot \epsilon \right. \\
& \left. + \mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \in [k_m, k_n]} \cdot (\mu(t) \cdot \mathcal{F}\{\phi\}(k_x, k_y, k_z) + (1 - \mu(t)) \cdot \epsilon) \right), \quad (17)
\end{aligned}$$

where $\mu(t)$ is a bump function fulfilling $\mu(0) = \mu(k_n - k_m) = 1$ and $\mu'(0) = -\mu'(k_n - k_m)$.

Replacing the Fourier transform with the general K -amplitude decomposition, the K-Flow is expressed in its general form as follows:

$$\Psi_{k_m+t} = \mathcal{F}^{-1} \left(\mathbb{I}_{k < k_m} \cdot \mathcal{F}\{\phi\}(k) + \left(1 - \mathbb{I}_{k \geq k_n} \right) \cdot \epsilon + \mathbb{I}_{k \in [k_m, k_n]} \cdot (\mu(t) \cdot \mathcal{F}\{\phi\}(k) + (1 - \mu(t)) \cdot \epsilon) \right),$$

where $\mathcal{F}\{\phi\}(k)$ is defined in the main text.

Experimental Implementation. In this paper's experiments, particularly in the Fourier and PCA versions of the K-Flow flow, we restrict the discrete values of k to $\{0, \frac{k_{\max}}{2}, k_{\max}\}$, with k_{\max} determined by resolution. We then extend k continuously using Equation 17.

Bump Function. We propose two types of bump functions: 1. Hard bump; 2. Soft bump. The **hard bump function** $\mu : [0, 1] \rightarrow \mathbf{R}^+$ satisfies the specific endpoint properties:

$$\mu(0) = \mu(1) = 1 \quad \text{and} \quad \mu'(0) = -\mu'(1). \quad (18)$$

Inspired by spline functions, such bump functions can be constructed using polynomials. For example, a quartic form used in our experiments is given by:

$$\mu(t) = 1 - 3t^2 + 2t^3. \quad (19)$$

For more examples, readers can explore modifications of the connection functions used in Meyer wavelets (Meyer, 1992).

In this paper, we utilize hard bump functions for constructing K-Flow flows with scaling discretization exceeding one component.

Soft Localization with Soft Bump Function. Consider that the scaling parameter is discretized to take values in an increasing sequence $\{k_i\}_{i=0}^n$. Consequently, the continuous k lies in the interval $k \in [k_0, k_n]$. Define

$$\psi_i := \mathbb{I}_{\sqrt{k_x^2+k_y^2+k_z^2} \in [k_i, k_{i+1}]}.$$

These ψ_i form a partition of unity for the K-Amplitude basis. The derivative of the soft bump function μ'_i is defined for each scaling component ψ_i (a.k.a. frequency band), expressed as:

$$\mu'_i(k; a_i, b) = \begin{cases} c \cdot \left(1 - \left(\frac{k-a_i}{b} \right)^2 \right)^n, & \text{if } |k - a_i| < b, \\ 0, & \text{if } |k - a_i| \geq b, \end{cases} \quad (20)$$

where $a_i = \frac{k_i + k_{i+1}}{2}$ and c is the normalization constant ensuring that the integral of the function over its compact support is 1. Note that hyper-parameter $b \leq k_n - k_0$ dictates the width or support region of the bump, while the degree n measures the sharpness of the bump. We retain b and n as hyperparameters. The bump function $\mu_i(k)$ is then obtained by integrating $\mu'_i(k)$, which is also a polynomial function.

It is evident that $\mu_i(k)$ satisfies:

$$\mu_i(k_0) = 0 \quad \text{and} \quad \mu_i(k_n) = 1.$$

Finally, conditioned on a sampled noise ϵ , the modified soft K-Flow flow at time $t \in [0, k_n - k_0]$ is expressed as:

$$\Psi_{k_0+t} = \mathcal{F}^{-1} \left(\sum_i \psi_i(k_x, k_y, k_z) \cdot \mu_i(k_0 + t) \cdot \mathcal{F}\{\phi\}(k_x, k_y, k_z) + \sum_i \psi_i(k_x, k_y, k_z) \cdot (1 - \mu_i(k_0 + t)) \cdot \epsilon \right). \quad (21)$$

1296 Through the application of this formula and a family of soft bump functions $\{\mu_i\}$, we can also
 1297 implement algorithm 1. In comparison to the hard bump functions, a K-Flow constructed with soft
 1298 bump functions assigns varying weights to each scale according to the scaling parameter k . Unlike
 1299 hard bump functions which strictly set other scales to zero for each stage of k , soft bump functions
 1300 provide a more gradual transition, allowing for multiple frequencies to flow concurrently, and the
 1301 relative weights are determined by the current scaling parameter k .

1302 **Comments on Haar and Meyer Wavelet K -amplitude.** One type of wavelet that offers both
 1303 frequency and spatial localization is the Meyer wavelet. The Meyer wavelet is originally defined in
 1304 the Fourier frequency domain, making it ideal for smooth frequency transitions.

1305 The 1D Meyer wavelet $\psi(t)$ and its scaling function $\phi(t)$ are defined via their Fourier transforms,
 1306 $\hat{\psi}(\omega)$ and $\hat{\phi}(\omega)$, respectively. The Meyer wavelet is constructed to ensure that the wavelet transform
 1307 will partition the frequency domain into octave bands.

1309 The Fourier transform of the scaling function $\hat{\phi}(\omega)$ is defined as:

$$\hat{\phi}(\omega) = \begin{cases} 1 & \text{if } |\omega| \leq \frac{2\pi}{3}, \\ \cos\left(\frac{\pi}{2}\nu\left(\frac{3|\omega|}{2\pi} - 1\right)\right) & \text{if } \frac{2\pi}{3} < |\omega| \leq \frac{4\pi}{3}, \\ 0 & \text{if } |\omega| > \frac{4\pi}{3}, \end{cases} \quad (22)$$

1315 where $\nu(t)$ is a smooth function defined as:

$$\nu(t) = \begin{cases} 0 & \text{if } t \leq 0, \\ t & \text{if } 0 < t < 1, \\ 1 & \text{if } t \geq 1. \end{cases} \quad (23)$$

1321 The Fourier transform of the Meyer wavelet $\hat{\psi}(\omega)$ is then defined as:

$$\hat{\psi}(\omega) = \begin{cases} \sin\left(\frac{\pi}{2}\nu\left(\frac{3|\omega|}{2\pi} - 1\right)\right) & \text{if } \frac{2\pi}{3} < |\omega| \leq \frac{4\pi}{3}, \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

1326 In other words, Meyer transformation can be seen as the Fourier transform with a spatial cutoff
 1327 window. Note that the scaling function and the wavelet function play different roles, where the
 1328 low-frequency content of data is obtained by convolving the signal with the scaling function.

1329 In the ablation section, we will employ a specific discretization of the Meyer wavelet to generate
 1330 our data. Additionally, we will explore the Haar wavelet method, which is implemented solely
 1331 through spatial convolution kernels and scaling operations. The Haar wavelet, being the simplest
 1332 form of wavelet, is particularly interesting because it uses piecewise constant functions to capture
 1333 local features at varying scales, providing a contrast to the smoother Meyer wavelet.

1335 D.5 IMPLEMENTATION DETAILS

1336 **Hyper-parameters.** In our experiments, we use the pretrained VAE from Stable Diffusion (Rom-
 1337 bach et al., 2022). The VAE encoder has a downsampling factor of 8 given an RGB pixel-based
 1338 image $\mathbf{x} \in \mathbb{R}^{h \times w \times 3}$, $\mathbf{z} = \mathcal{E}(\mathbf{x})$ has shape $\frac{h}{8} \times \frac{w}{8} \times 4$. All experiments are operated in the fixed
 1339 latent space.

1341 In Table S2, we provide training hyperparameters for the image generation tasks on the two datasets.
 1342 For implementing training algorithm Algorithm 1, the bump function is provided in eq. (19). For
 1343 the classifier-free sampling on the conditional generation task, the cfg-scale is set to be 1.5.

1344 D.6 RESOURCE REQUIREMENT AND TIME COMPLEXITY

1346 **Resources Requirement.** All experiments were conducted on NVIDIA H100 GPUs, with a total
 1347 computational budget of approximately 3,000 GPU-hours.

1348 **Time Complexity.** Our main focus is on comparing the computational complexity of the K -
 1349 amplitude flow with that of ordinary latent flows, we observe that during training, the additional

1350

Supplementary Table S2: Hyper-parameters of DiT network.

1351

1352

1353

1354

1355

1356

1357

1358

1359

1360

computational overhead introduced by the K -amplitude flow is minimal. From Algorithm 1, it is evident that the only additional computational step is the discrete inverse K -amplitude transform performed at each training iteration, while the remaining steps maintain the same complexity as the ordinary flow matching algorithm. For instance, when considering the Fourier transform, its computational complexity is $\mathcal{O}(N \log N)$, where N denotes the length of the flattened image vector in the latent space.

For inference, from Algorithm 3, compared to ordinary latent flow, the only additional step to perform the K amplitude flow is an inverse K amplitude transform to set up the initial noise for generation, and the remaining inference remains the same complexity. Thus, we expect a similar or slightly higher complexity than the ordinary latent flow during generation. Empirically, we test the averaged number of function evaluations (NFE) required for the adaptive solver to reach its prespecified numerical tolerance on the CelebA 256 dataset. In fact, our NFE is better than baseline latent flow (LFM):

1374

Supplementary Table S3: CelebA-HQ 256.

1375

1376

1377

1378

1379

1380

1381

1382

1383

1384

1385

Model	NFE \downarrow
LFM, ADM	85
LFM, DiT L/2	89
FM	128
K-Flow, DiT L/2 (Ours)	78

It is worth mentioning that when testing the FID, we apply the fixed-step ODE solver (“Euler”) with 50 steps. Thus, we also provide the average inference time of generating one CelebA sample on one H20 GPU:

1386

Supplementary Table S4: CelebA-HQ 256.

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

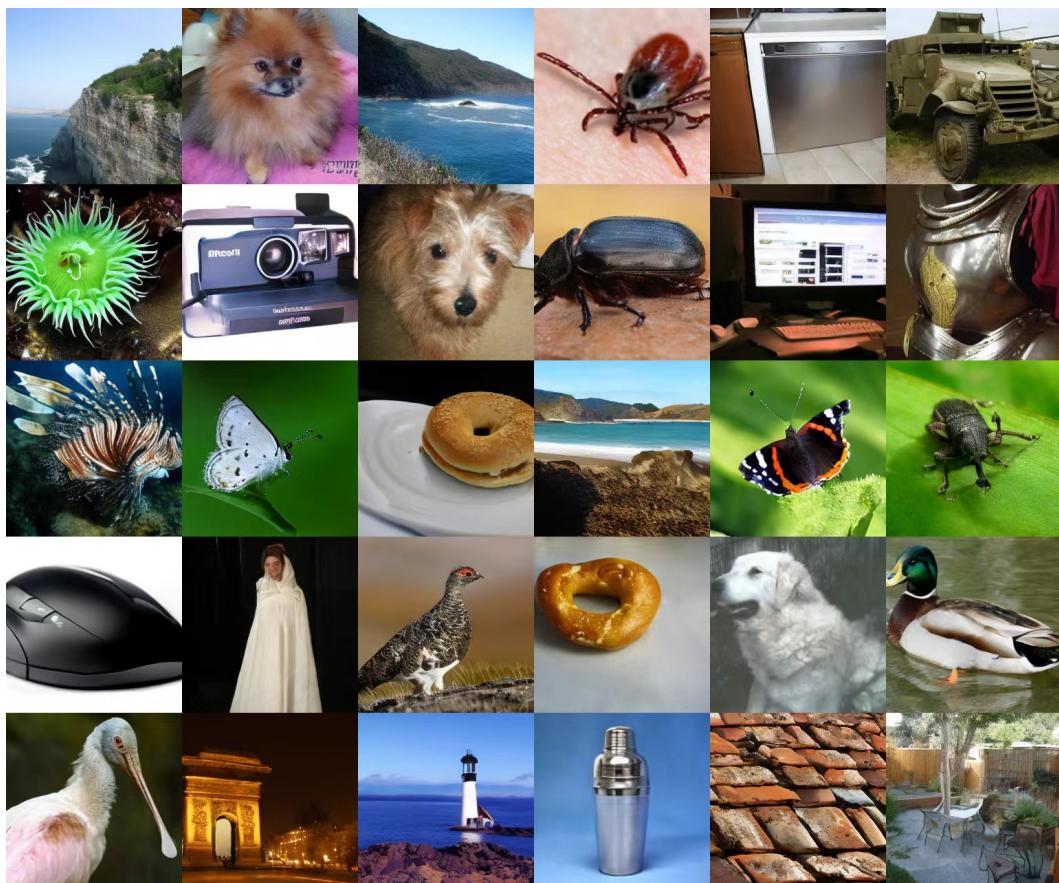
1401

1402

1403

Model	Time (s)
LFM, DiT L/2	0.583
K-Flow, DiT L/2 (Ours)	0.589

1404 **E ALGORITHMS**
14051406 In this section, we list three key algorithms.
14071408 **Algorithm 1** Training of K-Flow.
14091410 **Require:** Scaling parameter k with maximum k_{max} , K-Flow transform \mathcal{F} , inverse transform \mathcal{F}^{-1} ,
1411 noise distribution p , target distribution q
1412 Normalize k to be in $[0, 1]$: $k \leftarrow k/k_{max}$
1413 Initialize parameters θ of v_k
1414 **while** not converged **do**
1415 Sample scaling parameter $k \sim \mathcal{U}(0, 1)$
1416 Sample training example $\phi \sim q$, sample noise $\epsilon \sim p$
1417 Calculate current flow position Ψ_k according to K-Flow transform \mathcal{F} , \mathcal{F}^{-1} and Equation (10)
1418 Calculate the conditional vector field $\dot{\Psi}_k$ according to \mathcal{F} , \mathcal{F}^{-1} and Equation (11)
1419 Calculate the objective $\ell(\theta) = \|v_k(\Psi_k; \theta) - \dot{\Psi}_k\|_g^2$, following Equation (12)
1420 $\theta = \text{optimizer_step}(\ell(\theta))$
1421 **end while**1422 **Algorithm 2** Scaling-controllable Generation of K-Flow.
14231424 **Require:** Scaling parameter k , K -amplitude transform \mathcal{F} , inverse transform \mathcal{F}^{-1} , noise distribution
1425 p in the K -amplitude space, target distribution q
1426 Initialize pre-trained $v_k(\theta)$
1427 Sample one high-scaling noise $\epsilon_{\text{high}} \sim p$, sample two independent low-scaling noise $\epsilon_{\text{low}} \sim p$,
1428 $\tilde{\epsilon}_{\text{low}} \sim p$
1429 $\Psi_0 = \mathcal{F}^{-1}\{\epsilon_{\text{low}}, \epsilon_{\text{high}}\}$
1430 $\tilde{\Psi}_0 = \mathcal{F}^{-1}\{\tilde{\epsilon}_{\text{low}}, \epsilon_{\text{high}}\}$
1431 **for** $k \in [0, 1]$ **do**
1432 $\Psi_k \leftarrow \text{ODEstep}(v_k(\cdot, \theta), \Psi_0)$
1433 $\tilde{\Psi}_k \leftarrow \text{ODEstep}(v_k(\cdot, \theta), \tilde{\Psi}_0)$
1434 **end for**
1435
1436 **return** $\Psi_1, \tilde{\Psi}_1$ 1437
1438 **Algorithm 3** Class-conditional Generation of K-Flow with dropping.
14391440 **Require:** Pre-trained $v_k(\theta)$, conditioning class c , dropping time τ , noise distribution p , guidance
1441 parameter ω
1442 1: $\Psi_0 \sim p$
1443 2: **for** $k \in [0, \tau]$ **do**
1444 3: $\tilde{v}_k(\cdot) \leftarrow (1 - \omega)v_k^0(\cdot, \theta) + \omega u_k(\cdot, c, \theta)$ {guided velocity}
1445 4: $\Psi_\tau \leftarrow \text{ODEstep}(\tilde{v}_k(\cdot), \Psi_0)$
1446 5: **end for**
1447 6: **for** $k \in [\tau, 1]$ **do**
1448 7: $\Psi_1 \leftarrow \text{ODEstep}(v_k(\cdot, \theta), \Psi_\tau)$
1449 8: **end for**
1450 9:
1451 10: **return** Ψ_1

1458 F MORE RESULTS
14591460 F.1 MORE RESULTS ON UNCONDITIONAL GENERATION
14611462 We provide more results on the class-conditional generation using K-Flow in Figure S4.
14631464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492 Supplementary Figure S4: Non-curated samples of our reversing scaling variant on ImageNet (cfg = 1.5).
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511

1512
1513

F.2 UNCONDITIONAL GENERATION ON LSUN CHURCH

1514
1515
1516
1517
1518

We conducted unconditional generation experiments on LSUN Church Yu et al. (2015), with the resolution of 256×256 . The results are presented in Table S5. We test our K -amplitude flow with two and three scaling components using the db6 wavelet Karam (2012) K -amplitude transform, and we find that the three scaling components version achieves the best quantitative results in terms of FID and Recall.

1519
1520Supplementary Table S5: LSUN Church 256×256 .

1521

1522

1523

1524

1525

1526

1527

1528

1529

1530

1531

1532

1533

Results Table S5 summarizes the results on LSUN Church. We test our K-Flow with two and three scaling components using the db6 wavelet K -amplitude transform, and we find that the three scaling components version achieves the best quantitative results in terms of FID and Recall.

1534
1535
1536

F.3 CLASS-AWARE FID METRIC

1537
1538

We propose using the class-aware FID metric, defined as follows:

1539

1540

$$\text{FID}_{\text{class-conditional}} = \mathbb{E}_{c \sim p(c)} [\text{FID}(c)] \quad (25)$$

1541

where for each class c , the FID is calculated by:

1542

1543

$$\text{FID}(c) := \text{FID}(X_r^c, X_g^c) = \|\mu_r^c - \mu_g^c\|^2 + \text{Tr}(\Sigma_r^c + \Sigma_g^c - 2(\Sigma_r^c \Sigma_g^c)^{1/2}). \quad (26)$$

1544

1545

Here, X_r^c and X_g^c denote the real and generated data subsets for class c , respectively. Based on $\text{FID}(c)$, the Class-Dropping-Ratio (CDR) is defined by

1546

1547

$$\text{CDR} := \mathbb{E}_{c \sim p(c)} \left[\frac{\text{FID}_{\text{bef}}(c)}{\text{FID}_{\text{aft}}(c)} \right],$$

1548

1549

where FID_{bef} denotes the FID calculated for the flow model carried with the class condition for the whole process, and FID_{aft} denotes the FID calculated for the flow model carried with the class condition for only a subprocess (we keep the initial 30% of the inference time for the experiment). In practice, instead of computing the expectation over the entire class distribution $p(c)$, we randomly select 5 classes out of the total 1000 classes for evaluation.

1550

1551

1552

1553

1554

1555

1556

1557

1558

1559

1560

1561

1562

1563

1564

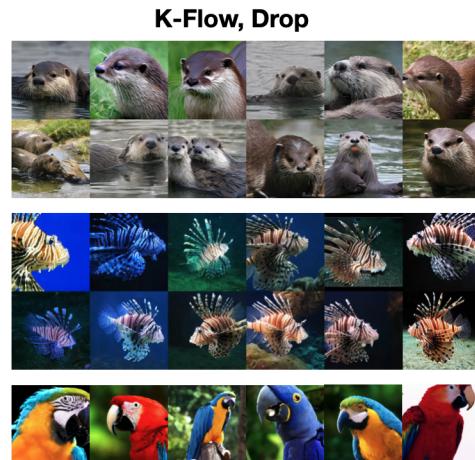
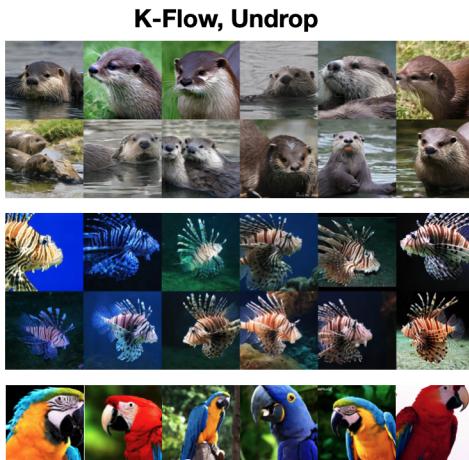
1565

1566
1567

F.4 ABLATION STUDIES ON CONTROLLABLE CLASS-CONDITIONAL GENERATION

1568
1569
1570
1571

In Section 4, we provide brief results on the controllable class-conditional generation over ImageNet. Here, we would like to give a more qualitative comparison between our model K-Flow and LFM.

1572
1573
1574
1575
1576
15771578
1579
1580
1581
1582
1583
1584
1585
1586

Supplementary Figure S5: Classifier-free guidance sampling of our Fourier-based K-Flow with a hyperparameter setting of $\text{cfg} = 3$. In the right columns, the class condition is omitted for the last 50% of the scaling steps during inference, using the same initial noise. It can be observed that as the cfg value increases and the duration of omitting the class condition decreases, the generated results appear nearly identical.

1590
1591
1592
1593
1594
1595
1596
15971598
1599
1600
1601

Supplementary Figure S6: Classifier-free guidance sampling of our wavelet-based K-Flow with a hyperparameter setting of $\text{cfg} = 2$. In the right columns, the class condition is dropped for the last 70% of the scaling steps during inference, using the same initial noise. It can be observed that after dropping, K-Flow still preserves the high-scaling contents.

1602
1603

F.5 ABLATION STUDY ON WAVELET BASE

1604
1605
1606
1607
1608
1609

From Table S6, we tested two additional wavelet base, the discrete Haar basis (Haar, 1911) and the discrete Meyer basis (Meyer, 1990) as a supplement of the Daubechies wavelet (db6, Karam (2012)) used in the main text. All three wavelets demonstrated comparable performance in terms of both the FID and Recall metrics.

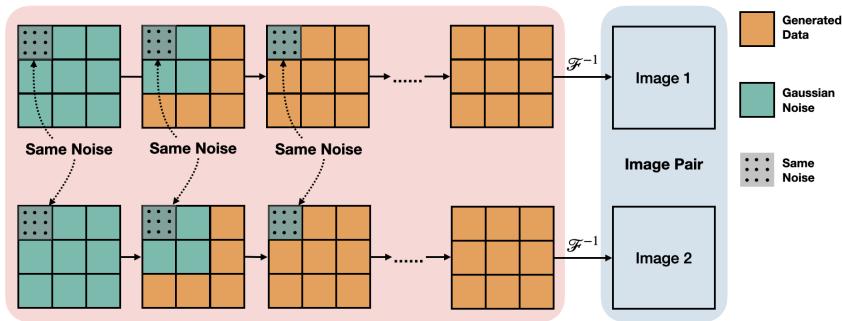
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619

Supplementary Table S6: CelebA-HQ 256.

Model	FID \downarrow	Recall \uparrow
CelebA-HQ 256		
K-Flow, Meyer-DiT L/2	5.01	0.47
K-Flow, Haar-DiT L/2	5.01	0.46
K-Flow, Db-DiT L/2 (three scales)	5.77	0.42

Supplementary Table S7: Conditional ImageNet 256.

Model	FID \downarrow	KSR \downarrow	Recall \uparrow
K-Flow, Wave (reverse)	23.06	-	0.58
+ cfg=1.5	5.1	-	0.46
K-Flow, Wave-DiT L/2 (Ours)	17.8	-	0.56
+ cfg=1.5	4.49	-	0.44
LFM (DiT L/2)	14.0	-	0.56
+ cfg=1.5	2.78	-	0.45

Supplementary Figure S7: Daubechies wavelet K -amplitude with more components trained on CelebA-256.

Supplementary Figure S8: Pipeline of scaling-controllable generation (high scaling).

F.6 ABLATION STUDY ON SCALING PARTITIONS

Although the quality of face generation appears similar to the naked eye, the model with three K -amplitude bands (the last row of Table S6) performed worse in terms of FID and Recall metrics. We provide the generated samples for qualitative evaluation in Figure S7.

Reversing the K -amplitude Scaling. In Table S7, we also tested a counterintuitive scaling order: from high to low. This means generating high-frequency details first and then filling in the low-frequency components during the generation process. We find that the model can still produce images normally (Figure S4), with a better diversity (Recall) but lower quality (FID) compared to the low-to-high scaling approach.

F.7 IMAGE SCALING-CONTROLLABLE GENERATION

Preserving Low Scaling, Modifying High Scaling. We need to highlight that in K-Flow, when modeling the flow from lower to higher scales, the noise at higher scales is used to predict the velocity at the lower scale. This is determined by the nature of ODE flow. To this end, we conduct a study by reversing the scaling direction in the Daubechies wavelet K-Flow, and the pipeline is illustrated in Figure S8. In such a reversed setup, we keep the low-scaling part the same noise while gradually denoising the high-scaling part.

The results are listed in Figure S9. According to the six pairs of results, we can observe that the low-scaling part stays the same, like the background of the image and the gender and color of the people, while the high-resolution details of facial expressions and outlook vary within each pair.

Remarks. Although the overall results are generally optimistic, some unexpected changes have been observed in the high-scaling parts. This may be attributed to two factors:

1. The compressed latent space may mix high and low content present in the original pixel space.
2. The loss Equation (12) may not be perfectly optimized, meaning that K-Flow localized vector field might not be perfectly confined to the low-scaling part. The second factor might be mitigated by training on larger datasets. Furthermore, by training a reversed K-Flow flow (from high to low), we observe that fixing the low-scaling noise enables unsupervised editing of detailed high-scaling content.



Supplementary Figure S9: Results of scaling-controllable generation. We display six pairs of images, where each pair of images preserves the low scaling and differs in the high scaling.



Supplementary Figure S10: LFM editing by Algorithm 2.

In Figure S9, we've tested the wavelet-based K-Flow and observed similar results with the Fourier-based K-Flow.

This insight further supports our model's capacity to decompose the generative process into distinct frequency bands, where specific frequency bands can be independently controlled. This separation aids in achieving more detailed and deliberate modifications to generated data, adding a layer of precision and flexibility to the generative framework.

F.8 IMAGE RESTORATION

In this section, we evaluate the performance of the K -amplitude flow on several image restoration tasks, including super-resolution and inpainting. These tasks typically involve reconstructing the high-frequency components of an image conditioned on the known low-frequency components. Unlike unsupervised editing based on different scales, the performance of this experiment can be

Supplementary Table S8: Performance comparison on image restoration. In this experiment, we pre-trained K flow with the same U-net architecture implemented (Martin et al., 2024).

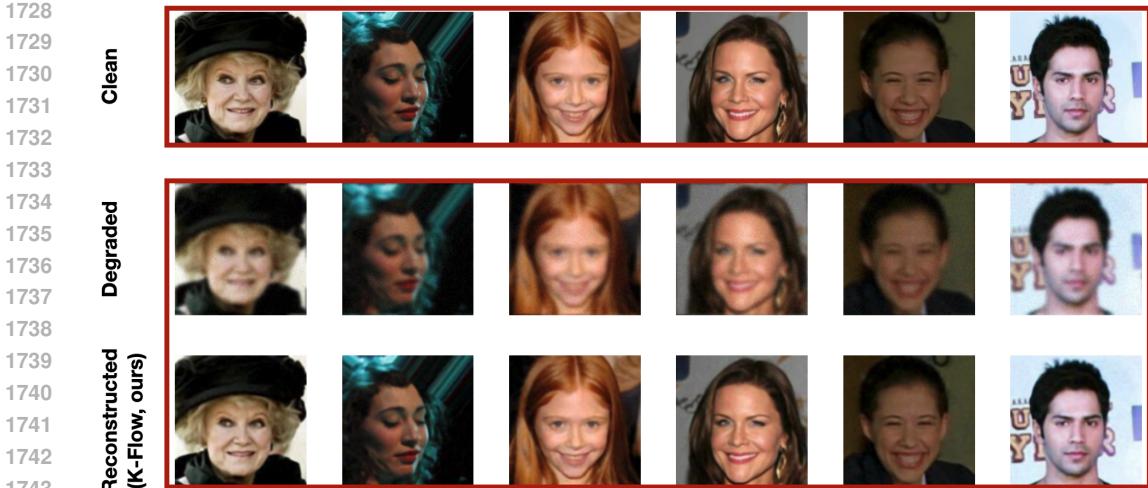
	Super-res.		Box inpaint.		Deblurring.	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Degraded	10.17	0.182	22.12	0.742	27.67	0.740
PnP-Diff (Zhu et al., 2023)	31.20	0.893	N/A	N/A	32.49	0.911
PnP-GS (Hurault et al., 2021)	30.69	0.889	N/A	N/A	33.65	0.924
OT-ODE Pokle et al. (2023)	31.05	0.902	28.84	0.914	32.63	0.915
D-Flow (Ben-Hamu et al., 2024)	29.17	0.833	25.30	0.805	31.07	0.877
Flow-Priors Zhang et al. (2024)	28.35	0.717	29.40	0.858	31.40	0.856
PnP-Flow (Martin et al., 2024)	31.49	0.907	30.59	0.943	34.51	0.940
K-Flow (ours)	32.51	0.934	30.49	0.943	35.89	1.034

quantitatively measured using reconstruction metrics such as Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM).

Datasets and Baselines Our method is benchmarked against standard diffusion and flow matching based restoration methods (see (Martin et al., 2024) for a detailed introduction). We evaluate all methods on CelebA dataset, with images resized to 128×128 .

Algorithm. Our training-free restoration method basically follows the efficient algorithm proposed in Martin et al. (2024) with two key changes adapted to K -amplitude:

1. The naive linear interpolation step is replaced by our scaling interpolation formula Equation (10).



Supplementary Figure S11: Visualization for image restoration using K-Flow.

2. Instead of starting restoration from pure noise, we start at $t = 0.5$, since our flow primarily denoises high-frequency components during the later period of time. From the K -amplitude perspective, this strategic initialization point provides a more informed starting state compared to conventional flow approaches, significantly reducing the inference computational overhead while maintaining restoration quality.

We provide the algorithm details in Algorithm 4.

Algorithm 4 PnP K-Flow.

Input: Pre-trained network v^θ by K-Flow, time sequence $(t_n)_n$ either finite with $t_n = n/N, N \in \mathbb{N}$ or infinite with $\lim_{n \rightarrow +\infty} t_n = 1$ and $t_0 = 0.3$, adaptive stepsizes $(\gamma_n)_n$.
Initialize: $x_0 \in \mathbb{R}^d$.
for $n = 0, 1, \dots$, **do**
 $z_n = x_n - \gamma_n \nabla F(x_n)$. ▷ Gradient step on the data-fidelity term
 \tilde{z}_n from z_n and noise ϵ through K -amplitude interpolation 10.
 $x_{n+1} = D_{t_n}(\tilde{z}_n)$ ▷ PnP step with restoration denoiser in (Martin et al., 2024)
return x_{n+1}

Results. We report benchmark results (following Martin et al. (2024)) for all methods across three restoration tasks, measuring average PSNR and SSIM on 100 test images, including super-resolution (with down sample rate $\times 2$), deblurring and Box inpainting problems. Results are averaged across 100 test images. From Table S8, we see that the K -amplitude flow achieves state-of-the-art (SOTA) quantitative results in the super-resolution task, deblurring task, and comparable results in inpainting tasks. In terms of time complexity, we only use 75 iterations in the super-resolution task, while PnP-flow's iteration number is set to 150. This superior performance without task-specific hyperparameter tuning can be attributed to our model's inherent frequency-aware design: both deblurring and super-resolution tasks primarily involve recovering high-frequency information (higher values of the scaling parameter k), which naturally aligns with the later stages of K-Flow's scaling-progressive generation process. From Figure S11, we can clearly see how K-Flow restores the high scaling components of a blurred picture.

F.9 MOLECULAR ASSEMBLY

We consider another scientific task: molecular assembly. The goal is to learn the trajectory on moving clusters of weakly-correlated molecular structures to the strongly-correlated structures.

1782
1783 Supplementary Table S9: K-Flow against seven generative models on COD-Cluster17 with 5K, 10K, and all
1784 samples. The best results are marked in **bold**.
1785

	COD-Cluster17-5K		COD-Cluster17-10K		COD-Cluster17-All	
	PM (atom) ↓	PM (center) ↓	PM (atom) ↓	PM (center) ↓	PM (atom) ↓	PM (center) ↓
GNN-MD	13.67 ± 0.06	13.80 ± 0.07	13.83 ± 0.06	13.90 ± 0.05	22.30 ± 12.04	14.51 ± 0.82
CrystalSDE-VE	15.52 ± 1.48	16.46 ± 0.99	17.25 ± 2.46	17.86 ± 1.11	17.28 ± 0.73	18.92 ± 0.03
CrystalSDE-VP	18.15 ± 3.02	19.15 ± 4.46	22.20 ± 3.29	21.39 ± 1.50	18.03 ± 4.56	20.02 ± 3.70
CrystalFlow-VE	14.87 ± 7.07	13.08 ± 4.51	16.41 ± 2.64	16.71 ± 2.35	12.80 ± 1.20	15.09 ± 0.34
CrystalFlow-VP	15.71 ± 2.69	17.10 ± 1.89	19.39 ± 4.37	16.01 ± 3.13	13.50 ± 0.44	13.28 ± 0.48
CrystalFlow-LERP	13.59 ± 0.09	13.26 ± 0.09	13.54 ± 0.03	13.20 ± 0.03	13.61 ± 0.00	13.28 ± 0.01
AssembleFlow	7.27 ± 0.04	6.13 ± 0.10	7.38 ± 0.03	6.21 ± 0.05	7.37 ± 0.01	6.21 ± 0.01
K-Flow (ours)	7.21 ± 0.12	6.11 ± 0.11	7.26 ± 0.06	6.12 ± 0.07	7.23 ± 0.01	6.07 ± 0.01

1794
1795 **Dataset and evaluation metrics.** We evaluate our method using the crystallization dataset
1796 COD-Cluster17 (Liu et al., 2024b), a curated subset of the Crystallography Open Database
1797 (COD)(Grazulis et al., 2009) containing 133K crystals. We consider three versions of COD-
1798 Cluster17 with 5K, 10K, and the full dataset. To assess the quality of the generated molecular
1799 assemblies, we employ *Packing Matching (PM)*(Chisholm & Motherwell, 2005), which quantifies
1800 how well the generated structures align with reference crystals in terms of spatial arrangement and
1801 packing density. Following (Liu et al., 2024b), we compute PM at both the atomic level (PM-atom)
1802 and the mass-center level (PM-center) (Chisholm & Motherwell, 2005).

1803 **Baselines.** We evaluate our approach against GNN-MD (Liu et al., 2024b), variations of Crys-
1804 talSDE and CrystalFlow (Liu et al., 2024b), and the state-of-the-art AssembleFlow (Guo et al.,
1805 2025). CrystalSDE-VE/VP model diffusion via stochastic differential equations, while CrystalFlow-
1806 VE/VP use flow matching, with VP focusing on variance preservation. CrystalFlow-LERP employs
1807 linear interpolation for efficiency. AssembleFlow (Guo et al., 2025) enhances rigidity modeling
1808 using an inertial frame.

1809 **Main results.** The main results in Table S9 show that K-Flow outperforms all baselines across three
1810 datasets. Building on AssembleFlow’s rigidity modeling, K-Flow decomposes molecular pairwise
1811 distances via spectral methods and projects geometric information from \mathbb{R}^3 and SO^3 accordingly.
1812 This approach achieves consistently superior packing matching performance.

1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835