



HAL
open science

Fast, faithful and photorealistic diffusion-based image super-resolution with enhanced Flow Map models

Maxence Noble, Gonzalo Iñaki Quintana, Benjamin Aubin, Clément Chadebec

► **To cite this version:**

Maxence Noble, Gonzalo Iñaki Quintana, Benjamin Aubin, Clément Chadebec. Fast, faithful and photorealistic diffusion-based image super-resolution with enhanced Flow Map models. 2026. <hal-05535533>

HAL Id: hal-05535533

<https://hal.science/hal-05535533v1>

Preprint submitted on 3 Mar 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Fast, faithful and photorealistic diffusion-based image super-resolution with enhanced Flow Map models

Maxence Noble^{1,2,*}, Gonzalo Iñaki Quintana¹, Benjamin Aubin¹ & Clément Chadebec¹

¹Jasper Research, Paris, France, ²CMAP, CNRS, Ecole polytechnique, Palaiseau, France



Figure 1: **Qualitative super-resolution results of the *Shortcut* variant of FlowMapSR for $\times 4$ and $\times 8$ upscaling.** A single Flow Map model is used for both upscaling factors, with 2 inference steps.

ABSTRACT

Diffusion-based image super-resolution (SR) has recently attracted significant attention by leveraging the expressive power of large pre-trained text-to-image diffusion models (DMs). A central practical challenge is resolving the trade-off between reconstruction faithfulness and photorealism. To address inference efficiency, many recent works have explored knowledge distillation strategies specifically tailored to SR, enabling one-step diffusion-based approaches. However, these teacher–student formulations are inherently constrained by information compression, which can degrade perceptual cues such as lifelike textures and depth of field, even with high overall perceptual quality. In parallel, self-distillation DMs, known as Flow Map models, have emerged as a promising alternative for image generation tasks, enabling fast inference while preserving the expressivity and training stability of standard DMs. Building on these developments, we propose FlowMapSR, a novel diffusion-based framework for image super-resolution explicitly designed for efficient inference. Beyond adapting Flow Map models to the SR task, we introduce two complementary enhancements: (i) positive–negative prompting guidance, based on a generalization of classifier free-guidance paradigm to Flow Map models, and (ii) adversarial fine-tuning using Low-Rank Adaptation (LoRA). Among the considered Flow Map formulations (*Eulerian*, *Lagrangian*, and *Shortcut*), we find that the *Shortcut* variant consistently achieves the best performance when combined with these enhancements. Extensive experiments show that FlowMapSR achieves a better balance between reconstruction faithfulness and photorealism than recent state-of-the-art methods for both $\times 4$ and $\times 8$ upscaling, while maintaining competitive inference time. Notably, a single model is used for both upscaling factors, without any scale-specific conditioning or degradation-guided mechanisms.

*Work done during internship at Jasper Research.

1 INTRODUCTION

Image super-resolution (SR) is a fundamental problem in computer vision, which aims to reconstruct a plausible high-resolution (HR) image from a low-resolution (LR) input. The task is inherently ill-posed, as SR exhibits a strong *one-to-many* ambiguity: the LR observation provides only partial structural information, leaving high-frequency details fundamentally underdetermined. Consequently, SR methods must navigate a delicate trade-off between reconstruction *faithfulness*—recovering details consistent with the underlying HR image—and *photorealism*—producing visually plausible textures for the human eye. This trade-off becomes increasingly pronounced as the upscaling factor increases, since the LR input becomes less informative and the space of plausible solutions expands, and is further complicated by unknown degradations such as blur, sensor noise, and compression artefacts.

Owing to these challenges, SR has long served as a driving force for advances in generative modeling. While early learning-based approaches were dominated by GAN-based methods (Wang et al., 2018; 2021) and Transformer architectures (Lu et al., 2022), recent research has increasingly focused on diffusion models (DMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021) and related flow-matching (FM) formulations (Lipman et al., 2022), which have demonstrated a strong capacity to synthesize fine details and realistic textures in image generation tasks. As a result, diffusion-based approaches now define the state-of-the-art for image SR. Existing diffusion-based SR approaches can be broadly divided into two categories. The first consists of inference-only methods, which leverage a pre-trained DM and apply advanced posterior sampling techniques to solve the SR inverse problem without need of ground truth LR-HR data pairs or additional training; see Janati et al. (2025) for a comprehensive overview. Although these methods achieve consistently strong results across a variety of SR scenarios, they typically require dozens of model evaluations per image, leading to prohibitive computational costs for large-scale applications. The second category specifically relies on LR-HR training data pairs (the LR inputs being synthetically obtained from available HR images) and focuses on training DMs explicitly for SR, either from scratch (Li et al., 2022; Saharia et al., 2022; Yue et al., 2023; 2024; Luo et al., 2023; Delbracio & Milanfar, 2023) or by fine-tuning large pre-trained backbones (Lin et al., 2024; Wang et al., 2024b; Wu et al., 2024b; Yang et al., 2024). Although these approaches deliver strong performance, they also rely on tens to hundreds of sampling steps at inference time.

Beyond the specific context of SR, reducing the number of inference steps has become an active research direction for DMs more broadly. Prominent strategies include the development of efficient samplers (Song et al., 2020; Lu et al., 2025), the design of straight or near-optimal transport paths (Liu et al., 2022; Esser et al., 2024; Wang et al., 2024a; Kim et al., 2024a), and various knowledge distillation techniques (Salimans & Ho, 2022; Sauer et al., 2024b;a; Yin et al., 2024; Chadebec et al., 2025a). While these approaches have proven highly effective for unconditional or text-to-image generation, they do not directly translate to SR, where generation is strongly conditioned on the input image whose informativeness depends on both the degradation process and the upscaling factor. To address this challenge, several one-step diffusion-based SR methods (i.e., requiring only a single inference step) have recently emerged, primarily relying on teacher–student distillation. These approaches notably include adversarial distillation (Zhang et al., 2024; Wu et al., 2025; Xie et al., 2024; He et al., 2024), score distillation (Wu et al., 2024a), and consistency techniques (Wang et al., 2024c), often combined with perceptual regularization losses. Despite their efficiency and strong empirical performance, those methods are fundamentally constrained by information loss between the teacher and student models, which may legitimately degrade photorealism, restrict texture diversity, and reduce robustness across different upscaling factors.

Recently, a new paradigm for fast diffusion inference, referred to as Flow Map models, has been introduced (Boffi et al., 2025a;b), with concrete instantiations proposed in the literature (Geng et al., 2025; Frans et al., 2025; Sabour et al., 2025). These models adopt a *self-distillation* paradigm: starting from a FM formulation, they learn a dynamical mapping between a source and a target distribution while enforcing self-consistency across different discretization scales through an additional training objective. This formulation enables stable training in high-dimensional spaces and supports inference with very few (potentially a single) model evaluations. Flow Map models have recently demonstrated state-of-the-art performance for image generation at resolution 512×512 (Sabour et al., 2025), and ongoing work has begun exploring their scalability to higher resolution (Zheng et al., 2025; Hu et al., 2025). Crucially, unlike distillation approaches that compress an iterative denoising process into the student network—thereby potentially incurring information loss—Flow Map models preserve the expressivity of large generative models through self-consistency, making them particularly appealing for the SR task.

In this work, we investigate the use of Flow Map models for image SR. Our objective is to combine the expressive power of large DMs with fast inference, while maintaining both faithfulness to the underlying HR content and strong perceptual quality. Our contributions are threefold:

- We introduce **FlowMapSR**, a simple and generic diffusion-based framework for SR built upon the three major Flow Map formulations (*Eulerian*, *Lagrangian*, *Shortcut*). To adapt Flow Map models to the specific challenges of SR, we propose a base implementation that guarantees perceptual quality along with two enhancements: positive–negative prompting guidance inspired by Zhang et al. (2024), and lightweight adversarial fine-tuning using LoRA modules (Hu et al., 2022).

- We empirically demonstrate that the FlowMapSR instance based on the *Shortcut* formulation (Frans et al., 2025) (equivalently the *Progressive* approach from Boffi et al. (2025a)) provides the best performance and scalability in high-resolution SR scenarios. With inference time comparable to existing one-step state-of-the-art SR methods, this version of FlowMapSR produces reconstructions that are in general more faithful to the HR reference while exhibiting improved photorealistic characteristics, including realistic textures, depth of field, and high-frequency details. We validate these results on both synthetic and real-world benchmarks, and demonstrate consistent performance across challenging upscaling scenarios ($\times 4$ and $\times 8$) using a single model, without additional degradation-specific guidance or conditioning.
- Finally, we present a comprehensive ablation study of the proposed framework. Our analysis highlights a clear trade-off between reconstruction faithfulness and perceptual realism, providing insights into how FlowMapSR can be tuned to different application requirements.

2 RELATED WORKS

Diffusion models have emerged as the dominant paradigm for image SR, thanks to their strong generative capabilities and their ability to overcome limitations of earlier GAN-based approaches Wang et al. (2018; 2021), that were prone to training instability and lack of photorealism. Early diffusion-based SR methods such as SR3 (Saharia et al., 2022) and SRDiff (Li et al., 2022) adapted the conditional DDPM framework (Ho et al., 2020) by explicitly conditioning the generation process on the LR input in pixel space. IDM (Gao et al., 2023) further extends this line of work by incorporating an implicit neural representation, enabling SR across arbitrary continuous upscaling factors. Thereafter, several methods extended diffusion-based SR to latent space by leveraging large pre-trained text-to-image backbones, such as Stable Diffusion (Rombach et al., 2022). Notable examples include DiffBIR (Lin et al., 2024), StableSR (Wang et al., 2024b), and SeeSR (Wu et al., 2024b), which exploit rich semantic and perceptual representations learned at scale to improve visual quality. Despite their scalability and improved semantic modeling, these approaches still rely on the standard iterative diffusion inference process, and consequently require a large number of inference steps—typically more than 50. To attenuate this issue, some methods such as ResShift (Yue et al., 2023; 2024), IR-SDE (Luo et al., 2023), and InDI (Delbracio & Milanfar, 2023) accelerated inference by designing shorter Markov chains or optimizing transition schedules, cutting the number of inference steps from hundreds to a dozen with minimal loss in reconstruction quality.

A complementary and increasingly active research direction investigates one-step (or few-step) diffusion-based formulations via distillation techniques. In this setting, a large pre-trained text-to-image DM acts as a teacher and guides the training of a lightweight student network, enabling SR with a drastically reduced number of model evaluations. Existing methods primarily differ in the choice of distillation objective and the regularization mechanisms used to preserve fidelity and perceptual quality. We may distinguish: **(a) consistency-based distillation**, in which the student is trained to match the teacher’s outputs through direct regression or bidirectional consistency constraints, see for instance SinSR (Wang et al., 2024c); **(b) score distillation**, which aims to minimize a predefined divergence between the teacher and student output distributions. OSEDiff (Wu et al., 2024a) follows this approach and incorporates the LPIPS perceptual loss (Zhang et al., 2018) to encourage the recovery of high-frequency details; **(c) adversarial distillation**, which explicitly aligns the student’s output distribution with that of the teacher using an auxiliary discriminator. For example, ADDSR (Xie et al., 2024) introduces a timestep-adaptive adversarial loss that dynamically balances teacher–student alignment across diffusion steps, leading to improved perceptual realism. Similarly, TAD-SR (He et al., 2024) combines a time-aware adversarial loss with high-frequency-enhanced score distillation to better recover fine textures. Along the same line, S3Diff (Zhang et al., 2024) enhances adversarial distillation by integrating LPIPS regularization, positive–negative prompting guidance, and LoRA-based degradation-aware conditioning, significantly improving robustness to real-world degradations. Finally, VPD-SR (Wu et al., 2025) augments adversarial distillation with semantic-aware supervision and high-frequency perceptual losses, leveraging pre-trained vision–language representations to further improve semantic alignment and texture fidelity.

3 PRELIMINARIES ON FLOW MAP MODELS

Consider two probability distributions π_0 (target) and π_1 (source) supported on \mathbb{R}^d , for some $d \geq 1$. We assume access to a training dataset $\mathcal{D} = \{(x_0^i, x_1^i)\}_{i=1}^N$, consisting of N paired samples drawn from a joint coupling of π_0 and π_1 , denoted by $\pi_{0,1}$. In this section, we study the problem of generating new samples from π_0 based on *unseen* samples from π_1 .

A natural approach to this conditional generative modeling problem is to learn a transport map $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that pushes π_1 to π_0 (such that $T(x_1) \sim \pi_0$ whenever $x_1 \sim \pi_1$) by leveraging the available samples from \mathcal{D} . In practice, this is commonly addressed within the Flow Matching (FM) (Lipman et al., 2022) setting, which we briefly recall in Section 3.1. Basically, FM models learn the drift of an Ordinary Differential Equation (ODE) whose solution transports π_1 to π_0 , using a neural network. The resulting dynamics (neural ODE) defines a valid candidate for T , by composing a sequence of local transport maps obtained by numerically integrating the neural ODE over small time intervals.

While this approach is conceptually simple and has proven effective across a wide range of data modalities, it comes with a significant computational cost at inference time. In particular, FM inference generally requires dozens of neural network evaluations to control time discretization error, which can make FM-based methods impractical in large-scale settings. Flow Map models (Boffi et al., 2025a;b) extend the FM framework by explicitly addressing this limitation. Instead of relying on ODE numerical integration, they introduce an additional self-distillation objective which aims to learn *exact* integrations of the FM-based neural ODE on larger time intervals, thus enabling efficient inference with substantially fewer evaluations. In Section 3.2, we review the main components of this recent approach. We then introduce in Section 3.3 additional elements on conditioned Flow Map models that will serve as the foundation of our FlowMapSR framework.

3.1 BACKGROUND ON FLOW MATCHING

Interpolant & ODE. FM can be seen as a particular instantiation of the Stochastic Interpolant (SI) framework (Albergo et al., 2023), whose paradigm is the following : (i) defining a continuous-time path of probability distributions $(p_t)_{t \in [0,1]}$ that interpolates between the boundary conditions $p_0 = \pi_0$ and $p_1 = \pi_1$, and (ii) constructing a marginal-preserving ODE by identifying a velocity field $v : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that the stochastic process $(X_t)_{t \in [0,1]}$ governed by

$$dX_t = v_t(X_t)dt, \quad X_1 \sim \pi_1, \quad (1)$$

verifies $X_t \sim p_t$ for any $t \in [0, 1]$. Then, one may generate samples from π_0 by simulating ODE (1) backward in time, starting from samples drawn from π_1 . In practice, the bridging distributions $(p_t)_{t \in [0,1]}$ are defined as the marginals of a *linear* interpolant $(I_t)_{t \in [0,1]}$ of the form $I_t = \alpha_t X_0 + \beta_t X_1$, where $(X_0, X_1) \sim \pi_{0,1}$ and $\alpha, \beta : [0, 1] \rightarrow [0, 1]$ are continuously differentiable functions satisfying boundary conditions $\alpha_0 = \beta_1 = 1$ and $\alpha_1 = \beta_0 = 0$. The standard FM setting is recovered by taking $\alpha_t = 1 - t$ and $\beta_t = t$. *Throughout the main paper, we restrict our attention to this standard formulation, and defer the general case to Section A.1.* In this setting, a valid—though generally intractable—expression for the velocity field is $v_t(x) = \mathbb{E}[v_{\text{FM}}^{\text{target}} | I_t = x]$, where $v_{\text{FM}}^{\text{target}} = X_1 - X_0$ corresponds to the conditional target velocity.

Training & Inference. If v_t were known, a valid transport map pushing π_1 to π_0 would be given for any $x \sim \pi_1$ by

$$T_{\text{FM}}(x) = x - \int_0^1 v_r(X_r^{1,x})dr, \quad (2)$$

where $(X_r^{t,x})_{r \in [0,t]}$ denotes for any $t \in [0, 1]$ the solution to ODE (1) on $[0, t]$, with terminal condition $X_t = x$. In practice, v_t has no closed form but it can still be efficiently learned by learning a parametric model v_t^θ (e.g., a neural network) via the mean-square regression problem

$$\arg \min_{\theta} \mathbb{E}[\mathcal{L}_{\text{FM}}(v_t^\theta, I_t)], \quad \mathcal{L}_{\text{FM}}(v_t^\theta, I_t) = \|v_t^\theta(I_t) - v_{\text{FM}}^{\text{target}}\|^2, \quad (3)$$

where the expectation is taken over $(X_0, X_1) \sim \pi_{0,1}$ (approximated by samples from \mathcal{D}) and $t \sim q(t)$, q being a predefined time distribution on $[0, 1]$. At inference time, replacing v_t with its learned approximation v_t^θ yields a neural ODE whose associated transport map remains intractable and must itself be approximated numerically. The standard approach relies on discretizing the integral in (2) on small time intervals (backward in time) with numerical integration methods, starting from $t = 1$. Whether using the first-order Euler scheme or high-order ODE solvers such as Heun or Runge–Kutta methods, this inference procedure induces a large number of neural network evaluations, which can imply a prohibitive cost for large scale applications. Finally, the FM interpolation framework allows one to derive an *approximate* x_0 -prediction given a state (x_t, t) of the ODE trajectory, if we assume that the target velocity field is straight and constant :

$$\hat{x}_0(x_t) = x_t - t v_t^\theta(x_t). \quad (4)$$

While this approximation is accurate for small time values (where it coincides with a first-order Euler expansion around $t = 0$), it becomes generally loose at larger times for arbitrary couplings $\pi_{0,1}$, as it fails to capture the curvature of the neural ODE trajectory. However, in the special case where $\pi_{0,1}$ is the optimal transport plan associated with the quadratic cost between π_0 and π_1 , the x_0 -prediction (4) holds exactly for all $t \in [0, 1]$ (Pooladian et al., 2023; Tong et al., 2023). This observation has motivated a substantial body of work on few-step (and theoretically single-step) FM models (Liu, 2022; Lee et al., 2024; Kim et al., 2024a; Kornilov et al., 2024), which provide an alternative to Flow Map approaches. A detailed study of these methods is beyond the scope of this paper.

FM enhancing via CFG. When data samples are associated with an additional class information \mathbf{c} (e.g., a text label), the natural extension of FM consists in learning a class-conditional velocity field $v_t(x|\mathbf{c}) = \mathbb{E}[v_{\text{FM}}^{\text{target}} | I_t = x, \mathbf{c}]$, which replaces v_t in ODE (1). In image generation, however, it has been empirically observed that combining $v_t(x|\mathbf{c})$ with its unconditional counterpart $v_t(x|\emptyset)$ (that is learned in standard FM) significantly improves sample quality.

This popular trick, known as *Classifier-Free Guidance* (CFG) (Ho & Salimans, 2022), amounts to substitute $v_t(x|\mathbf{c})$ at inference by the following linear combination

$$v_t^{\text{cfg}}(x|\mathbf{c}) = w v_t(x|\mathbf{c}) + (1 - w)v_t(x|\emptyset), \quad (5)$$

where $w \in (1, \infty)$ denotes the guidance strength, recovering the canonical setting in the limit $w \rightarrow 1$. In practice, increasing w enhances perceptual quality but often reduces faithfulness to realistic content. The conditional velocity $v_t(x|\mathbf{c})$ is learned via a class-conditional neural network $v_t^\theta(x|\mathbf{c})$ by minimizing a conditional variant of (3),

$$\arg \min_{\theta} \mathbb{E}[\mathcal{L}_{\text{cFM}}(v_t^\theta, \mathbf{I}_t, \mathbf{c})], \quad \mathcal{L}_{\text{cFM}}(v_t^\theta, \mathbf{I}_t, \mathbf{c}) = \|v_t^\theta(\mathbf{I}_t|\mathbf{c}) - v_{\text{FM}}^{\text{target}}\|^2, \quad (6)$$

where the expectation is taken over $(X_0, X_1, \mathbf{c}) \sim \pi_{0,1}$ and $t \sim q(t)$. To expose v^θ to class-unconditional samples, a common practice is to drop the class condition \mathbf{c} in (6) with a prescribed probability during training.

3.2 BACKGROUND ON FLOW MAP MODELS

Definition & parameterization. Closely related to consistency-based DMs (Kim et al., 2024b; Zheng et al., 2024; Heek et al., 2024; Li & He, 2025; Kim et al., 2025), Flow Map models (Boffi et al., 2025b;a) and related approaches (Geng et al., 2025; Sabour et al., 2025; Frans et al., 2025) extend the FM framework to alleviate its high inference cost. Their key idea is to directly approximate the *integrated* velocity field over large time intervals, thereby enabling efficient inference with only a few steps. Formally, these methods aim to estimate the so-called *Flow Map* $X : [0, 1]^2 \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ associated with the ODE (1). For any pair of timesteps $(s, t) \in [0, 1]^2$ such that $s \leq t$ and any point $x \in \mathbb{R}^d$, the Flow Map is defined as

$$X_{s,t}(x) = x - \int_s^t v_r(X_r^{t,x}) dr. \quad (7)$$

In particular, the transport map $x \mapsto X_{0,1}(x) = \mathbb{T}_{\text{FM}}(x)$ can be obtained with a single evaluation of the Flow Map, making this formulation especially well suited for one-step inference. Following prior work (Boffi et al., 2025a; Frans et al., 2025), we consider in the rest of the paper an Euler-like parameterization of the Flow Map given by

$$X_{s,t}(x) = x - (t - s)u_{s,t}(x), \quad u_{s,t}(x) = \frac{1}{t - s} \int_s^t v_r(X_r^{t,x}) dr, \quad (8)$$

where $u_{s,t}$ denotes the *average* velocity field on time interval $[s, t]$ associated to ODE (1) conditioned on terminal point $X_t = x$, that we now aim to estimate with a neural network $u_{s,t}^\theta$. Note that the instantaneous velocity field $v_t(x)$ corresponds to a zeroth-order approximation of $u_{s,t}(x)$ in the limit regime $s \rightarrow t$, i.e., $u_{t,t}(x) = v_t(x)$, which establishes a straightforward bridge between FM and Flow Map frameworks. Since we already know how to learn v_t , see (3), the remaining challenge is to estimate $u_{s,t}$ for $s < t$.

Self-distillation losses. To this end, prior works introduce three consistency-based characterizations of the Flow Map, leading to so-called *Self-Distillation* (SD) objective functions. Given timesteps $s < t$ along with the FM interpolant \mathbf{I}_t , the SD losses take the general form

$$\mathcal{L}_{\text{SD}}(u_{s,t}^\theta, \mathbf{I}_t) = \|u_{s,t}^\theta(\mathbf{I}_t) - \text{sg}(u_{\text{SD}}^{\text{target}})\|^2, \quad (9)$$

where sg is the stopgradient operator, and $u_{\text{SD}}^{\text{target}}$ depends on the chosen characterization:

1. *Lagrangian* SD (LSD) loss (Boffi et al., 2025a, Eq. (13))

$$u_{\text{LSD}}^{\text{target}} = v_{\text{FM}}^{\text{target}} + (t - s)\partial_s u_{s,t}^\theta(\mathbf{I}_t)$$

2. *Eulerian* SD (ESD) loss (Geng et al., 2025, Eq. (9)), or equivalently (Boffi et al., 2025a, Eq. (14))

$$u_{\text{ESD}}^{\text{target}} = v_{\text{FM}}^{\text{target}} - (t - s)(\nabla u_{s,t}^\theta(\mathbf{I}_t)v_{\text{FM}}^{\text{target}} + \partial_t u_{s,t}^\theta(\mathbf{I}_t))$$

3. *Shortcut* SD (SSD) loss (Frans et al., 2025, Eq. (5)), or equivalently (Boffi et al., 2025a, Eq. (15))

$$u_{\text{SSD}}^{\text{target}} = \frac{1}{2}u_{r,t}^\theta(\mathbf{I}_t) + \frac{1}{2}u_{s,r}^\theta\left(\mathbf{I}_t - \frac{t-s}{2}u_{r,t}^\theta(\mathbf{I}_t)\right), \quad r = \frac{s+t}{2}$$

For completeness, we provide detailed derivations of these objectives in Section A.1. We note that SD losses are more expensive to evaluate than the standard FM loss: the LSD and ESD losses indeed require computing the derivatives of u^θ with respect to time and space (which can be efficiently implemented via jvp operations at the cost of an additional backward pass), while the SSD loss requires two extra evaluations of u^θ .

Training & Inference. Learning the Flow Map amounts to simply combining the flow matching loss \mathcal{L}_{FM} and the self-distillation loss \mathcal{L}_{SD} leading to the optimization problem

$$\arg \min_{\theta} \mathbb{E}[\mathcal{L}_{\text{FM-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t)], \quad \mathcal{L}_{\text{FM-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t) = \begin{cases} \mathcal{L}_{\text{FM}}(u_{t,t}^{\theta}, \mathbf{I}_t) & \text{if } s = t, \\ \mathcal{L}_{\text{SD}}(u_{s,t}^{\theta}, \mathbf{I}_t) & \text{else,} \end{cases} \quad (10)$$

where the expectation is taken over $(X_0, X_1) \sim \pi_{0,1}$ and $(s, t) \sim \bar{q}(s, t)$, \bar{q} being a predefined joint distribution over $\{(s, t) \in [0, 1]^2 : s \leq t\}$. Due to the Euler-like parameterization given in (8), inference with a Flow Map model is straightforward. Given a time discretization $\{t_k\}_{k=0}^K$ of $[0, 1]$ of size $(K + 1)$, with $K \geq 1$, $t_0 = 0$ and $t_K = 1$, and an initial sample $x_K \sim \pi_1$, one can recursively compute the iterates $\{x_k\}_{k=0}^K$ via the update

$$x_k = x_{k+1} - \delta_k u_{t_k, t_{k+1}}^{\theta}(x_{k+1}), \quad (11)$$

where $\delta_k = t_{k+1} - t_k$ denotes the k -th step-size. In particular, x_0 is approximately sampled from π_0 if u^{θ} accurately approximates u . Importantly, this formulation avoids the discretization error inherent to standard FM, allowing the number of inference steps K to remain very small, i.e., $K \in \{1, 2, 4\}$ (Boffi et al., 2025a; Sabour et al., 2025). As in the FM setting, Flow Map models admit direct x_0 -prediction at any intermediate step of the ODE trajectory; indeed, given a current state (x_t, t) , the corresponding Flow Map prediction is given by

$$\hat{x}_0(x_t) = x_t - t u_{0,t}^{\theta}(x_t). \quad (12)$$

Unlike the FM prediction in (4), the Flow Map prediction (12) is *exact* for any given timestep $t \in [0, 1]$.

3.3 GENERALIZATION OF CFG TO THE FLOW MAP SETTING

Assume now that each training sample $(x_0, x_1) \in \mathcal{D}$ is provided with an additional knowledge \mathbf{c} (e.g., a class conditioning), that we wish to exploit to improve sample quality. In the FM setting, the CFG paradigm discussed in Section 3.1 has been shown to bring substantial gains in image generation tasks, which naturally motivates its use within Flow Map models, especially in the context of SR. However, its application is not straightforward, as CFG is originally formulated for instantaneous velocity fields, whereas Flow Map models operate on averaged velocity fields over time intervals. Building on the insights of Geng et al. (2025) in the ESD setting, we propose to integrate CFG directly into Flow Map models at *training stage*, by exploiting the connection between FM and Flow Map formulations across all three SD settings. This leads to new CFG-aware expressions for the LSD and SSD objectives. To promote model flexibility and reduce the risk of overfitting, we further propose to *randomly* sample the guidance scale w during training from the interval $(1, w_{\max})$, where $w_{\max} > 1$ is treated as a hyperparameter¹. This design prevents specialization to a single guidance strength while still retaining the benefits of CFG at inference. Below, we detail the derivation of the Flow Map training objectives in the CFG setting.

Flow matching loss. Given a guidance scale $w \in (1, w_{\max})$, we rewrite the loss \mathcal{L}_{cFM} in (6) to directly introduce CFG at the velocity level, yielding the following objective

$$\mathcal{L}_{\text{cfg-FM}}(u_{t,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w) = \|u_{t,t}^{\theta}(\mathbf{I}_t | \mathbf{c}) - \text{sg}(v_{\text{cfg-FM}}^{\text{target}})\|^2, \quad v_{\text{cfg-FM}}^{\text{target}} = w v_{\text{FM}}^{\text{target}} + (1 - w) u_{t,t}^{\theta}(\mathbf{I}_t | \emptyset). \quad (13)$$

This loss naturally extends the standard CFG formula (5) to the Flow Map setting, by leveraging the identity $v_t = u_{t,t}$. Compared to the unconditional setting, each loss evaluation now requires one extra model evaluation.

Self-distillation losses. In the same spirit as in the FM setting, we derive the following CFG-based SD loss

$$\mathcal{L}_{\text{cfg-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w) = \|u_{s,t}^{\theta}(\mathbf{I}_t | \mathbf{c}) - \text{sg}(u_{\text{cSD}}^{\text{target}})\|^2, \quad (14)$$

where $u_{\text{cSD}}^{\text{target}}$ still depends on the chosen characterization:

1. CFG-based *Lagrangian* SD (cfg-LSD) loss

$$u_{\text{cfg-LSD}}^{\text{target}} = v_{\text{cfg-FM}}^{\text{target}} + (t - s) \partial_s u_{s,t}^{\theta}(\mathbf{I}_t | \mathbf{c}), \quad v_{\text{cfg-FM}}^{\text{target}} = w v_{\text{FM}}^{\text{target}} + (1 - w) u_{s,s}^{\theta}(\mathbf{I}_t - (t - s) u_{s,t}^{\theta}(\mathbf{I}_t | \emptyset) | \emptyset),$$

2. CFG-based *Eulerian* SD (cfg-ESD) loss, already given in (Geng et al., 2025, Eq. (17))

$$u_{\text{cfg-ESD}}^{\text{target}} = v_{\text{cfg-FM}}^{\text{target}} - (t - s) \left(\nabla u_{s,t}^{\theta}(\mathbf{I}_t | \mathbf{c}) v_{\text{cfg-FM}}^{\text{target}} + \partial_t u_{s,t}^{\theta}(\mathbf{I}_t | \mathbf{c}) \right), \quad v_{\text{cfg-FM}}^{\text{target}} = w v_{\text{FM}}^{\text{target}} + (1 - w) u_{t,t}^{\theta}(\mathbf{I}_t | \emptyset),$$

3. CFG-based *Shortcut* SD (cfg-SSD) loss

$$u_{\text{cfg-SSD}}^{\text{target}} = \frac{1}{2} u_{r,t}^{\theta}(\mathbf{I}_t | \mathbf{c}) + \frac{1}{2} u_{s,r}^{\theta}(\hat{\mathbf{I}}_s | \mathbf{c}), \quad \hat{\mathbf{I}}_s = \mathbf{I}_t - \frac{t - s}{2} u_{r,t}^{\theta}(\mathbf{I}_t | \mathbf{c}), \quad r = \frac{s + t}{2}.$$

Detailed derivations of these objectives are provided in Section A.2. Relative to the unconditional setting, computing the SD loss incurs two additional model evaluations in the Lagrangian setting, one in the Eulerian setting, and none in the Shortcut setting.

¹In contrast, Geng et al. (2025) adopt a deterministic training procedure with a fixed guidance scale.

Training & Inference. When combining both objectives previously established, we obtain the CFG-based objective

$$\arg \min_{\theta} \mathbb{E}[\mathcal{L}_{\text{cfg-FM-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w)], \mathcal{L}_{\text{cfg-FM-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w) = \begin{cases} \mathcal{L}_{\text{cfg-FM}}(u_{t,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w) & \text{if } s = t, \\ \mathcal{L}_{\text{cfg-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t, \mathbf{c}, w) & \text{else,} \end{cases}$$

where the expectation is taken over $(X_0, X_1, \mathbf{c}) \sim \pi_{0,1}$, $(s, t) \sim \bar{q}(s, t)$ and $w \sim \text{U}(1, w_{\max})$. Analogously to standard CFG practice, we randomly drop the class condition \mathbf{c} in (13) and (14) with a probability of 10%, ensuring that the drop is applied consistently across all evaluations of u^{θ} within a given setting. When this occurs, we additionally set $w = 1$, thereby recovering the original *unconditional* FM formulation for (13) and the corresponding *unconditional* SD formulation for (14). We observed in early experiments that this design choice significantly improves the quality of the generated samples. One could further incorporate additional evaluations of the instantaneous velocity $u_{t,t}^{\theta}(\cdot|\mathbf{c})$ or $u_{s,s}^{\theta}(\cdot|\mathbf{c})$ when computing $v_{\text{FM}}^{\text{target}}$, at the cost of introducing an additional hyperparameter, as discussed in (Geng et al., 2025, Appendix B.1) for the ESD setting. To avoid extra hyperparameter tuning, we do not explore this extension and leave it for future work.

Importantly, the proposed CFG-based Flow Map formulation avoids the inference-time cost increase typically associated with traditional CFG. In particular, it incurs the same computational cost as the unconditional setting, allowing us to benefit from CFG without additional overhead. Given an initial sample $x_K^{\mathbf{c}} \sim \pi_1$ with class condition \mathbf{c} , the unconditional update rule (11) naturally becomes

$$x_k^{\mathbf{c}} = x_{k+1}^{\mathbf{c}} - \delta_k u_{t_k, t_{k+1}}^{\theta}(x_{k+1}^{\mathbf{c}}|\mathbf{c}). \quad (15)$$

Similarly, the x_0 -prediction defined in (12) extends to the conditional setting as

$$\hat{x}_0(x_t|\mathbf{c}) = x_t - t u_{0,t}^{\theta}(x_t|\mathbf{c}). \quad (16)$$

4 METHODOLOGY

We propose to address the SR task within the Flow Map framework. In this setting, each paired sample $(x_0, x_1) \sim \pi_{0,1}$ consists of a low-resolution (LR) image $x_1 \in \mathbb{R}^{3 \times w \times h}$ and its high-resolution (HR) counterpart $x_0 \in \mathbb{R}^{3 \times W \times H}$, with a fixed aspect ratio satisfying $H/h = W/w \geq 1$. We refer to this ratio as the SR ‘‘upsampling factor’’ and denote it by s_{up} , while its inverse $s_{\text{down}} = 1/s_{\text{up}}$ corresponds to the ‘‘downscaling factor’’. In Section 4.1, we introduce a complete methodology tailored to the SR setting, which we term *FlowMapSR*. To further enhance both reconstruction faithfulness and perceptual quality, we extend FlowMapSR with a positive–negative prompting strategy inspired by Zhang et al. (2024), described in Section 4.2, as well as an additional adversarial fine-tuning stage detailed in Section 4.3.

4.1 BASE IMPLEMENTATION OF FLOWMAPSR

Data. Each training pair in our dataset \mathcal{D} consists of a ground-truth HR image x_0 and a corresponding LR image x_1 , synthetically generated by applying a sequence of random degradations to x_0 , including downscaling by a factor s_{down} . We build on the widely used degradation pipeline of Wang et al. (2021), which we adapt to better reflect real-world LR images similar to those shown in Figure 4. In particular, we reduce the strength of noise components (e.g., Gaussian noise, gray noise, and blur) to emphasize the effects of downscaling and compression; implementation details are provided in Section B. To improve robustness across diverse SR scenarios and promote generalization, we further adopt the continuous SR strategy of Gao et al. (2023) by sampling $s_{\text{down}} \sim \text{U}(0.1, 1)$ independently for each training pair.

The resulting LR image is then resized back to the spatial resolution of x_0 using a randomly selected interpolation mode (nearest, area, bilinear, or bicubic), followed by additional blur and pixel-value clamping. This procedure yields an LR source x_1 with the same spatial dimensions as x_0 , which is required by the Flow Map framework, while implicitly defining a continuous upscaling factor $s_{\text{up}} \sim \text{U}(1, 10)$. Although this setup theoretically enables SR up to $\times 10$, we restrict our main evaluation to $\times 4$ and $\times 8$ upscaling, to specifically emphasize the benefits of FlowMapSR approach in challenging SR scenarios. Finally, to enable large-scale training and inference, we operate in latent space by encoding image pairs with a non-trainable variational encoder. As a result, both training and inference of the Flow Map models described in Section 3 are conducted in latent space, with a final decoding step, denoted by **Dec**, applied at inference time to recover the HR image.

Time-dependent regularization with perceptual loss. Following prior work (Wu et al., 2024a; Zhang et al., 2024), we augment the standard Flow Map training objective (10) with an additional LPIPS perceptual loss (Zhang et al., 2018), which aligns the decoded HR prediction with the ground-truth HR image. This regularization significantly accelerates convergence and improves the recovery of meaningful image features at inference. The resulting objective is

$$\mathcal{L}_{\text{FM-SD}}^{\text{reg}}(u_{s,t}^{\theta}, \mathbf{I}_t) = \mathcal{L}_{\text{FM-SD}}(u_{s,t}^{\theta}, \mathbf{I}_t) + \lambda_s^{\text{LPIPS}} \text{LPIPS}(\hat{X}_0, X_0), \text{ where } \hat{X}_0 = \text{Dec}(\mathbf{I}_t - t u_{s,t}^{\theta}(\mathbf{I}_t)),$$

where $\lambda^{\text{LPIPS}} : [0, 1] \rightarrow [0, \infty)$ is a time-dependent weight and \hat{X}_0 denotes an approximate pixel-space x_0 -prediction from \mathbf{I}_t . Unlike previous approaches that use a constant λ^{LPIPS} , we deliberately design this weight to be *strictly decreasing* with

time, reflecting the accuracy of direct x_0 prediction from I_t under our Flow Map model: (a) when $s = t$ (FM case), \hat{X}_0 corresponds to the decoded approximation in (4), whose accuracy decreases as s increases; accordingly, λ^{LPIPS} is set higher for small s and lower for large s ; (b) when $s < t$ (SD case), \hat{X}_0 coincides with the exact prediction in (12) only for $s = 0$, but remains reliable for small s , motivating a similar time-dependent weighting. The specific form of λ^{LPIPS} used in our experiments is provided in Section B.

Dynamic loss weighting. To ensure meaningful self-consistency, the FM-SD loss must be evaluated over a broad range of timestep pairs (s, t) . In practice, this can induce strong time-dependent variance that hampers stable optimization. To mitigate this issue, we follow Boffi et al. (2025a), inspired by Karras et al. (2024), and jointly learn a *dynamic weighting* function $\lambda_{s,t}^\psi$, parameterized as a lightweight neural network described in Section B, alongside the Flow Map model $u_{s,t}^\theta$. This leads to the reweighted objective

$$\mathcal{L}(\theta, \psi) = \mathbb{E}[\mathcal{L}(u_{s,t}^\theta, I_t; \psi)], \quad \mathcal{L}(u_{s,t}^\theta, I_t; \psi) = \exp(-\lambda_{s,t}^\psi) \mathcal{L}_{\text{FM-SD}}^{\text{reg}}(u_{s,t}^\theta, I_t) + \lambda_{s,t}^\psi. \quad (17)$$

At optimality, this formulation encourages loss gradients with respect to θ to have comparable magnitudes across different (s, t) values (Karras et al., 2024, Eq. (26)), thereby stabilizing training.

Discrete timestep sampling. As described in the following section, our experiments rely on models with discrete time inputs, which precludes the use of continuous-time sampling strategies for designing \bar{q} (Boffi et al., 2025a). Instead, inspired by Frans et al. (2025), we sample timestep pairs (s, t) by first selecting a discretization level d of the interval $[0, 1]$. Each level d is drawn from the set $0, 1, \dots, d_{\text{max}}$, where $d_{\text{max}} \geq 1$ is a hyperparameter. This induces a grid of $2^d + 1$ uniformly spaced timesteps on $[0, 1]$, allowing up to $K_d = 2^d$ inference steps. In this formulation, $d = 0$ corresponds to one-step inference, $d = 1$ to two-step inference, and so on. To prevent the self-distillation loss from overwhelming the FM loss, we adopt a non-uniform allocation between FM and SD training samples by fixing the probability of sampling boundary cases with $s = t$. Following prior work, we set this probability to 75% in all experiments, resulting in 75% FM samples and 25% SD samples. Details on how this scheme is adapted to each SD setting are provided in Section B.

FM training warm-start. Finally, to prevent the self-distillation objective from dominating early training dynamics, we initialize each Flow Map model by training *without* the SD component in the FM-SD loss for a predefined number of iterations. In our experiments, we first train our model for 5,000 gradient steps using the FM objective, followed by an additional 5,000 gradient steps using the FM-SD loss. Further implementation details are provided in Section B.

4.2 ENHANCING FLOWMAPSR VIA POSITIVE-NEGATIVE PROMPTING GUIDANCE

Our base implementation of FlowMapSR, described in Section 4.1, is unconditional in the sense that it takes only the LR image as input and predicts the corresponding HR image at inference using the update rule (11). While this formulation effectively captures meaningful structures from the LR input, see Figure 2, and renders them in a visually plausible manner, the resulting images often lack sharpness and contrast, and therefore fall short in terms of photorealism. To address this limitation, we propose to incorporate CFG into FlowMapSR by adapting the positive–negative prompting strategy introduced in Zhang et al. (2024) to our framework, as detailed below. Specifically, we consider the following two generic text prompts:

- a negative prompt \mathbf{c}^{neg} : “oil painting, cartoon, blur, dirty, messy, low quality, deformation, low resolution, oversmooth”
- a positive prompt \mathbf{c}^{pos} : “a high-resolution, 8K, ultra-realistic image with sharp focus, vibrant colors, and natural lighting”

The objective is to enhance the performance of FlowMapSR conditioned on \mathbf{c}^{pos} by explicitly modeling poor visual quality through \mathbf{c}^{neg} , and enforcing a strong separation between the two. This strategy provides an effective guidance signal while avoiding reliance on external label extractors, which may be inaccurate or brittle.

Following the CFG formalism introduced in Section 3.3, we treat \mathbf{c}^{pos} as the *target* conditioning (replacing \mathbf{c}) and \mathbf{c}^{neg} as the *unconditional* conditioning (replacing \emptyset). To explicitly learn the notion of degraded quality, we introduce a key modification to the CFG–FM–SD loss related to the SR setting: when the class condition \mathbf{c}^{pos} is dropped and \mathbf{c}^{neg} is used instead, we replace the HR target x_0 with a synthetic low-quality version $x_0^{\text{neg}} \neq x_1$. This target is constructed to be less degraded than x_1 , ensuring that the model does not learn a trivial noise trajectory. Concretely, x_0^{neg} is obtained by applying the same degradation pipeline used to generate x_1 , except that the downscaling factor is sampled as $s_{\text{down}}^{\text{neg}} \sim \mathcal{U}(s_{\text{down}}, 1)$, where s_{down} denotes the downscaling factor associated with x_1 . This design guides the unconditional (negative) model toward generating low-quality versions of the HR image, while still preserving informative structure relative to the input x_1 through the difference in upscaling factors. At inference time, we exclusively use the Flow Map model conditioned on \mathbf{c}^{pos} to generate HR images via the update rule (15). As a result, each text-conditioned Flow Map model (LSD, ESD, or SSD) is enhanced with CFG so as to improve perceptual quality relative to its unconditional counterpart. In early experiments, however, we observe that this strategy is beneficial only for the *Shortcut* formulation, as illustrated in Figure 2. In contrast, the *Lagrangian* and *Eulerian* formulations exhibit noticeable artefacts when combined with CFG; qualitative examples illustrating this failure are provided in Section C.



Figure 2: **Qualitative comparison of positive-negative CFG enhancement applied to LSD-, ESD-, and SSD-based FlowMapSR for $\times 4$ upscaling.** The HR reference image exhibits sharp edges and rich high-frequency details. The base FlowMapSR model (“Without CFG”) successfully recovers structural details but lacks perceptual sharpness. Applying positive–negative CFG improves visual realism in the *Shortcut* formulation (e.g., more natural hair textures), while it introduces noticeable degradations and artefacts in the *Lagrangian* and *Eulerian* formulations.

We attribute the poor performance of our *Lagrangian* and *Eulerian* CFG-based formulations to unstable training dynamics, likely amplified by the use of JVP operators, which are challenging to handle in practice. This instability is consistent with observations reported in prior unconditional implementations, see (Boffi et al., 2025a, Figure 4), and could potentially be mitigated through specialized training strategies, such as those proposed in (Sabour et al., 2025, Section 3.4). In contrast, our *Shortcut* CFG-based formulation exhibits stable training behavior without requiring additional heuristics, which may explain its superior empirical performance. In the following, we denote by $\mathcal{L}(u_{s,t}^\theta, I_t, \mathbf{c}, w; \psi)$ the CFG-based adaptation of the previously established loss $\mathcal{L}(u_{s,t}^\theta, I_t; \psi)$ given in (17) with positive-negative conditioning \mathbf{c} and guidance scale w , where $\mathcal{L}_{\text{FM-SD}}$ is now replaced by $\mathcal{L}_{\text{cfg-FM-SD}}$.

4.3 ENHANCING FLOWMAPSR VIA LORA-BASED ADVERSARIAL FINE-TUNING

Although positive–negative CFG improves the performance of the *Shortcut*-based formulation of FlowMapSR, we still observe residual artefacts in the generated HR images, which are attributable to the guidance mechanism itself (a phenomenon that is also prone to occur in standard image generation tasks). To further mitigate these artefacts and strengthen the alignment between the model outputs and the ground-truth HR distribution, we introduce an additional fine-tuning stage that augments the CFG-based training objective with an adversarial loss, treating the Flow Map model as the generator.

Inspired by recent work Sabour et al. (2025), we adopt the Relativistic Pairing GAN (RPGAN) framework (Jolicoeur-Martineau, 2018), which has been shown to be more robust to mode collapse than standard GAN formulations (Sun et al., 2020). The adversarial training is performed in latent space. Given a discriminator \mathbf{D}^ϕ parameterized by a neural network (described in Section B), we define the following adversarial objective, with separate generator (G-loss) and discriminator (D-loss) terms:

$$\mathcal{L}^{\text{adv}}(u_{s,t}^\theta, I_t, \mathbf{c}, w; \psi, \phi) = \begin{cases} \text{Softplus} \left(\mathbf{D}^\phi(\hat{Z}_0^{\text{adv}}) - \mathbf{D}^\phi(Z_0) \right) + \lambda^{\text{adv}} \mathcal{L}(u_{s,t}^\theta, I_t, \mathbf{c}, w; \psi) & \text{(G-loss)} \\ \text{Softplus} \left(\mathbf{D}^\phi(Z_0) - \mathbf{D}^\phi(\text{sg}(\hat{Z}_0^{\text{adv}})) \right) & \text{(D-loss)}, \end{cases}$$

where $\lambda^{\text{adv}} > 0$ controls the contribution of the original loss, Z_0 denotes the latent encoding of X_0 , and \hat{Z}_0^{adv} is the latent x_0 -prediction produced by the FlowMapSR model from X_1 . To fully exploit the expressivity of the Flow Map model, we compute \hat{Z}_0^{adv} using a two-step prediction (rather than single-step), following (15). In particular, the intermediate state at $t = 1/2$ is detached to avoid the computation of higher-order derivatives, in line with Shen et al. (2025). The parameter λ^{adv} balances adversarial alignment with preservation of the pre-trained FlowMapSR behavior, as excessive adversarial pressure may induce divergence from the original model. As in standard GAN settings, we consider reaching adversarial convergence when we have equality between the “fake” score $\mathbf{D}^\phi(\hat{Z}_0^{\text{adv}})$ and the “real” score $\mathbf{D}^\phi(Z_0)$.

Table 1: **Comparison of inference efficiency between FlowMapSR and competing SR methods.** We report the average inference time measured on LR RealSR inputs for $\times 4$ upscaling (input resolution 128×128) using a single H100 GPU. For FlowMapSR, we evaluate three variants with 1, 2, and 4 inference steps. FlowMapSR-1 achieves inference time comparable to OSEDiff-1, while FlowMapSR-2 and FlowMapSR-4 remain competitive with S3Diff-1.

Method	Real-ESRGAN	OSEDiff-1	S3Diff-1	FlowMapSR-1	FlowMapSR-2	FlowMapSR-4
Inference time (s)	0.05	0.14	0.40	0.14	0.22	0.40

Owing to the relative nature of RPGAN, which promotes stable training dynamics, we alternate optimization of the G-loss and D-loss at each training step. We found it unnecessary to include additional R_1 or R_2 regularization terms in the discriminator loss; however, pretraining the discriminator for a small number of iterations was observed to accelerate convergence, and we therefore adopt this strategy in all experiments.

Finally, to keep the fine-tuning stage lightweight while preserving the knowledge acquired by the CFG-based Flow Map model, we restrict adversarial training to Low-Rank Adaptation (LoRA) layers (Hu et al., 2022) within the generator, following the strategy from Chadebec et al. (2025a). This choice reduces computational cost during training and enables precise control of the LoRA scaling factor at inference time, allowing us to modulate the impact of the adversarial refinement.

5 EXPERIMENTS

In this section, we evaluate FlowMapSR on synthetic and real-world LR inputs for $\times 4$ and $\times 8$ super-resolution. We first describe the experimental setup in Section 5.1, then present a quantitative comparison with competing SR methods in Section 5.2, complemented by representative qualitative examples (with additional results in Section C, including $\times 2$ super-resolution). These visual comparisons emphasize the recovery of lifelike textures (e.g., fur, human skin, metals) and photo-realistic attributes (e.g., depth of field, contrast, color coherence). Finally, in Section 5.3, we study the key hyperparameters of FlowMapSR that control the trade-off between reconstruction faithfulness and perceptual quality, including the number of inference steps, the LoRA scaling factor at inference, and the maximum positive-negative CFG scale.

5.1 EXPERIMENTAL SETUP

FlowMapSR framework. Throughout this section, our default Flow Map model corresponds to the *Shortcut* variant of FlowMapSR, augmented with positive-negative CFG (Section 4.2) and adversarial fine-tuning (Section 4.3). We do not include results for the *Eulerian* and *Lagrangian* variants, as they consistently underperform on the SR task compared to the *Shortcut* formulation (see Figure 2). Unless stated otherwise, all reported results are obtained by using the guidance strength $w_{\max} = 3.5$, setting the LoRA scale as 1.5 and applying 2 inference steps (denoted by *FlowMapSR-2*). The model adopts a UNet architecture following SDXL (Podell et al., 2024), initialized from the image restoration weights of Chadebec et al. (2025b). Training is conducted on a single H100 GPU using HR targets of up to 1024×1024 resolution; additional training details and hyperparameter settings are provided in Section B. Notably, the same model is used for both $\times 4$ and $\times 8$ upscaling, without any scale-specific conditioning. At inference time, LR inputs are first resized in pixel space using bicubic interpolation according to the target scale, then processed in latent space by the Flow Map model, and finally decoded to obtain the HR output. When memory constraints arise due to large target resolutions, inference is performed after decoding on pixel-space tiles, followed by Gaussian blending.

Testing details. For synthetic evaluation with available ground truth, we generate LR-HR pairs from the DIV2K validation (DIV2K-Val) set (Agustsson & Timofte, 2017), RealSR (Cai et al., 2019), and DRealSR (Wei et al., 2020), using the degradation pipeline described in Section 4.1. We recall that this pipeline is a lighter variant of the one proposed in Wang et al. (2021), with reduced additional noise, see Section B for more details. Performance on synthetic datasets is assessed using both reference metrics—PSNR, SSIM (Wang et al., 2004), LPIPS (Zhang et al., 2018), DISTS (Ding et al., 2020), and FID (Heusel et al., 2017)—which measure fidelity to the HR ground truth, and non-reference metrics—NIQE (Zhang et al., 2015), MANIQA (Yang et al., 2022), MUSIQ (Ke et al., 2021), and CLIPIQA (Wang et al., 2023)—which better reflect perceptual quality. We additionally present extensive qualitative results on DIV2K-Val and RealSR. For real-world evaluation, we use the RealSet65 dataset (Yue et al., 2023) and report qualitative comparisons. FlowMapSR is compared against the GAN-based method Real-ESRGAN (Wang et al., 2021) and state-of-the-art single-step diffusion-based methods OSEDiff (Wu et al., 2024a) and S3Diff (Zhang et al., 2024) (respectively denoted OSEDiff-1 and S3Diff-1 to emphasize their single-step nature). For fairness, we rely on publicly available code and pre-trained checkpoints to generate HR outputs on the same test sets. While FlowMapSR differs substantially from prior distillation-based approaches in training budget (under 100M parameters for OSEDiff and S3Diff versus over 2B for FlowMapSR), its inference time remains competitive for the default FlowMapSR-2 model, as reported in Table 1.

Table 2: **Quantitative comparison between FlowMapSR and competing SR methods ($\times 4$ and $\times 8$ upscaling).** The best and second best results are highlighted in **bold** and underlined, respectively.

Dataset (s_{up})	Method	Reference metrics					Non-reference metrics			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow	MUSIQ \uparrow	CLIPQA \uparrow
DIV2K-Val ($\times 4$)	Real-ESRGAN	23.03	0.6880	0.2739	0.1276	21.97	3.758	0.3536	0.6112	0.4960
	OSDiff-1	21.89	<u>0.6500</u>	0.2625	0.1107	19.23	3.373	0.4363	<u>0.6777</u>	0.5974
	S3Diff-1	20.87	0.6223	0.2475	<u>0.1093</u>	<u>15.39</u>	3.198	<u>0.4223</u>	0.6780	0.5893
	FlowMapSR-2	<u>22.12</u>	0.6412	<u>0.2554</u>	0.0995	13.05	3.036	0.4010	0.6431	0.5419
RealSR ($\times 4$)	Real-ESRGAN	22.28	0.7047	<u>0.2743</u>	0.1939	126.70	5.770	0.4253	0.6497	0.4575
	OSDiff-1	21.15	0.6610	<u>0.2752</u>	<u>0.1907</u>	106.10	5.180	0.4936	0.6925	0.5288
	S3Diff-1	<u>21.75</u>	<u>0.6762</u>	0.2585	0.1775	94.32	5.661	0.5075	0.6986	<u>0.5511</u>
	FlowMapSR-2	20.42	0.5965	0.3011	0.1964	<u>95.16</u>	<u>5.322</u>	0.5362	<u>0.6975</u>	0.5901
DrealSR ($\times 4$)	Real-ESRGAN	25.22	0.7771	0.2700	<u>0.1875</u>	140.31	6.954	0.3876	0.6001	0.4581
	OSDiff-1	23.54	0.7173	<u>0.2884</u>	0.1894	114.69	<u>5.745</u>	0.4758	0.6565	0.5634
	S3Diff-1	<u>24.36</u>	<u>0.7200</u>	0.2886	0.1818	105.48	5.734	0.4862	0.6720	<u>0.5771</u>
	FlowMapSR-2	23.15	0.6688	0.3056	0.1937	<u>108.74</u>	5.779	0.4863	<u>0.6600</u>	0.5894
DIV2K-Val ($\times 8$)	Real-ESRGAN	20.55	0.6061	0.3940	0.1882	<u>40.98</u>	5.209	0.2505	0.4612	0.3317
	OSDiff-1	19.68	0.5567	0.3909	<u>0.1667</u>	43.41	3.394	0.4080	0.6511	0.5759
	S3Diff-1	19.30	0.5386	0.3743	<u>0.1762</u>	47.49	3.069	0.3468	<u>0.6227</u>	<u>0.5712</u>
	FlowMapSR-2	<u>19.90</u>	<u>0.5671</u>	<u>0.3767</u>	0.1469	28.25	<u>3.197</u>	<u>0.3793</u>	0.6135	0.5008
RealSR ($\times 8$)	Real-ESRGAN	19.72	0.6014	0.3536	0.2327	179.34	6.294	0.2594	0.4787	0.2865
	OSDiff-1	18.86	<u>0.5294</u>	0.4174	0.2526	180.05	<u>4.592</u>	<u>0.4396</u>	0.6454	0.5011
	S3Diff-1	18.93	0.5224	<u>0.3828</u>	<u>0.2335</u>	<u>172.07</u>	4.208	0.4298	0.6546	0.5217
	FlowMapSR-2	17.96	0.4878	0.4190	0.2576	163.35	5.173	0.5042	0.6773	0.5746
DrealSR ($\times 8$)	Real-ESRGAN	22.46	0.6898	0.3378	0.2244	<u>178.19</u>	7.246	0.2612	0.4630	0.3252
	OSDiff-1	21.11	0.5989	0.4190	0.2520	180.56	4.763	<u>0.4216</u>	0.6153	0.5201
	S3Diff-1	<u>21.17</u>	0.5631	0.4150	<u>0.2405</u>	185.93	4.390	0.4099	0.6321	0.5359
	FlowMapSR-2	20.70	<u>0.5858</u>	<u>0.3944</u>	0.2481	170.81	5.743	0.4472	0.6365	0.5483

5.2 MAIN RESULTS

We report quantitative results for $\times 4$ and $\times 8$ upscaling in Table 2, from which we draw the following observations.

(a) On DIV2K-Val, FlowMapSR achieves strong performance on reference-based metrics, particularly perceptual measures such as LPIPS, DISTS, and FID, compared to competing methods. This high fidelity to the DIV2K-Val ground-truth images is consistent with the qualitative results shown in Figure 3, see respectively the rows 1-3 for $\times 4$ upscaling and the rows 4-7 for $\times 8$ upscaling. In particular, FlowMapSR exhibits faithful recovery of textures and details (e.g., human skin and attributes, animal fur), geometric structures, lighting, environmental context, and depth of field. While these qualitative examples demonstrate strong photorealism—often surpassing competing methods—this advantage is not always reflected in non-reference metrics, where OSDiff-1 tends to achieve higher scores. Visually, this may be explained by the stronger sharpness and contrast produced by OSDiff-1, which benefit such metrics but come at the expense of reference fidelity. We attribute the slightly lower non-reference scores of FlowMapSR in this setting to occasional boundary artefacts introduced by Gaussian tiling, see Section C, which is not required for processing RealSR and DRealSR images. Improving this tiling strategy is left for future work.

(b) On RealSR and DRealSR, the trend is reversed: FlowMapSR achieves the best average performance on non-reference metrics for $\times 4$ and $\times 8$ upscaling on both datasets, while still maintaining competitive DISTS and FID scores, especially at higher scaling factor. Although Real-ESRGAN attains strong reference-based scores on these datasets, this should be interpreted with caution. Visual inspection in Figure 3 (row 4 for $\times 4$, and row 8 for $\times 8$) reveals that Real-ESRGAN often produces visibly low-quality outputs. Despite being counter-intuitive, high reference-based scores and poor visual quality are not contradictory here; both reflect the limited visual quality of the RealSR and DrealSR ground truth, thereby revealing the limitations of reference metrics in this case. In contrast, while FlowMapSR may appear less faithful based on reference metrics, it consistently produces higher-quality images, with more realistic textures and convincing depth of field areas, where S3Diff comparatively struggles.

Finally, we present qualitative comparisons on real-world inputs from RealSet65 in Figure 4, where we show SR results for selected regions of the LR images. In this fully real-world setting, FlowMapSR consistently produces the most photorealistic and faithful outputs for both $\times 4$ and $\times 8$ upscaling, particularly in terms of texture realism and depth of field effects. Overall, FlowMapSR demonstrates strong and consistent performance across both synthetic and real-world scenarios, using a single model without any upscaling-specific conditioning. While this behavior is only partially captured by quantitative metrics, it is clearly reflected in qualitative results. Additional comparisons on DIV2K-Val, RealSR, and RealSet65 for both $\times 4$ and $\times 8$ upscaling are provided in Section C.

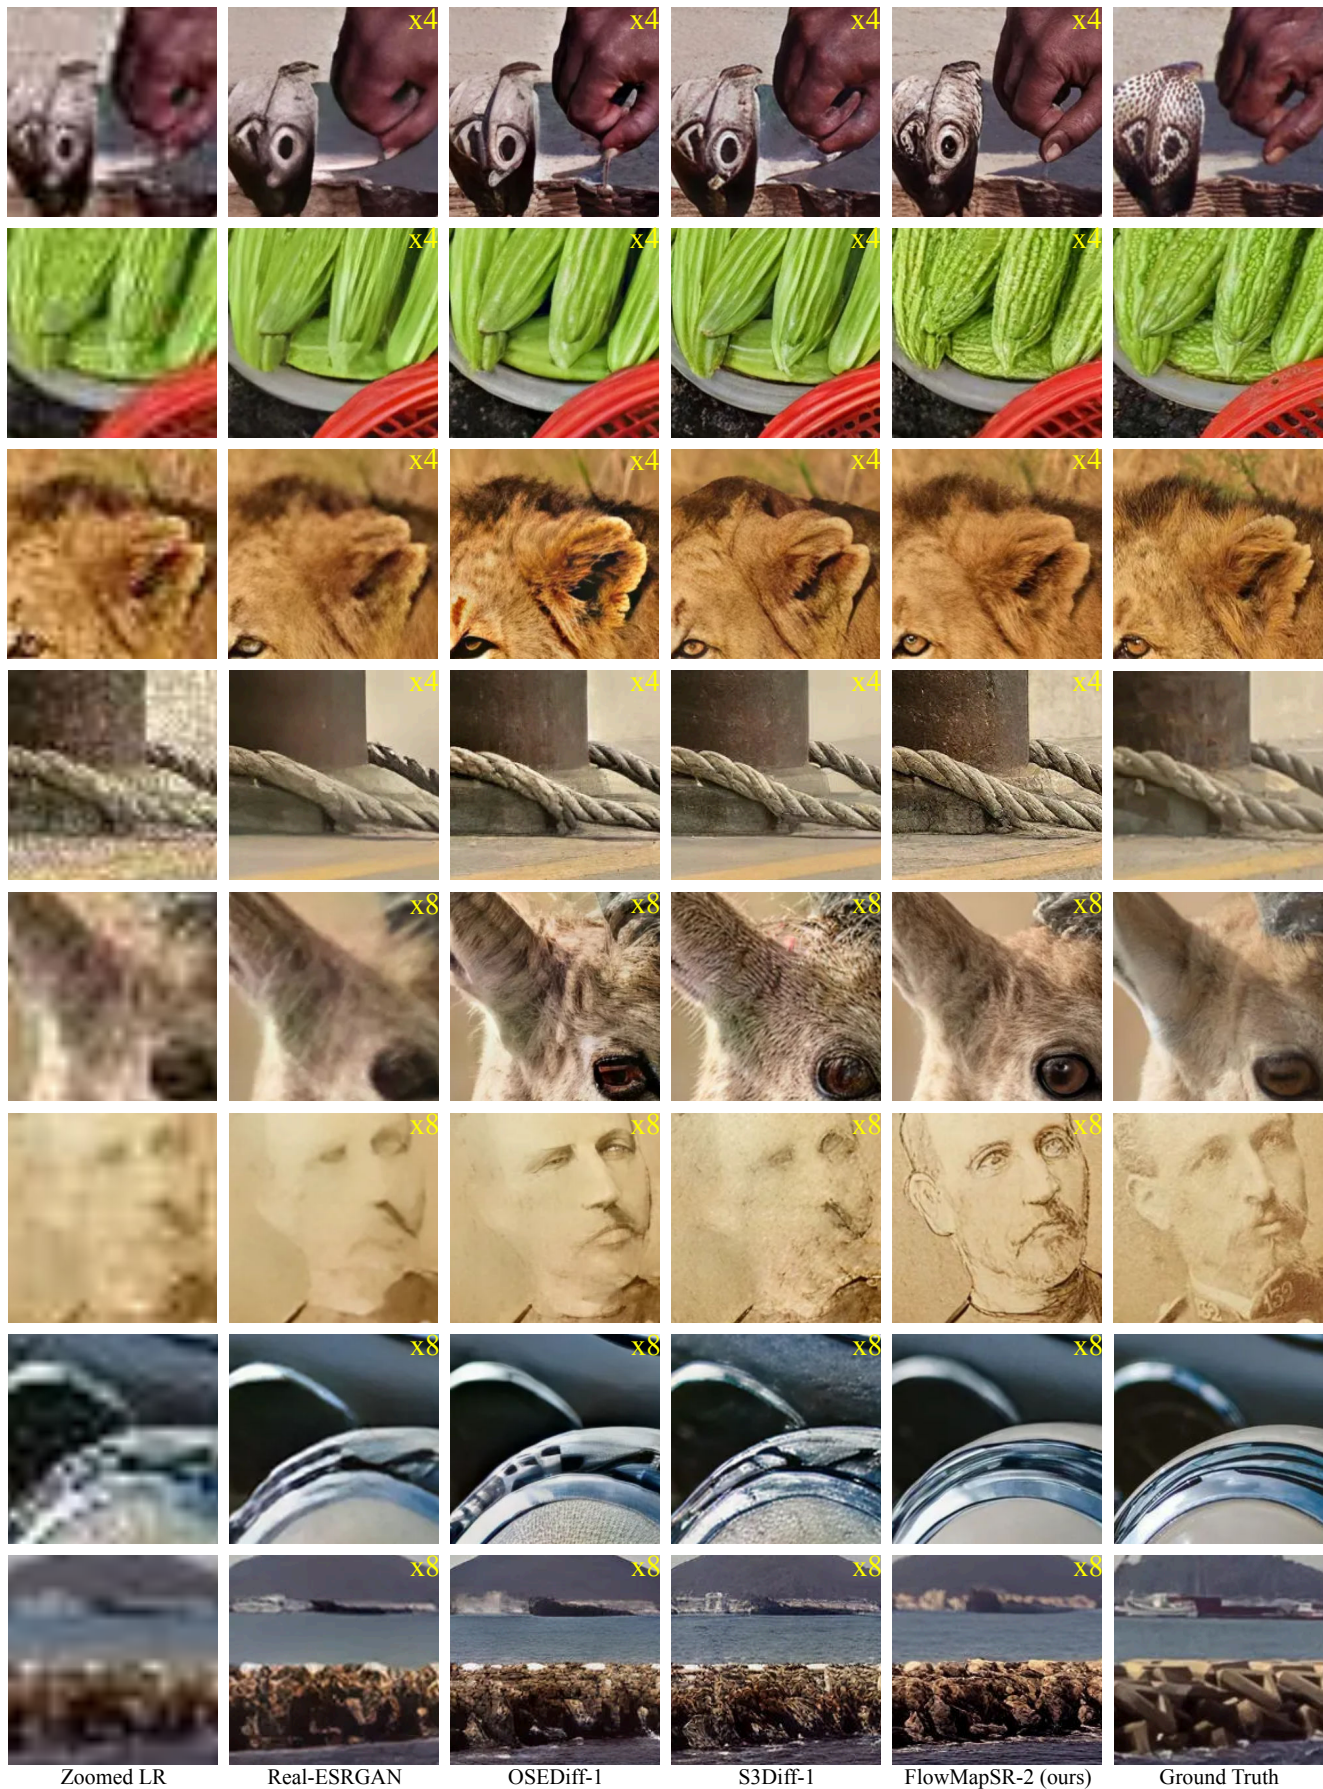


Figure 3: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 4$ (rows 1–4) and $\times 8$ (rows 5–8) upscaling.** Rows 1–3 and 5–7 show examples from DIV2K-Val, while rows 4 and 8 correspond to RealSR samples.



Figure 4: **Qualitative comparison between FlowMapSR and competing diffusion-based SR methods on real-world LR inputs for $\times 4$ and $\times 8$ upscaling.** The two LR images are taken from the RealSet65 dataset. For each image, a **green** region is selected for $\times 4$ upscaling and a **blue** region for $\times 8$ upscaling. FlowMapSR-2 consistently achieves the best balance between reconstruction faithfulness and photorealism.

5.3 ABLATION STUDIES

We analyze the influence of three key hyperparameters of FlowMapSR: the number of inference steps, the LoRA scaling factor applied at inference, and the maximum CFG scale used during training. For each setting, we report quantitative results on DIV2K-Val and RealSR for $\times 4$ and $\times 8$ upscaling, complemented by qualitative comparisons for the first two ablations. Additional visual results are provided in Section C.

Impact of the number of inference steps. Our Flow Map formulation allows the number of inference steps (NFE) to be chosen as a power of two; we evaluate 1, 2, and 4 steps. As shown in Table 3, using a single step consistently yields higher fidelity to the reference images, while increasing NFE improves model expressivity and perceptual quality, as reflected by non-reference metrics. These results highlight a clear trade-off: increasing NFE is not always beneficial. Qualitative examples in Figure 5 illustrate that larger NFE can degrade depth-of-field consistency (rows 1 and 3) and introduce overly shiny or glassy textures (rows 2 and 4), whereas a single step, while more faithful, lacks visual refinement. Overall, using two inference steps, as in our default configuration, provides an effective balance between faithfulness and perceptual quality.

Impact of the LoRA scale at inference. A similar trade-off is observed with respect to the LoRA scaling factor, which controls the influence of the adversarially fine-tuned parameters at inference. Increasing the LoRA scale (e.g., from 1 to 2) enhances sharpness and contrast, as illustrated in Figure 6. However, this comes at the cost of reduced reconstruction faithfulness and, in some cases, degraded photorealism, particularly in depth-of-field rendering (see rows 1 and 2). This effect becomes more pronounced for larger values (e.g., LoRA scale 3), which lead to clearly suboptimal results. The quantitative results in Table 4 confirm this behavior, with higher reference-based scores at lower LoRA scales and improved non-reference metrics at larger scales. In practice, the default setting with a LoRA scale of 1.5 offers a good compromise between fidelity and perceptual enhancement.

Table 3: **Quantitative comparison of FlowMapSR with varying number of inference steps (NFE) ($\times 4$ and $\times 8$ upscaling).** All FlowMapSR variants use $w_{\max} = 3.5$ and a LoRA scale of 1.5, consistent with the FlowMapSR-2 configuration. The best and second best results are highlighted in **bold** and underlined, respectively. **Red** cells denote the default FlowMapSR-2 configuration used in Table 2 and Figure 3, while **blue** cells mark the best results when also considering competing methods reported in Table 2.

Dataset (s_{sup})	NFE	Reference metrics					Non-reference metrics			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow	MUSIQ \uparrow	CLIPQA \uparrow
DIV2K-Val ($\times 4$)	1	22.89	0.6703	0.2464	0.0986	14.03	3.750	0.3810	0.6248	0.4995
	2	<u>22.12</u>	<u>0.6412</u>	<u>0.2554</u>	<u>0.0995</u>	13.05	<u>3.036</u>	<u>0.4010</u>	<u>0.6431</u>	<u>0.5419</u>
	4	20.94	0.5881	0.3033	0.1126	<u>13.55</u>	2.775	0.4134	0.6575	0.5848
RealSR ($\times 4$)	1	21.74	0.6517	0.2636	0.1790	90.54	5.630	0.4949	0.6764	0.5196
	2	<u>20.42</u>	<u>0.5965</u>	<u>0.3011</u>	<u>0.1964</u>	95.16	5.322	<u>0.5362</u>	<u>0.6975</u>	<u>0.5901</u>
	4	19.15	0.5322	0.3598	0.2174	103.86	<u>5.342</u>	0.5607	0.7095	0.6204
DIV2K-Val ($\times 8$)	1	20.75	0.5981	0.3568	0.1499	27.97	3.857	0.3537	0.5914	0.4648
	2	<u>19.90</u>	<u>0.5671</u>	<u>0.3767</u>	0.1469	<u>28.25</u>	<u>3.197</u>	<u>0.3793</u>	<u>0.6135</u>	<u>0.5008</u>
	4	18.85	0.5223	0.4249	0.1525	29.82	3.077	0.3987	0.6311	0.5449
RealSR ($\times 8$)	1	19.28	0.5486	0.3886	0.2458	166.73	5.758	0.4536	0.6495	0.5178
	2	<u>17.96</u>	<u>0.4878</u>	<u>0.4190</u>	<u>0.2576</u>	163.35	5.173	<u>0.5042</u>	<u>0.6773</u>	<u>0.5746</u>
	4	16.67	0.4179	0.4773	0.2755	<u>164.99</u>	<u>5.535</u>	0.5309	0.6970	0.6127




Figure 5: **Qualitative comparison of FlowMapSR with varying numbers of inference steps (NFE) for $\times 4$ (rows 1–2) and $\times 8$ (rows 3–4) upscaling.** The LR inputs are drawn from the DIV2K-Val dataset. Increasing NFE improves perceptual quality, while fewer steps yield higher faithfulness to the reference, illustrating the trade-off quantified in Table 3.

Table 4: **Quantitative comparison of FlowMapSR with varying LoRA scale ($\times 4$ and $\times 8$ upscaling).** All FlowMapSR variants use $w_{\max} = 3.5$ with 2 inference steps, consistent with the FlowMapSR-2 configuration. The best and second best results are highlighted in **bold** and underline, respectively. **Red** cells correspond to the default version of FlowMapSR-2 used in Table 2 and Figure 3.

Dataset (s_{up})	LoRA scale	Reference metrics					Non-reference metrics			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow	MUSIQ \uparrow	CLIPQA \uparrow
DIV2K-Val ($\times 4$)	1	22.29	0.6520	<u>0.2571</u>	0.0958	12.92	3.170	0.3865	0.6309	0.5138
	1.5	<u>22.12</u>	<u>0.6412</u>	0.2554	<u>0.0995</u>	<u>13.05</u>	<u>3.036</u>	0.4010	<u>0.6431</u>	0.5419
	2	21.83	0.6224	0.2675	0.1093	13.43	2.928	0.4059	0.6456	<u>0.5649</u>
	3	20.21	0.5232	0.3928	0.1800	19.20	3.569	0.4149	0.6323	0.6486
RealSR ($\times 4$)	1	20.50	0.6084	0.2891	0.1907	92.76	5.252	<u>0.5285</u>	<u>0.6961</u>	0.5753
	1.5	<u>20.42</u>	<u>0.5965</u>	<u>0.3011</u>	<u>0.1964</u>	<u>95.16</u>	<u>5.322</u>	0.5362	0.6975	<u>0.5901</u>
	2	20.22	0.5753	0.3274	0.2109	103.50	5.659	0.5271	0.6860	0.5967
	3	19.39	0.5280	0.4140	0.2526	137.74	7.134	0.4991	0.6662	0.6009
DIV2K-Val ($\times 8$)	1	20.08	0.5782	<u>0.3792</u>	0.1415	27.86	3.313	0.3671	0.6021	0.4721
	1.5	<u>19.90</u>	<u>0.5671</u>	0.3767	<u>0.1469</u>	<u>28.25</u>	<u>3.197</u>	<u>0.3793</u>	<u>0.6135</u>	0.5008
	2	19.62	0.5457	0.3888	0.1645	30.45	3.022	0.3788	0.6115	<u>0.5367</u>
	3	18.71	0.4670	0.4987	0.2392	43.63	3.592	0.3938	0.6136	0.6600
RealSR ($\times 8$)	1	18.07	0.5019	0.4041	0.2480	158.81	5.163	0.4989	0.6785	0.5634
	1.5	<u>17.96</u>	<u>0.4878</u>	<u>0.4190</u>	<u>0.2576</u>	<u>163.35</u>	<u>5.173</u>	0.5042	<u>0.6773</u>	0.5746
	2	17.80	0.4638	0.4563	0.2797	182.77	5.537	0.4791	0.6580	<u>0.5814</u>
	3	17.76	0.4410	0.5492	0.3148	229.96	6.784	0.4491	0.6226	0.5954

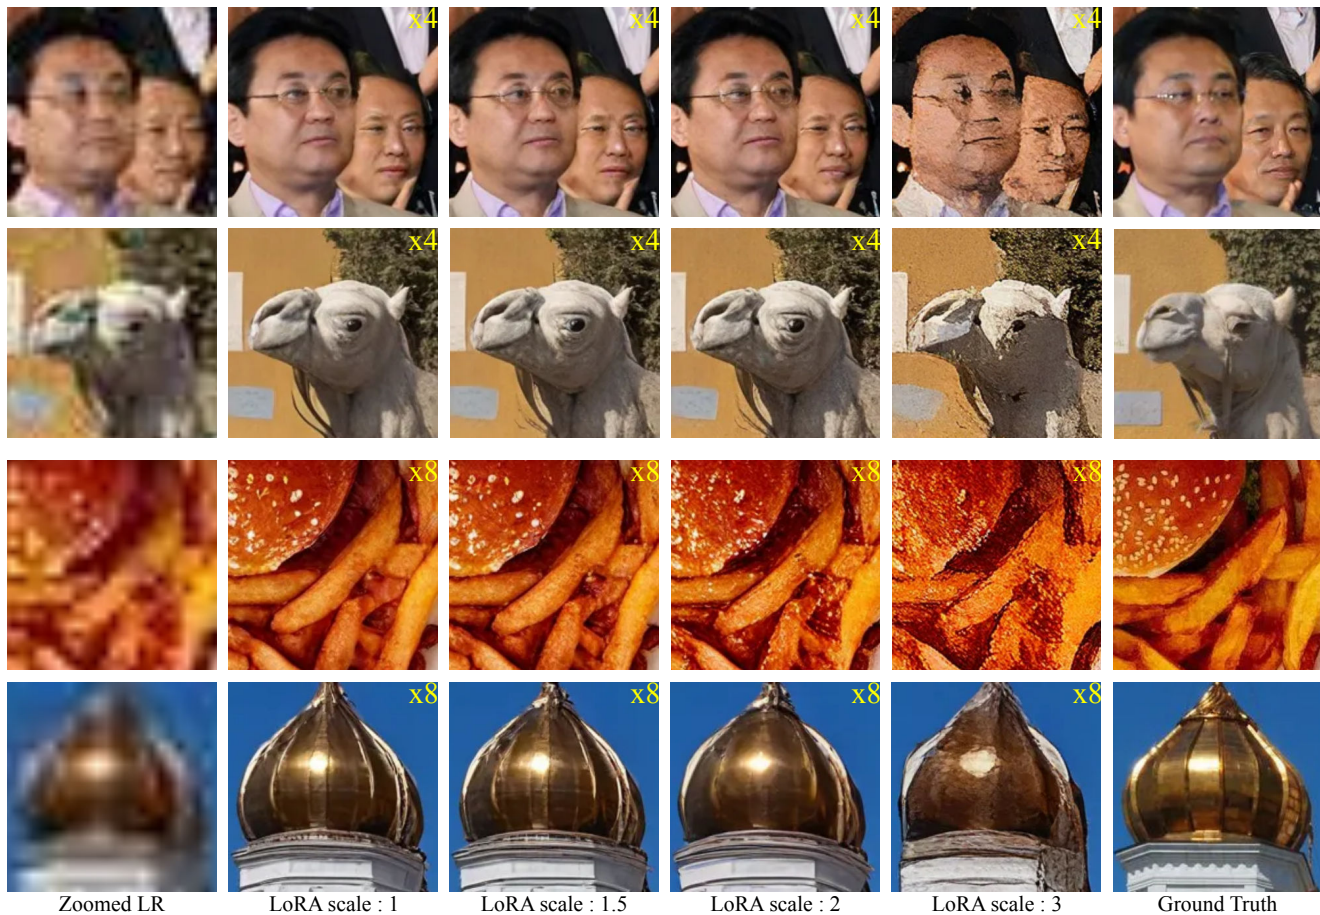


Figure 6: **Qualitative comparison of FlowMapSR with varying LoRA scale for $\times 4$ (rows 1–2) and $\times 8$ (rows 3–4) upscaling.** Inputs are from the DIV2K-Val dataset. Smaller LoRA scales better preserve fidelity to the reference, while larger scales favor sharpness (eventually degrading quality with LoRA scale set to 3), consistent with the trade-off in Table 4.

Table 5: **Quantitative comparison of FlowMapSR with maximal positive-negative CFG scale w_{\max} ($\times 4$ and $\times 8$ upscaling).** The best and second best results are highlighted in **bold** and underlined, respectively. **Blue** cells mark the best results when also considering competing methods reported in Table 2. All FlowMapSR variants use 2 inference steps, consistent with the FlowMapSR-2 configuration, but are obtained **without adversarial fine-tuning**.

Dataset (s_{up})	Max. CFG scale	Reference metrics					Non-reference metrics			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow	MUSIQ \uparrow	CLIPQA \uparrow
DIV2K-Val (x4)	2	23.15	0.6759	0.2568	0.1015	11.93	3.489	0.3484	0.6042	0.4605
	2.5	22.98	0.6675	0.2578	0.0987	<u>12.08</u>	<u>3.314</u>	0.3584	<u>0.6068</u>	0.4794
	3	23.48	0.6913	0.2658	0.1052	12.50	4.052	0.3344	0.5822	0.4415
	3.5	22.12	0.6413	0.2554	<u>0.0995</u>	13.05	3.036	<u>0.4008</u>	0.6433	0.5415
	4	20.37	0.5530	0.3602	0.1612	15.34	5.072	0.4230	<u>0.6068</u>	0.6195
RealSR (x4)	2	<u>22.32</u>	<u>0.6627</u>	<u>0.2402</u>	0.1692	88.92	4.911	0.4473	0.6630	0.4860
	2.5	22.08	0.6545	0.2450	0.1688	87.15	4.677	0.4634	0.6691	0.5106
	3	22.77	0.6924	0.2180	0.1582	81.01	5.278	0.4435	0.6602	0.4755
	3.5	20.43	0.5968	0.3009	0.1964	94.71	5.331	<u>0.5368</u>	0.6976	<u>0.5910</u>
	4	18.75	0.5003	0.4191	0.2469	107.96	8.727	0.5513	<u>0.6777</u>	0.6371
DIV2K-Val (x8)	2	20.27	0.5807	0.3634	0.1420	27.59	3.331	0.3601	<u>0.6118</u>	0.4737
	2.5	<u>20.45</u>	<u>0.5818</u>	0.3771	0.1494	28.57	<u>3.270</u>	0.3582	<u>0.5926</u>	0.4622
	3	20.87	0.6091	<u>0.3726</u>	0.1499	28.75	4.158	0.3237	0.5629	0.4183
	3.5	19.90	0.5671	0.3767	<u>0.1469</u>	<u>28.31</u>	3.199	<u>0.3794</u>	0.6135	<u>0.5006</u>
	4	19.53	0.5334	0.4639	0.1988	30.71	4.553	0.3907	0.5634	0.5510
RealSR (x8)	2	18.78	<u>0.5241</u>	<u>0.3695</u>	<u>0.2350</u>	151.44	4.524	0.4542	0.6673	0.5238
	2.5	18.98	0.5168	<u>0.3835</u>	0.2404	155.66	4.474	0.4540	0.6628	0.5383
	3	19.40	0.5656	0.3465	0.2280	<u>154.32</u>	5.131	0.4342	0.6590	0.4978
	3.5	17.96	0.4877	0.4188	0.2570	162.22	5.178	<u>0.5034</u>	0.6776	<u>0.5749</u>
	4	17.64	0.4316	0.5353	0.2996	175.57	8.143	0.5140	0.6408	0.6261

Impact of the positive–negative CFG strength. Finally, we observe a similar trade-off when studying the impact of the CFG scale used to train FlowMapSR (**without adversarial fine-tuning**), as shown in Table 5. Increasing the maximum guidance strength w_{\max} consistently improves perceptual quality, as indicated by non-reference metrics, with $w_{\max} = 4$ yielding CLIPQA scores that surpass competing methods. This gain, however, comes at the cost of reconstruction faithfulness: reference-based metrics degrade at higher guidance strengths, reflecting weaker alignment with the HR ground truth. Lower CFG scales favor fidelity, though the optimal value is not straightforward: $w_{\max} = 2$ performs best in terms of FID and DISTS, while PSNR and SSIM peak at $w_{\max} = 3$. Overall, we adopt $w_{\max} = 3.5$ as a balanced compromise between perceptual quality and faithfulness. Although adversarial fine-tuning does not change the quantitative metrics for this setting (see Table 2), it leads to clear visual improvements in the generated HR images (see Figure 16 in Section C).

6 CONCLUSION

In this paper, we introduced FlowMapSR, a diffusion-based framework for image super-resolution built on Flow Map generative models and designed for fast inference. Unlike prior diffusion-based SR approaches based on teacher–student distillation, FlowMapSR directly trains a large, expressive model while maintaining efficiency. The framework supports all three Flow Map formulations (*Lagrangian*, *Eulerian*, and *Shortcut*) and incorporates two SR-specific enhancements: a unified CFG-based positive–negative prompting strategy and a lightweight LoRA-based adversarial fine-tuning stage for controllable perceptual refinement at inference.

Empirically, the *Shortcut* formulation best accommodates these enhancements, while the *Lagrangian* and *Eulerian* variants suffer from unstable training. We conjecture that this limitation could be alleviated through finer-grained loss control (Sabour et al., 2025), which we leave for future work. Using the *Shortcut* variant as our default configuration, FlowMapSR achieves a strong balance between reconstruction faithfulness and photorealism, outperforming competing diffusion-based SR methods in qualitative realism. It produces more lifelike textures (especially, for human subjects), improved depth-of-field rendering, and fewer overly sharp or glossy artefacts, all using a single model without scale- or degradation-specific conditioning. While training is more expensive than distillation-based approaches, inference time remains competitive. Finally, our ablations studies show that the fundamental SR trade-off is explicitly controlled by FlowMapSR hyperparameters: increasing the number of inference steps, CFG strength, or LoRA scale improves perceptual quality, while lower values favor fidelity.

A current limitation of our approach is the reliance on Gaussian tiling at very high resolutions, which can introduce mild blurring at the boundaries of generated images; such artefacts could likely be alleviated through more advanced tiling strategies. Additionally, our procedure may exhibit noticeable color shifts, a phenomenon commonly observed in diffusion models (Deck & Bischoff, 2023). In practice, FlowMapSR could be combined with GPU-like post-processing techniques to correct these shifts, as explored for instance in Feydy et al. (2019). We hope our results motivate further exploration of Flow Map models for other image-to-image translation tasks, such as object removal and relighting, depth and normal estimation.

REFERENCES

- Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 126–135, 2017.
- Michael S. Albergo, Nicholas M. Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions, 2023. URL <https://arxiv.org/abs/2303.08797>.
- Nicholas M. Boffi, Michael S. Albergo, and Eric Vanden-Eijnden. How to build a consistency model: Learning flow maps via self-distillation, 2025a. URL <https://arxiv.org/abs/2505.18825>.
- Nicholas Matthew Boffi, Michael Samuel Albergo, and Eric Vanden-Eijnden. Flow map matching with stochastic interpolants: A mathematical framework for consistency models. *Transactions on Machine Learning Research*, 2025b. ISSN 2835-8856. URL <https://openreview.net/forum?id=cqDH0e6ak2>.
- Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 3086–3095, 2019.
- Clément Chadebec, Onur Tasar, Eyal Benaroch, and Benjamin Aubin. Flash diffusion: Accelerating any conditional diffusion model for few steps image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 15686–15695, 2025a.
- Clément Chadebec, Onur Tasar, Sanjeev Sreetharan, and Benjamin Aubin. Lbm: Latent bridge matching for fast image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 29086–29098, October 2025b.
- Katherine Deck and Tobias Bischoff. Easing color shifts in score-based diffusion models. *arXiv preprint arXiv:2306.15832*, 2023.
- Mauricio Delbracio and Peyman Milanfar. Inversion by direct iteration: An alternative to denoising diffusion for image restoration. *arXiv preprint arXiv:2303.11435*, 2023.
- Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2567–2581, 2020.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trounev, and Gabriel Peyré. Interpolating between optimal transport and mmd using sinkhorn divergences. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 2681–2690, 2019.
- Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=0lzB6LnXcS>.
- Sicheng Gao, Xuhui Liu, Bohan Zeng, Sheng Xu, Yanjing Li, Xiaoyan Luo, Jianzhuang Liu, Xiantong Zhen, and Baochang Zhang. Implicit diffusion models for continuous super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10021–10030, 2023.
- Zhengyang Geng, Mingyang Deng, Xingjian Bai, J. Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling, 2025. URL <https://arxiv.org/abs/2505.13447>.
- Xiao He, Huaao Tang, Zhijun Tu, Junchao Zhang, Kun Cheng, Hanting Chen, Yong Guo, Mingrui Zhu, Nannan Wang, Xinbo Gao, et al. One step diffusion-based super-resolution with time-aware distillation. *arXiv preprint arXiv:2408.07476*, 2024.
- Jonathan Heek, Emiel Hoogeboom, and Tim Salimans. Multistep consistency models, 2024. URL <https://arxiv.org/abs/2403.06807>.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022. URL <https://arxiv.org/abs/2207.12598>.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- Zheyuan Hu, Chieh-Hsin Lai, Yuki Mitsufuji, and Stefano Ermon. Cmt: Mid-training for efficient learning of consistency, mean flow, and flow map models. *arXiv preprint arXiv:2509.24526*, 2025.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Yazid Janati, Eric Moulines, Jimmy Olsson, and Alain Oliviero-Durmus. Bridging diffusion posterior sampling and monte carlo methods: a survey. *Philosophical Transactions A*, 383(2299):20240331, 2025.
- A Jolicoeur-Martineau. The relativistic discriminator: A key element missing from standard gan. *arXiv preprint arXiv:1807.00734*, 2018.
- Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 24174–24184, 2024.
- Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5148–5157, 2021.
- Beomsu Kim, Yu-Guan Hsieh, Michal Klein, Marco Cuturi, Jong Chul Ye, Bahjat Kawar, and James Thornton. Simple reflow: Improved techniques for fast flow models. *arXiv preprint arXiv:2410.07815*, 2024a.
- Beomsu Kim, Jaemin Kim, Jeongsol Kim, and Jong Chul Ye. Generalized consistency trajectory models for image manipulation, 2025. URL <https://arxiv.org/abs/2403.12510>.
- Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsufuji, and Stefano Ermon. Consistency trajectory models: Learning probability flow ODE trajectory of diffusion. In *The Twelfth International Conference on Learning Representations*, 2024b. URL <https://openreview.net/forum?id=ymjI8feDTD>.
- Nikita Kornilov, Petr Mokrov, Alexander Gasnikov, and Aleksandr Korotin. Optimal flow matching: Learning straight trajectories in just one step. *Advances in Neural Information Processing Systems*, 37:104180–104204, 2024.
- Sangyun Lee, Zinan Lin, and Giulia Fanti. Improving the training of rectified flows. *Advances in neural information processing systems*, 37:63082–63109, 2024.
- Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing*, 479:47–59, 2022.
- Liangchen Li and Jiajun He. Bidirectional consistency models, 2025. URL <https://arxiv.org/abs/2403.18035>.
- Xinqi Lin, Jingwen He, Ziyang Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *European conference on computer vision*, pp. 430–448. Springer, 2024.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022.
- Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. *Machine Intelligence Research*, pp. 1–22, 2025.

- Zhisheng Lu, Juncheng Li, Hong Liu, Chaoyan Huang, Linlin Zhang, and Tiejong Zeng. Transformer for single image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 457–466, 2022.
- Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration with mean-reverting stochastic differential equations. *arXiv preprint arXiv:2301.11699*, 2023.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=di52zR8xgf>.
- Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky TQ Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Amirmojtaba Sabour, Sanja Fidler, and Karsten Kreis. Align your flow: Scaling continuous-time flow map distillation, 2025. URL <https://arxiv.org/abs/2506.14603>.
- Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4713–4726, 2022.
- Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022.
- Axel Sauer, Frederic Boesel, Tim Dockhorn, Andreas Blattmann, Patrick Esser, and Robin Rombach. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, 2024a.
- Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *European Conference on Computer Vision*, pp. 87–103. Springer, 2024b.
- Dazhong Shen, Guanglu Song, Yi Zhang, Bingqi Ma, Lujundong Li, Dongzhi Jiang, Zhuofan Zong, and Yu Liu. Adt: Tuning diffusion models with adversarial supervision, 2025. URL <https://arxiv.org/abs/2504.11423>.
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *The Ninth International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=PXTIG12RRHS>.
- Ruoyu Sun, Tiantian Fang, and Alexander Schwing. Towards a better global loss landscape of gans. *Advances in Neural Information Processing Systems*, 33:10186–10198, 2020.
- Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Hugué, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.
- Fu-Yun Wang, Ling Yang, Zhaoyang Huang, Mengdi Wang, and Hongsheng Li. Rectified diffusion: Straightness is not your need in rectified flow. *arXiv preprint arXiv:2410.07303*, 2024a.
- Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 2555–2563, 2023.
- Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 132(12):5929–5949, 2024b.

- Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pp. 0–0, 2018.
- Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *International Conference on Computer Vision Workshops (ICCVW)*, 2021.
- Yufei Wang, Wenhan Yang, Xinyuan Chen, Yaohui Wang, Lanqing Guo, Lap-Pui Chau, Ziwei Liu, Yu Qiao, Alex C Kot, and Bihan Wen. Sinsr: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 25796–25805, 2024c.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Pengxu Wei, Ziwei Xie, Hannan Lu, Zongyuan Zhan, Qixiang Ye, Wangmeng Zuo, and Liang Lin. Component divide-and-conquer for real-world image super-resolution. In *European conference on computer vision*, pp. 101–117. Springer, 2020.
- Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems*, 37:92529–92553, 2024a.
- Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 25456–25467, 2024b.
- Xue Wu, Jingwei Xin, Zhijun Tu, Jie Hu, Jie Li, Nannan Wang, and Xinbo Gao. One-step diffusion-based real-world image super-resolution with visual perception distillation, 2025. URL <https://arxiv.org/abs/2506.02605>.
- Rui Xie, Chen Zhao, Kai Zhang, Zhenyu Zhang, Jun Zhou, Jian Yang, and Ying Tai. Addsr: Accelerating diffusion-based blind super-resolution with adversarial diffusion distillation. *arXiv preprint arXiv:2404.01717*, 2024.
- Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1191–1200, 2022.
- Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *European conference on computer vision*, pp. 74–91. Springer, 2024.
- Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6613–6623, 2024.
- Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Resshift: Efficient diffusion model for image super-resolution by residual shifting. *Advances in Neural Information Processing Systems*, 36:13294–13307, 2023.
- Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Efficient diffusion model for image restoration by residual shifting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024.
- Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015.
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.
- Jianbin Zheng, Minghui Hu, Zhongyi Fan, Chaoyue Wang, Changxing Ding, Dacheng Tao, and Tat-Jen Cham. Trajectory consistency distillation. *CoRR*, 2024.
- Kaiwen Zheng, Yuji Wang, Qianli Ma, Huayu Chen, Jintao Zhang, Yogesh Balaji, Jianfei Chen, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Large scale diffusion distillation via score-regularized continuous-time consistency. *arXiv preprint arXiv:2510.08431*, 2025.

ORGANIZATION OF THE SUPPLEMENTARY

The appendix is organized as follows. Section A first presents complementary theoretical material related to the Flow Map models introduced in Section 3. Section B then provides additional implementation details regarding the data, model architecture, training procedure, and inference setup used in our experiments. Finally, Section C reports further qualitative results for $\times 4$ and $\times 8$ super-resolution, including comparisons with competing methods on additional synthetic and real-world samples, as well as extended analyses of the ablation studies discussed in Section 5.

A THEORETICAL DETAILS

A.1 BACKGROUND ON FLOW MATCHING AND FLOW MAP MODELS

In the case where the interpolant I_t admits the general form $I_t = \alpha_t X_0 + \beta_t X_1$, a velocity field candidate for the generative ODE (1) is given by $v_t(x) = \mathbb{E}[v_{\text{FM}}^{\text{target}} | I_t = x]$ where $v_{\text{FM}}^{\text{target}} = \dot{I}_t = \dot{\alpha}_t X_0 + \dot{\beta}_t X_1$ is the conditional target velocity. Following Boffi et al. (2025a), the loss target in the SD optimization problem (9) related to this general SI writes as follows

1. *Lagrangian* SD (LSD) loss

$$u_{\text{LSD}}^{\text{target}} = u_{s,s}^{\theta}(I_t - (t-s)u_{s,t}^{\theta}(I_t)) + (t-s)\partial_s u_{s,t}^{\theta}(I_t), \quad (18)$$

2. *Eulerian* SD (ESD) loss

$$u_{\text{ESD}}^{\text{target}} = u_{t,t}^{\theta}(I_t) - (t-s)(\nabla u_{s,t}^{\theta}(I_t)u_{t,t}^{\theta}(I_t) + \partial_t u_{s,t}^{\theta}(I_t)), \quad (19)$$

3. *Shortcut* SD (SSD) loss

$$u_{\text{SSD}}^{\text{target}} = \frac{t-r}{t-s}u_{r,t}^{\theta}(I_t) + \frac{r-s}{t-s}u_{s,r}^{\theta}(I_t - (t-r)u_{r,t}^{\theta}(I_t)), \quad r \sim q(r|s,t), \quad (20)$$

where $q(r|t,s)$ is a time distribution on $\{r \in [0,1] : s < r < t\}$. These expressions are derived from straightforward computations on the Flow Map (8), summarized below as an informal counterpart to the proofs in Boffi et al. (2025a).

Proof. Lagrangian setting. By computing the derivative of the Flow Map (8) evaluated in the trajectory $(X_t)_{t \in [0,T]}$ with respect to time variable s , we have

$$\frac{dX_{s,t}(X_t)}{ds} = u_{s,t}(X_t) - (t-s)\partial_s u_{s,t}(X_t),$$

By definition of the Flow Map given in (7), we also have that $dX_{s,t}(X_t)/ds = v_s(X_{s,t}(X_t))$. By combining these two results, we obtain the *Lagrangian characterization*

$$\begin{aligned} u_{s,t}(X_t) &= v_s(X_{s,t}(X_t)) + (t-s)\partial_s u_{s,t}(X_t) \\ &= u_{s,s}(X_{s,t}(X_t)) + (t-s)\partial_s u_{s,t}(X_t). \end{aligned}$$

After replacing u by u^{θ} in the previous equality, we obtain the LSD target (18) from the right term.

Eulerian setting. By computing the derivative of the Flow Map (8) evaluated in the trajectory $(X_t)_{t \in [0,T]}$ with respect to time variable t , we have

$$\begin{aligned} \frac{dX_{s,t}(X_t)}{dt} &= \frac{dX_t}{dt} - u_{s,t}(X_t) - (t-s)\frac{du_{s,t}(X_t)}{dt} \\ &= v_t(X_t) - u_{s,t}(X_t) - (t-s)(\nabla u_{s,t}(X_t)v_t(X_t) + \partial_t u_{s,t}(X_t)). \end{aligned}$$

By definition of the Flow Map given in (7), we also have that $dX_{s,t}(X_t)/dt = v_t(X_t) - v_t(X_t) = 0$. By combining these two results, we obtain the *Eulerian characterization*

$$\begin{aligned} u_{s,t}(X_t) &= v_t(X_t) - (t-s)(\nabla u_{s,t}(X_t)v_t(X_t) + \partial_t u_{s,t}(X_t)) \\ &= u_{t,t}(X_t) - (t-s)(\nabla u_{s,t}(X_t)u_{t,t}(X_t) + \partial_t u_{s,t}(X_t)). \end{aligned}$$

After replacing u by u^{θ} in the previous equality, we obtain the ESD target (19) from the right term.

Shortcut setting. By definition of the Flow Map given in (7), we have for any time $r \in (s,t)$

$$X_{s,t}(X_t) = X_{s,r}(X_{r,t}(X_t)).$$

When substituting the parametrization (8) in the previous equality, we obtain the *Shortcut characterization*

$$\begin{aligned} X_t - (t-s)u_{s,t}(X_t) &= \{X_t - (t-r)u_{r,t}(X_t)\} - (r-s)u_{s,r}(X_{r,t}(X_t)) \\ \iff u_{s,t}(X_t) &= \frac{t-r}{t-s}u_{r,t}(X_t) + \frac{r-s}{t-s}u_{s,r}(X_{r,t}(X_t)). \end{aligned}$$

After replacing u by u^{θ} in the previous equality, we obtain the SSD target (20) from the right term. \square

Simplifications with the standard interpolant. Now assume that we have $I_t = (1-t)X_0 + tX_1$. To obtain the objective functions presented in Section 3.2, we proceed to the following simplifications: (a) taking inspiration from the conditional trick in FM, we replace each occurrence of the *marginal* velocity field ($u_{s,s}^\theta$ in LSD loss, $u_{t,t}^\theta$ in ESD loss) by the *conditional* velocity field, exactly defined as $v_{\text{FM}}^{\text{target}} = X_1 - X_0$, which has the computational advantage to be constant over the whole interpolant, (b) following Frans et al. (2025), we deterministically set $q(r|s, t) = \delta_{(s+t)/2}(r)$ in SSD formulation.

A.2 GENERALIZATION OF CFG TO THE FLOW MAP SETTING

To incorporate CFG into Flow Map models, we suggest to replace the target velocity field v_t by its CFG counterpart (5), where the class-unconditional velocity $v_t(\cdot|\emptyset)$ is legitimately substituted by $u_{t,t}(\cdot|\emptyset)$. This naturally leads to the CFG-based loss for the FM setting detailed in (13), which is unchanged in the general SI setting (recalling that we now have $v_{\text{FM}}^{\text{target}} = \dot{I}_t$). With this paradigm, the SD targets defined in Section A.1 naturally turn into the following CFG-based targets

1. CFG-based *Lagrangian* SD (cfg-LSD) loss

$$\begin{aligned} u_{\text{cfg-LSD}}^{\text{target}} &= v_{\text{cfg-FM}}^{\text{target}} + (t-s)\partial_s u_{s,t}^\theta(I_t|\mathbf{c}), \\ v_{\text{cfg-FM}}^{\text{target}} &= w u_{s,s}^\theta(I_t - (t-s)u_{s,t}^\theta(I_t|\mathbf{c}) | \mathbf{c}) + (1-w)u_{s,s}^\theta(I_t - (t-s)u_{s,t}^\theta(I_t|\emptyset) | \emptyset), \end{aligned}$$

2. CFG-based *Eulerian* SD (cfg-ESD) loss

$$\begin{aligned} u_{\text{cfg-ESD}}^{\text{target}} &= v_{\text{cfg-FM}}^{\text{target}} - (t-s)\left(\nabla u_{s,t}^\theta(I_t|\mathbf{c})v_{\text{cfg-FM}}^{\text{target}} + \partial_t u_{s,t}^\theta(I_t|\mathbf{c})\right), \\ v_{\text{cfg-FM}}^{\text{target}} &= w u_{t,t}^\theta(I_t|\mathbf{c}) + (1-w)u_{t,t}^\theta(I_t|\emptyset), \end{aligned}$$

3. CFG-based *Shortcut* SD (cfg-SSD) loss

$$u_{\text{cfg-SSD}}^{\text{target}} = \frac{t-r}{t-s}u_{r,t}^\theta(I_t|\mathbf{c}) + \frac{r-s}{t-s}u_{s,r}^\theta(I_t - (t-r)u_{r,t}^\theta(I_t|\mathbf{c}) | \mathbf{c}), \quad r \sim q(r|t, s).$$

Simplifications with the standard interpolant. Now assume that we have $I_t = (1-t)X_0 + tX_1$. To obtain the CFG-based SD objectives presented in Section 3.3, we apply the same simplifications as in the unconditional setting: in particular, we replace each occurrence of the *class-conditional marginal* velocity field in LSD and ESD settings by its conditional counterpart $v_{\text{FM}}^{\text{target}} = X_1 - X_0$, thus avoiding extra model evaluations, and set $q(r|s, t) = \delta_{(s+t)/2}(r)$ in SSD setting.

B IMPLEMENTATION DETAILS

B.1 GENERAL SETTINGS

Data. We assemble a collection of publicly available, free-to-use images to construct the LR–HR training dataset \mathcal{D} . To generate synthetic LR images from HR ground truth, we adapt the widely used degradation pipeline of Wang et al. (2021), applying the following operations in sequence:

1. *Moderate blurring*: applied with 80% probability;
2. *Downscaling* (core operation): applied systematically using a random scale $s_{\text{down}} \sim \text{U}(0.1, 1)$ and a randomly chosen interpolation mode (nearest, area, bilinear, or bicubic);
3. *Addition of low-level noise*: Gaussian noise with 50% probability, otherwise Poisson noise;
4. *JPEG compression*: applied systematically with the default compression level from Wang et al. (2021).

To match the spatial resolution of the HR image and define a valid interpolant for the Flow Map formulation, the degraded output is resized back using a random interpolation mode, followed by additional blurring and pixel-value clamping. Both HR and LR images are then rescaled to the range $[-1, 1]$. Since we operate in latent space to leverage large pre-trained diffusion models, both HR and synthetically generated LR images are embedded using the pre-trained variational encoder from (Podell et al., 2024).

To support training across multiple resolutions, we adopt a bucketing strategy (Podell et al., 2024) that accommodates varying aspect ratios and image sizes. The specific configuration depends on the scalability of each Flow Map variant:

- *SSD setting*: pixel budgets of $[512^2, 1024^2]$ sampled with probabilities $[0.2, 0.8]$ and batch sizes $[4, 2]$, respectively, to fit within a single H100 GPU; aspect ratios are sampled in $[1/4, 4]$.

- *LSD and ESD settings*: pixel budgets of $[256^2, 512^2]$ with the same sampling probabilities and batch sizes as in the SDD setting. Higher resolutions could not be accommodated due to memory constraints during training, which could not be alleviated by multi-GPU sharding owing to software limitations imposed by the use of the jvp package. Aspect ratios are again sampled in $[1/4, 4]$.

Architectures. In the FlowMapSR framework described in Section 4, we employ three neural network components:

(A) *Flow Map model* $u_{s,t}^\theta$. This component is implemented as a UNet with an architecture inspired by the SDXL text-to-image diffusion model (Podell et al., 2024). Conditioning is handled via projection-based embeddings. Specifically, to evaluate the model at a timestep pair (s, t) , we apply positional encodings to each input, using 256 channels per timestep by default. In addition, conditioning on positive or negative text prompts is achieved through a frozen CLIP text encoder, producing embeddings of dimension 1280. Overall, the resulting model contains approximately 2.5B parameters.

(B) *Dynamic weighting function* $\lambda_{s,t}^\psi$. This component is implemented as a lightweight neural network. It consists of positional embeddings of s and t (256 channels each), whose outputs are concatenated into a 512-dimensional vector and passed through a two-layer MLP ($512 \rightarrow 1$ with SiLU activation, followed by $1 \rightarrow 1$). The final layer is initialized to zero using Kaiming uniform initialization, ensuring that the initial weighting is constant and equal to 1 across all timestep pairs. As a result, this network is extremely compact, with fewer than 1k parameters.

(C) *Discriminator* \mathbf{D}^ϕ . It follows a patch-based architecture in the style of Pix2Pix (Isola et al., 2017). It comprises four Conv2d layers with 64, 128, 256, and 512 channels, respectively. Group normalization is applied after each layer except the first, followed by a SiLU activation. A final Conv2d layer with a single output channel produces a patch-wise output of spatial size $(w/8 - 2, h/8 - 2)$ for an input of size (w, h) . The discriminator contains approximately 2.8M parameters.

B.2 TRAINING PROCEDURE

Initialization of unconditional FlowMapSR. For unconditional models, we initialize the network using the Latent Bridge Matching (LBM) image deblurring checkpoint from Chadebec et al. (2025b), which provides a strong initialization for image restoration. Although the LBM interpolant is defined as a Brownian bridge—i.e., a *stochastic* variant of the standard FM interpolant with added Gaussian noise scaled by $\sigma_t = \sigma\sqrt{t(1-t)}$ —this initialization remains appropriate in our setting, since the noise level is negligible ($\sigma = 5 \times 10^{-3}$). As the LBM model is unconditional, we load only the shared parameters, leaving class-embedding weights randomly initialized.

Initialization of CFG-based FlowMapSR. CFG-based models are initialized from their corresponding unconditional FlowMapSR counterparts, previously trained without guidance. As before, only shared parameters are loaded, while class-embedding weights are initialized randomly.

Time-dependent regularization with perceptual loss. We set the time-dependent perceptual weight as $\lambda_s^{\text{LPIPS}} = 5 \times \exp(-4s)$, which we found to balance effective pixel-level regularization with minimal bias toward the FM-SD objective. To reduce memory usage, we follow the random cropping strategy of Chadebec et al. (2025b), computing the LPIPS loss on randomly sampled patches from X_0 and \hat{X}_0 . In all experiments, the maximum patch size is fixed to 512^2 .

Discrete timestep sampling. Let $\{t_k\}_{k=0}^K$ denote a discretization of $[0, 1]$ into K uniform sub-intervals. We set $d_{\max} = 7$, corresponding to a finest discretization of $2^7 + 1 = 129$ timesteps. The joint sampling distribution $\bar{q}(s, t)$ is defined as follows, depending on the loss type:

- *FM loss.* With 75% probability, we set $d = d_{\max}$, sample k uniformly from $\{0, \dots, K_d\}$, and take $s = t = t_k$.
- *SD loss (LSD/ESD).* With 25% probability, we sample d uniformly from $\{0, \dots, d_{\max}\}$, then sample k uniformly from $\{1, \dots, K_d\}$ and k' uniformly from $\{0, \dots, k\}$, setting $(s, t) = (t_{k'}, t_k)$.
- *SD loss (SSD).* With 25% probability, we sample d uniformly from $\{0, \dots, d_{\max} - 1\}$, sample k uniformly from $\{1, \dots, K_d\}$, and set $(s, t) = (t_{k-1}, t_k)$. While more restrictive than the LSD/ESD setting, this choice simplifies the definition of the midpoint $r = (s + t)/2$ and matches the protocol of Frans et al. (2025).

Positive–negative prompting guidance & LoRA-based adversarial fine-tuning. Following standard CFG practice, we sample the positive prompt \mathbf{c}^{pos} with 90% probability and the negative prompt \mathbf{c}^{neg} with 10% probability. We set $\lambda^{\text{adv}} = 0.1$ and fix the LoRA rank to 64 across all experiments.

Optimization details. Training instances of FlowMapSR begins with a pure FM warm-start phase of 5k iterations, followed by 5k iterations of unconditional self-distillation and an additional 3k iterations with CFG. Adversarial fine-tuning then proceeds for 2.5k to 4k iterations, depending on RPGAN convergence. The effective batch size is fixed to 256 using gradient accumulation. We use AdamW with default hyperparameters, setting the learning rate to 4×10^{-5} for the Flow Map model and 10^{-4} for the discriminator. Since the initial LBM checkpoint is already well aligned with our framework, we do not employ EMA, which we found to degrade performance.

B.3 UNSUCCESSFUL ENHANCEMENT STRATEGIES

In preliminary experiments, we explored several alternative strategies to further enhance SR performance. However, none of the approaches described below led to measurable improvements over our final configuration.

- (a) *Conditioning on the upscaling factor s_{up} .* We investigated explicitly conditioning the Flow Map model on the target upscaling factor using positional encodings with appropriate scaling, but it did not yield any noticeable performance gains.
- (b) *Fourier coefficient matching loss.* Inspired by (Wu et al., 2025), we evaluated the use of Fourier coefficient matching as an additional pixel-level regularization term, using the same coefficient settings as for the LPIPS loss. In our experiments, this approach did not improve output sharpness or visual quality.
- (c) *Pixel-space RPGAN fine-tuning.* To exhaustively explore adversarial training options, we also experimented with aligning pixel-space distributions after CFG training. To mitigate memory constraints, we adopted the same random cropping strategy as for LPIPS regularization, evaluating the discriminator on randomly sampled patches from \hat{X}_0^{adv} and X_0 , with a maximum patch size of 512^2 . This approach consistently produced inferior results compared to our default latent-space adversarial fine-tuning. Representative qualitative examples are shown in Figure 16.

C ADDITIONAL RESULTS

Additional comparison with competing methods. In this section, we complement the main results of Section 5. To demonstrate the consistency of FlowMapSR across upscaling factors, we report quantitative results on DIV2K-Val for $\times 2$ SR in Table 6, with qualitative examples in Figure 9. These results align with those observed for $\times 4$ and $\times 8$ upscaling in Section 5. We further present qualitative comparisons for $\times 4$ upscaling on DIV2K-Val (Figure 10) and RealSR (Figure 11), as well as for $\times 8$ upscaling on DIV2K-Val (Figure 12) and RealSR (Figure 13). Real-world LR inputs from the RealSet65 dataset are also included in Figure 14 for both $\times 4$ and $\times 8$ SR. Across all settings, FlowMapSR consistently recovers fine-grained structures aligned with HR references, produces lifelike textures, particularly for human and animal details (e.g., rows 5–8 in Figure 10, rows 3–4 in Figure 14, and rows 2, 3, and 7 in Figure 12), and avoids visually disruptive artifacts. Moreover, FlowMapSR more faithfully preserves realistic depth of field, where competing methods often fail (e.g., row 7 in Figure 10, rows 2 and 7 in Figure 11, and row 8 in Figure 12).

Ablation studies. We further provide complementary qualitative analyses to illustrate the effect of key design choices in FlowMapSR. In Figure 7, we compare *deterministic* CFG training with the *stochastic* CFG strategy proposed in this work, observing reduced sharpness and increased artifacts under deterministic guidance. In Figure 15, we report additional examples for the *Lagrangian* and *Eulerian* FlowMapSR variants, with and without CFG: the *Lagrangian* formulation deteriorates rapidly as CFG strength increases, while the *Eulerian* variant remains more stable but struggles to produce visually compelling results. In Figure 16, we analyze the impact of adversarial fine-tuning in the *Shortcut* formulation, comparing latent-space and pixel-space RPGAN approaches under both unconditional and CFG-based training. Latent-space adversarial fine-tuning consistently yields superior visual quality, and its benefits are significantly amplified when combined with CFG. These observations motivate our final design choice, where adversarial fine-tuning is applied jointly with CFG rather than in isolation. Finally, we provide additional qualitative examples for the ablation studies on the number of inference steps (Figure 17) and the LoRA scale (Figure 18), which further confirm the trends discussed in Section 5.




Figure 7: **Qualitative comparison of CFG-enhanced Shortcut variants of FlowMapSR for $\times 4$ upscaling (deterministic vs stochastic).** We compare several deterministic choices of the guidance scale w with the default stochastic CFG configuration ($w_{max} = 3.5$). Adversarial fine-tuning is not applied in these comparisons. The LR images are taken from the Div2K-Val dataset.




Figure 8: **Examples of tiling failure in FlowMapSR for $\times 4$ upscaling.** The LR inputs are drawn from the DIV2K-Val dataset. For HR images, we display the same zoomed region (blue for FlowMapSR, green for ground truth).

Visual limitation of FlowMapSR. We observe in some cases that the Gaussian tiling strategy used in FlowMapSR to generate HR images with resolution exceeding 1024×1024 introduces visible artefacts, manifested as blurry rectangular regions near image boundaries, see Figure 8. We conjecture that this limitation may contribute to the inferior performance observed on non-reference evaluation metrics in Table 2, as such artefacts degrade image sharpness and perceptual realism. Addressing and improving the tiling procedure is left for future work.

Table 6: **Quantitative comparison between FlowMapSR and competing SR methods ($\times 2$ upscaling).** The best and second best results are highlighted in **bold** and underlined, respectively.

Dataset (s_{up})	Method	Reference metrics					Non-reference metrics			
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	DISTS \downarrow	FID \downarrow	NIQE \downarrow	MANIQA \uparrow	MUSIQ \uparrow	CLIPQA \uparrow
DIV2K-Val ($\times 2$)	Real-ESRGAN	25.16	0.7829	0.2123	0.1050	14.75	3.855	0.4122	0.6450	0.5515
	OSEDiff-1	22.72	0.6931	<u>0.2363</u>	0.1063	17.78	<u>3.386</u>	0.4533	0.6894	0.6183
	S3Diff-1	22.39	0.6748	0.2316	<u>0.1034</u>	<u>10.97</u>	3.715	0.3700	0.6175	0.5045
	FlowMapSR-2	<u>23.82</u>	<u>0.6989</u>	0.1999	0.0770	7.16	2.904	0.3892	0.6390	0.5356




Figure 9: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 2$ upscaling on DIV2K-Val.** Visually, FlowMapSR more faithfully recovers fine details, textures and depth of field while maintaining high overall image quality with no noticeable artifacts.

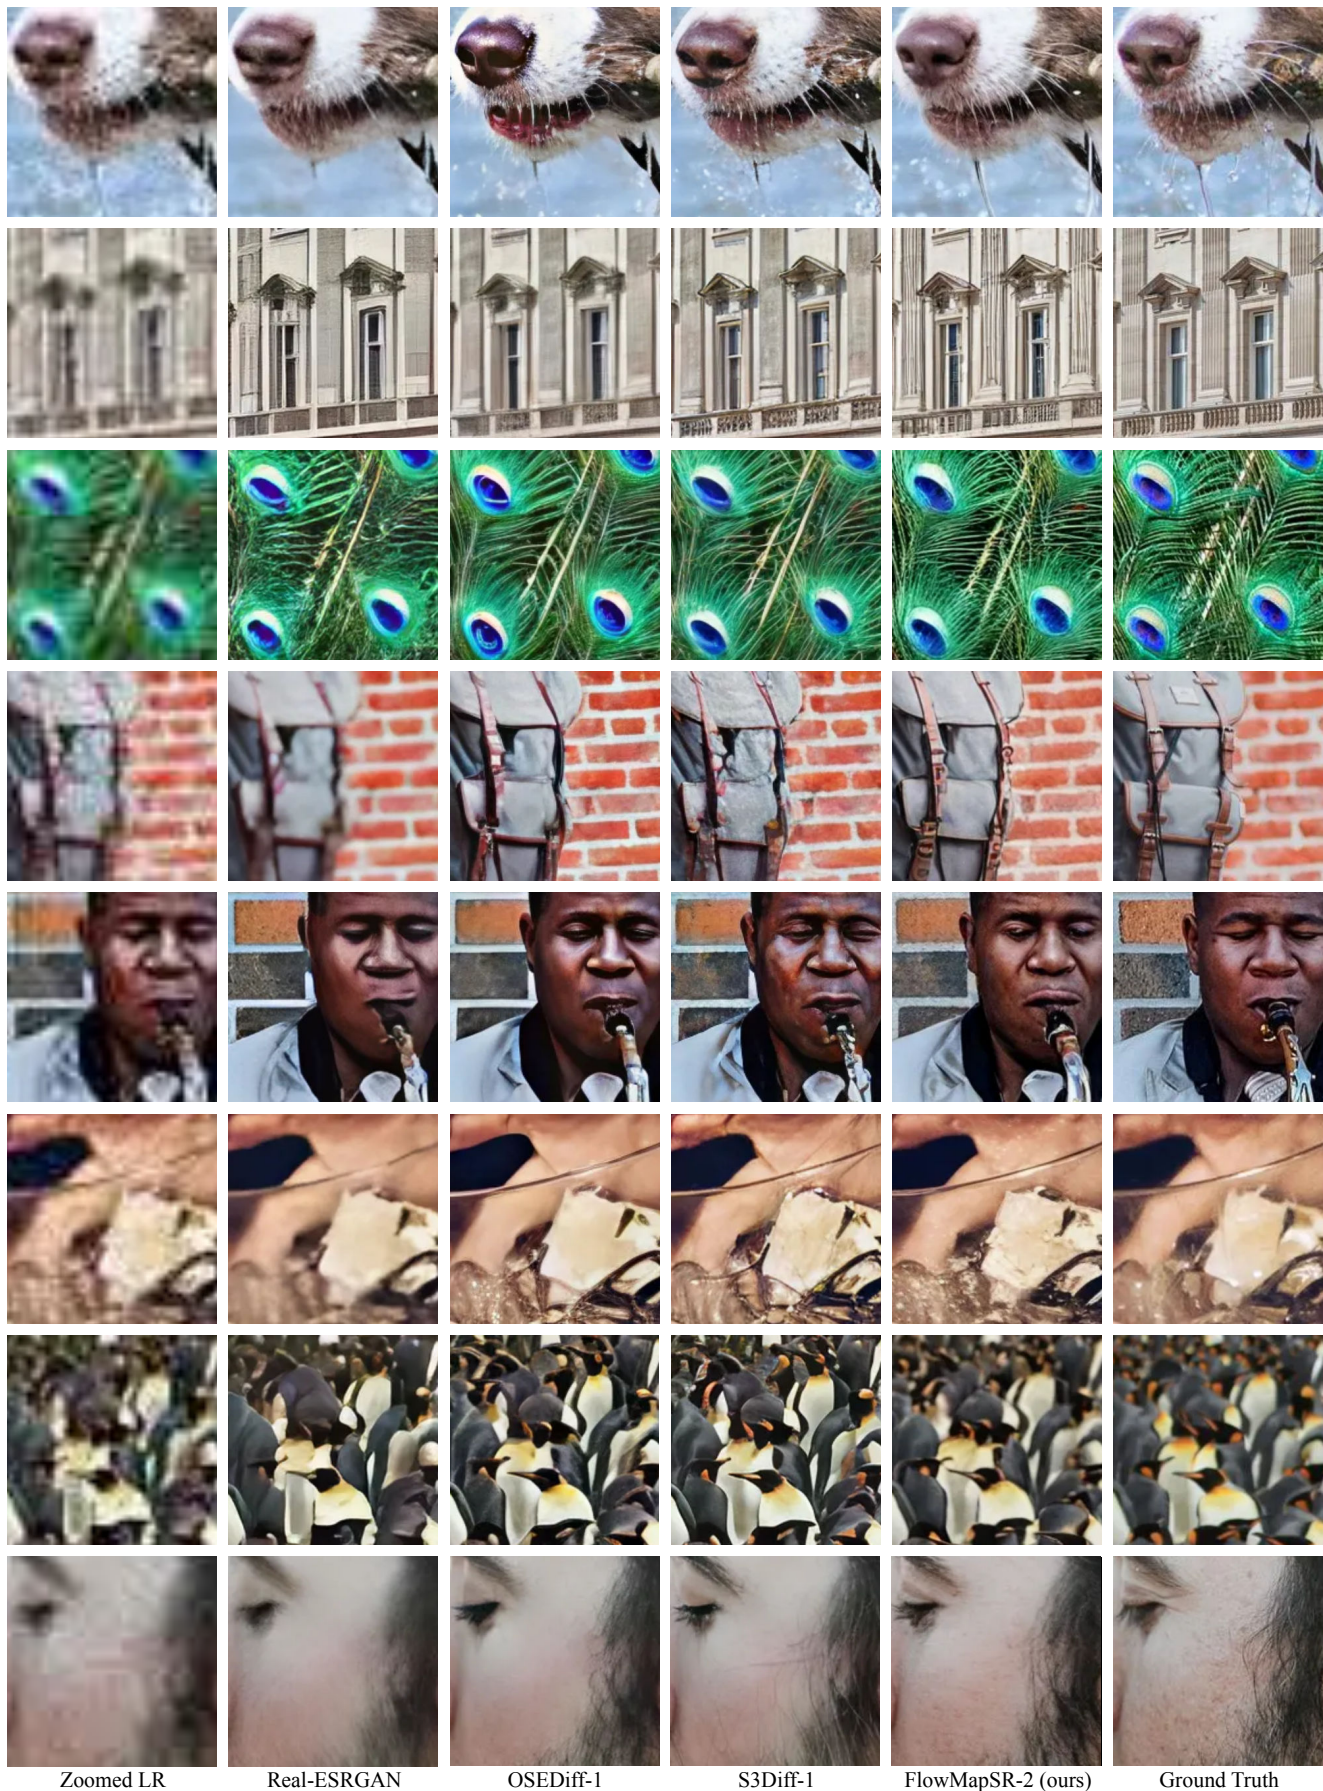


Figure 10: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 4$ upscaling on DIV2K-Val.** This is complementary to Figure 3.



Figure 11: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 4$ upscaling on RealSR.** This is complementary to Figure 3.

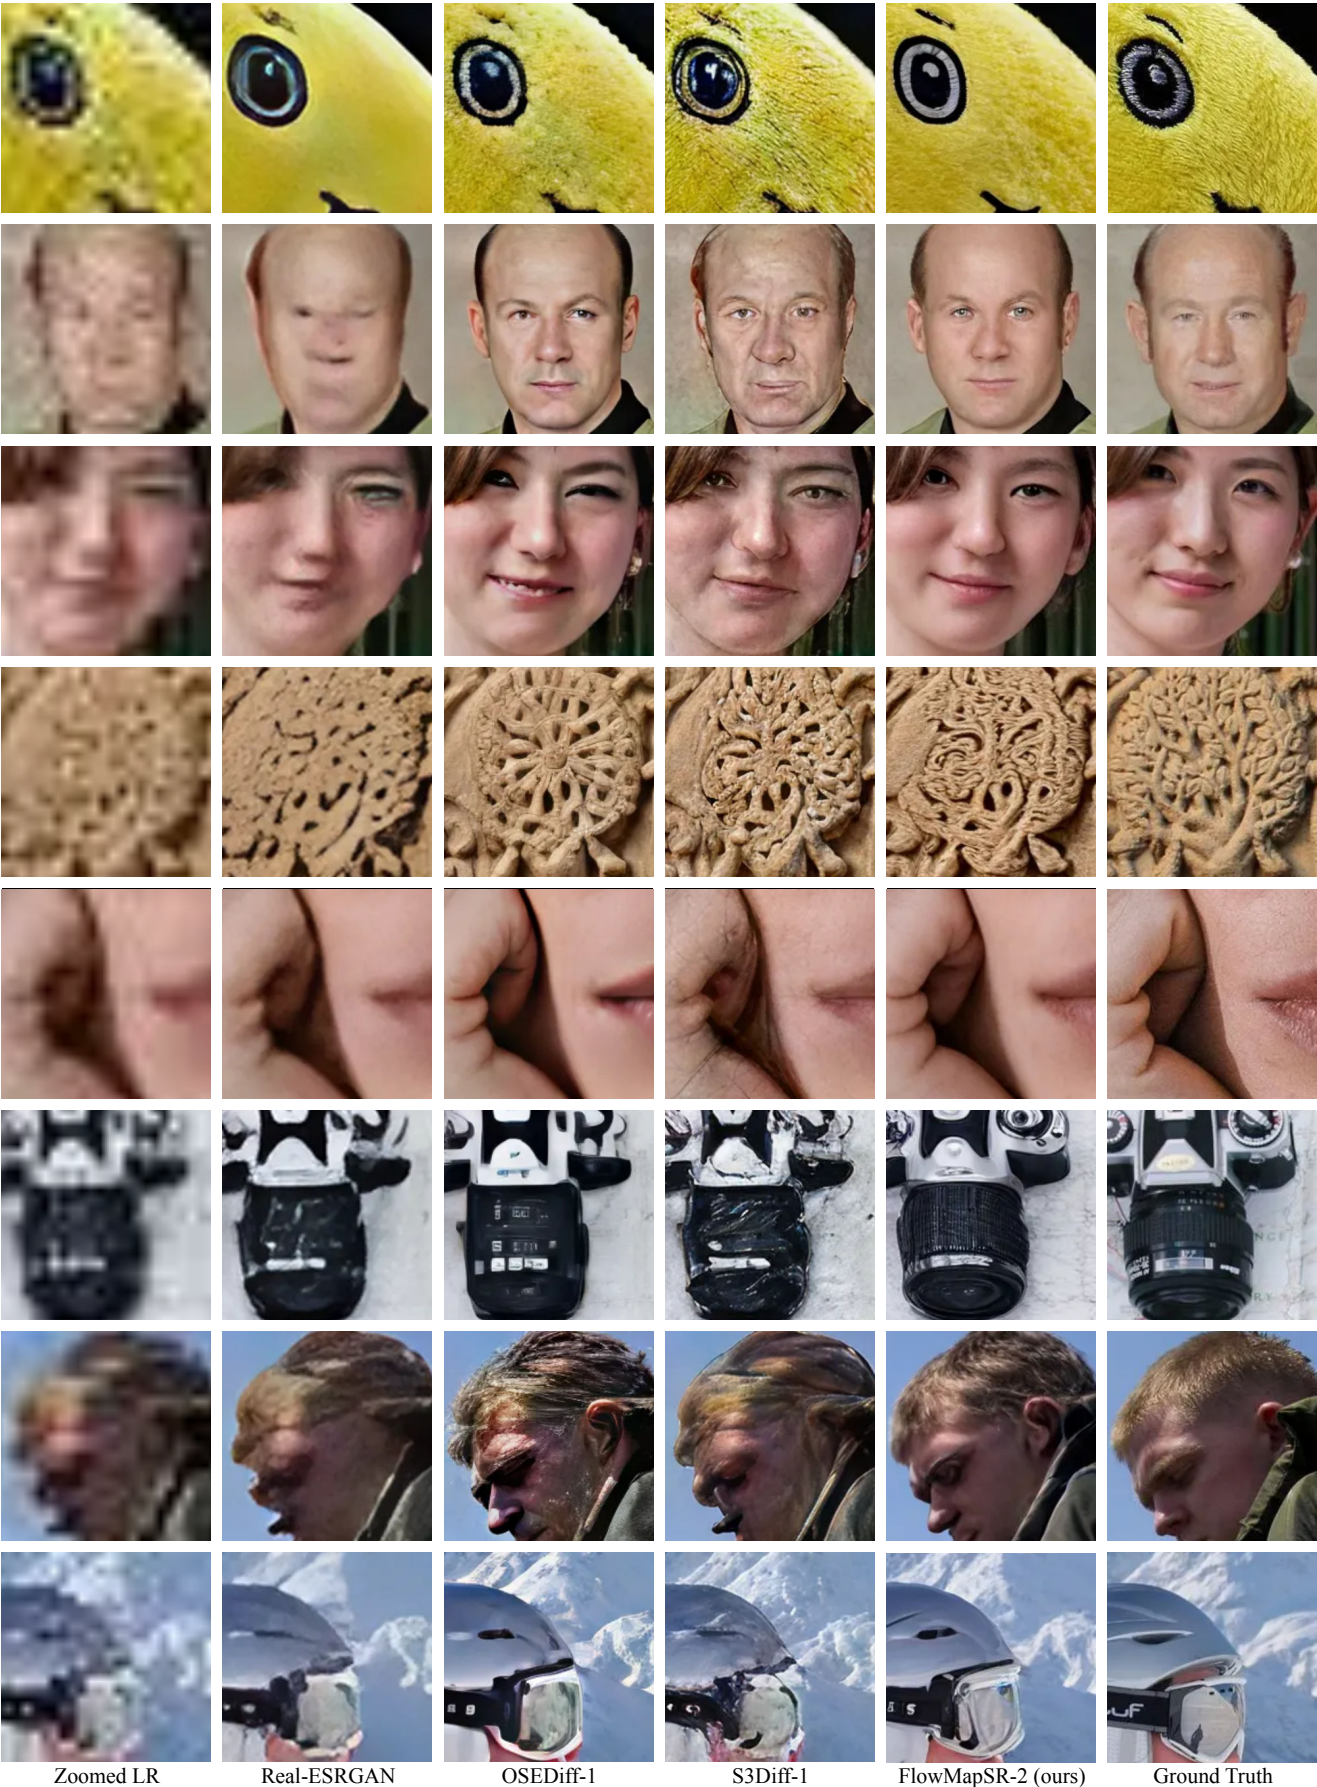


Figure 12: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 8$ upscaling on DIV2K-Val.** This is complementary to Figure 3.




Figure 13: **Qualitative comparison between FlowMapSR and competing SR methods for $\times 8$ upscaling on RealSR.** This is complementary to Figure 3.




Figure 14: **Qualitative comparison between FlowMapSR and competing diffusion-based SR methods on real-world LR inputs for $\times 4$ and $\times 8$ upscaling.** The LR images are taken from the RealSet65 dataset. This is complementary to Figure 4.




Figure 15: **Qualitative comparison of positive-negative CFG enhancement applied to LSD- and ESD-based FlowMapSR for $\times 4$ and $\times 8$ upscaling.** The LR images are taken from the Div2K-Val dataset. This is complementary to Figure 2.



Figure 16: **Qualitative comparison of adversarially finetuned Shortcut variants of FlowMapSR for $\times 4$ upscaling.** The green setting ($w_{\max} = 3.5$ & latent RPGAN) is the default configuration used in our main experiments. The LR images are taken from the Div2K-Val dataset. For each LR–HR pair, increasing the CFG scale alone (first row) already leads to noticeably sharper results, reflecting improved perceptual quality. When adversarial fine-tuning is additionally applied, texture fidelity and sharpness further improve. This effect is especially pronounced when fine-tuning is performed in latent space (third row), whereas pixel-space fine-tuning yields only marginal gains (second row). Overall, combining CFG with latent adversarial fine-tuning produces superior visual results compared to using solely CFG or adversarial fine-tuning on the standard FlowMapSR model.




Figure 17: Qualitative comparison of FlowMapSR with varying numbers of inference steps (NFE) for $\times 4$ (rows 1–3) and $\times 8$ (rows 4–6) upscaling. The LR inputs are drawn from the DIV2K-Val dataset. This is complementary to Figure 5.

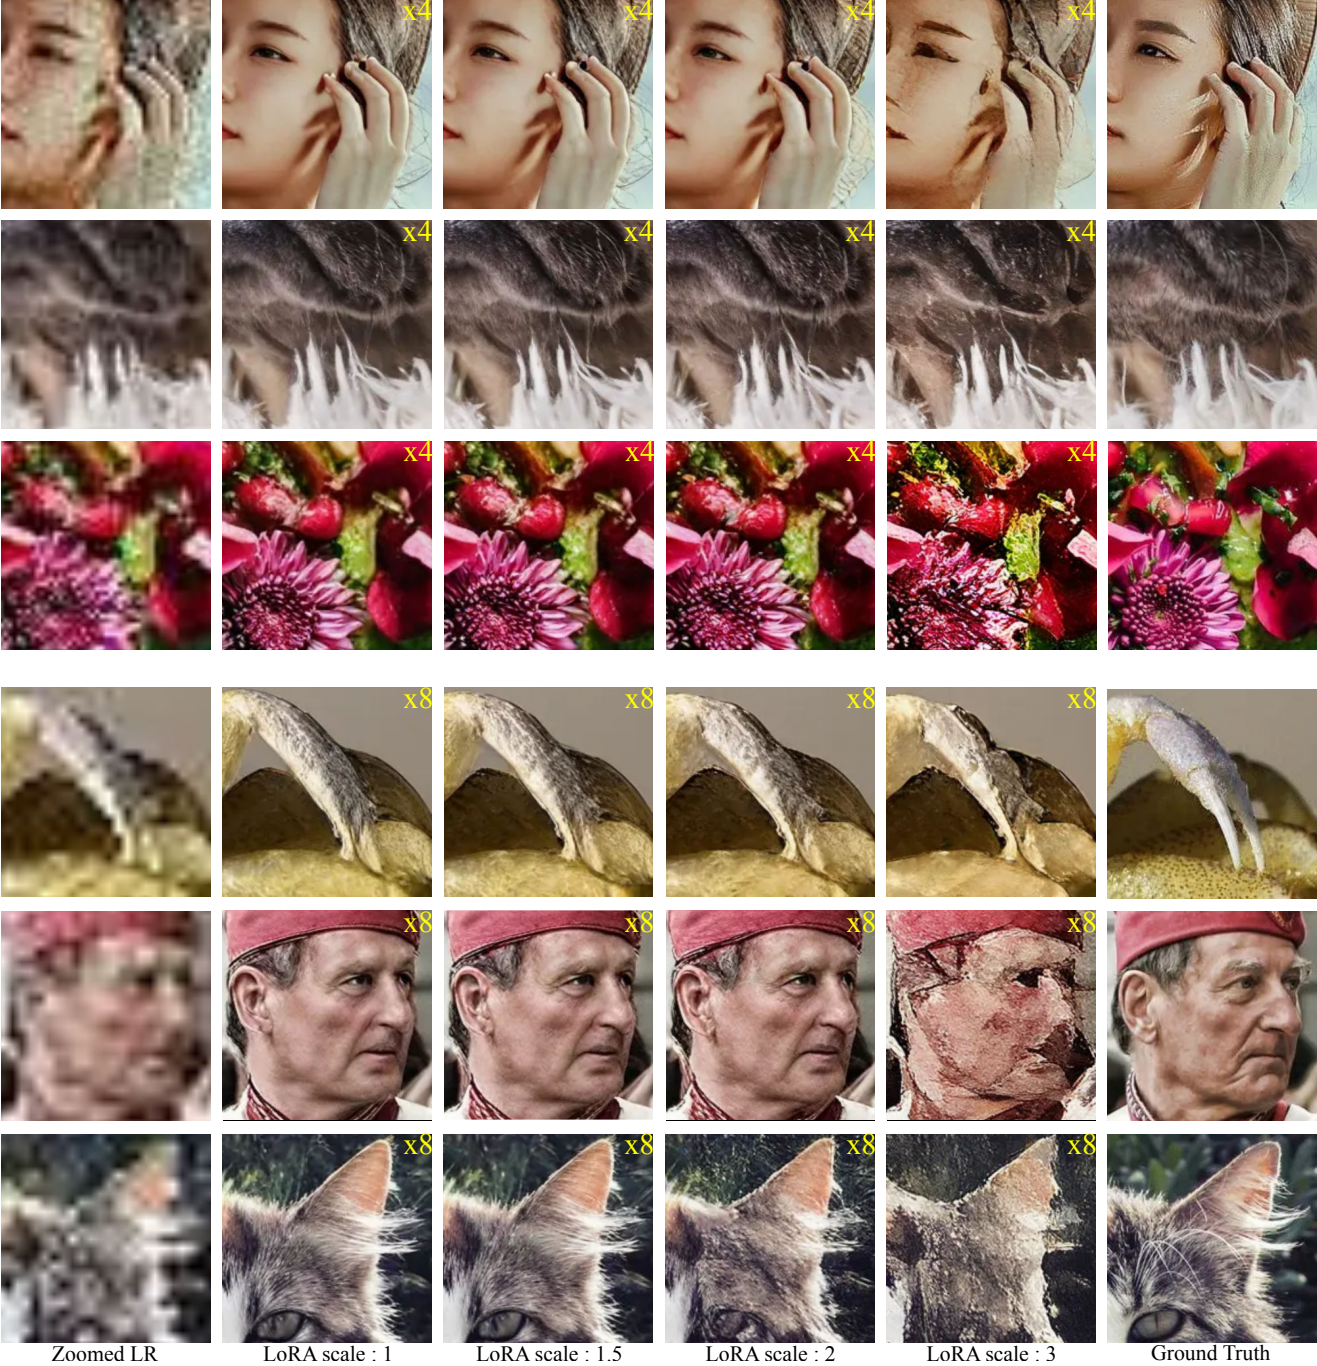


Figure 18: **Qualitative comparison of FlowMapSR with varying LoRA scale for $\times 4$ (rows 1–3) and $\times 8$ (rows 4–6) upscaling.** The LR inputs are drawn from the DIV2K-Val dataset. This is complementary to Figure 6.