

Cross-Space Distillation: Teaching One-Step Students with Modern Diffusion Teachers

Anh Nguyen^{1*}, Ngan Nguyen^{1*}, Duc Vu^{1*}
 Trung Dao², Viet Nguyen³, Quan Dao⁴
 Kien Nguyen¹, Chi Tran¹, Phong Nguyen¹, Khoi Nguyen¹, Cuong Pham¹
 Dimitris Metaxas⁴, Vishal M. Patel³, and Anh Tran¹

¹Qualcomm AI Research[†]
³Johns Hopkins University

²University of Wisconsin–Madison
⁴Rutgers University




Fig. 1: When Latents Don’t Match. (A) Existing distribution-based distillation methods rely on a **Shared-Space constraint**, assuming Teacher and Student share the same latent resolution and VAE. This prevents transferring knowledge from high-resolution teachers (*e.g.*, 1024^2) to compact students (*e.g.*, 512^2), as their latent tensors are inherently incompatible. (B) We formalize this setting as **Cross-Space Distillation** and introduce **Bridge** ($\mathcal{B}_\phi$), a lightweight module that maps student latents z_S to teacher-compatible latents $\hat{z}_T = \mathcal{B}_\phi(z_S)$, enabling standard one-step distillation under Cross-Resolution and Cross-VAE mismatch without modifying the student backbone.

Abstract. Modern one-step diffusion models achieve impressive quality through distribution-based timestep distillation. Yet, they rely on a critical assumption: Teacher and Student must inhabit the same latent space. This Shared-Space constraint prevents knowledge transfer from modern high-capacity Teachers (*e.g.*, SD 3.5 and Flux) into compact, deployment-friendly Students such as SD 1.5, whose latent resolution and VAE parameterization differ from the Teacher. We formalize this overlooked regime as **Cross-Space Distillation**, where Teacher and Student differ in both latent resolution and VAE space. To enable distillation under this mismatch, we introduce the **Bridge** ($\mathcal{B}_\phi$), a lightweight

* Equal contribution.

[†] Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.

latent interface that maps Student latents into the Teacher space without modifying the Student backbone. Bridge combines a frozen Student VAE decoder as a spatial prior with a compact learnable projector, and is trained with latent reconstruction and attention fidelity objectives for stable Teacher-space alignment. Across diverse modern Teachers, Bridge enables substantial gains for compact one-step Students; for example, it improves SD 1.5 from **5.4** to **9.4** HPSv3 while preserving one-step inference, low latency, and broad ecosystem compatibility. These results show that heterogeneous large Teachers can be distilled into efficient, deployable backbones through a lightweight latent-space interface.

Keywords: Diffusion Distillation · One-Step Generation

1 Introduction

State-of-the-art text-to-image diffusion models, such as Stable Diffusion 3.5 [6] and Flux [12, 13], achieve remarkable visual fidelity at high resolutions. However, their large backbones and multi-step sampling impose substantial computational and memory costs, limiting deployment on consumer hardware.

To reduce sampling overhead, recent works distill multi-step teachers into single-step generators via latent consistency distillation [15, 43, 51], variational score distillation [5, 26, 27], or adversarial distillation [32, 33]. Among these approaches, the latter two, which are both grounded in *distribution-based distillation*, have become central to modern one and few-step high-fidelity models. However, they generally preserve the teacher’s latent structure and assume a shared latent space, causing the distilled models to inherit large parameter counts or high latent resolutions of the teacher, which limits on-device practicality.

Another line of work compresses large generators through network pruning, structural compression, or lightweight architectures [1, 2, 9, 16, 50, 53]. These methods face three major challenges: pruning often hits a hard sparsity threshold beyond which quality drops, designing compact architectures requires extensive retraining and manual tuning specific to a teacher.

Rather than building new backbones from scratch, we propose reusing compact, widely supported models like Stable Diffusion 1.5 and transferring the advanced generative quality of modern teachers via distillation. This approach is compact, easily integrates with existing ecosystems, is model-agnostic, and benefits from pretrained weights to stabilize and accelerate optimization.

Transferring knowledge to such students introduces four key challenges: **Cross-Resolution** (from a teacher at 1024×1024 to a student at 512×512), **Cross-VAE** (different latent spaces), **Cross-Architecture** (transformers vs. UNets), and **Cross-Mechanism** (flow matching vs. standard diffusion). Empirically, the architecture and mechanism gaps are comparatively minor; the main barriers are resolution and latent-space mismatch, which prevent standard distillation objectives from being applied directly.

We address these two constraints via **Cross-Space Distillation**, introducing **Bridge**, a lightweight module that maps the student latent into the teacher

latent space. This enables standard distillation objectives without modifying the student architecture. Bridge uses a compact single-stage design with the student’s VAE frozen for spatial alignment, incorporating an Attention Fidelity loss to preserve fine structure and high-quality reconstruction.

Cross-Space Distillation. We study distillation where the Teacher and Student do *not* share the same latent space:

1. **Cross-Resolution:** The latent grids have different spatial shapes, *e.g.*, $z_S \in \mathbb{R}^{h \times w \times C_S}$ and $z_T \in \mathbb{R}^{H \times W \times C_T}$ with $(h, w) \neq (H, W)$.
2. **Cross-VAE:** The autoencoders differ, so latent spaces are not aligned between $\mathcal{Z}_S$ and $\mathcal{Z}_T$.

Consequence: Standard distillation objectives are not directly applicable between z_S and z_T without an explicit alignment mapping.

This design preserves backbone compatibility while unlocking high-fidelity outputs. Experiments show that Bridge substantially improves SD 1.5 through Cross-Resolution and Cross-VAE distillation, effectively transferring the teacher’s visual priors. Beyond distillation, Bridge is modular, enabling seamless high-resolution synthesis upgrades for existing architectures.

In short, our contributions are as follows:

1. **Cross-Space Distillation.** We formalize a new distillation setting where Teacher and Student differ simultaneously in latent resolution and VAE space, enabling knowledge transfer across mismatched latent spaces.
2. **Bridge.** We introduce the Bridge ($\mathcal{B}_\phi$), a lightweight module that maps Student latents into the Teacher space. It combines (i) an architectural prior that leverages the Student’s spatial decoder as a scaffold for a compact Teacher-space projector, and (ii) an Attention Fidelity objective preserving structural and fine-grained detail via a reverse-KL on attention distributions.
3. **Practical Effectiveness.** Experiments across modern teacher–compact student pairs show substantial improvements in generation quality, pushing compact students close to teacher-level performance while keeping the student backbone unchanged and significantly lighter than the teacher.

2 Preliminary

2.1 Diffusion and Flow

Diffusion models [8, 28, 38] and flow matching models [4, 19, 20, 40] provide closely related continuous-time formulations for generative modeling. Both define a forward corruption process that progressively transforms data into Gaussian-like noise and learn a reverse-time generative procedure that maps noise back to the data distribution. During inference, samples are generated by numerically solving the learned reverse dynamics with a discretized solver, typically requiring multiple network evaluations [11, 21, 37].

Diffusion. Diffusion models define a fixed forward Gaussian noising process that gradually perturbs clean samples $x_0 \sim p_{\text{data}}$ over T time steps, characterized by

$$x_t = \alpha_t x_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (1)$$

where the coefficients α_t and σ_t are predefined noise schedules [8]. A neural network, parametrized by θ , is trained to approximate the reverse-time dynamics by gradually removing noise from x_t . Common parameterizations of the reverse process include ϵ -prediction, x_0 -prediction, and v -prediction [31].

Flow matching trains a continuous normalizing flow (CNF) by regressing the conditional velocity field of a chosen probability path [19]. A common Gaussian probability path uses linear interpolation between data and Gaussian noise,

$$x_t = (1 - t)x_0 + t\epsilon, \quad t \in (0, 1). \quad (2)$$

Instead of learning a reverse-time stochastic process, flow matching directly trains a neural network to predict the conditional velocity field $v_\theta(x_t, t)$ that transports samples along this path [19]. Related formulations include rectified flow [20] and conditional/OT flow matching objectives [40].

Connection. Diffusion and flow matching can be interpreted within a unified continuous-time view where a Gaussian probability path transports data samples to a Gaussian distribution [19, 38]. At endpoints, different parameterizations are related under closed-form transformations at Gaussian endpoints, allowing flow matching velocities to be expressed in terms of diffusion-model outputs [31, 36]. In our work, we convert outputs to a consistent x_0 -estimate in a shared latent space, enabling Teacher supervision after applying our alignment module.

2.2 One-Step Distribution-Based Distillation

To reduce inference latency to a single network function evaluation, one-step distillation trains a fast generator $\mathcal{S}_\theta$ to match the sample distribution of a multi-step Teacher [45]. In this work, we focus on *distribution-based* objectives that align the Student distribution with a Teacher-implied target distribution, rather than *trajectory-based* distillation methods that supervise intermediate sampling states or enforce step-to-step consistency, such as consistency distillation [14, 22].

Variational score distillation (VSD) is a distribution-based recipe that updates the Student by regressing the difference between Student and Teacher scores evaluated on the same noised sample [26, 41, 45–47]. Let $x = \mathcal{S}_\theta(\xi, c)$ denote a sample generated by the student model $\mathcal{S}_\theta$ conditioned on c from Gaussian noise ξ . We denote x_t as noised sample of x at time t obtained by the forward diffusion process with $\epsilon \sim \mathcal{N}(0, I)$. The generator update takes the form

$$\nabla_\theta \mathbb{E}_{t, \xi, \epsilon} [(s_\psi(x_t, t, c) - s_\eta(x_t, t, c)) \nabla_\theta \mathcal{S}_\theta(\xi, c)], \quad (3)$$

where s_ψ is the frozen Teacher score model and s_η approximates the intractable Student score. In practice, s_η is commonly implemented by an auxiliary score

network trained online on Student samples, with alternating updates between $\mathcal{S}_\theta$ and the score estimator [45]. Besides, this distillation often operates entirely in latent space, assuming the Teacher and Student share the same space.

Adversarial distillation. Several works introduce a discriminator to align few-step Student outputs with those of a multi-step Teacher by providing an additional distribution-level training signal [32, 34]. This objective is often combined with VSD-style distribution matching to further reduce the gap between a one-step Student and a multi-step Teacher [32, 45].

In our work. These objectives require Teacher and Student to have comparable representations, such as the input to the Teacher score model or any discriminator. Our method satisfies this under Cross-Resolution and Cross-VAE mismatch by aligning Student states into a Teacher-compatible latent space, enabling standard VSD and adversarial objectives with minimal modification.

3 Distillation Enabled by the Bridge

3.1 Cross-Space Distillation

One-step distribution-based distillation methods [5, 26, 27, 32, 33, 46, 47] reduce the *temporal* cost of diffusion generation by decreasing the number of sampling steps. However, these methods typically assume that the teacher and student share the same latent space, thereby requiring supervision to be applied directly within that common representation.

This **Shared-Space constraint** becomes restrictive in practical large-scale distillation scenarios. Modern state-of-the-art teacher models often operate at higher latent resolution and channels and rely on different VAE architectures, whereas compact student models typically employ lower latent resolutions and distinct VAE designs. Under such representational mismatch, standard distillation objectives cannot be directly applied to the teacher and student latent variables without additional alignment.

We refer to this regime as **Cross-Space Distillation**, in which supervision must be conducted across latent spaces that differ in both **spatial resolution** and **underlying VAE representations**. Addressing these discrepancies requires learning an explicit alignment mapping to enable distillation between teacher and student representations.

3.2 Problem Setting and Notation

We denote the frozen Teacher denoiser by $\mathcal{T}$ and the Student denoiser by $\mathcal{S}$. When referring to autoencoders, we use $(\mathcal{E}_T, \mathcal{D}_T)$ and $(\mathcal{E}_S, \mathcal{D}_S)$ for the Teacher and Student VAE encoder/decoder, respectively. In this paper, we use the term *space* to refer to a model’s *latent representation space* induced by its VAE parameterization and latent grid resolution.

$$\hat{z}_T = \mathcal{B}_\phi(z_S), \quad \hat{z}_T \in \mathcal{Z}_T. \quad (4)$$




Fig. 2: Bridge for Cross-Space Distillation. **Left. Bridge Training:** given an image x , we encode it with both VAEs to obtain paired latents (z_S, z_T) . The Bridge $\mathcal{B}_\phi$ maps $z_S \rightarrow \hat{z}_T$ using a **frozen spatial prior** from a frozen prefix of the Student decoder $\mathcal{D}_S^{(n)}$ to expand the latent grid, followed by a **learnable projector** g_ϕ that outputs a Teacher-compatible latent. We train $\mathcal{B}_\phi$ with latent reconstruction $\mathcal{L}_{rec}$ and attention fidelity $\mathcal{L}_{attn}$ computed from a frozen Teacher denoiser by matching attention responses induced by $\hat{z}_T$ and z_T . **Top-right. Distillation Enabled by Bridge:** after training, $\mathcal{B}_\phi$ converts Student outputs into the Teacher latent space so standard one-step distillation losses such as $\mathcal{L}_{VSD}$ and $\mathcal{L}_{GAN}$ could apply, bypassing Cross-Resolution and Cross-VAE mismatch. **Bottom-right. Inference with Resolution Upgrade:** a low-resolution Student sample can be mapped by $\mathcal{B}_\phi$ and decoded with the Teacher decoder to synthesize at the Teacher resolution without changing the Student denoiser.

A useful Bridge should meet three criteria: (1) High precision in reconstructing the Teacher latent $\hat{z}_T$, (2) Efficiency in memory and computation, adding negligible overhead to the Student network, and (3) Distillation-friendliness, enabling effective knowledge transfer using standard distribution-based distillation methods. To satisfy these requirements, we introduce new designs across Bridge architecture (Sec. 3.3), training objectives (Sec. 3.4), and describe how the trained Bridge is used for distillation and inference in Sec. 3.4.

3.3 Architecture with Rich Priors

The key challenge in distillation under latent mismatch is that z_S and z_T differ in spatial shape and latent parameterization. We therefore factor $\mathcal{B}_\phi$ into two stages. The first stage aligns spatial resolution from $h \times w$ to $H \times W$. The second stage aligns features into the Teacher latent representation.

Spatial prior. For spatial alignment, one could learn a latent upsampler jointly with the Bridge, but this requires the model to learn spatial upsampling behavior from scratch. Instead, we reuse the early decoding blocks of the Student VAE decoder, which already implement latent upsampling and are trained with a massive amount of data. Let $\mathcal{D}_S^{(n)}$ denote the first n decoding blocks of the Student VAE decoder. We keep them frozen and choose n such that their output has spatial resolution $H \times W$, matching the Teacher’s latent resolution. Given

$z_S \in \mathbb{R}^{h \times w \times C_S}$, we compute:

$$f_{\text{prior}} = \mathcal{D}_S^{(n)}(z_S), \quad f_{\text{prior}} \in \mathbb{R}^{H \times W \times C_{\text{prior}}}, \quad (5)$$

where C_{prior} is the channel dimension of this decoder feature. This stage performs the spatial expansion without introducing additional trainable parameters and provides a strong scaffold for the next stage.

Projection head. After spatial alignment, f_{prior} has the correct spatial resolution but remains in the Student decoder feature space rather than the Teacher latent representation. We therefore learn a compact projection head g_ϕ that maps f_{prior} to a Teacher compatible latent:

$$\hat{z}_T = \mathcal{B}_\phi(z_S) = g_\phi(f_{\text{prior}}), \quad \hat{z}_T \in \mathbb{R}^{H \times W \times C_T}. \quad (6)$$

Overall, our Bridge is: $z_S \xrightarrow{\mathcal{D}_S^{(n)}} f_{\text{prior}} \xrightarrow{g_\phi} \hat{z}_T$, which isolates spatial alignment in the frozen decoder prefix while using the trainable module g_ϕ for channel and semantic alignment. We optimize g_ϕ with the objectives in Sec. 3.4.

3.4 Training Objectives

We optimize the *Bridge* using two complementary objectives: the first term provides a regression signal in the Teacher latent space, while the second term adds Teacher-based supervision by matching internal attention responses.

Latent Reconstruction ($\mathcal{L}_{\text{rec}}$). Given paired latents (z_S, z_T) extracted from the Student and Teacher encoders, the Bridge maps the Student latent to the Teacher latent space as $\hat{z}_T = \mathcal{B}_\phi(z_S)$. We supervise the Bridge with an ℓ_1 reconstruction loss:

$$\mathcal{L}_{\text{rec}} = \|z_T - \hat{z}_T\|_1. \quad (7)$$

While ℓ_2 is a common choice, we observe that it is often inefficient and less effective for latent-space reconstruction in practice, as noted in recent distillation/acceleration works [18, 25, 49]. In our setting, the ℓ_1 objective provides better training stability and helps preserve salient structure in the Teacher latent space.

Attention Fidelity ($\mathcal{L}_{\text{attn}}$). While latent reconstruction $\mathcal{L}_{\text{rec}}$ is necessary for alignment, it is often *insufficient*: small latent discrepancies can produce perceptible semantic errors after decoding. Pixel-wise losses may tolerate slight misalignments, thus losing *fine-grained structure*. Adding adversarial losses can sharpen details but is sensitive to hyperparameters and reduces training stability. We therefore seek a supervision signal that is **(i)** stable, **(ii)** globally aware, and **(iii)** defined directly in Teacher space.

To this end, inspired by [23], we align the Teacher denoiser’s self-attention *distributions*. Attention captures long-range dependencies, so mismatches between $\hat{z}_T$ and z_T are amplified in the Teacher’s internal responses; a compatible latent should induce similar attention patterns under identical diffusion conditions. We extract attention maps $P^l(z; t, c) \in [0, 1]^{N \times N}$ from a fixed set of Teacher layers

$l \in \text{Layers}$, where N is the number of tokens. Each P^l is obtained via row-wise softmax over key positions from the Teacher’s query and key activations. Following MiniLLM [7], we use reverse KL divergence, which focuses on the dominant attention mass and is empirically more stable than forward KL. We minimize:

$$\mathcal{L}_{\text{attn}} = \sum_{l \in \text{Layers}} \tau^2 \text{KL}(P^l(\hat{z}_T; t, c) \parallel P^l(z_T; t, c)), \quad (8)$$

where $\tau > 0$ is a temperature for numerical stability, and we normalize attention logits before applying the softmax.

Final Objective ($\mathcal{L}_{\text{final}}$). We train the Bridge by minimizing a weighted sum of latent reconstruction and attention fidelity:

$$\mathcal{L}_{\text{final}} = \alpha \mathcal{L}_{\text{rec}} + \beta \mathcal{L}_{\text{attn}}, \quad (9)$$

where α, β are hyperparameters and set to $\alpha = \beta = 1$ in our experiments.

Bridge Usage. After training the Bridge, we attach the frozen module to the learnable Student to form an augmented Student that supports one-step distillation and one-step inference from the Teacher. This procedure follows the prior work summarized in Sec. 2.2 and is illustrated in Fig. 2, in the top-right and bottom-right panels.

4 Experiments

4.1 Experimental Setup

Student Models ($\mathcal{S}$). We initialize the Student from one-step diffusion variants: SD 1.5 backbone from DMD2 [46, 47] and SD 2.1 backbone from SiD-LSG [52]. These models serve as representative legacy generators operating at 512×512 resolution.

Teacher Models ($\mathcal{T}$). We evaluate our method under distillation settings using five state-of-the-art teachers operating at 1024×1024 resolution, including SDXL [29], Kolors [39], PixArt- Σ [3], FLUX.2 [klein] 4B [13], and SD 3.5 Medium [6]. These models span different backbone architectures (U-Net, DiT, and MMDiT), and employ different latent configurations with varying VAE designs and channel dimensions (4, 16, and 32 channels).

Bridge ($\mathcal{B}_\phi$). Our Bridge module $\mathcal{B}_\phi$ is a lightweight alignment network with approximately 5M trainable parameters. As described in Sec. 3.3, $\mathcal{B}_\phi$ is factorized into a frozen spatial prior and a learnable projection head. The spatial prior is implemented by the first n decoding blocks of the Student VAE decoder, denoted $\mathcal{D}_S^{(n)}$, which expands the Student latent grid to the Teacher spatial resolution. In all experiments, we set $n = 1$ and freeze $\mathcal{D}_S^{(1)}$. The learnable projection head g_ϕ is implemented with SwinIR [17] and maps the intermediate decoder features to a Teacher-compatible latent $\hat{z}_T$. We train $\mathcal{B}_\phi$ with the objectives in Sec. 3.4. For

Table 1: Quantitative results. A: modern multi-step Teachers. **B:** compact one-step Students enabled by *Cross-Space Distillation*. Despite large gaps in resolution, VAE space, architecture, and generative mechanism, Bridge transfers knowledge from diverse modern Teachers into the same lightweight Student backbones while preserving one-step inference. In **B**, the second column lists the Teacher used for each individual distillation run; initialized and merged rows summarize the corresponding Student family. **Paras** are parameter counts in billions (B), and **NFE** denotes inference-time network evaluations. **Pink** marks initialized Students, **pale green** marks individual Bridge-distilled Students, and **stronger green** marks checkpoints obtained by averaging all five distilled Students in the same family. Across teacher families and metrics, Bridge **consistently improves** the initialized Students, and the merged checkpoints provide a simple post-training route to combine knowledge from multiple Teacher sources. Higher values are better for all metrics.

Teacher		Paras	NFE	Metrics				
Architect.	Model	(B)		HPSv3	HPSv2	IR	MPS	DPG
A. Multi-Step Teachers								
Large U-Net	SDXL [29]	2.57	50	9.25	28.36	0.66	13.72	74.00
	Kolors [39]	2.57	50	10.59	30.94	0.88	14.12	76.52
DiT	PixArt- Σ -1024 [3]	0.61	20	9.62	30.39	0.92	14.14	80.00
MM-DiT	FLUX.2-klein-4B [13]	4.00	50	10.05	28.90	0.80	14.00	83.20
	SD 3.5 Medium [6]	2.50	50	10.86	30.02	0.96	14.13	84.50
B. Cross-Space Distillation								
SD 1.5 Student, DMD2 [46] (<i>init.</i>)		0.86	1	5.37	21.90	-0.29	10.42	59.85
Large U-Net	SDXL [29]	0.86	1	9.04	27.37	0.30	11.93	63.00
	Kolors [39]	0.86	1	9.33	27.38	0.34	11.66	64.46
DiT	PixArt- Σ -1024 [3]	0.86	1	8.65	26.43	0.30	11.94	63.43
MM-DiT	FLUX.2-klein-4B [13]	0.86	1	9.49	26.64	0.40	12.35	64.06
	SD 3.5 Medium [6]	0.86	1	9.42	28.30	0.62	12.96	65.75
Merged (All 5 Teachers)		0.86	1	10.53	29.07	0.65	12.62	66.67
SD 2.1 Student, SiD [52] (<i>init.</i>)		0.86	1	6.42	23.74	0.10	11.29	61.33
Large U-Net	SDXL [29]	0.86	1	8.40	26.73	0.23	11.88	67.03
	Kolors [39]	0.86	1	8.52	26.57	0.20	11.91	68.34
DiT	PixArt- Σ -1024 [3]	0.86	1	8.42	28.29	0.44	12.20	68.51
MM-DiT	FLUX.2-klein-4B [13]	0.86	1	8.74	28.31	0.33	12.03	66.00
	SD 3.5 Medium [6]	0.86	1	8.64	29.11	0.45	12.92	68.52
Merged (All 5 Teachers)		0.86	1	9.75	30.00	0.74	12.60	68.50

Attention Fidelity, we use noising level $t = 1$ and temperature $\tau = 3.0$. Ablations on architecture and objectives are provided in the supplementary material.

Distillation Details. We follow the distillation protocol of DMD2 [46]. All models are trained for 20 hours on $8 \times$ NVIDIA H100 GPUs with 80GB memory using the AdamW optimizer. Additional implementation details are provided in the supplementary material.

Training Data. To maximize *training efficiency* and *data accessibility*, we avoid relying on large-scale real-image datasets. Instead, we construct a static synthetic dataset of approximately 2M images by sampling text prompts from **JourneyDB** and **LAION** [35] and generating images with the corresponding teacher model. We use this synthetic dataset to train $\mathcal{B}_\phi$ during space-bridge training. During distillation, we recycle the same dataset to compute an auxiliary adversarial objective, following the setup of DMD2.

Metrics. We evaluate Cross-Space Distillation using five widely adopted metrics: HPSv3 [24], HPSv2 [42], ImageReward [44], MPS [48], and DPG Bench [10]. HPSv3 and HPSv2 assess human preference alignment on 12,000 and 3,200 prompts, respectively, while ImageReward is computed on 100 prompts. We further report MPS on the HPSv2 and ImageReward prompt sets to measure multi-dimensional scores. DPG Bench evaluates text-image alignment on 1,065 prompts. For all metrics, higher scores indicate better performance.

4.2 Main Results

Cross-Architecture and Cross-Mechanism Distillation. As demonstrated in Tab. 1, our framework successfully facilitates distillation across architectures and diffusion frameworks, such as from DiT-based, flow matching models (*e.g.*, SD 3.5 and FLUX.2) to UNet-based, noise-prediction models (*e.g.*, SD 1.5). This versatility stems from two key design choices. First, our proposed Bridge explicitly neutralizes the Cross-Resolution and Cross-VAE constraints. Second, we adopt a mechanism-agnostic formulation by reparameterizing the student’s prediction back to x_0 before sending it to the teacher, thereby eliminating discretization inconsistencies between the diffusion and flow matching formulations. Furthermore, we deliberately avoid layer-wise matching; supervision is imposed only at the prediction level, thereby eliminating structural constraints between the DiT and UNet backbones. Notably, even without architecture-based alignment, our method still achieves strong performance, demonstrating that once resolution and VAE mismatches are properly addressed, effective Cross-Space distillation becomes not only feasible but robust.

Quantitative Results. Tab. 1 summarizes performance under five complementary metrics. Across all teacher models, Cross-Space Distillation improves the one-step Students on these metrics, suggesting that the gains are not confined to a single evaluator. For instance, using SD 3.5 Medium as the Teacher increases the SD 1.5 Student from 5.37 to 9.42 on HPSv3 and from -0.29 to 0.62 on ImageReward, with matching improvements on HPSv2, MPS, and DPG Bench. Overall, these results suggest higher perceived quality and stronger prompt adherence while preserving the same one-step inference budget.

Qualitative Results. We report qualitative examples in Fig. 3. Across a diverse set of prompts, our approach produces images with clearer textures, stronger structural consistency, and more faithful local details than the initialized baseline. In contrast, the baseline more frequently exhibits softness, broken fine struc-




Fig. 3: Qualitative Comparison of Cross-Space Distillation. We show samples from the initialized one-step student SD 1.5 DMD2 at 512×512 and from one-step students distilled with our method using different 1024×1024 teachers. Columns labeled **+Teacher** denote the student distilled from that teacher.

ture, and occasional anatomical artifacts. These examples visually support the quantitative gains by showing that Bridge enabled distillation improves perceptual sharpness and coherence while preserving one-step generation.

4.3 Post-Training Improvement via Model Merging

As an optional post-training refinement, we investigate **model merging** across students distilled from different teachers. This experiment is *not* part of the core Bridge design; rather, it serves as a practical extension that can further improve the final student **without changing the model architecture, parameter count, or inference cost**. Since all distilled checkpoints within each student family share the same backbone and parameterization, they can be combined through **direct parameter averaging**.

Given a set of distilled one-step student checkpoints $\mathcal{S} = \{\theta_1, \theta_2, \dots, \theta_{|\mathcal{S}|}\}$, we form a merged checkpoint by taking the element-wise mean of the parameters:

$$\theta_{\text{merged}} = \frac{1}{|\mathcal{S}|} \sum_{i=1}^{|\mathcal{S}|} \theta_i. \quad (10)$$

Table 2: Pruning versus Bridge enabled distillation. We prune SD 3.5 Medium to 30% sparsity with OBS-Diff [53] and use it to initialize a one-step Student, then distill from the full SD 3.5 Medium Teacher using the same setup as Sec. 4.1. Bridge enabled distillation yields higher preference scores and cleaner details, while pruning based initialization more often produces blur and artifacts.

Setting	Params (B)	HPSv3	HPSv2	IR	MPS	DPG
Prune (30%)	1.5	2.76	21.94	-0.36	9.54	55.24
Ours	0.86	9.42	28.30	0.62	12.96	65.75

In our experiments, we merge the five students obtained by distilling from the five different teacher models in the main paper. This simple procedure consolidates multiple teacher-specific distilled checkpoints into a **single student checkpoint** while preserving the same inference-time NFE and model size.

Tab. 1 reports the merged results as highlighted rows under both SD 1.5 and SD 2.1 student families. Parameter averaging yields a stronger *overall* metric profile than the initialized student and is often competitive with, or better than, the individual distilled checkpoints on HPSv3, HPSv2, and ImageReward. For example, the merged SD 1.5 checkpoint improves *HPSv3 from 5.37 to 10.53* and *HPSv2 from 21.90 to 29.07*, while the merged SD 2.1 checkpoint improves *HPSv3 from 6.42 to 9.75* and *HPSv2 from 23.74 to 30.00*. These results suggest that model merging can serve as a lightweight post-training enhancement for Cross-Space Distillation.

This model merging presents an *unique benefit* of our technique. Normally, it is hard to combine the image-generation capabilities of state-of-the-art models like Flux, SD 3.5, or Kolors, given their different network architectures and latent spaces. However, our technique distills them into student models with the same backbone and latent space, and combining knowledge across these student models is trivial. It provides a pathway to **unify generation capabilities from multiple model sources**. With more teacher models and/or stronger distillation techniques beyond DMD2, we expect an even stronger merged student model that can outperform every single teacher model.

4.4 Additional Analyses

Student Initialization. We compare two Student initializations: one-step SD 1.5 DMD2 checkpoint and standard multi-step SD 1.5 Teacher weights. With SD 3.5 Medium as the Teacher (Tab. 3), DMD2-init is slightly better in HPSv2, while SD 1.5 is slightly better on other metrics. This suggests initialization has minimal influence on performance in our framework. To maximize training speed and efficiency, we opt to initialize DMD2 as students.

Comparison with Network Pruning. To benchmark against parameter reduction, we compare the Bridge with network pruning. Using OBS-Diff [53], we prune SD 3.5 Medium to 30% sparsity (2.5B $\rightarrow$ 1.5B) and use it to initialize

Table 3: Student initialization. We compare initializing the Student from a one-step SD 1.5 DMD2 checkpoint versus standard multi-step SD 1.5 weights, with both distilled from the same SD 3.5 Medium Teacher. Results are similar across metrics, so we use DMD2 initialization for subsequent experiments due to better training efficiency.

Setting	Iteration	HPSv3	HPSv2	IR	MPS	DPG
SD 1.5 Init.	9k	9.22	28.44	0.66	13.08	66.89
SD 1.5 DMD2 Init.	3k	9.42	28.30	0.62	12.96	65.75



Fig. 4: Qualitative comparisons. (Left) Pruning vs. Bridge distillation. We compare a pruned baseline (1.5B) against our Bridge-enabled compact student (0.6B), generated with identical prompts and $\text{NFE} = 1$. **(Right) Resolution upgrade at inference time.** A 512×512 sample from the same one-step student is mapped by **Bridge** into the teacher latent space and decoded with the teacher decoder to produce a 1024×1024 output.

the one-step Student $\mathcal{S}$, then distill from the full SD 3.5 Medium Teacher following Sec. 4. As shown in Tab. 2 and Fig. 4, with the same $\text{NFE}=1$ budget, our approach yields notably sharper outputs and stronger preferences, while pruned-model distillation tends to introduce blur and artifacts.

Resolution Upgrade with Bridge. Once trained, our Bridge $\mathcal{B}_\phi$ can also serve as a plug-and-play inference module that upgrades a low-resolution generator to higher-resolution outputs without retraining the generator itself. As shown in the right pane of Fig. 4, the Bridge synthesizes semantically consistent high-frequency details from low-resolution latents. This demonstrates that $\mathcal{B}_\phi$ generalizes to unseen inputs and provides a robust, cross-resolution mapping.

5 Related Work

Diffusion and Flow. *Diffusion models* have become the dominant paradigm for high-quality image generation. Early works such as DDPM [8] and score-based generative models [38] introduced a stochastic forward corruption process and a learned reverse denoising process for generative modeling. Subsequent improvements have significantly accelerated sampling through more efficient solvers [11, 21, 37] and improved training formulations [31].

Recently, *flow matching models* [19,20,40] provide an alternative continuous-time formulation that directly learns the velocity field transporting samples from noise to data. Modern text-to-image systems such as SD 3.5 [6] and Flux [12,13] build on these advances with large-scale backbones and high-resolution latent representations, achieving strong visual fidelity and prompt alignment. However, these models typically require multiple network evaluations during inference, resulting in substantial computational overhead.

Distillation for Fast Diffusion Sampling. To reduce inference cost, a large body of work distills multi-step diffusion models into fast generators with significantly fewer sampling steps. Distribution-based distillation methods train a Student generator to match the sample distribution implied by a pretrained Teacher model. Representative approaches include variational score distillation and related methods [26,27,41,45–47], which use score differences between Teacher and Student predictions as training signals. Other approaches explore consistency-style distillation [14,22] or trajectory-based methods that directly approximate multi-step sampling dynamics.

Several works further combine distribution matching with adversarial supervision to improve perceptual quality and reduce the gap between fast generators and their multi-step Teachers [32,33]. While these approaches substantially reduce the number of inference steps, they typically assume that the Teacher and Student operate within the same latent representation and spatial resolution. Consequently, the distilled model often inherits the architectural scale or latent structure of the Teacher, limiting its suitability for lightweight deployment.

In this paper, we study **Cross-Space Distillation**, introducing a lightweight alignment module that maps Student representations into the Teacher latent space, enabling standard distillation objectives without modifying the Student architecture. This allows compact backbones such as Stable Diffusion 1.5 [30] to inherit the generative capability of modern high-capacity models while maintaining ecosystem compatibility.

6 Conclusion

We identify a common but restrictive assumption in distribution-based one-step distillation: Teacher and Student are expected to share the same latent representation. We relax this assumption through **Cross-Space Distillation** and introduce **Bridge** ($\mathcal{B}_\phi$), a lightweight latent interface that maps Student latents into the Teacher space, making standard one-step distillation objectives applicable under latent-resolution and VAE mismatch. Across modern Teachers and compact Students, Bridge yields strong preference gains while preserving one-step inference, the original Student backbone, and compatibility with widely used deployment ecosystems. Beyond training, Bridge can also be reused at inference time to map low-resolution Student samples onto the Teacher latent grid for higher-resolution decoding. We hope this work motivates broader exploration of distillation and model reuse across heterogeneous latent representations, and inspires future research on alignment interfaces that unlock knowledge transfer across mismatched latent spaces.

References

1. Cai, F., Guo, Y., Li, J., Li, W., Chen, J., Fang, X.: Fastflux: Pruning flux with block-wise replacement and sandwich training. arXiv preprint arXiv:2506.10035 (2025)
2. Chen, J., Hu, D., Huang, X., Coskun, H., Sahni, A., Gupta, A., Goyal, A., Lahiri, D., Singh, R., Idelbayev, Y., et al.: Snapgen: Taming high-resolution text-to-image models for mobile devices with efficient architectures and training. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 7997–8008 (2025)
3. Chen, J., Ge, C., Xie, E., Wu, Y., Yao, L., Ren, X., Wang, Z., Luo, P., Lu, H., Li, Z.: Pixart- σ : Weak-to-strong training of diffusion transformer for 4k text-to-image generation. In: European Conference on Computer Vision (ECCV) (2024), <https://arxiv.org/abs/2403.04692>
4. Dao, Q., Metaxas, D.: Mpdit: Multi-patch global-to-local transformer architecture for efficient flow matching and diffusion model. arXiv preprint arXiv:2603.26357 (2026)
5. Dao, T., Nguyen, T.H., Le, T., Vu, D., Nguyen, K., Pham, C., Tran, A.: Swift-brush v2: Make your one-step diffusion model better than its teacher. In: European Conference on Computer Vision. pp. 176–192. Springer (2024)
6. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first international conference on machine learning
7. Gu, Y., Dong, L., Wei, F., Huang, M.: Minillm: Knowledge distillation of large language models. arXiv preprint arXiv:2306.08543 (2023)
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
9. Hu, D., Gupta, A., Gabidolla, M., Sahni, A., Coskun, H., Li, Y., Idelbayev, Y., Mahmood, A., Lebedev, A., Lahiri, D., et al.: Snapgen++: Unleashing diffusion transformers for efficient high-fidelity image generation on edge devices. arXiv preprint arXiv:2601.08303 (2026)
10. Hu, X., Wang, R., Fang, Y., et al.: Ella: Equip diffusion models with llm for enhanced semantic alignment. arXiv preprint arXiv:2403.05135 (2024)
11. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2022), <https://arxiv.org/abs/2206.00364>
12. Labs, B.F.: Announcing FLUX.1. <https://blackforestlabs.ai/announcing-flux-1> (2024), accessed: 2026-03-04
13. Labs, B.F.: FLUX.2: Frontier Visual Intelligence. <https://bfl.ai/blog/flux-2> (2025), accessed: 2026-03-04
14. Lee, S., Xu, Y., Geffner, T., Fantì, G., Kreis, K., Vahdat, A., Nie, W.: Truncated consistency models. In: *International Conference on Learning Representations (ICLR)* (2025), <https://openreview.net/pdf/bb8f3dceac43037618899ff56c90995c5e08e978.pdf>
15. Li, J., Feng, W., Chen, W., Wang, W.Y.: Reward guided latent consistency distillation. arXiv preprint arXiv:2403.11027 (2024). <https://doi.org/10.48550/arXiv.2403.11027>
16. Li, Y., Wang, H., Jin, Q., Hu, J., Chemerys, P., Fu, Y., Wang, Y., Tulyakov, S., Ren, J.: Snapfusion: Text-to-image diffusion model on mobile devices within two seconds. *Advances in Neural Information Processing Systems* **36**, 20662–20678 (2023)

17. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1833–1844 (2021)
18. Lin, S., Wang, A., Yang, X.: Sdxl-lightning: Progressive adversarial diffusion distillation. arXiv preprint arXiv:2402.13929 (2024). <https://doi.org/10.48550/arXiv.2402.13929>, <https://arxiv.org/abs/2402.13929>
19. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: International Conference on Learning Representations (ICLR) (2023), <https://arxiv.org/abs/2210.02747>
20. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022), <https://arxiv.org/abs/2209.03003>
21. Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In: Advances in Neural Information Processing Systems (NeurIPS) (2022), <https://arxiv.org/abs/2206.00927>
22. Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference. arXiv preprint arXiv:2310.04378 (2023), <https://arxiv.org/abs/2310.04378>
23. Ma, J., Peng, Q., Guo, X., Chen, C., Lu, H., Yang, Z.: X2i: Seamless integration of multimodal understanding into diffusion transformer via attention distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16733–16744 (October 2025)
24. Ma, Y., Wu, X., Sun, K., Li, H.: Hpsv3: Towards wide-spectrum human preference score. arXiv preprint arXiv:2508.03789 (2025)
25. Nguyen, A., Nguyen, V., Vu, D., Dao, T., Tran, C., Tran, T., Tran, A.: Improved training technique for shortcut models. arXiv preprint arXiv:2510.21250 (2025). <https://doi.org/10.48550/arXiv.2510.21250>, <https://arxiv.org/abs/2510.21250>, accepted at NeurIPS 2025
26. Nguyen, T.H., Tran, A.: Swiftbrush: One-step text-to-image diffusion model with variational score distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7807–7816 (2024)
27. Nguyen, V., Nguyen, A., Dao, T., Nguyen, K., Pham, C., Tran, T., Tran, A.: Supercharged one-step text-to-image diffusion models with negative prompts. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18004–18013 (2025)
28. Nguyen, V., Vu, G., Thanh, T.N., Than, K., Tran, T.: On inference stability for diffusion models. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 14449–14456 (2024)
29. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=di52zR8xgf>
30. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10684–10695 (June 2022)
31. Salimans, T., Ho, J.: Progressive distillation for fast sampling of diffusion models. arXiv preprint arXiv:2202.00512 (2022), <https://arxiv.org/abs/2202.00512>

32. Sauer, A., Boesel, F., Dockhorn, T., Blattmann, A., Esser, P., Rombach, R.: Fast high-resolution image synthesis with latent adversarial diffusion distillation. In: SIGGRAPH Asia 2024 Conference Papers. pp. 1–11 (2024)
33. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. In: European Conference on Computer Vision. pp. 87–103. Springer (2024)
34. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. In: European Conference on Computer Vision. pp. 87–103. Springer (2024)
35. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in neural information processing systems* **35**, 25278–25294 (2022)
36. Schusterbauer, J., Gui, M., Fundel, F., Ommer, B.: Diff2flow: Training flow matching models via diffusion model alignment. *arXiv preprint arXiv:2506.02221* (2025), <https://arxiv.org/abs/2506.02221>
37. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020), <https://arxiv.org/abs/2010.02502>
38. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (ICLR) (2021), <https://arxiv.org/abs/2011.13456>
39. Team, K.: Kolores: Effective training of diffusion model for photorealistic text-to-image synthesis. *arXiv preprint* (2024), <https://github.com/Kwai-Kolores/Kolores>
40. Tong, A., Fatras, K., Malkin, N., Huguier, G., Zhang, Y., Rector-Brooks, J., Wolf, G., Bengio, Y.: Improving and generalizing flow-based generative models with minibatch optimal transport. *arXiv preprint arXiv:2302.00482* (2023), <https://arxiv.org/abs/2302.00482>
41. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023), <https://arxiv.org/abs/2305.16213>
42. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341* (2023)
43. Xie, Q., Liao, Z., Deng, Z., Chen, C., Lu, H.: Tlcm: Training-efficient latent consistency model for image generation with 2-8 steps. *arXiv preprint arXiv:2406.05768* (2024). <https://doi.org/10.48550/arXiv.2406.05768>
44. Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., Dong, Y.: Imagereward: Learning and evaluating human preferences for text-to-image generation. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2023)
45. Xu, Y., Nie, W., Vahdat, A.: One-step diffusion models with f -divergence distribution matching. *arXiv preprint arXiv:2502.15681* (2025), <https://arxiv.org/abs/2502.15681>
46. Yin, T., Gharbi, M., Park, T., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T.: Improved distribution matching distillation for fast image synthesis. In: *NeurIPS* (2024)
47. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: *CVPR* (2024)
48. Zhang, .: Learning multi-dimensional human preference for text-to-image generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2024)

49. Zhang, Y., Hooi, B.: Hipa: Enabling one-step text-to-image diffusion models via high-frequency-promoting adaptation. arXiv preprint arXiv:2311.18158 (2023). <https://doi.org/10.48550/arXiv.2311.18158>, <https://arxiv.org/abs/2311.18158>
50. Zhao, Y., Xu, Y., Xiao, Z., Jia, H., Hou, T.: Mobilediffusion: Instant text-to-image generation on mobile devices. In: European Conference on Computer Vision. pp. 225–242. Springer (2024)
51. Zheng, J., Hu, M., Fan, Z., Wang, C., Ding, C., Tao, D., Cham, T.J.: Trajectory consistency distillation: Improved latent consistency distillation by semi-linear consistency function with trajectory mapping. arXiv preprint arXiv:2402.19159 (2024). <https://doi.org/10.48550/arXiv.2402.19159>
52. Zhou, M., Wang, Z., Zheng, H., Huang, H.: Guided score identity distillation for data-free one-step text-to-image generation (2024), <https://arxiv.org/abs/2406.01561>
53. Zhu, J., Wang, H., Su, M., Wang, Z., Wang, H.: Obs-diff: Accurate pruning for diffusion models in one-shot. arXiv preprint arXiv:2510.06751 (2025), <https://arxiv.org/abs/2510.06751>