

Curvature-Adaptive Consistency Flow Matching: Autonomous Trajectory Optimization via Reinforcement Learning

Songtao Tian^{1*}, Guhan Chen^{1*}, Bohan Li¹, Jingyi Ma¹, and Zixiong Yu^{2**}

¹Tsinghua University, Beijing, China

²Huawei Large Model Data Technology Lab, Shenzhen, China

Abstract. Consistency distillation has significantly accelerated diffusion-model inference, but its sampling dynamics remain underexplored. We reveal an asymmetry: although Logit-Normal sampling priors work well for standard iterative generation, consistency distillation exhibits a different difficulty profile (*e.g.*, U-shaped), with optimization bottlenecks concentrated at the boundary stages rather than intermediate steps. To address the limitations of static sampling under evolving learning demands, we propose Curvature-Adaptive Consistency Flow Matching (CACFM). By formulating distillation as a dynamic decision process, CACFM uses a lightweight reinforcement learning agent to probe Probability Flow ODE trajectories and construct an efficiency-oriented curriculum that prioritizes critical regions without manual scheduling. Combined with Flow-adapted DMD and adversarial consistency objectives, our RL-based scheduler achieves state-of-the-art results on large-scale models such as FLUX and SDXL, mitigating structural deformities and preserving high-frequency details in extreme few-step regimes.

1 Introduction

In recent years, artificial intelligence (AI) systems have become a central pillar of modern computation. Although recent theoretical studies have advanced our understanding of these systems [16, 17], this understanding remains incomplete [3, 55]. Modern AI models have demonstrated strong capabilities in language understanding [43, 54], visual synthesis [21, 23], video generation [10, 12, 20], and multimodal reasoning [52, 56]. These advances have been driven by progress in architectures [8, 53], data construction [9, 34, 49], and algorithm design [39, 47, 57].

Within generative modeling, diffusion models and their recent extension to Flow Matching (FM) frameworks have emerged as a powerful paradigm across multiple domains [25, 40, 42]. Notably, modern large-scale transformer-based architectures [43], such as FLUX [15] and SD3 [5] utilize the Flow Matching

* Equal contribution. Work by Guhan Chen was done during an internship at Huawei.

** Corresponding author: yuzx19@tsinghua.org.cn



Fig. 1: 1024×1024 samples produced by our 4-step generator distilled from FLUX.1-dev.

formulation, which relies on deterministic Probability Flow ODEs (PF-ODEs) to transform noise into data [25]. While this formulation provides flexible control over the trade-off between computational cost and output quality, it inherently requires numerous sequential evaluations of computationally intensive neural networks, creating a significant bottleneck for real-time applications.

To mitigate this computational burden, extensive research has focused on acceleration techniques, ranging from advanced ODE solvers [13, 28] to distillation-based approaches like Consistency Models (CMs [41]) and Phased Consistency Models (PCMs [46]). However, existing methods exhibit a distinct static nature in their sampling strategies: they predominantly rely on pre-defined discretization schemes or heuristic priors, such as basic uniform sampling or directly adopting the Logit-Normal distribution [5] widely utilized in Rectified Flow training.

However, we contend that high-dimensional generative flows in consistency distillation exhibit pronounced Geometric Heterogeneity, where curvature and optimization difficulty vary drastically across timesteps. A uniform allocation of computational budget fails to resolve high-curvature regions. Furthermore, directly adopting static heuristics like the Logit-Normal distribution lacks the adaptive flexibility to accommodate the unique optimization dynamics of consistency distillation, especially as optimization focal points shift during training.

To validate this perspective, we compute the Oracle Consistency Error along a converged teacher trajectory to reveal the true difficulty distribution. Building on this, we propose Curvature-Adaptive Consistency Flow Matching (CACFM; see Fig. 1 for visual examples), which reformulates consistency distillation as a dynamic decision process. Instead of manually designing sampling schedules, we deploy a lightweight Reinforcement Learning (RL) agent to dynamically probe the trajectory, identifying and prioritizing sub-trajectories with higher optimization returns. This leads to the following fundamental discovery:

1. **U-Shaped Difficulty Profile:** Contrary to the common observation in standard iterative generation training that intermediate steps are the most challenging (modeled by Logit-Normal sampling), the consistency error analysis reveals that in consistency distillation tasks, the boundary stages are more challenging than the intermediate phases (Fig. 2).
2. **Emergent Curriculum Learning:** Beyond learning the aforementioned curve, the agent develops a coarse-to-fine strategy, progressing from establishing global structures in early training to refining high-frequency details in later stages (Fig. 6).




Fig. 2: Comparing sampling distributions along the FM trajectory. (a) The Logit-Normal prior assumes that intermediate steps are the most important. (b) The U-shaped distribution identifies boundary stages (initialization and final refinement) as the primary bottlenecks, and our agent learns this pattern (see Fig. 6 for experimental results).

While curvature-adaptive sampling effectively allocates the training budget to resolve trajectory bottlenecks, extreme few-step generation still suffers from distributional shift at the manifold level. To further enhance fidelity, we integrate Flow-adapted Distribution Matching Distillation (DMD), inspired by [51], and adversarial consistency as complementary training objectives; the core algorithmic novelty of CACFM remains the RL-based temporal scheduler. Extensive experiments on FLUX and SDXL architectures demonstrate that CACFM establishes new state-of-the-art (SOTA) performance. By intelligently allocating the computational budget to regions of high geometric complexity, our method achieves superior FID and aesthetic quality, proving that *learning where to look* is as important as *learning how to match*. Our contributions are threefold:

- **Discovery of Geometric Bottlenecks:** We reveal that consistency distillation exhibits a boundary-dominated (*e.g.*, U-shaped) difficulty profile, contrary to the intermediate-focused priors of standard iterative generation.
- **Curvature-Adaptive Optimization via RL:** We propose CACFM, the first framework that dynamically adapts to the geometric heterogeneity of Flow Matching trajectories, replacing static heuristics with an agent.
- **Adaptive Scheduler with Complementary Distillation Losses:** CACFM leverages an RL scheduler and complementary distillation losses, outperforming existing acceleration methods on large-scale benchmarks.

2 Related Work

Flow Matching and Static Consistency Distillation. FM [1, 25, 26] has become a prominent framework for continuous-time generative modeling, using PF-ODEs to transform noise into data. Recent architectures such as FLUX [15] and SD3 [5] achieve remarkable quality but require expensive iterative sampling. To accelerate inference, CMs [41] and Latent CMs (LCMs [30]) map trajectories directly to solutions, while PCMs [46] split trajectories into sub-segments.

The Static Prior Limitation: Existing distillation methods rely on static sampling priors, from uniform sampling in CMs to Logit-Normal distributions [5], and thus cannot adapt to varying flow curvature. While adaptive sampling has been explored for data-level selection in broader training contexts [33, 50], CACFM extends this idea to trajectory-level policy learning over PF-ODE sub-trajectories.

Efficient Distillation. Recent works improve few-step generation from complementary angles. SANA-Sprint [4] combines consistency distillation with architectural and adversarial stabilizers, while ADM [29] uses adversarial distribution matching for image and video synthesis. Beyond diffusion, efficient student training has also been explored across diverse distillation settings [18, 19, 35].

Teacher Matching vs. Trajectory Allocation: These methods mainly improve *how to match* the teacher distribution. In contrast, CACFM focuses on *where to learn* along the trajectory, making our scheduler orthogonal to these loss-, teacher-, and architecture-centric advances.

RL in Generative Models. RL has been widely applied to diffusion models, mainly for downstream alignment. Methods such as DDPO [44] and DPOK [7] frame the denoising process as a Markov Decision Process to optimize objectives such as human preference scores (HPS) or image compressibility. More broadly, feedback-driven policy learning has also been explored for adaptive solver control and rendering-aware generation outside diffusion distillation [22, 31].

Structural vs. Alignment Optimization: Unlike these approaches, which use RL to optimize generation objectives (what to generate), CACFM employs RL to optimize training structure (how to learn). We treat sub-trajectory selection as a decision process to discover geometric bottlenecks in the flow.

3 Preliminaries

FM and Geometric Complexity. FM models define a continuous-time probability path $p_t(\mathbf{x})$ that transports the data distribution $p_0(\mathbf{x})$ to a simple prior $p_1(\mathbf{x}) = \mathcal{N}(0, \mathbf{I})$. Unlike standard diffusion formulations based on stochastic differential equations, FM is defined through ODE dynamics: $d\mathbf{x}_t/dt = \mathbf{v}(\mathbf{x}_t, t)$, $t \in [0, 1]$, where $\mathbf{v} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ is a time-dependent vector field that induces the flow map $\phi_t(\mathbf{x})$, transporting data to noise during the forward process.

Specifically within the context of Optimal Transport Flow Matching, the forward process is defined by a linear interpolation between a target data sample

$\mathbf{x}_0 \sim p_0$ and standard Gaussian noise $\mathbf{x}_1 \sim p_1$: $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$. This conditional path induces a constant conditional vector field $\mathbf{u}_t(\mathbf{x}|\mathbf{x}_0, \mathbf{x}_1) = \mathbf{x}_1 - \mathbf{x}_0$. The objective of FM is to train a neural network $\mathbf{v}_\phi(\mathbf{x}, t)$ to regress this target vector field by minimizing the expected mean squared error:

$$\mathcal{L}_{\text{FM}}(\phi) = \mathbb{E}_{t, \mathbf{x}_0, \mathbf{x}_1} [\|\mathbf{v}_\phi(\mathbf{x}_t, t) - (\mathbf{x}_1 - \mathbf{x}_0)\|^2].$$

Once trained, the generation (sampling) process involves solving the learned ODE backward from an initial noise sample $\mathbf{x}_1 \sim \mathcal{N}(0, \mathbf{I})$ to $t = 0$. The learned vector field $\mathbf{v}_\phi$ often exhibits geometric heterogeneity: it varies rapidly in high-curvature regions of the generative trajectory while remaining smooth in others. Consequently, accurately solving the learned ODE using standard numerical solvers requires many evaluations to traverse these geometrically complex regions.

Consistency Distillation and the Sampling Bias. CMs [41] accelerate sampling by learning to map any point $\mathbf{x}_t$ on the PF-ODE trajectory directly to its origin. To mitigate the difficulty of learning long-range mappings, PCMs [46] partition the trajectory into M sub-segments using boundaries $\{s_m\}_{m=0}^M$.

Ideally, a model learns a local consistency function $\mathbf{f}_\theta^m$ for each sub-trajectory $[s_m, s_{m+1}]$. The training objective utilizes a discretized grid $\{t_n\}_{n=0}^N$ and a one-step ODE solver estimate $\hat{\mathbf{x}}_{t_n}^\phi$, formulated as:

$$\mathcal{L}^{\text{PCM}}(\theta; \phi) = \mathbb{E}[\lambda(t_n) d(\mathbf{f}_\theta^m(\mathbf{x}_{t_{n+1}}, t_{n+1}), \mathbf{f}_{\theta^-}^m(\hat{\mathbf{x}}_{t_n}^\phi, t_n))] \quad (1)$$

where $d(\cdot, \cdot)$ denotes a distance metric, expectations are taken over $\mathbb{P}(m)$, $\mathbb{P}(n|m)$ and $\mathbb{P}(\mathbf{x}_{t_{n+1}}|n, m)$, *i.e.*, uniform distributions of sub-trajectory indices m , interval indices n , and distribution $\mathbb{P}_{t_{n+1}}$. The term θ^- denotes the exponential moving average of the parameters, updated via $\theta^- \leftarrow \mu\theta^- + (1-\mu)\theta$.

The key bottleneck in Eq. (1) is the sampling distribution $\mathbb{P}(m)$. Existing methods sample the sub-trajectory index m either uniformly or via a fixed Logit-Normal distribution, implicitly assuming that consistency difficulty is either homogeneous (Uniform) or concentrated in the middle (Logit-Normal). We argue that this static assumption is flawed. As we will show in Fig. 6, the true "consistency bottlenecks" are often located at the boundaries (*e.g.*, U-shaped).

4 Methodology: Curvature-Adaptive Optimization

The central hypothesis of this work is that the PF-ODE exhibits non-uniform geometric complexity: boundary regions require precise structural alignment, and optimization bottlenecks shift dynamically throughout distillation.

Static sampling strategies (Uniform, Logit-Normal, or even U-shaped priors) operate blindly, failing to adapt to this heterogeneity. They waste computational budget on linear, low-curvature segments while under-fitting highly curved turns.

To address this, we propose **CACFM** (Fig. 3), a framework that reformulates consistency distillation as a curvature-adaptive decision process. We employ an RL agent as a probe, using consistency error as a curvature proxy to dynamically detect high-yield trajectory segments and prioritize them during training.




Fig. 3: Training paradigm of CACFM.

4.1 RL as a Geometric Probe

We formulate the sub-trajectory selection as a Markov Decision Process (MDP). Unlike prior works that use RL for content alignment (*e.g.*, optimizing aesthetic scores), our agent optimizes the *training dynamics* itself.

Semantic Stage Abstraction: A key challenge in applying RL to trajectory optimization is the potential explosion of the state space. To resolve this, we introduce a semantic stage abstraction. By partitioning the continuous trajectory $t \in [0, 1]$ into a small number of coarse geometric stages (e.g., $M = 4$), we map the process into distinct semantic phases of generation: *Initialization*, *Structural Formation*, *Texture Filling*, and *Final Refinement*.

- **State Space $\mathcal{S}$ (Relative Difficulty Profile):** To make informed decisions, the agent must perceive the relative difficulty across stages. We define the state s_t as the ordinal ranking of the consistency losses $\{\mathcal{L}^{(1)}, \dots, \mathcal{L}^{(M)}\}$ associated with each sub-trajectory. *Why this works:* By limiting $M = 4$, the state space size is $|\mathcal{S}| = 4! = 24$. This compact representation transforms an intractable continuous control problem into a highly sample-efficient tabular Q-learning task, converging rapidly without overhead.
- **Action Space $\mathcal{A}$ (Curvature Focus):** At each step, the agent selects an action $a_t \in \{1, \dots, M\}$ for the next training iteration. This enables the model to allocate resources directly to the identified bottleneck.
- **Reward Function $\mathcal{R}$ (Baseline Advantage):** Directly using the raw loss difference as a reward is problematic because different sub-trajectories exhibit inherent variations in optimization difficulty (*i.e.*, some segments naturally have higher loss values). To address this, we adopt a baseline-based advantage strategy. We maintain an exponential moving average (EMA) of the consistency loss for each sub-trajectory m , denoted as $B_t(m)$. The reward r_t is defined as the improvement of the current consistency loss $\mathcal{L}_{\text{con}}$ over its historical baseline, scaled by a factor λ_r :

$$r_t = \lambda_r \cdot (B_t(a_t) - \mathcal{L}_{\text{con}}), \quad (2)$$

where the baseline is updated as $B_{t+1}(a_t) \leftarrow \beta B_t(a_t) + (1 - \beta)\mathcal{L}_{\text{con}}$ with $\lambda_r = 100$ and $\beta = 0.95$. This formulation encourages the agent to select

sub-trajectories that yield higher optimization benefits, effectively realizing a dynamic curriculum learning strategy.

- **Policy Update:** We use tabular Q-learning rather than a neural policy. After observing the next state s_{t+1} , the Q-table is updated with $\alpha = 0.1$ and $\gamma = 0.9$:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)],$$

We adopt an ϵ -greedy exploration policy, with ϵ linearly decayed from 1.0 to 0.1 during the first 20,000 training steps. This state-dependent update enables temporal credit assignment beyond myopic loss-proportional sampling.

4.2 Hybrid Objective for High-Fidelity Distillation

While RL optimizes *where* to learn, we must ensure the model learns *effectively* at the chosen locations. Relying solely on the consistency loss is prone to mode collapse or blurry artifacts in few-step regimes. To address this, we integrate DMD and Adversarial Consistency into the objective.

DMD. Inspired by [51], we incorporate a DMD objective to explicitly align the student’s generated distribution with the teacher’s target distribution. Specifically, we minimize the Kullback-Leibler (KL) divergence between the student’s marginal distribution $\mathbb{P}_{\theta}^{\text{student}}(\mathbf{x}_0)$ and the teacher’s distribution $\mathbb{P}_{\phi}^{\text{teacher}}(\mathbf{x}_0)$:

$$\mathcal{L}^{\text{DMD}} = D_{\text{KL}}(\mathbb{P}_{\theta}^{\text{student}}(\mathbf{x}_0) \parallel \mathbb{P}_{\phi}^{\text{teacher}}(\mathbf{x}_0)).$$

Following the derivation in [51], the gradient of this objective is approximated via the score difference evaluated at perturbed samples. Let $\mathbf{y} = \mathbf{f}_{\theta}(\mathbf{x}_t, t)$ denote the clean sample predicted by the student. To provide a stable optimization signal, we introduce a re-noising step to compute the gradient update:

$$\nabla_{\theta} \mathcal{L}^{\text{DMD}} \approx \mathbb{E}[(\mathbf{s}^{\text{teacher}}(\mathbf{x}_{\tau}) - \mathbf{s}^{\text{student}}(\mathbf{x}_{\tau})) \nabla_{\theta} \mathbf{f}_{\theta}(\mathbf{x}_t, t)], \quad (3)$$

where $\tau \sim \mathcal{U}(0, 1)$ is a re-noising timestep, $\epsilon \sim \mathcal{N}(0, I)$, and $\mathbf{x}_{\tau} = (1 - \tau)\mathbf{y} + \tau\epsilon$ is the re-noised sample (using the Flow Matching noise schedule). Here $\mathbf{s}^{\text{teacher}}$ and $\mathbf{s}^{\text{student}}$ are the score functions evaluated at the noisy state $\mathbf{x}_{\tau}$. Specifically, the computation involves three steps: 1) The student predicts a clean sample $\mathbf{y}$. 2) $\mathbf{y}$ is re-noised to $\mathbf{x}_{\tau}$. 3) The scores are computed at $\mathbf{x}_{\tau}$. For the student score $\mathbf{s}^{\text{student}}(\mathbf{x}_{\tau})$, we efficiently approximate it using the student model itself via the vector field relationship, avoiding the need for an extra discriminator network.

Score Estimation in FM. Applying the DMD loss requires evaluating the score function $\mathbf{s}_t(\mathbf{x}) = \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$. While traditional diffusion models predict the noise ϵ to estimate the score, our FM teacher predicts the vector field $\mathbf{v}_{\theta}$. Under the optimal transport probability path $p_t(\mathbf{x}|\mathbf{x}_0) = \mathcal{N}((1 - t)\mathbf{x}_0, t^2\mathbf{I})$, we leverage Tweedie’s formula to analytically derive the score function directly from the velocity output: $\mathbf{s}_{\theta}(\mathbf{x}_t, t) = -[\mathbf{x}_t + (1 - t)\mathbf{v}_{\theta}]/t$.

This rigorous transformation allows us to seamlessly compute the score difference for the DMD loss using exclusively the velocity outputs. The complete step-by-step mathematical derivation is provided in Appendix A.

Algorithm 1 Curvature-Adaptive Consistency Flow Matching (CACFM)

-
- 1: **Input:** Data $\mathcal{D}$, Teacher ϕ , Student θ , Solver Ψ , Sub-trajectory partitions $\{s_j\}$.
 - 2: **Initialize:** Target $\theta^- \leftarrow \theta$, Q-table $Q(s, a)$.
 - 3: **repeat**
 - 4: Sample batch $(\mathbf{z}, \mathbf{c}) \sim \mathcal{D}$.
 - 5: **Observe Geometry:** Construct state s_t by ranking current losses.
 - 6: **Select Curvature Focus:** Choose sub-trajectory $m \in \{1, \dots, M\}$ via ϵ -greedy policy based on $Q(s_t, \cdot)$.
 - 7: Sample t_{n+1} within segment m . Generate $\mathbf{x}_{t_{n+1}}$.
 - 8: Compute Teacher Target: $\mathbf{x}_{t_n}^\phi \leftarrow \text{SolverStep}(\Psi, \mathbf{x}_{t_{n+1}})$.
 - 9: Compute Student Prediction: $\tilde{\mathbf{x}}_{s_m} = \mathbf{f}_\theta^m(\mathbf{x}_{t_{n+1}}, t_{n+1})$.
 - 10: **Compute Hybrid Loss:**
 - 11: $\mathcal{L}_{\text{total}} = \mathcal{L}^{\text{PCM}} + 0.1\mathcal{L}^{\text{adv}} + 0.5\mathcal{L}^{\text{DMD}}$
 - 12: **Update Model:** $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{total}}$.
 - 13: Update Target: $\theta^- \leftarrow \mu\theta^- + (1 - \mu)\theta$.
 - 14: **Update Policy:**
 - 15: Compute Reward r_t (2). Update $Q(s_t, m)$ via Q-learning.
 - 16: **until** convergence
-

Adversarial Consistency Loss. To enhance distribution consistency in few-step generation, we introduce an adversarial loss defined as:

$$\mathcal{L}^{\text{adv}}(\theta, \theta^-; \phi, m) = \text{ReLU}(1 + \mathcal{D}(\tilde{\mathbf{x}}_s, \mathbf{c})) + \text{ReLU}(1 - \mathcal{D}(\hat{\mathbf{x}}_s, \mathbf{c}))$$

where $\tilde{\mathbf{x}}_s = \mathbf{f}_\theta^m(\mathbf{x}_{t_{n+1}}, t_{n+1}) + \epsilon_1$, $\hat{\mathbf{x}}_s = \mathbf{f}_{\theta^-}(\mathbf{x}_{t_n}^\phi, t_n) + \epsilon_2$ with ϵ_1 and ϵ_2 denoting noise perturbations. Here $\mathcal{D}$ is a discriminator conditioned on prompts $\mathbf{c}$.

This adversarial objective acts as a robust regularizer, ensuring that the consistency mapping strictly resides on the natural image manifold. Furthermore, this constraint performs a critical stabilizing function: as our DMD objective utilizes the student’s own score estimates (*i.e.*, self-distillation), the discriminator effectively prevents the model from drifting into degenerate solutions, counteracting the potential instability of bootstrapping.

4.3 Algorithm and Analysis of the Optimization Mechanism

The training procedure of **CACFM** is summarized in Algorithm 1 and illustrated in Fig. 3. Our framework operates in a dynamic closed loop: the RL agent perceives the geometric bottlenecks (State), selects the high-yield segment (Action), and the student model updates its flow via the hybrid loss (Consistency + Adv + DMD). This update alters the trajectory’s geometry, reshaping the consistency error profile (Reward) and guiding the agent’s next decision.

Crucially, to reduce variance from adversarial and distillation objectives, the reward r_t (2) is evaluated solely based on the structural consistency loss $\mathcal{L}_{\text{PCM}}$, while the full hybrid loss $\mathcal{L}_{\text{total}}$ is used for back-propagation. Unless otherwise stated, we set $\lambda_{\text{adv}} = 0.1$, $\lambda_{\text{DMD}} = 0.5$, and the PCM consistency weight to 1.0.

While prior works like [6] employed Policy Gradient methods primarily for *alignment* (optimizing a specific reward model or discriminator score), our

CACFM framework utilizes Q-learning for **structural discovery**. Our agent does not merely optimize the output; it probes the PF-ODE trajectory to identify and resolve geometric bottlenecks that hinder consistent distillation.

By integrating reinforcement learning with the geometric properties of Flow Matching, our framework transcends standard acceleration techniques. The proposed approach yields three critical advantages:

- **Geometry-Awareness:** Unlike static strategies that function blindly, our agent can adapt to training characteristics, automatically allocating computational budget to high-error, high-curvature trajectory segments.
- **Emergent Curriculum:** The reward-guided optimization naturally induces a *coarse-to-fine* training curriculum. Our empirical analysis reveals that the agent autonomously shifts its focus from global structure initialization at the trajectory boundaries to high-frequency detail refinement, mirroring human learning processes without manual scheduling.
- **Holistic Stability:** By combining the exploration capability of RL with the strict constraints of Adversarial and DMD losses, CACFM avoids the "reward hacking" often seen in RL, ensuring that the accelerated trajectory remains faithful to the teacher’s original distribution.

Computational Efficiency. Crucially, the computational footprint introduced by our RL agent is negligible. Since the state space is low-dimensional (discrete loss rankings) and the Q-learning update is computationally lightweight, the additional cost is virtually non-existent compared to the gradient backpropagation through the massive Transformer backbones (*e.g.*, FLUX with 12B parameters). Thus, CACFM achieves acceleration without incurring extra training latency.

5 Experiments

5.1 Experimental Setup

Dataset Configuration. We conduct our image generation experiments using the large-scale LAION dataset [38] for model training. For evaluation purposes, we utilize a randomly selected 15,000-image subset from the CC3M dataset [2], ensuring a representative sample for performance assessment.

Model Architecture. All text-to-image generation experiments are built upon FLUX [15] and Stable Diffusion XL (SDXL [32]) as our base architectures. This choice provides a robust foundation for comparing various acceleration techniques across different flow formulations (*i.e.*, Rectified Flow vs. DDPM).

Evaluation Protocol. We employ a comprehensive suite of metrics to rigorously assess both generation quality and prompt adherence:

- Image Quality: FID [11] for distribution similarity.
- Human Preference: HPSv2 [48] and PickScore [14].
- Aesthetic Quality: LAION Aesthetic Score [38].

All metrics are uniformly computed on the 15K CC3M validation split to ensure strict, consistent, and fair benchmarking across all compared methods.

RL Hyperparameter Configuration. To strike an optimal balance between routing granularity and computational overhead, we partition the generative trajectory into $M = 4$ discrete segments for the RL agent. Consequently, the state space cardinality is limited to $|S| = M! = 24$. We validate this choice by varying M in Section 5.5 and measure the computational overhead of the $M = 4$ setting in Section 5.4. The results show that this compact state space enables rapid policy convergence with negligible overhead compared to backbone training.

5.2 Comparative Analysis

Baseline Methods. We compare CACFM against different SOTA approaches:

- Consistency Models: LCM [30], Hyper-SD [36], PCM [46], TDD [45], TCD [58].
- Flow-based Methods: InstaFlow (Insta. [27]).
- Adversarial Distillation: Lightning (Light. [24]), Turbo [37], Schnell [15].

Qualitative Results. Figures 4 and 5 present comprehensive visual comparisons. As illustrated, CACFM demonstrates superior capability in handling complex geometries. While baseline methods (*e.g.*, Turbo, PCM) frequently suffer from structural collapse or severe texture blurring under few-step regimes, which are indicative of under-optimized high-curvature segments, CACFM maintains:

- **Structural Integrity:** Significantly lower structural artifacts, suggesting the RL agent successfully reinforced the learning of difficult boundary conditions.
- **High-Frequency Detail:** Better preservation of textures, attributed to the emergent curriculum that prioritizes refinement in later training stages.

Quantitative Results.

1) FID Performance: Tables 1 and 2 present the FID scores evaluated on the CC3M dataset. Our CACFM achieves state-of-the-art results across all configurations. Notably, the performance gap is most pronounced in the extreme low-step regime (4-Step), where CACFM outperforms FLUX-schnell by a substantial margin of over 2 FID points. This empirical evidence confirms our





Fig. 4: Qualitative comparisons with LoRA-based approaches on FLUX architecture.

Table 1: FID Comparison on CC3M (FLUX). Best results are **bolded**.

Methods	FID ($\downarrow$)		
	4-Step	8-Step	16-Step
Turbo [37]	58.22	46.24	43.66
Hyper-SD [36]	45.35	43.86	43.40
TDD [45]	45.81	41.65	41.12
Schnell [15]	41.19	40.47	39.89
CACFM(Ours)	39.19	36.96	37.62

Table 2: FID Comparison on CC3M (SDXL). Best results are **Bold**.

Steps	FID ($\downarrow$) Comparison across Different Methods								
	Light. [24]	Turbo [37]	LCM [30]	Hyper-SD [36]	PCM [46]	Insta. [27]	TDD [45]	TCD [58]	Ours
4	37.49	52.90	45.57	39.43	37.26	38.13	41.75	46.40	35.29
8	38.28	65.25	43.67	41.63	39.30	35.60	46.00	49.51	34.42
16	40.22	77.13	43.33	44.12	40.47	34.43	51.22	54.68	33.49

**Fig. 5:** Qualitative comparisons with LoRA-based approaches on SDXL architecture.

hypothesis: when the inference budget is constrained, curvature-adaptive training is essential to resolve geometric bottlenecks that uniform sampling misses.

- Consistency Model Comparison: CACFM outperforms LCM, PCM, and TCD, validating the benefit of dynamic trajectory optimization.
- SOTA Benchmarking: Our approach rivals or surpasses Adversarial Distillation (Turbo) and Rectified-Flow (Insta.) methods, setting a new efficiency standard.

2) Human-Centric Evaluation: Tables 3 and 4 report metrics assessing perceptual quality and text alignment.

- Consistent Performance: CACFM achieves top-tier performance (consistently ranking 1st or 2nd) across the vast majority of configurations.
- Robustness: Our approach maintains stable improvements across increasing steps. Crucially, the high Aesthetic and PickScores suggest that the emergent curriculum effectively transitions focus to fine-grained details (Phase 3) in later training stages, directly boosting perceptual fidelity.

Inference Scalability. As presented in Tables 1 and 2, CACFM exhibits strong zero-shot scalability. Although the RL policy is trained with a fixed granularity of $M = 4$ segments, it consistently improves generation quality when extrapolated to more inference steps ($N = 8$ & 16). For instance, on the FLUX backbone, the FID score substantially improves from 39.19 (4-step) to 36.96 (8-step). This suggests that RL-guided optimization better regularizes the probability flow trajectory, allowing users to trade computation for perceptual quality at inference time. In contrast, conventional distillation baselines frequently overfit to their specific training step counts, failing to capitalize on increased inference budgets.

Table 3: Aesthetic evaluation on FLUX. Best results are **Bold**.

Methods	Step 4			Step 8			Step 16		
	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore
Turbo	0.1293	5.5666	17.1633	0.1372	5.7197	17.0438	0.2658	5.7665	21.3205
Hyper-SD	0.2626	5.6734	21.2479	0.2747	5.7120	21.3677	0.2757	5.7617	21.4478
TDD	0.2454	5.6165	20.7101	0.2610	5.6576	20.9471	0.2609	5.6383	20.8932
Schnell	0.2713	5.5204	21.3134	0.2696	5.4914	21.2159	0.2671	5.4802	21.1154
Ours ^(rank)	0.2638 ⁽²⁾	5.7955⁽¹⁾	21.3160⁽¹⁾	0.2683 ⁽³⁾	5.7961⁽¹⁾	21.3818⁽¹⁾	0.2754 ⁽²⁾	5.7967⁽¹⁾	21.4520⁽¹⁾

Table 4: Aesthetic evaluation on SDXL: Our method consistently achieves top-tier performance, often ranking among the top three or even attaining the best results. Best results are **Bold**, and the superscript indicates the rank of our method.

Methods	Step 4			Step 8			Step 16		
	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore
Lightning	0.2666	5.7135	21.3174	0.2721	5.8287	21.2698	0.2660	5.8600	21.0609
Turbo	0.2587	5.3267	20.6089	0.2471	5.2256	20.2537	0.2393	5.1566	20.0273
Hyper-SD	0.2855	5.8806	21.2701	0.2862	5.9284	21.4551	0.2898	5.9372	21.4797
LCM	0.2431	5.4165	20.8978	0.2493	5.4680	20.9471	0.2473	5.4863	20.8295
PCM	0.2663	5.6441	21.0911	0.2731	5.7256	21.1061	0.2704	5.7573	20.9734
InstaFlow	0.2472	5.5360	21.0710	0.2522	5.5813	21.1527	0.2560	5.6136	21.1934
TDD	0.2609	5.7519	20.9910	0.2602	5.8571	20.8012	0.2511	5.8673	20.4932
TCD	0.2576	5.5966	20.7705	0.2543	5.6558	20.5408	0.2450	5.6267	20.2597
Ours ^(rank)	0.2764 ⁽²⁾	5.8931⁽¹⁾	21.2241 ⁽³⁾	0.2875⁽¹⁾	5.9423⁽¹⁾	21.3321 ⁽²⁾	0.2932⁽¹⁾	5.9823⁽¹⁾	21.5532⁽¹⁾

5.3 Analysis: Uncovering Geometric Heterogeneity

This section presents the experimental results that systematically validate the core observations outlined at the beginning of the paper.

U-Shaped difficulty profile. We computed the *Oracle Consistency Error* (proxy for ground truth difficulty) across the trajectory using a converged SDXL teacher model. Fig. 6 (Left) reveals a profound geometric insight:

- The ground truth error follows a "U-shaped" profile, indicating that the *boundary conditions* initialization ($t \rightarrow 0$) and final refinement ($t \rightarrow 1$) are the actual geometric bottlenecks.
- Our RL agent’s learned policy (Red Line) perfectly correlates with this ground truth ($\rho > 0.95$), automatically shifting focus to the boundaries.

The widely used Logit-Normal prior is geometrically misaligned for consistency distillation. CACFM outperforms it by simply "learning where to look."

Emergent Curriculum Learning. Fig. 6 (Right) visualizes the temporal evolution of the agent’s policy. We observe an emergent coarse-to-fine curriculum:

- *Early Training (0–20k steps)*: The agent prioritizes the $t \rightarrow 0$ phase (Phase 0), focusing heavily on global structural formation.
- *Late Training (20k+ steps)*: The focus autonomously shifts towards $t \rightarrow 1$ (Phase 3), dedicating capacity to refining high-frequency details and textures. Crucially, this sophisticated behavior emerges purely from reward-driven optimization, entirely without any manual heuristic scheduling.




Fig. 6: Analysis of the Learned Policy. *Left:* The agent’s selection frequency (Red) perfectly aligns with the U-shaped optimization difficulty profile (Blue Bars), validating its geometry-awareness. *Right:* The selection probability heatmap shows an emergent curriculum, shifting focus from initialization to refinement (Phase 0 to 3) over time.

Table 5: Wall-Clock Time Efficiency Comparison. We compare the FID scores of PCM and our CACFM at fixed training intervals. Although CACFM incurs a slight overhead per iteration ($\sim 1.18\times$ relative time per step), the RL-guided trajectory selection accelerates convergence significantly.

Methods	Relative Time per Step	FID ($\downarrow$) on CC3M over Time		
		12 Hours	24 Hours	36 Hours
PCM	1.00 $\times$	46.52	43.10	41.65
CACFM (Ours)	1.18 $\times$	44.15	40.82	39.19

5.4 Training Efficiency and Wall-Clock Time Analysis.

A potential concern regarding the introduction of the RL agent and discriminator is the computational overhead per training iteration. To evaluate the true training efficiency, we compare CACFM against a strong baseline, PCM [46], under a strict *fixed wall-clock time* budget, rather than a fixed number of iterations. This comparison accounts for the additional forward/backward passes introduced by our auxiliary modules. We conducted training on 8 accelerators with 80GB device memory each. As reported in Table 5, we measured the FID on the CC3M validation set at different training durations (12h, 24h, and 36h).

Results and Discussion: The empirical results in Table 5 demonstrate that CACFM consistently maintains a superior FID-to-time ratio compared to PCM across all checkpoints. Notably, CACFM achieves an FID of 40.82 within just 24 hours, effectively surpassing the performance of PCM at 36 hours (FID 41.65).

While CACFM introduces a marginal computational overhead of approximately 18% per iteration due to the RL agent and DMD loss updates, the resulting gains in *data efficiency* are substantial. The RL agent effectively acts as a dynamic curriculum learner, prioritizing high-error sub-trajectories and avoiding redundant training on easy segments. This accelerates the overall convergence rate, making CACFM more time-efficient in practical training scenarios despite the increased model complexity.

Table 6: Ablation Study: adversarial vs. DMD loss. Avg Rank ($\downarrow$) denotes the mean ranking across all 9 evaluation dimensions.

Methods	Step 4			Step 8			Step 16			Avg Rank ($\downarrow$)
	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore	HPS	Aesthetic	PickScore	
Naive CFM (Uniform)	0.236	5.432	20.69	0.264	5.602	21.12	0.272	5.646	21.16	4.11
CFM Logit-Normal	0.232	5.458	20.70	0.265	5.598	20.21	0.270	5.637	21.11	4.89
CFM Loss-Aware	0.236	5.436	20.77	0.263	5.588	21.16	0.270	5.659	21.17	4.00
CACFM w/o DMD	0.239	5.485	20.78	0.263	5.594	21.16	0.270	5.623	21.18	3.78
CACFM w/o RL	0.235	5.430	20.71	0.265	5.610	21.20	0.271	5.658	21.23	3.22
CACFM (Ours)	0.276	5.893	21.22	0.288	5.942	21.33	0.293	5.982	21.55	1.00

Table 7: Comparison of dynamic adaptive sampling strategies on FLUX 4-step generation. Loss-Aware and MAB adapt online but treat stages myopically, while CACFM uses Q-learning over ranking states.

Sampling Strategy	HPS	Aesthetic	PickScore
Loss-Aware (EMA)	0.236	5.436	20.77
MAB (ϵ -greedy)	0.242	5.515	20.84
CACFM (Ours)	0.276	5.893	21.22

Table 8: Hyperparameter ablation on the number of semantic stages M under the 4-step setting. $M = 4$ achieves the best FID by balancing trajectory granularity and exploration cost.

Stages	State Space	FID ($\downarrow$)
$M = 3$	$3! = 6$	39.85
$M = 4$	$4! = 24$	39.19
$M = 6$	$6! = 720$	39.51

5.5 Ablation Study

Static vs. Adaptive Geometric Priors. To isolate the benefit of adaptive RL, we compare CACFM against the following static sampling strategies in Table 6.

- **Naive CFM (Uniform)** and **Logit-Normal CFM** serve as baselines for fixed geometric assumptions.
- **CFM Loss-Aware** is a dynamic loss-proportional sampler that updates stage probabilities using the EMA of per-stage consistency losses; it lacks the state transition modeling and temporal credit assignment of Q-learning.

The results show that CACFM outperforms both uniform and heuristic-based methods. Notably, CFM Logit-Normal (which prioritizes intermediate steps) underperforms relative to Naive CFM, further validating our finding that the true difficulty profile is U-shaped, not bell-shaped. Compared with dynamic but myopic samplers such as Loss-Aware and MAB (Table 7), CACFM further benefits from state-dependent long-horizon scheduling.

Impact of Hybrid Loss. We evaluate the contribution of the DMD loss component. As shown in the "w/o DMD loss" rows, removing distribution alignment leads to a noticeable drop in aesthetic scores. This indicates that while RL effectively optimizes the structural consistency (fidelity), the DMD loss is crucial for aligning the student's fine-grained texture distribution with the teacher's, validating the synergy of our hybrid objective.

Sensitivity to the Number of Semantic Stages. The number of semantic stages M determines the granularity at which the flow trajectory is ranked and routed by the

Q-table policy. A smaller M gives a compact state space and fast exploration, but it can merge intervals whose consistency losses follow different dynamics. A larger M provides a finer ordering of local bottlenecks, but the induced permutation state space grows factorially and can make early exploration less reliable.

We therefore validate this design choice by varying the number of semantic stages while keeping the 4-step sampling setting fixed. As shown in Table 8, $M = 4$ achieves the best balance between routing resolution and exploration cost: $M = 3$ is too coarse and merges distinct bottlenecks, whereas $M = 6$ enlarges the state space to 720 states and slows early Q-table exploration. This trend supports using a moderate discretization of the trajectory rather than treating more stages as monotonically better.

6 Conclusions and Limitations

Conclusions. We presented CACFM, a framework that reformulates consistency distillation as geometric trajectory optimization. By employing an RL agent to dynamically target high-curvature bottlenecks and integrating a unified Flow-DMD objective, CACFM induces an emergent training curriculum that moves beyond static sampling priors. This adaptive strategy improves trajectory learning, mitigates structural artifacts, and achieves state-of-the-art few-step generation on large-scale FLUX and SDXL backbones.

Limitations. As a distillation approach, CACFM’s generative quality is inherently bounded by the teacher model’s vector field manifold and by the expressiveness of the student model used to approximate it. Moreover, although adaptive trajectory routing improves few-step generation, extreme compression to single-step inference remains challenging. Highly curved or topologically complex flows cannot always be faithfully approximated by a single linear transition, which may lead to minor artifacts or loss of fine-grained details in difficult samples.

Acknowledgements

We thank the members of Huawei Large Model Data Technology Lab for their valuable suggestions and support throughout this work. The authors also benefited from the academic environment and exchanges at the Department of Statistics and Data Science, Tsinghua University, and are grateful to the faculty members and fellow students for their guidance, insights, and helpful discussions. We also thank the anonymous reviewers and area chairs for their insightful and constructive comments, which helped improve the quality of this work.

References

1. Albergo, M.S., Vanden-Eijnden, E.: Building normalizing flows with stochastic interpolants. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=li7qeBbCR1t>, accessed: 27 Jun 2026

2. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In: Conference on Computer Vision and Pattern Recognition. pp. 3558–3568 (2021), <https://ieeexplore.ieee.org/document/9578388>, accessed: 27 Jun 2026
3. Chen, G., Li, Y., Lin, Q.: On the impacts of the random initialization in the neural tangent kernel theory. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024), <https://openreview.net/forum?id=ni3Ud2BV3G>, accessed: 27 Jun 2026
4. Chen, J., Xue, S., Zhao, Y., Yu, J., Paul, S., Chen, J., Cai, H., Han, S., Xie, E.: Sana-sprint: One-step diffusion with continuous-time consistency distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16185–16195 (October 2025), <https://doi.org/10.1109/ICCV51701.2025.01502>, accessed: 27 Jun 2026
5. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Rombach, R.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first International Conference on Machine Learning (2024), <https://openreview.net/forum?id=FPnUhsQJ5B>, accessed: 27 Jun 2026
6. Fan, Y., Lee, K.: Optimizing DDPM sampling with shortcut fine-tuning. In: Proceedings of the 40th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 9623–9639. PMLR (23–29 Jul 2023), <https://proceedings.mlr.press/v202/fan23b.html>, accessed: 27 Jun 2026
7. Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., Lee, K.: Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. In: Advances in Neural Information Processing Systems. vol. 36, pp. 79858–79885. Curran Associates, Inc. (2023), <https://arxiv.org/abs/2305.16381>, accessed: 27 Jun 2026
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016), https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html, accessed: 27 Jun 2026
9. He, Y., Zhang, C., Chen, F., Cao, J.: Cinematte: Background matting for virtual production and beyond. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Findings. pp. 8725–8735 (June 2026), <https://arxiv.org/abs/2605.18328>, accessed: 27 Jun 2026
10. He, Y., Zhang, W., Deng, J., Cong, Y.: Prior-knowledge-free video frame interpolation with bidirectional regularized implicit neural representations. In: MultiMedia Modeling. pp. 112–126. Springer Nature Switzerland, Cham (2024), https://link.springer.com/chapter/10.1007/978-3-031-53311-2_9, accessed: 27 Jun 2026
11. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in Neural Information Processing Systems **30** (2017), <https://dl.acm.org/doi/10.5555/3295222.3295408>, accessed: 27 Jun 2026
12. Hu, Y., Li, Z., Chen, Z., Huang, Q., Fu, Z., Xu, M., Nie, L.: Refine: Composed video retrieval via shared and differential semantics enhancement. ACM Trans. Multimedia Comput. Commun. Appl. (Feb 2026), <https://doi.org/10.1145/3796712>, accessed: 27 Jun 2026

13. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* **35**, 26565–26577 (2022), <https://dl.acm.org/doi/10.5555/3600270.3602196>, accessed: 27 Jun 2026
14. Kirstain, Y., Polyak, A., Singer, U., Matiana, S., Penna, J., Levy, O.: Pick-a-pic: An open dataset of user preferences for text-to-image generation. In: *Thirty-seventh Conference on Neural Information Processing Systems* (2023), <https://openreview.net/forum?id=G5RwHpBUv0>, accessed: 27 Jun 2026
15. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux> (2024), accessed: 27 Jun 2026
16. Lai, J., Yu, Z., Tian, S., Lin, Q.: Generalization ability of wide residual networks (2023), <https://arxiv.org/abs/2305.18506>, accessed: 27 Jun 2026
17. Li, Y., Yu, Z., Chen, G., Lin, Q.: On the eigenvalue decay rates of a class of neural-network related kernel functions defined on general domains. *Journal of Machine Learning Research* **25**(82), 1–47 (2024), <http://jmlr.org/papers/v25/23-0866.html>, accessed: 27 Jun 2026
18. Li, Y., Yang, C., Dong, J., Yao, Z., Xu, H., Dong, Z., Zeng, H., An, Z., Tian, Y.: Ammkd: Adaptive multimodal multi-teacher distillation for lightweight vision-language models (2025), <https://arxiv.org/abs/2509.00039>, accessed: 27 Jun 2026
19. Li, Y., Yang, C., Zeng, H., Dong, Z., An, Z., Xu, Y., Tian, Y., Wu, H.: Frequency-aligned knowledge distillation for lightweight spatiotemporal forecasting. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7262–7272 (2025), <https://doi.org/10.1109/ICCV51701.2025.00682>, accessed: 27 Jun 2026
20. Li, Z., Hu, Y., Chen, Z., Huang, Q., Qiu, G., Fu, Z., Liu, M.: Retrack: Evidence-driven dual-stream directional anchor calibration network for composed video retrieval. In: *AAAI*. vol. 40, pp. 23373–23381 (2026), <https://arxiv.org/abs/2604.17898>, accessed: 27 Jun 2026
21. Liang, G., Wang, Z., Hu, J., Zhou, H., Xue, Z., Zhang, J., Xu, D., Yu, Q.: Render-in-the-loop: Vector graphics generation via visual self-feedback (2026), <https://arxiv.org/abs/2604.20730>, accessed: 27 Jun 2026
22. Liang, G., Wang, Z., Wang, C., Hu, J., Zhou, H., Liu, J., Zhang, J., Xu, D., Yu, Q.: Vanim: Rendering-aware sparse state modeling for structure-preserving vector animation (2026), <https://arxiv.org/abs/2605.01517>, accessed: 27 Jun 2026
23. Lin, J., Zhang, J., Jin, G., Song, W., Liu, T., Lu, G.: 3d plant root skeleton detection and extraction. In: *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 3011–3017. IEEE (2025), <https://arxiv.org/abs/2508.08094>, accessed: 27 Jun 2026
24. Lin, S., Wang, A., Yang, X.: Sd-xl-lightning: Progressive adversarial diffusion distillation (2024), <https://arxiv.org/abs/2402.13929>, accessed: 27 Jun 2026
25. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=PqvMRDCJT9t>, accessed: 27 Jun 2026
26. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=XVjTT1nw5z>, accessed: 27 Jun 2026

27. Liu, X., Zhang, X., Ma, J., Peng, J., Liu, Q.: Instaflo: One step is enough for high-quality diffusion-based text-to-image generation. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=1k4yZbbDqX>, accessed: 27 Jun 2026
28. 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 (2022), https://openreview.net/forum?id=2uAaGw1P_V, accessed: 27 Jun 2026
29. Lu, Y., Ren, Y., Xia, X., Lin, S., Wang, X., Xiao, X., Ma, A.J., Xie, X., Lai, J.H.: Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16818–16829 (October 2025), <https://doi.org/10.1109/ICCV51701.2025.01562>, accessed: 27 Jun 2026
30. Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference (2023), <https://arxiv.org/abs/2310.04378>, accessed: 27 Jun 2026
31. Mai, J., Liao, B., Zhao, Z., Zeng, Y., Li, H., Civera, J., Wu, T., Zhou, Y., Liu, P.: Neural predictor-corrector: Solving homotopy problems with reinforcement learning (2026), <https://arxiv.org/abs/2602.03086>, accessed: 27 Jun 2026
32. 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>, accessed: 27 Jun 2026
33. Rao, J., Liu, X., Deng, H., Lin, Z., Yu, Z., Wei, J., Meng, X., Zhang, M.: Dynamic sampling that adapts: Self-aware iterative data persistent optimization for mathematical reasoning. In: Findings of the Association for Computational Linguistics: ACL 2026. pp. 28337–28350. Association for Computational Linguistics (Jul 2026), <https://aclanthology.org/2026.findings-acl.1412/>, accessed: 27 Jun 2026
34. Rao, J., Liu, X., Yan, H., Shen, J., Mo, H., Dong, Y., Yan, Z., Wang, Z., Lin, Z., Meng, X., Yu, Z., Deng, L., Wei, J., Wang, Y., Zhang, M.: A data-centric perspective on the lifecycle of large language models. TechRxiv **2025**(1220) (2025), <https://www.techrxiv.org/doi/abs/10.36227/techrxiv.176620610.03288677/v1>, accessed: 27 Jun 2026
35. Rao, J., Meng, X., Ding, L., Qi, S., Liu, X., Zhang, M., Tao, D.: Parameter-efficient and student-friendly knowledge distillation. IEEE Transactions on Multimedia **26**, 4230–4241 (2024). <https://doi.org/10.1109/TMM.2023.3321480>, accessed: 27 Jun 2026
36. Ren, Y., Xia, X., Lu, Y., Zhang, J., Wu, J., Xie, P., WANG, X., Xiao, X.: HyperSD: Trajectory segmented consistency model for efficient image synthesis. In: The Thirty-eighth Annual Conference on Neural Information Processing Systems (2024), <https://openreview.net/forum?id=05Xb0oi0x3>, accessed: 27 Jun 2026
37. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation (2023), <https://arxiv.org/abs/2311.17042>, accessed: 27 Jun 2026
38. 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. In: Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2022), <https://openreview.net/forum?id=M3Y74vmsMcY>, accessed: 27 Jun 2026

39. Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y.K., Wu, Y., Guo, D.: Deepseekmath: Pushing the limits of mathematical reasoning in open language models (2024), <https://arxiv.org/abs/2402.03300>, accessed: 27 Jun 2026
40. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: Proceedings of the 32nd International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 37, pp. 2256–2265. PMLR, Lille, France (07–09 Jul 2015), <https://proceedings.mlr.press/v37/sohl-dickstein15.html>, accessed: 27 Jun 2026
41. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. In: Proceedings of the 40th International Conference on Machine Learning. pp. 32211–32252 (2023), <https://arxiv.org/abs/2303.01469>, accessed: 27 Jun 2026
42. 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 (2021), <https://arxiv.org/abs/2011.13456>, accessed: 27 Jun 2026
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017), https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf, accessed: 27 Jun 2026
44. Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Purushwalkam, S., Ermon, S., Xiong, C., Joty, S., Naik, N.: Diffusion model alignment using direct preference optimization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8228–8238 (June 2024), <https://arxiv.org/abs/2311.12908>, accessed: 27 Jun 2026
45. Wang, C., Guo, Z., Duan, Y., Li, H., Chen, N., Tang, X., Hu, Y.: Target-driven distillation: Consistency distillation with target timestep selection and decoupled guidance. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 7619–7627 (2025), <https://arxiv.org/abs/2409.01347>, accessed: 27 Jun 2026
46. Wang, F.Y., Huang, Z., Bergman, A., Shen, D., Gao, P., Lingelbach, M., Sun, K., Bian, W., Song, G., Liu, Y., et al.: Phased consistency models. Advances in neural information processing systems **37**, 83951–84009 (2024), <https://arxiv.org/abs/2405.18407>, accessed: 27 Jun 2026
47. Wu, W., Li, Y., Chen, G., Wang, L., Chen, H.: Tool-augmented policy optimization: Synergizing reasoning and adaptive tool use with reinforcement learning (2026), <https://arxiv.org/abs/2510.07038>, accessed: 27 Jun 2026
48. 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 (2023), <https://arxiv.org/abs/2306.09341>, accessed: 27 Jun 2026
49. Xiao, C., Dou, J., Lin, Z., Ke, Z., Hou, L.: From points to coalitions: Hierarchical contrastive shapley values for prioritizing data samples. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 40, pp. 15995–16003 (2026), <https://ojs.aaai.org/index.php/AAAI/article/view/38633>, accessed: 27 Jun 2026
50. Yao, J., Li, C., Xiao, C.: Swift sampler: Efficient learning of sampler by 10 parameters. In: Advances in Neural Information Processing Systems. vol. 37, pp. 59030–59053. Curran Associates, Inc. (2024), <https://arxiv.org/abs/2410.05578>, accessed: 27 Jun 2026

51. Yin, T., Gharbi, M., Zhang, R., Shechtman, E., Durand, F., Freeman, W.T., Park, T.: One-step diffusion with distribution matching distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6613–6623 (2024), <https://arxiv.org/abs/2311.18828>, accessed: 27 Jun 2026
52. Yu, B., Liu, D., Shi, H., Chang, G., Wei, J., Sun, L., Tian, S., Bu, L.: Sam-wav2lip++: Enhancing behavioral realism in synthetic agents through audio-driven speech and action refinement. In: 2024 IEEE International Conference on Systems, Man, and Cybernetics (SMC). pp. 2999–3006 (2024), <https://ieeexplore.ieee.org/abstract/document/10832087>, accessed: 27 Jun 2026
53. Yu, Z., Chen, G., Lai, J., Li, B., Tian, S.: Branch scaling manifests as implicit architectural regularization for improving generalization in overparameterized resnets. In: Forty-third International Conference on Machine Learning (2026), <https://openreview.net/forum?id=iWV5o07Bxv>, accessed: 27 Jun 2026
54. Yu, Z., Rao, J., Chen, G., Tian, S., Li, B., Wei, J., Zhang, M., Meng, X.: MathAgent: Adversarial evolution of constraint graphs for mathematical reasoning data synthesis. In: Findings of the Association for Computational Linguistics: ACL 2026. pp. 28298–28321. Association for Computational Linguistics (Jul 2026), <https://aclanthology.org/2026.findings-acl.1410/>, accessed: 27 Jun 2026
55. Yu, Z., Tian, S., Chen, G.: Divergence of neural tangent kernel in classification problems. In: The Thirteenth International Conference on Learning Representations (2025), <https://openreview.net/forum?id=VEJzjAvaIy>, accessed: 27 Jun 2026
56. Zhang, J., Lu, G.: Vision-language embodiment for monocular depth estimation. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 29479–29489 (2025), <https://arxiv.org/abs/2503.16535>, accessed: 27 Jun 2026
57. Zhao, Z., Yang, H., Liao, B., Zeng, Y., Yan, S., Gu, Y., Liu, P., Zhou, Y., Li, H., Civera, J.: Advances in global solvers for 3d vision (2026), <https://arxiv.org/abs/2602.14662>, accessed: 27 Jun 2026
58. 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 (2024), <https://arxiv.org/abs/2402.19159>, accessed: 27 Jun 2026