Novelty Assessment Report

Paper: BézierFlow: Learning Bézier Stochastic Interpolant Schedulers for Few-Step Generation

PDF URL: https://openreview.net/pdf?id=PCuDI32xhQ

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

We introduce BézierFlow, a lightweight training approach for few-step generation with pretrained diffusion and flow models. BézierFlow achieves a 2–3× performance improvement for sampling with $\leq$ 10 NFEs while requiring only 15 minutes of training. Recent lightweight training approaches have shown promise by learning optimal timesteps, but their scope remains restricted to ODE discretizations. To broaden this scope, we propose learning the optimal transformation of the sampling trajectory by parameterizing stochastic interpolant (SI) schedulers. The main challenge lies in designing a parameterization that satisfies critical desiderata, including boundary conditions, differentiability, and monotonicity of the SNR. To effectively meet these requirements, we represent scheduler functions as Bézier functions, where control points naturally enforce these properties. This reduces the problem to learning an ordered set of points in the time range, while the interpretation of the points changes from ODE timesteps to Bézier control points. Across a range of pretrained diffusion and flow models, BézierFlow consistently outperforms prior timestep-learning methods, demonstrating the effectiveness of expanding the search space from discrete timesteps to Bézier-based trajectory transformations.

Core Task Landscape

This paper addresses: **Few-Step Generation with Pretrained Diffusion and Flow Models**

A total of **50 papers** were analyzed and organized into a taxonomy with **20 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Distillation-Based Acceleration Methods**
- **Sampling Trajectory Optimization**
- **Architectural and Training Innovations**
- **Parallel and Non-Autoregressive Sampling**
- **Domain-Specific Applications**
- **Conditional Generation and Inverse Problems**

Complete Taxonomy Tree

- Few-Step Generation with Pretrained Diffusion and Flow Models Survey Taxonomy
- Distillation-Based Acceleration Methods
  - Consistency Distillation Approaches (7 papers)
  - [10] Latent consistency models: Synthesizing high-resolution images with few-step inference (Luo, 2023) View paper
  - [12] Sana-sprint: One-step diffusion with continuous-time consistency distillation (Chen Jun-song, 2025) View paper
  - [14] CMT: Mid-Training for Efficient Learning of Consistency, Mean Flow, and Flow Map Models (Hu, 2025) View paper
  - [17] SCott: Accelerating Diffusion Models with Stochastic Consistency Distillation (Liu HongJian, 2025) View paper
  - [32] Flow map matching with stochastic interpolants: A mathematical framework for consistency models (Boffi, 2024) View paper
  - [48] Align Your Flow: Scaling Continuous-Time Flow Map Distillation (Sabour, 2025) View paper
  - [49] How to build a consistency model: Learning flow maps via self-distillation (Boffi, 2025) View paper
  - Trajectory-Preserving Distillation (3 papers)
  - [2] On distillation of guided diffusion models (Chenlin Meng, 2023) View paper
  - [8] Flash diffusion: Accelerating any conditional diffusion model for few steps image generation (Chadebec, 2025) View paper
  - [37] Toward Theoretical Insights into Diffusion Trajectory Distillation via Operator Merging (Gao Wei-guo, 2025) View paper
  - Trajectory-Reformulation Distillation (2 papers)
  - [38] Hyper-sd: Trajectory segmented consistency model for efficient image synthesis (Ren Yuxi, 2024) View paper
  - [43] Simple ReFlow: Improved Techniques for Fast Flow Models (Kim Beom-Su, 2024) View paper
  - Hybrid Distillation Frameworks (8 papers)
  - [4] Adversarial diffusion distillation (Sauer, 2024) View paper
  - [20] SwiftVideo: A Unified Framework for Few-Step Video Generation through Trajectory-Distribution Alignment (Sun, 2025) View paper
  - [24] AnimateDiff-Lightning: Cross-Model Diffusion Distillation (Lin, 2024) View paper
  - [25] Fast High-Resolution Image Synthesis with Latent Adversarial Diffusion Distillation (Axel Sauer, 2024) View paper
  - [29] Learning Few-Step Diffusion Models by Trajectory Distribution Matching (Luo, 2025) View paper
  - [33] Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis (LU Yanzuo, 2025) View paper
  - [35] Self-Corrected Flow Distillation for Consistent One-Step and Few-Step Image Generation (Quan Dao, 2025) View paper
  - [36] Nitrofusion: High-fidelity single-step diffusion through dynamic adversarial training (Dar-Yen Chen, 2025) View paper
  - Distribution-Matching Distillation (2 papers)

- [28] One-step Diffusion Models with Bregman Density Ratio Matching (Zhu Yuanzhi, 2025) View paper
  - [30] Joint Distillation for Fast Likelihood Evaluation and Sampling in Flow-based Models (Xinyue Ai, 2025) View paper
  - GAN-Augmented Distillation (1 papers)
  - [21] Ufogen: You forward once large scale text-to-image generation via diffusion gans (Yanwu Xu, 2024) View paper
- Sampling Trajectory Optimization
  - Scheduler and Timestep Learning (2 papers)
  - [19] Entropic Time Schedulers for Generative Diffusion Models (Dejan Stancevic, 2025) View paper
  - [34] Differentiable Solver Search for Fast Diffusion Sampling (Wang, 2025) View paper
  - Trajectory Transformation Learning ★ (1 papers)
  - [0] BézierFlow: Learning Bézier Stochastic Interpolant Schedulers for Few-Step Generation (Anon et al., 2026) View paper
  - Guidance Optimization Techniques (2 papers)
  - [7] Applying Guidance in a Limited Interval Improves Sample and Distribution Quality in Diffusion Models (KynkÃ¤Ã¤nniemi, 2024) View paper
  - [16] Model-agnostic human preference inversion in diffusion models (Kim, 2024) View paper
- Architectural and Training Innovations
  - Direct Few-Step Training Methods (2 papers)
  - [3] One step diffusion via shortcut models (Frans, 2024) View paper
  - [5] Directly denoising diffusion models (Zhang Dan, 2024) View paper
  - Unified Multi-Step and Few-Step Frameworks (1 papers)
  - [6] Unified Continuous Generative Models (Sun Peng, 2025) View paper
  - Design Space and Preconditioning Studies (1 papers)
  - [1] Elucidating the design space of diffusion-based generative models (Karras, 2022) View paper
  - Flow Matching and Interpolant Methods (4 papers)
  - [11] Improving and generalizing flow-based generative models with minibatch optimal transport (Tong, 2023) View paper
  - [13] Local flow matching generative models (Xu Chen, 2024) View paper
  - [22] Faster Inference of Flow-Based Generative Models via Improved Data-Noise Coupling (Aram Davtyan, 2025) View paper
  - [45] Towards high-order mean flow generative models: Feasibility, expressivity, and provably efficient criteria (Cao Yang, 2025) View paper
  - Restoration and Corruption-Robust Training (1 papers)
  - [31] Restoration Score Distillation: From Corrupted Diffusion Pretraining to One-Step High-Quality Generation (Zhang, 2025) View paper
- Parallel and Non-Autoregressive Sampling (3 papers)
  - [15] Accelerating Parallel Sampling of Diffusion Models (Tang Zhiwei, 2024) View paper
  - [44] Inference Acceleration of Autoregressive Normalizing Flows by Selective Jacobi Decoding (Zhang Jia-ru, 2025) View paper
  - [50] Accelerating Diffusion Models with Parallel Sampling: Inference at Sub-Linear Time Complexity (Chen Hao-xuan, 2024) View paper
- Domain-Specific Applications
  - Text-to-Speech and Audio Synthesis (3 papers)
  - [9] F5-TTS: A Fairytaler that Fakes Fluent and Faithful Speech with Flow Matching (Chen Yu-Shen, 2024) View paper
  - [41] Guided conditioning with predictive network on score-based diffusion model for speech enhancement (Dail Kim, 2024) View paper
  - [46] ASD-Diffusion: Anomalous Sound Detection with Diffusion Models (Zhang, 2024) View paper
  - Protein and Molecular Design (2 papers)
  - [26] ProtFlow: Fast Protein Sequence Design via Flow Matching on Compressed Protein Language Model Embeddings (Zhu Yi-heng, 2025) View paper
  - [40] ReQFlow: Rectified Quaternion Flow for Efficient and High-Quality Protein Backbone Generation (Wang ZiChong, 2025) View paper
  - Medical Imaging and Restoration (1 papers)
  - [23] Few-shot CBCT-based synthetic CT generation with denoising diffusion probabilistic model. (Ping Lin Yeap, 2025) View paper
  - Language Modeling with Diffusion (1 papers)
  - [42] Seed Diffusion: A Large-Scale Diffusion Language Model with High-Speed Inference (Song Yuxuan, 2025) View paper
- Conditional Generation and Inverse Problems (4 papers)
  - [18] Diffusion models as plug-and-play priors (Graikos, 2022) View paper
  - [27] Upsample Guidance: Scale Up Diffusion Models without Training (Hwang, 2024) View paper
  - [39] An Efficient Sparse Blocks Inference Method for Image Editing Based on Diffusion Models (Zhuochao Yang, 2024) View paper
  - [47] LD-RPS: Zero-Shot Unified Image Restoration via Latent Diffusion Recurrent Posterior Sampling (Wang Yong, 2025) View paper

Narrative

Core task: few-step generation with pretrained diffusion and flow models. The field has organized itself around several complementary strategies for accelerating sampling without retraining large models from scratch. Distillation-Based Acceleration Methods compress multi-step trajectories into fewer function evaluations by training student models to mimic teacher outputs, often leveraging adversarial losses or consistency objectives as seen in works like Adversarial Diffusion Distillation[4] and Latent Consistency Models[10]. Sampling Trajectory Optimization refines the paths taken through latent space—whether by learning better couplings, adjusting time schedules, or transforming entire trajectories—while Architectural and Training Innovations explore modifications to network design and loss formulations that inherently support faster inference. Parallel and Non-Autoregressive Sampling investigates methods that generate multiple timesteps or tokens simultaneously, Domain-Specific Applications tailor acceleration techniques to modalities such as video or protein design, and Conditional Generation and Inverse Problems address guidance and task-specific constraints during few-step synthesis.

Within Sampling Trajectory Optimization, a particularly active line of work focuses on trajectory transformation learning: rather than distilling individual steps, these methods learn to map or rectify entire sampling paths to achieve straighter, more efficient flows. BezierFlow[0] exemplifies this approach by parameterizing trajectories with Bézier curves, enabling smooth interpolation and direct optimization of the end-to-end path. This contrasts with approaches like Shortcut Models[3], which construct piecewise shortcuts between intermediate states, and Directly Denoising[5], which emphasizes single-step jumps from noise to data. Meanwhile, works such as Flow Map Matching[32] and Self Corrected Flow[35] explore alternative geometric structures and iterative refinement strategies for trajectory design. BezierFlow[0] sits naturally among these trajectory-centric methods, sharing their emphasis on global path structure while offering a distinct parametric framework that balances expressiveness with computational tractability.

Related Works in Same Category

No sibling papers were found in the same taxonomy leaf. A taxonomy-subtopic-level comparison will be produced instead.

Taxonomy-Level Summary

All three subtopics address optimization of the sampling process in diffusion and flow models, but target different aspects. Trajectory Transformation Learning focuses on continuous, differentiable transformations of entire sampling paths. Guidance Optimization Techniques specifically optimizes how guidance signals are applied across steps. Scheduler and Timestep Learning learns discrete timestep sequences or scheduler functions.

Similarities: - All aim to improve sampling efficiency and quality in pretrained diffusion/flow models - All involve learning or optimizing aspects of the sampling trajectory rather than model architecture - All operate within the few-step generation paradigm to reduce computational cost

Differences: - Trajectory Transformation Learning uses continuous, differentiable representations of entire trajectories, while Scheduler and Timestep Learning focuses on discrete timestep selection - Guidance Optimization Techniques specifically targets guidance signal application, while the other two are guidance-agnostic - Trajectory Transformation Learning parameterizes transformations of paths, while Scheduler and Timestep Learning learns sequences or functions for step selection - The boundary between Trajectory Transformation and Scheduler Learning is whether the method operates on continuous trajectory transformations versus discrete timestep optimization

Suggested Search Directions: - Investigate whether trajectory transformation methods can incorporate guidance-specific optimization - Explore hybrid approaches that combine continuous trajectory transformations with learned discrete schedulers - Examine whether guidance optimization techniques benefit from trajectory-level versus step-level parameterization

Sibling Subtopics

- **Guidance Optimization Techniques** (leaves: 1, papers: 2)
- Scope: Methods that optimize the application of guidance signals across sampling steps to improve quality and efficiency.
- Exclude: Methods that learn schedulers without guidance-specific optimization belong to Scheduler and Timestep Learning.
- **Scheduler and Timestep Learning** (leaves: 1, papers: 2)
- Scope: Techniques that learn optimal timestep sequences or scheduler functions to improve sampling efficiency.
- Exclude: Methods that learn trajectory transformations beyond discrete timesteps belong to Trajectory Transformation Learning.

Contributions Analysis

Overall novelty summary. BézierFlow proposes learning optimal stochastic interpolant schedulers via Bézier-based parameterization to accelerate few-step generation with pretrained diffusion and flow models. The paper resides in the 'Trajectory Transformation Learning' leaf under 'Sampling Trajectory Optimization', a sparse subcategory containing only this work among the 50 papers surveyed. This positioning reflects a relatively unexplored research direction: rather than learning discrete timesteps or distilling multi-step teachers, the method parameterizes continuous trajectory transformations using differentiable Bézier functions. The sparsity of this leaf suggests the approach occupies a niche within the broader few-step generation landscape.

The taxonomy reveals that most acceleration efforts concentrate on distillation-based methods (21 papers across six subcategories) or domain-specific applications (8 papers). Within 'Sampling Trajectory Optimization', the sibling leaf 'Scheduler and Timestep Learning' contains two papers focused on discrete timestep selection, while 'Guidance Optimization Techniques' addresses guidance signal scheduling. BézierFlow diverges by targeting continuous trajectory transformations rather than discrete schedules or guidance modulation. The taxonomy narrative mentions related geometric approaches like Flow Map Matching and Self Corrected Flow, which explore alternative trajectory structures but are not classified in the same leaf, suggesting methodological distinctions in how trajectory optimization is formulated.

Among 29 candidates examined, the contribution-level analysis shows varied novelty signals. The core idea of learning stochastic interpolant schedulers examined 9 candidates with no clear refutations, suggesting limited prior work on this specific formulation. The Bézier parameterization examined 10 candidates, again with no refutations, indicating the representation choice appears distinctive. However, the BézierFlow training framework examined 10 candidates and found 1 refutable match, suggesting some overlap with existing lightweight training approaches. The limited search scope (29 papers, not exhaustive) means these findings reflect top-K semantic matches rather than comprehensive coverage of all related work.

Given the sparse taxonomy leaf and limited refutations across most contributions, BézierFlow appears to introduce a relatively novel trajectory transformation approach within the examined literature. The single refutation for the training framework likely reflects overlap with general lightweight training paradigms rather than the specific Bézier-based scheduler parameterization. However, the analysis is constrained by the 29-candidate search scope and may not capture all relevant prior work in trajectory optimization or scheduler learning. The taxonomy structure suggests the method occupies a distinct but narrow niche within the broader few-step generation field.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Learning optimal stochastic interpolant schedulers for few-step generation

Description: The authors propose optimizing the sampling trajectories themselves by learning stochastic interpolant schedulers, which govern the geometry of the sampling path. This broadens the scope beyond learning discrete ODE timesteps to learning continuous trajectory transformations that preserve endpoint distributions.

This contribution was assessed against **9 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Stochastic Interpolants via Conditional Dependent Coupling

URL: View paper

Brief Assessment

Conditional Dependent Coupling[58] focuses on multi-stage generation with conditional coupling between stages for high-resolution image synthesis, not on learning optimal sampling trajectory transformations or schedulers for few-step generation as in the original paper.

2. Generative AI for Molecular Simulations

URL: View paper

Brief Assessment

Molecular Simulations[55] applies stochastic interpolants to molecular dynamics and temperature transformations in statistical mechanics, not to few-step generation in diffusion/flow models for image synthesis.

3. Consistent3D: Towards Consistent High-Fidelity Text-to-3D Generation with Deterministic Sampling Prior

URL: View paper

Brief Assessment

Consistent3D[54] focuses on text-to-3D generation using SDE/ODE trajectory sampling for score distillation, not on learning stochastic interpolant schedulers for few-step generation in diffusion/flow models.

4. Building normalizing flows with stochastic interpolants

URL: View paper

Brief Assessment

Stochastic Interpolants[51] focuses on building normalizing flows between probability densities using a fixed interpolant formulation (Eq. 5), not on learning optimal schedulers for few-step generation. The candidate establishes the theoretical framework for stochastic interpolants but does not address trajectory optimization or scheduler learning for accelerated sampling.

5. Dynamic-TreeRPO: Breaking the Independent Trajectory Bottleneck with Structured Sampling

URL: View paper

Brief Assessment

Dynamic-TreeRPO[53] focuses on RL-based optimization for text-to-image generation with tree-structured sampling strategies, not on learning continuous trajectory transformations or stochastic interpolant schedulers for generative models.

6. Lipschitz-guided design of interpolation schedules in generative models

URL: View paper

Brief Assessment

Lipschitz Guided Design[52] focuses on optimizing interpolation schedules based on numerical properties (Lipschitz constants) rather than learning them through data-driven training. The original paper learns schedulers via teacher-forcing with neural network outputs, while the candidate derives schedules analytically through mathematical optimization criteria.

7. Probabilistic Forecasting with Stochastic Interpolants and FÃ¶llmer Processes

URL: View paper

Brief Assessment

Follmer Processes[57] focuses on probabilistic forecasting by mapping point mass measures to target distributions for time-series prediction, not on optimizing sampling trajectories for few-step generation in pretrained diffusion/flow models. The candidate addresses a fundamentally different problem domain (forecasting future states) rather than accelerating generative model sampling.

8. Likely Interpolants of Generative Models

URL: View paper

Brief Assessment

Likely Interpolants[56] focuses on computing interpolation curves constrained to data distributions for generative models, not on learning optimal sampling trajectory transformations or schedulers for few-step generation as in the original paper.

9. Flow map matching with stochastic interpolants: A mathematical framework for consistency models

URL: View paper

Brief Assessment

Flow Map Matching[32] focuses on learning two-time flow maps for consistency models, not on optimizing stochastic interpolant schedulers that govern sampling trajectory geometry as in the original paper.

Contribution 2: Bézier-based parameterization of SI schedulers

Description: The authors introduce a Bézier-based parameterization for stochastic interpolant schedulers that naturally satisfies boundary conditions, differentiability, and monotonicity of the signal-to-noise ratio. This reduces the problem to learning an ordered set of control points while ensuring the scheduler functions remain smooth and well-defined.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Modelling Additive Manufacturing Processes via Graph-Conditioned Diffusion Models

URL: View paper

Brief Assessment

Additive Manufacturing[67] focuses on modeling 3D printing processes using graph-conditioned diffusion models, not on scheduler parameterization for generative models. The provided context contains only declaration/metadata pages without technical content about Bézier functions or stochastic interpolants.

2. BÃ©zierFlow: BÃ©zier Stochastic Interpolant Schedulers for Few-Step Generation

URL: View paper

Brief Assessment

BezierFlow[61] proposes the same Bézier-based parameterization for stochastic interpolant schedulers. The candidate paper is the original paper itself, so no refutation is possible.

3. Bezier Distillation

URL: View paper

Brief Assessment

Bezier Distillation[65] applies Bézier curves to distillation in rectified flow models, focusing on multi-teacher knowledge distillation and error accumulation reduction. The original paper uses Bézier parameterization for stochastic interpolant schedulers in diffusion/flow models to satisfy boundary conditions, differentiability, and SNR monotonicity—a fundamentally different technical application.

4. Implicit BÃ©zier Motion Model for Precise Spatial and Temporal Control

URL: View paper

Brief Assessment

Implicit Bezier Motion[64] applies Bézier curves to character animation motion control, not to stochastic interpolant schedulers in diffusion/flow models. The domains and technical objectives are fundamentally different.

5. Draw Step by Step: Reconstructing CAD Construction Sequences from Point Clouds via Multimodal Diffusion.

URL: View paper

Brief Assessment

CAD Construction Sequences[59] focuses on reconstructing CAD models from point clouds using multimodal diffusion for design sequences, not on Bézier parameterization for stochastic interpolant schedulers in generative models.

6. Research on Ship Automatic Berthing Algorithm Based on Flow Matching and Velocity Matching

URL: View paper

Brief Assessment

Ship Automatic Berthing[60] applies Bézier curves to ship berthing path planning in maritime navigation, not to stochastic interpolant scheduler parameterization in generative models. The domains and technical objectives are entirely different.

7. CLDM-Palm: A controllable latent diffusion model for high-fidelity palmprint generation based on BÃ©zier curves

URL: View paper

Brief Assessment

CLDM-Palm[66] focuses on palmprint generation using Bézier curves as control conditions for image synthesis, not on parameterizing stochastic interpolant schedulers for diffusion model sampling trajectories.

8. An adaptive compressor characteristic map method based on the BÃ©zier curve

URL: View paper

Brief Assessment

Compressor Characteristic Map[63] applies Bézier curves to compressor performance modeling in gas turbine engines, not to stochastic interpolant schedulers or diffusion/flow models for generative AI.

9. SketchRefiner: Text-Guided Sketch Refinement Through Latent Diffusion Models

URL: View paper

Brief Assessment

SketchRefiner[62] uses Bézier curves for sketch representation in a graphics application, not for parameterizing stochastic interpolant schedulers in diffusion models. The technical domains are entirely different.

10. Is Your Diffusion Model Actually Denoising?

URL: View paper

Brief Assessment

Actually Denoising[68] focuses on measuring schedule deviation in diffusion models to assess non-denoising behavior, not on parameterizing schedulers using Bézier functions for few-step generation optimization.

Contribution 3: BézierFlow lightweight training framework

Description: The authors develop BézierFlow, a complete lightweight training framework that combines the optimization of sampling trajectories with Bézier-based continuous parameterization. The method achieves 2–3× performance improvement for sampling with ≤10 NFEs while requiring only 15 minutes of training.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. HDFlow: Hierarchical Diffusion-Flow Planning for Long-horizon Robotic Assembly

URL: View paper

Brief Assessment

HDFlow[72] focuses on hierarchical planning for robotic assembly using diffusion and rectified flow models, not on lightweight training frameworks for few-step generation in pretrained diffusion models. The candidate addresses a different problem domain (robotic manipulation) with different technical objectives (long-horizon planning vs. sampling acceleration).

2. Flash diffusion: Accelerating any conditional diffusion model for few steps image generation

URL: View paper

Brief Assessment

Flash Diffusion[8] focuses on distilling pretrained diffusion models using adversarial and distribution matching losses for few-step generation, not on learning Bézier-based continuous parameterization of sampling trajectories. The methods address different aspects of acceleration.

3. ACE-Step: A Step Towards Music Generation Foundation Model

URL: View paper

Brief Assessment

ACE Step[70] focuses on music generation using diffusion models with deep compression autoencoders and linear transformers, not on lightweight training frameworks for optimizing sampling trajectories in pretrained diffusion models.

4. Fast Autoregressive Models for Continuous Latent Generation

URL: View paper

Brief Assessment

Fast Autoregressive[73] focuses on replacing diffusion heads with shortcut heads for autoregressive image generation in continuous token spaces, not on learning optimal sampling trajectories or Bézier-based schedulers for pretrained diffusion/flow models.

5. Learning Few-Step Diffusion Models by Trajectory Distribution Matching

URL: View paper

Brief Assessment

Trajectory Distribution Matching[29] focuses on distribution-level trajectory matching for diffusion models using score distillation objectives, while BézierFlow optimizes sampling trajectories through Bézier-based continuous parameterization of stochastic interpolant schedulers. The candidate does not address Bézier parameterization or stochastic interpolant schedulers.

6. S4S: Solving for a Diffusion Model Solver

URL: View paper

Prior Art Analysis

S4S[71] demonstrates that prior work exists for lightweight training frameworks that optimize sampling trajectories in pretrained diffusion models for few-step generation. Both papers propose learning-based methods that require minimal training time (S4S[71] is described as 'lightweight, data-free' while BézierFlow requires '15 minutes of training') and achieve similar performance improvements in the few-NFE regime. S4S[71] directly optimizes a solver to match a teacher solver's output, which is conceptually similar to BézierFlow's approach of learning optimal sampling trajectories through teacher-forcing. The candidate paper's existence and publication timeline suggest that the concept of lightweight training for few-step generation was not novel to BézierFlow.

Evidence

Evidence 1 - **Rationale**: Both papers propose lightweight training methods that learn to optimize solvers/trajectories for few-step generation in pretrained diffusion models, demonstrating that S4S[71] established this approach prior to BézierFlow. - **Original**: we introduce b´ezierflow, a lightweight training approach for few-step generation with pretrained diffusion and flow models. b ´ezierflow achieves a 2-3x performance improvement for sampling with≤10 nfes while requiring only 15 minutes of training. - **Candidate**: we propose a new method that learns a good solver for the dm, which we call solving for the solver (s4s). s4s directly optimizes a solver to obtain good generation quality by learning to match the output of a strong teacher solver.

Evidence 2 - **Rationale**: S4S[71] already proposed optimizing both the solver and discretization schedule, which overlaps with BézierFlow's claim of broadening scope beyond timestep learning to trajectory optimization. - **Original**: recent lightweight training approaches have shown promise by learning optimal timesteps, but their scope remains restricted to ode discretizations. to broaden this scope, we propose learning the optimal transformation of the sampling trajectory by parameterizing stochastic interpolant (si) scheduler... - **Candidate**: our method is lightweight, data-free, and can be plugged in black-box on top of any discretization schedule or architecture to improve performance. building on top of this, we also propose s4s-alt, which optimizes both the solver and the discretization schedule.

Evidence 3 - **Rationale**: Both papers claim to improve upon traditional ODE solvers through lightweight training methods, with S4S[71] demonstrating this capability before BézierFlow. - **Original**: b ´ezierflow consistently outperforms prior timestep-learning methods, demonstrating the effectiveness of expanding the search space from discrete timesteps to b´ezier-based trajectory transformations. - **Candidate**: in all settings, s4s uniformly improves the sample quality relative to traditional ode solvers. moreover, our method is lightweight, data-free, and can be plugged in black-box on top of any discretization schedule or architecture to improve performance.

---

7. AdaDiff: Adaptive Step Selection for Fast Diffusion Models

URL: View paper

Brief Assessment

AdaDiff[69] focuses on learning instance-specific step usage policies for diffusion models, not on optimizing sampling trajectories through Bézier-based continuous parameterization of stochastic interpolant schedulers.

---

8. Revisiting Diffusion Models: From Generative Pre-training to One-Step Generation

URL: View paper

Brief Assessment

Generative Pre Training[76] focuses on GAN-based one-step generation from pretrained diffusion models, not on learning Bézier-parameterized sampling trajectories for few-step generation (≤10 NFEs). The technical approaches are fundamentally different.

---

9. Dip-go: A diffusion pruner via few-step gradient optimization

URL: View paper

Brief Assessment

Dip Go[74] focuses on structural pruning of diffusion models to reduce computational overhead during inference by removing redundant blocks, not on learning optimal sampling trajectories or schedulers for few-step generation as in BézierFlow.

---

10. Latent Diffusion Model without Variational Autoencoder

URL: View paper

Brief Assessment

Latent Diffusion Without VAE[75] focuses on replacing VAE-based latent spaces with self-supervised representations (DINO features) for diffusion models, not on learning optimal sampling trajectories or schedulers for few-step generation.

Appendix: Text Similarity Detection

Textual similarity detection checked 29 papers and found 3 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. BÃ©zierFlow: BÃ©zier Stochastic Interpolant Schedulers for Few-Step Generation

Detected in: Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] BézierFlow: Learning Bézier Stochastic Interpolant Schedulers for Few-Step Generation View paper
- [1] Elucidating the design space of diffusion-based generative models View paper
- [2] On distillation of guided diffusion models View paper
- [3] One step diffusion via shortcut models View paper
- [4] Adversarial diffusion distillation View paper
- [5] Directly denoising diffusion models View paper
- [6] Unified Continuous Generative Models View paper
- [7] Applying Guidance in a Limited Interval Improves Sample and Distribution Quality in Diffusion Models View paper
- [8] Flash diffusion: Accelerating any conditional diffusion model for few steps image generation View paper
- [9] F5-TTS: A Fairytaler that Fakes Fluent and Faithful Speech with Flow Matching View paper
- [10] Latent consistency models: Synthesizing high-resolution images with few-step inference View paper

- [11] Improving and generalizing flow-based generative models with minibatch optimal transport View paper
- [12] Sana-sprint: One-step diffusion with continuous-time consistency distillation View paper
- [13] Local flow matching generative models View paper
- [14] CMT: Mid-Training for Efficient Learning of Consistency, Mean Flow, and Flow Map Models View paper
- [15] Accelerating Parallel Sampling of Diffusion Models View paper
- [16] Model-agnostic human preference inversion in diffusion models View paper
- [17] SCott: Accelerating Diffusion Models with Stochastic Consistency Distillation View paper
- [18] Diffusion models as plug-and-play priors View paper
- [19] Entropic Time Schedulers for Generative Diffusion Models View paper
- [20] SwiftVideo: A Unified Framework for Few-Step Video Generation through Trajectory-Distribution Alignment View paper
- [21] Ufogen: You forward once large scale text-to-image generation via diffusion gans View paper
- [22] Faster Inference of Flow-Based Generative Models via Improved Data-Noise Coupling View paper
- [23] Few-shot CBCT-based synthetic CT generation with denoising diffusion probabilistic model. View paper
- [24] AnimateDiff-Lightning: Cross-Model Diffusion Distillation View paper
- [25] Fast High-Resolution Image Synthesis with Latent Adversarial Diffusion Distillation View paper
- [26] ProtFlow: Fast Protein Sequence Design via Flow Matching on Compressed Protein Language Model Embeddings View paper
- [27] Upsample Guidance: Scale Up Diffusion Models without Training View paper
- [28] One-step Diffusion Models with Bregman Density Ratio Matching View paper
- [29] Learning Few-Step Diffusion Models by Trajectory Distribution Matching View paper
- [30] Joint Distillation for Fast Likelihood Evaluation and Sampling in Flow-based Models View paper
- [31] Restoration Score Distillation: From Corrupted Diffusion Pretraining to One-Step High-Quality Generation View paper
- [32] Flow map matching with stochastic interpolants: A mathematical framework for consistency models View paper
- [33] Adversarial distribution matching for diffusion distillation towards efficient image and video synthesis View paper
- [34] Differentiable Solver Search for Fast Diffusion Sampling View paper
- [35] Self-Corrected Flow Distillation for Consistent One-Step and Few-Step Image Generation View paper
- [36] Nitrofusion: High-fidelity single-step diffusion through dynamic adversarial training View paper
- [37] Toward Theoretical Insights into Diffusion Trajectory Distillation via Operator Merging View paper
- [38] Hyper-sd: Trajectory segmented consistency model for efficient image synthesis View paper
- [39] An Efficient Sparse Blocks Inference Method for Image Editing Based on Diffusion Models View paper
- [40] ReQFlow: Rectified Quaternion Flow for Efficient and High-Quality Protein Backbone Generation View paper
- [41] Guided conditioning with predictive network on score-based diffusion model for speech enhancement View paper
- [42] Seed Diffusion: A Large-Scale Diffusion Language Model with High-Speed Inference View paper
- [43] Simple ReFlow: Improved Techniques for Fast Flow Models View paper
- [44] Inference Acceleration of Autoregressive Normalizing Flows by Selective Jacobi Decoding View paper
- [45] Towards high-order mean flow generative models: Feasibility, expressivity, and provably efficient criteria View paper
- [46] ASD-Diffusion: Anomalous Sound Detection with Diffusion Models View paper
- [47] LD-RPS: Zero-Shot Unified Image Restoration via Latent Diffusion Recurrent Posterior Sampling View paper
- [48] Align Your Flow: Scaling Continuous-Time Flow Map Distillation View paper
- [49] How to build a consistency model: Learning flow maps via self-distillation View paper
- [50] Accelerating Diffusion Models with Parallel Sampling: Inference at Sub-Linear Time Complexity View paper
- [51] Building normalizing flows with stochastic interpolants View paper
- [52] Lipschitz-guided design of interpolation schedules in generative models View paper
- [53] Dynamic-TreeRPO: Breaking the Independent Trajectory Bottleneck with Structured Sampling View paper
- [54] Consistent3D: Towards Consistent High-Fidelity Text-to-3D Generation with Deterministic Sampling Prior View paper
- [55] Generative AI for Molecular Simulations View paper
- [56] Likely Interpolants of Generative Models View paper
- [57] Probabilistic Forecasting with Stochastic Interpolants and FÃ¶llmer Processes View paper
- [58] Stochastic Interpolants via Conditional Dependent Coupling View paper
- [59] Draw Step by Step: Reconstructing CAD Construction Sequences from Point Clouds via Multimodal Diffusion. View paper
- [60] Research on Ship Automatic Berthing Algorithm Based on Flow Matching and Velocity Matching View paper
- [61] BÃ©zierFlow: BÃ©zier Stochastic Interpolant Schedulers for Few-Step Generation View paper
- [62] SketchRefiner: Text-Guided Sketch Refinement Through Latent Diffusion Models View paper
- [63] An adaptive compressor characteristic map method based on the BÃ©zier curve View paper
- [64] Implicit BÃ©zier Motion Model for Precise Spatial and Temporal Control View paper
- [65] Bezier Distillation View paper
- [66] CLDM-Palm: A controllable latent diffusion model for high-fidelity palmprint generation based on BÃ©zier curves View paper
- [67] Modelling Additive Manufacturing Processes via Graph-Conditioned Diffusion Models View paper
- [68] Is Your Diffusion Model Actually Denoising? View paper
- [69] AdaDiff: Adaptive Step Selection for Fast Diffusion Models View paper
- [70] ACE-Step: A Step Towards Music Generation Foundation Model View paper
- [71] S4S: Solving for a Diffusion Model Solver View paper
- [72] HDFlow: Hierarchical Diffusion-Flow Planning for Long-horizon Robotic Assembly View paper
- [73] Fast Autoregressive Models for Continuous Latent Generation View paper
- [74] Dip-go: A diffusion pruner via few-step gradient optimization View paper
- [75] Latent Diffusion Model without Variational Autoencoder View paper
- [76] Revisiting Diffusion Models: From Generative Pre-training to One-Step Generation View paper