

Novelty Assessment Report

Paper: GLASS Flows: Efficient Inference for Reward Alignment of Flow and Diffusion Models

PDF URL: <https://openreview.net/pdf?id=vH7OAPZ2dR>

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-29

Abstract

The performance of flow matching and diffusion models can be greatly improved at inference time using reward adaptation algorithms, yet efficiency remains a major limitation. While several algorithms were proposed, we demonstrate that a common bottleneck is the sampling method these algorithms rely on: many algorithms require to sample Markov transitions via SDE sampling, which is significantly less efficient and often less performant than ODE sampling. To remove this bottleneck, we introduce GLASS Flows, a new sampling paradigm that simulates a "flow matching model within a flow matching model" to sample Markov transitions. As we show in this work, this "inner" flow matching model can be retrieved from any pre-trained model without any re-training, effectively combining the efficiency of ODEs with the stochastic evolution of SDEs. On large-scale text-to-image models, we show that GLASS Flows eliminate the trade-off between stochastic evolution and efficiency. GLASS Flows improve state-of-the-art performance in text-to-image generation, making it a simple, drop-in solution for inference-time scaling of flow and diffusion models.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Reward Alignment of Flow and Diffusion Models**

A total of **50 papers** were analyzed and organized into a taxonomy with **19 categories**.

Taxonomy Overview

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

- **Inference-Time Alignment Methods**
- **Training-Based Alignment Methods**
- **Hybrid and Unified Alignment Frameworks**
- **Reward Hacking and Over-Optimization Mitigation**
- **Specialized Alignment Applications**
- **Accelerated and Efficient Alignment**
- **Survey and Review Papers**
- **Auxiliary and Related Methods**

Complete Taxonomy Tree

- Reward Alignment of Flow and Diffusion Models Survey Taxonomy
- Inference-Time Alignment Methods
 - Gradient-Based Guidance Approaches (6 papers)
 - [2] Inference-time alignment control for diffusion models with reinforcement learning guidance (Qiu Zijie, 2025) [View paper](#)
 - [4] Reward-guided controlled generation for inference-time alignment in diffusion models: Tutorial and review (Masatoshi Uehara, 2025) [View paper](#)
 - [8] Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review (Uehara, 2025) [View paper](#)
 - [9] A reward-directed diffusion framework for generative design optimization (Keramati, 2025) [View paper](#)
 - [36] CoDe: Blockwise Control for Denoising Diffusion Models (Singh, 2025) [View paper](#)
 - [41] Unified Control for Inference-Time Guidance of Denoising Diffusion Models (Maurya Goyal, 2025) [View paper](#)
 - Sampling and Search-Based Alignment (6 papers)
 - [29] Test-Time Alignment of Discrete Diffusion Models with Sequential Monte Carlo (C Pani, 2025) [View paper](#)
 - [30] Tree Reward-Aligned Search for TReASURe in Masked Diffusion Language Models (Li Ming, 2025) [View paper](#)
 - [37] Diffusion Tree Sampling: Scalable inference-time alignment of diffusion models (Jain Vineet, 2025) [View paper](#)
 - [44] Test-time Alignment of Diffusion Models without Reward Over-optimization (Kim Sun-woo, 2025) [View paper](#)
 - [47] TTSnap: Test-Time Scaling of Diffusion Models via Noise-Aware Pruning (Qingtao Yu, 2025) [View paper](#)
 - Prompt and Embedding Optimization (2 papers)
 - [38] Test-Time Alignment of Text-to-Image Diffusion Models via Null-Text Embedding Optimisation (Taehoon Kim, 2025) [View paper](#)
 - [45] Plug-and-Play Prompt Refinement via Latent Feedback for Diffusion Model Alignment (Lee, 2025) [View paper](#)
- Training-Based Alignment Methods
 - Reinforcement Learning Fine-Tuning
 - Policy Gradient and Actor-Critic Methods (4 papers)
 - [6] Flow-grpo: Training flow matching models via online rl (Liu Jie, 2025) [View paper](#)
 - [22] DanceGRPO: Unleashing GRPO on Visual Generation (Xue, 2025) [View paper](#)
 - [26] Advantage weighted matching: Aligning rl with pretraining in diffusion models (Ge, 2025) [View paper](#)
 - [33] Fine-tuning Flow Matching Generative Models with Intermediate Feedback (Fan Jiajun, 2025) [View paper](#)
 - Value-Based RL Approaches (2 papers)
 - [12] VARD: Efficient and Dense Fine-Tuning for Diffusion Models with Value-based RL (Zhuang, 2025) [View paper](#)

- [42] Aligning diffusion behaviors with q-functions for efficient continuous control (Chen Huayu, 2024) [View paper](#)
- Sparse Reward and Credit Assignment (1 papers)
 - [3] Towards better alignment: Training diffusion models with reinforcement learning against sparse rewards (Zijing Hu, 2025) [View paper](#)
- Preference Optimization Methods
- Direct Preference Optimization (2 papers)
 - [11] Preference-Based Alignment of Discrete Diffusion Models (Wells, 2025) [View paper](#)
 - [31] Balanceddp: Adaptive multi-metric alignment (Tamboli, 2025) [View paper](#)
- Multi-Preference and Distributional Alignment (4 papers)
 - [1] Calibrated multi-preference optimization for aligning diffusion models (Kyungmin Lee, 2025) [View paper](#)
 - [10] DRAGON: Distributional Rewards Optimize Diffusion Generative Models (Bai, 2025) [View paper](#)
 - [24] Regularized conditional diffusion model for multi-task preference alignment (Yu, 2024) [View paper](#)
 - [46] Diffusion Blend: Inference-Time Multi-Preference Alignment for Diffusion Models (Cheng Min, 2025) [View paper](#)
- Reward-Weighted and Flow Matching Alignment (2 papers)
- [25] Online reward-weighted fine-tuning of flow matching with wasserstein regularization (Fan Jiajun, 2025) [View paper](#)
- [48] Diffusion Alignment as Variational Expectation-Maximization (Lee Jae-Woo, 2025) [View paper](#)
- GFlowNet-Based Alignment (2 papers)
- [13] Improving gflownets for text-to-image diffusion alignment (Zhang, 2024) [View paper](#)
- [28] Efficient diversity-preserving diffusion alignment via gradient-informed gflownets (Liu Zhen, 2024) [View paper](#)
- Hybrid and Unified Alignment Frameworks ★ (3 papers)
 - [0] GLASS Flows: Efficient Inference for Reward Alignment of Flow and Diffusion Models (Anon et al., 2026) [View paper](#)
 - [14] GLASS Flows: Transition Sampling for Alignment of Flow and Diffusion Models (Holderrrieth, 2025) [View paper](#)
 - [27] ReALM-GEN: Real-World Constrained and Preference-Aligned Flow-and Diffusion-based Generative Models (P Giampouras, 2026) [View paper](#)
- Reward Hacking and Over-Optimization Mitigation (3 papers)
 - [7] ImageReFL: Balancing Quality and Diversity in Human-Aligned Diffusion Models (Nakhodnov, 2025) [View paper](#)
 - [17] Elucidating optimal reward-diversity tradeoffs in text-to-image diffusion models (Rohit Jena, 2025) [View paper](#)
 - [35] MIRA: Towards Mitigating Reward Hacking in Inference-Time Alignment of T2I Diffusion Models (Singh, 2025) [View paper](#)
- Specialized Alignment Applications
 - Video and Temporal Alignment (1 papers)
 - [50] Align-A-Video: Deterministic Reward Tuning of Image Diffusion Models for Consistent Video Editing (Shengzhi Wang, 2025) [View paper](#)
 - 3D and Multi-Modal Alignment (1 papers)
 - [19] Nabla-R2D3: Effective and Efficient 3D Diffusion Alignment with 2D Rewards (Liu Qing-ming, 2025) [View paper](#)
 - Robotics and Control Applications (2 papers)
 - [20] TeViR: Text-to-Video Reward with Diffusion Models for Efficient Reinforcement Learning (Chen YuHui, 2025) [View paper](#)
 - [21] Genflowrl: Shaping rewards with generative object-centric flow in visual reinforcement learning (Yu, 2025) [View paper](#)
 - Scientific and Domain-Specific Alignment (1 papers)
 - [23] Align-DA: Align Score-based Atmospheric Data Assimilation with Multiple Preferences (Sun, 2025) [View paper](#)
- Accelerated and Efficient Alignment (1 papers)
 - [18] Reward-Instruct: A Reward-Centric Approach to Fast Photo-Realistic Image Generation (Luo, 2025) [View paper](#)
- Survey and Review Papers (2 papers)
 - [15] Preference Alignment on Diffusion Model: A Comprehensive Survey for Image Generation and Editing (Wu Sihao, 2025) [View paper](#)
 - [16] Alignment and Safety of Diffusion Models via Reinforcement Learning and Reward Modeling: A Survey (Preeti Lamba, 2025) [View paper](#)
- Auxiliary and Related Methods (6 papers)
 - [5] Diffusion-rainbowpa: Improvements integrated preference alignment for diffusion-based text-to-image generation (H Sun, 2025) [View paper](#)
 - [32] Continuous Alignment of Multi-Target Preferences via Instructed Diffusion Model (Y Zhao, 2025) [View paper](#)
 - [34] Algorithms for optimal adaptation of diffusion models to reward functions (KD Dvijotham, 2023) [View paper](#)
 - [39] Solution Augmentation for ARC Problems Using GFlowNet: A Probabilistic Exploration Approach (S Hwang, 2024) [View paper](#)
 - [40] Reward-based modulation of task-switching performance: a diffusion model analysis. (Timo Weber, 2024) [View paper](#)
 - [43] Multi-agent Auto-Bidding with Latent Graph Diffusion Models (Huh, 2025) [View paper](#)

Narrative

Core task: Reward alignment of flow and diffusion models. The field has organized itself around several complementary strategies for steering generative models toward desired behaviors. Inference-Time Alignment Methods adjust sampling procedures without retraining, offering flexibility at the cost of computational overhead during generation. Training-Based Alignment Methods modify model parameters through reinforcement learning or preference optimization, embedding reward signals directly into the generative process. Hybrid and Unified Alignment Frameworks combine both paradigms, seeking to balance training efficiency with inference-time adaptability. Additional branches address Reward Hacking and Over-Optimization Mitigation to prevent degenerate solutions, Specialized Alignment Applications targeting domains like video or molecular design, and Accelerated and Efficient Alignment techniques that reduce computational burdens. Survey and Review Papers synthesize emerging best practices, while Auxiliary and Related Methods explore connections to broader generative modeling principles.

Recent work reveals tension between sample quality and computational cost, with some studies favoring lightweight inference-time guidance (Inference-Time RL Guidance[2], Training-Free Alignment[49]) and others advocating for parameter updates that internalize reward structure (Flow GRPO[6], ImageReFL[7]). A particularly active line explores how to handle sparse or multi-objective rewards without collapsing diversity (Sparse Reward Alignment[3], Reward-Diversity Tradeoffs[17]). GLASS Flows[0] sits within the Hybrid and Unified Alignment Frameworks branch, closely related to GLASS Transition Sampling[14] and ReALM-GEN[27], emphasizing a principled integration of flow-matching dynamics with reward-driven adjustments. Compared to purely inference-based approaches like Inference-Time RL Guidance[2], GLASS Flows[0] offers tighter coupling between training and sampling, while differing from purely training-centric methods by retaining inference-time flexibility. This positioning reflects ongoing efforts to unify the strengths of both paradigms without incurring prohibitive over-optimization or computational expense.

Related Works in Same Category

The following **2 sibling papers** share the same taxonomy leaf node with the original paper:

1. GLASS Flows: Transition Sampling for Alignment of Flow and Diffusion Models

Authors: Holderrieth, Peter, Singer, Uriel, Peter Holderrieth, et al. (18 authors total) | **Year/Venue:** 2025 | **URL:** [\[1\]](#)

Abstract

The performance of flow matching and diffusion models can be greatly improved at inference time using reward alignment algorithms, yet efficiency remains a major limitation. While several algorithms were proposed, we demonstrate that a common bottleneck is the sampling method these algorithms rely on: many algorithms require to sample Markov transitions via SDE sampling, which is significantly less efficient and often less performant than ODE sampling. To remove this bottleneck, we introduce GLA...

△ Similarity Notice

This paper is highly similar to the original paper; it may be a variant or near-duplicate. Please manually verify.

2. ReALM-GEN: Real-World Constrained and Preference-Aligned Flow-and Diffusion-based Generative Models

Authors: P Giampouras, M Mardani, Y Li, G Daras | **Year/Venue:** 2026 | **URL:** [\[2\]](#)

Abstract

Flow-based models are already trained to generate samples from a probability distribution $p_0(x)$. Aligned generative and reward-based alignment of flow- and diffusion-based models.

Relationship Analysis

Both papers belong to the Hybrid and Unified Alignment Frameworks category, addressing the integration of training and inference-time strategies for reward alignment of flow and diffusion models. The original paper (GLASS Flows) presents a specific technical method for efficient transition sampling via ODEs to enable inference-time alignment algorithms like SMC and guidance, while the candidate paper is a workshop proposal that broadly surveys the entire landscape of post-training alignment methods, including inverse problems, conditional generation, steering, and fine-tuning approaches. The key difference is that GLASS Flows contributes a novel algorithmic solution for a specific bottleneck in alignment methods, whereas the workshop proposal provides a comprehensive framework for organizing and understanding the diverse approaches to alignment across multiple application domains.

Contributions Analysis

Overall novelty summary. The paper introduces GLASS Flows, a sampling paradigm that simulates a ‘flow matching model within a flow matching model’ to enable efficient Markov transitions for reward adaptation algorithms. It resides in the Hybrid and Unified Alignment Frameworks leaf, which contains only three papers total, indicating a relatively sparse research direction. This leaf focuses on methods that combine training-time and inference-time strategies, distinguishing it from purely inference-based guidance or purely training-based fine-tuning approaches that dominate other branches of the taxonomy.

The taxonomy reveals substantial activity in neighboring areas: Inference-Time Alignment Methods includes six gradient-based guidance papers and six sampling-based alignment papers, while Training-Based Alignment Methods spans multiple subtopics with over twenty papers across RL fine-tuning and preference optimization. GLASS Flows bridges these domains by addressing a bottleneck in inference-time reward adaptation—specifically, the inefficiency of SDE sampling—while maintaining compatibility with pre-trained models. The scope notes clarify that hybrid frameworks must integrate both paradigms, whereas purely inference-based methods (e.g., gradient guidance) or purely training-based methods (e.g., policy gradient fine-tuning) belong elsewhere.

Among the three contributions analyzed, none were clearly refuted by the twenty-nine candidates examined. The first contribution (GLASS Flows sampling paradigm) examined ten candidates with zero refutable overlaps; the second (efficient ODE-based transition sampling) examined nine candidates with zero refutations; the third (application to inference-time reward alignment) examined ten candidates with zero refutations. This suggests that within the limited search scope—primarily top-K semantic matches and citation expansion—no prior work directly anticipates the specific combination of flow-within-flow sampling and ODE-based Markov transitions.

Based on the limited literature search of twenty-nine candidates, the work appears to occupy a distinct position within the sparse hybrid alignment space. The analysis does not cover exhaustive exploration of all inference-time or training-based methods, nor does it examine unpublished or domain-specific variants. The contribution-level statistics indicate no immediate prior work overlap among examined candidates, though the small size of the hybrid frameworks leaf and the modest search scope leave open the possibility of related techniques in adjacent branches or [3].

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

Contribution 1: GLASS Flows sampling paradigm for Markov transitions

Description: The authors propose GLASS Flows, a method that constructs an inner flow matching model to sample Markov transitions from pre-trained flow and diffusion models without retraining. This approach combines the efficiency of ODEs with the stochastic evolution characteristic of SDEs by using sufficient statistics to transform pre-trained models.

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. Flow matching with general discrete paths: A kinetic-optimal perspective

URL: [\[4\]](#)

Brief Assessment

Kinetic-Optimal Paths[59] focuses on discrete-space flow matching with continuous-time Markov chains for discrete data (text, materials, images), while GLASS Flows addresses continuous-space flow/diffusion models for sampling Markov transitions without retraining. The domains and technical approaches are fundamentally different.

2. Generator matching: Generative modeling with arbitrary markov processes

URL: [\[5\]](#)

Brief Assessment

Generator Matching[54] focuses on learning generators for arbitrary Markov processes through a unified framework, while GLASS Flows specifically addresses efficient sampling of pre-defined Markov transitions from pre-trained models without retraining. The candidate’s approach involves training generative models via generator matching objectives, whereas GLASS Flows transforms existing models for inference-time use.

3. Flow Matching: Markov kernels, stochastic processes and transport plans

URL: [\[6\]](#)

Brief Assessment

Flow Matching Kernels[56] focuses on mathematical foundations of flow matching via transport plans, Markov kernels, and stochastic processes. It does not propose a specific sampling method for Markov transitions from pre-trained models without retraining, which is the core novelty of GLASS Flows.

4. Jarzynski Reweighting and Sampling Dynamics for Training Energy-Based Models: Theoretical Analysis of Different Transition Kernels

URL: [\[1\]\[2\]](#)

Brief Assessment

Jarzynski Reweighting[57] focuses on theoretical analysis of Jarzynski reweighting for training energy-based models in flow-based diffusion and restricted Boltzmann machines. It does not address the construction of inner flow matching models for sampling Markov transitions from pre-trained models without retraining, which is the core novelty of GLASS Flows.

5. Flow network based generative models for non-iterative diverse candidate generation

URL: [\[1\]\[2\]](#)

Brief Assessment

Flow Network Candidates[51] focuses on learning flow networks for sequential generation where probability is proportional to rewards, not on sampling Markov transitions from pre-trained flow/diffusion models. The candidate addresses a different problem of converting energy functions to generative distributions through flow consistency equations.

6. Flow marching for a generative PDE foundation model

URL: [\[1\]\[2\]](#)

Brief Assessment

Flow Marching PDE[52] focuses on PDE-governed spatiotemporal dynamics with a bridge parameter k for neural operator learning, not on general flow/diffusion model alignment or reward-based inference-time scaling.

7. Stability of Schrödinger bridges and Sinkhorn semigroups for log-concave models

URL: [\[1\]\[2\]](#)

Brief Assessment

Schrodinger Bridge Stability[55] focuses on entropic optimal transport and Schrödinger bridges with log-concave models, not on flow matching models for sampling Markov transitions in generative models.

8. Bayesian structure learning with generative flow networks

URL: [\[1\]\[2\]](#)

Brief Assessment

Bayesian GFlowNets[58] focuses on structure learning for Bayesian networks using generative flow networks to approximate posterior distributions over DAGs. This is fundamentally different from GLASS Flows, which addresses sampling Markov transitions in pre-trained flow/diffusion models for reward alignment without retraining.

9. Generative flow networks for discrete probabilistic modeling

URL: [\[1\]\[2\]](#)

Brief Assessment

GFlowNets Discrete[53] focuses on learning generative flow networks for discrete probabilistic modeling using sequential construction policies, not on sampling Markov transitions from pre-trained flow/diffusion models via an inner flow matching framework as in GLASS Flows.

10. Videoflow: A flow-based generative model for video

URL: [\[1\]\[2\]](#)

Brief Assessment

VideoFlow[60] focuses on video prediction using normalizing flows for multi-frame sequences, not on sampling Markov transitions from pre-trained flow/diffusion models for reward alignment. The technical approaches and application domains are fundamentally different.

Contribution 2: Efficient transition sampling via ODEs without SDE bottleneck

Description: The method eliminates the common bottleneck in reward alignment algorithms by enabling efficient sampling of Markov transitions using ODEs rather than slower SDE sampling. This is achieved by retrieving an inner flow matching model from pre-trained models without additional training.

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. A training-free conditional diffusion model for learning stochastic dynamical systems

URL: [\[1\]\[2\]](#)

Brief Assessment

Training-Free Stochastic[77] focuses on learning stochastic flow maps of SDEs using a training-free conditional diffusion model, not on reward alignment algorithms or eliminating SDE bottlenecks in such contexts.

2. Fast sampling of diffusion models with exponential integrator

URL: [\[1\]\[2\]](#)

Brief Assessment

Exponential Integrator[69] focuses on discretizing the probability flow ODE for faster sampling of diffusion models, not on sampling Markov transitions for reward alignment. The candidate addresses numerical integration methods for ODEs, while the original paper introduces a method to sample stochastic transitions using an inner flow matching model for inference-time reward alignment.

3. The Effect of Stochasticity in Score-Based Diffusion Sampling: a KL Divergence Analysis

URL: [\[1\]\[2\]](#)

Brief Assessment

KL Divergence Analysis[76] focuses on analyzing stochasticity effects in score-based diffusion sampling through KL divergence bounds, not on developing methods for efficient transition sampling via ODEs for reward alignment algorithms.

4. Stochastic Transport Maps in Diffusion Models and Sampling

URL: [\[11\]](#)

Brief Assessment

Stochastic Transport Maps[75] focuses on theoretical foundations for constructing transport maps between probability distributions using diffusion processes and proving existence/uniqueness of SDEs. The candidate does not address the specific bottleneck of reward alignment algorithms requiring efficient Markov transition sampling.

5. An Ordinary Differential Equation Sampler with Stochastic Start for Diffusion Bridge Models

URL: [\[11\]](#)

Brief Assessment

Stochastic Start ODE[73] focuses on diffusion bridge models for conditional image generation tasks (restoration/translation), not on reward alignment algorithms for flow matching models. The technical contexts are fundamentally different.

6. New algorithms for sampling and diffusion models

URL: [\[11\]](#)

Brief Assessment

New Sampling Algorithms[68] focuses on reverse diffusion processes and sampling from known distributions using SDE theory, not on eliminating SDE bottlenecks in reward alignment by retrieving inner flow matching models from pre-trained models.

7. AdjointDEIS: Efficient gradients for diffusion models

URL: [\[11\]](#)

Brief Assessment

AdjointDEIS[71] focuses on computing gradients through diffusion sampling processes for optimization tasks, not on sampling Markov transitions as an alternative to SDE sampling for reward alignment algorithms.

8. Sa-solver: Stochastic adams solver for fast sampling of diffusion models

URL: [\[11\]](#)

Brief Assessment

SA-Solver[72] focuses on stochastic sampling methods for diffusion SDEs with variance control, not on eliminating SDE sampling bottlenecks in reward alignment algorithms through flow matching models.

9. On the mathematics of diffusion models

URL: [\[11\]](#)

Brief Assessment

Mathematics of Diffusion[74] focuses on deriving fundamental differential equations for diffusion models and analyzing reverse-diffusion processes theoretically, not on developing efficient transition sampling methods for reward alignment algorithms.

Contribution 3: Application to inference-time reward alignment with state-of-the-art performance

Description: The authors demonstrate that GLASS Flows, when combined with Feynman-Kac Steering, achieve state-of-the-art performance improvements in text-to-image generation. The method serves as a plug-in solution for inference-time reward alignment algorithms that previously relied on inefficient SDE sampling.

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. MIRA: Towards Mitigating Reward Hacking in Inference-Time Alignment of T2I Diffusion Models

URL: [\[11\]](#)

Brief Assessment

MIRA[35] focuses on mitigating reward hacking in noise optimization for text-to-image diffusion models through image-space KL regularization, whereas GLASS Flows addresses efficient transition sampling for general flow/diffusion models via ODE-based methods. These are distinct technical approaches to different problems in the inference-time alignment space.

2. Reno: Enhancing one-step text-to-image models through reward-based noise optimization

URL: [\[11\]](#)

Brief Assessment

Reno[67] focuses on optimizing initial noise for one-step text-to-image models using reward models, while GLASS Flows addresses efficient transition sampling for flow/diffusion models via ODEs. These are complementary inference-time techniques with different mechanisms and applications.

3. End-to-end Learning of Sparse Interventions on Activations to Steer Generation

URL: [\[11\]](#)

Brief Assessment

Sparse Activation Interventions[64] focuses on steering model activations through learned sparse interventions for alignment tasks, not on efficient transition sampling methods for flow/diffusion models used in reward alignment algorithms.

4. Test-time Alignment of Diffusion Models without Reward Over-optimization

URL: [\[11\]](#)

Brief Assessment

Test-Time Without Over-optimization[44] focuses on test-time alignment using Sequential Monte Carlo (SMC) sampling to avoid reward over-optimization, while the original paper applies GLASS Flows with Feynman-Kac Steering for inference-time reward alignment. These are distinct technical approaches to related but different problems in the reward alignment space.

5. Dreamsync: Aligning text-to-image generation with image understanding feedback

URL: [\[11\]](#)

Brief Assessment

DreamSync[61] focuses on training-time alignment of text-to-image models using VLM feedback and LoRA finetuning, not inference-time reward alignment methods like those addressed by GLASS Flows.

6. Parrot: Pareto-optimal multi-reward reinforcement learning framework for text-to-image generation

URL: [\[11\]](#)

Brief Assessment

Parrot[65] focuses on multi-reward RL fine-tuning for text-to-image generation using Pareto-optimal selection, not on inference-time reward alignment algorithms that rely on SDE sampling transitions as addressed by GLASS Flows.

7. Connections between reinforcement learning with feedback, test-time scaling, and diffusion guidance: An anthology

URL: [\[12\]](#)

Brief Assessment

Feedback Test-Time Anthology[63] focuses on theoretical connections between RL-based feedback methods and test-time scaling for autoregressive and diffusion models, not on practical inference-time reward alignment methods like Feynman-Kac Steering combined with novel transition sampling techniques.

8. Controlling Latent Diffusion Using Latent CLIP

URL: [\[13\]](#)

Brief Assessment

Latent CLIP Control[66] focuses on operating CLIP models in latent space for text-to-image generation tasks, not on inference-time reward alignment algorithms or flow matching models. The technical approaches are fundamentally different.

9. Reward-Instruct: A Reward-Centric Approach to Fast Photo-Realistic Image Generation

URL: [\[14\]](#)

Brief Assessment

Reward-Instruct[18] focuses on training-time reward optimization for few-step generation without diffusion distillation losses, whereas the original paper addresses inference-time reward alignment using GLASS Flows as a plug-in solution for existing algorithms like Feynman-Kac Steering.

10. RL for consistency models: Reward guided text-to-image generation with fast inference

URL: [\[15\]](#)

Brief Assessment

RL Consistency Models[62] focuses on fine-tuning consistency models for reward optimization, not on inference-time reward alignment for flow matching models using steering methods like GLASS Flows.

Appendix: Text Similarity Detection

Textual similarity detection checked 32 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. GLASS Flows: Transition Sampling for Alignment of Flow and Diffusion Models

Detected in: Core Task (sibling)

△ **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] GLASS Flows: Efficient Inference for Reward Alignment of Flow and Diffusion Models [View paper](#)
 - [1] Calibrated multi-preference optimization for aligning diffusion models [View paper](#)
 - [2] Inference-time alignment control for diffusion models with reinforcement learning guidance [View paper](#)
 - [3] Towards better alignment: Training diffusion models with reinforcement learning against sparse rewards [View paper](#)
 - [4] Reward-guided controlled generation for inference-time alignment in diffusion models: Tutorial and review [View paper](#)
 - [5] Diffusion-rainbowpa: Improvements integrated preference alignment for diffusion-based text-to-image generation [View paper](#)
 - [6] Flow-grpo: Training flow matching models via online rl [View paper](#)
 - [7] ImageReFL: Balancing Quality and Diversity in Human-Aligned Diffusion Models [View paper](#)
 - [8] Inference-time alignment in diffusion models with reward-guided generation: Tutorial and review [View paper](#)
 - [9] A reward-directed diffusion framework for generative design optimization [View paper](#)
 - [10] DRAGON: Distributional Rewards Optimize Diffusion Generative Models [View paper](#)
 - [11] Preference-Based Alignment of Discrete Diffusion Models [View paper](#)
 - [12] VARD: Efficient and Dense Fine-Tuning for Diffusion Models with Value-based RL [View paper](#)
 - [13] Improving gflownets for text-to-image diffusion alignment [View paper](#)
 - [14] GLASS Flows: Transition Sampling for Alignment of Flow and Diffusion Models [View paper](#)
 - [15] Preference Alignment on Diffusion Model: A Comprehensive Survey for Image Generation and Editing [View paper](#)
 - [16] Alignment and Safety of Diffusion Models via Reinforcement Learning and Reward Modeling: A Survey [View paper](#)
 - [17] Elucidating optimal reward-diversity tradeoffs in text-to-image diffusion models [View paper](#)
 - [18] Reward-Instruct: A Reward-Centric Approach to Fast Photo-Realistic Image Generation [View paper](#)
 - [19] Nabla-R2D3: Effective and Efficient 3D Diffusion Alignment with 2D Rewards [View paper](#)
 - [20] TeViR: Text-to-Video Reward with Diffusion Models for Efficient Reinforcement Learning [View paper](#)
 - [21] Genflowrl: Shaping rewards with generative object-centric flow in visual reinforcement learning [View paper](#)
 - [22] DanceGRPO: Unleashing GRPO on Visual Generation [View paper](#)
 - [23] Align-DA: Align Score-based Atmospheric Data Assimilation with Multiple Preferences [View paper](#)
 - [24] Regularized conditional diffusion model for multi-task preference alignment [View paper](#)
 - [25] Online reward-weighted fine-tuning of flow matching with wasserstein regularization [View paper](#)
 - [26] Advantage weighted matching: Aligning rl with pretraining in diffusion models [View paper](#)
 - [27] ReALM-GEN: Real-World Constrained and Preference-Aligned Flow-and Diffusion-based Generative Models [View paper](#)
 - [28] Efficient diversity-preserving diffusion alignment via gradient-informed gflownets [View paper](#)

- [29] Test-Time Alignment of Discrete Diffusion Models with Sequential Monte Carlo [View paper](#)
- [30] Tree Reward-Aligned Search for TReASURe in Masked Diffusion Language Models [View paper](#)
- [31] Balanceddp0: Adaptive multi-metric alignment [View paper](#)
- [32] Continuous Alignment of Multi-Target Preferences via Instructed Diffusion Model [View paper](#)
- [33] Fine-tuning Flow Matching Generative Models with Intermediate Feedback [View paper](#)
- [34] Algorithms for optimal adaptation of diffusion models to reward functions [View paper](#)
- [35] MIRA: Towards Mitigating Reward Hacking in Inference-Time Alignment of T2I Diffusion Models [View paper](#)
- [36] CoDe: Blockwise Control for Denoising Diffusion Models [View paper](#)
- [37] Diffusion Tree Sampling: Scalable inference-time alignment of diffusion models [View paper](#)
- [38] Test-Time Alignment of Text-to-Image Diffusion Models via Null-Text Embedding Optimisation [View paper](#)
- [39] Solution Augmentation for ARC Problems Using GFlowNet: A Probabilistic Exploration Approach [View paper](#)
- [40] Reward-based modulation of task-switching performance: a diffusion model analysis. [View paper](#)
- [41] Unified Control for Inference-Time Guidance of Denoising Diffusion Models [View paper](#)
- [42] Aligning diffusion behaviors with q-functions for efficient continuous control [View paper](#)
- [43] Multi-agent Auto-Bidding with Latent Graph Diffusion Models [View paper](#)
- [44] Test-time Alignment of Diffusion Models without Reward Over-optimization [View paper](#)
- [45] Plug-and-Play Prompt Refinement via Latent Feedback for Diffusion Model Alignment [View paper](#)
- [46] Diffusion Blend: Inference-Time Multi-Preference Alignment for Diffusion Models [View paper](#)
- [47] TTSnap: Test-Time Scaling of Diffusion Models via Noise-Aware Pruning [View paper](#)
- [48] Diffusion Alignment as Variational Expectation-Maximization [View paper](#)
- [49] Alignment without Over-optimization: Training-Free Solution for Diffusion Models [View paper](#)
- [50] Align-A-Video: Deterministic Reward Tuning of Image Diffusion Models for Consistent Video Editing [View paper](#)
- [51] Flow network based generative models for non-iterative diverse candidate generation [View paper](#)
- [52] Flow marching for a generative PDE foundation model [View paper](#)
- [53] Generative flow networks for discrete probabilistic modeling [View paper](#)
- [54] Generator matching: Generative modeling with arbitrary markov processes [View paper](#)
- [55] Stability of Schrödinger bridges and Sinkhorn semigroups for log-concave models [View paper](#)
- [56] Flow Matching: Markov kernels, stochastic processes and transport plans [View paper](#)
- [57] Jarzynski Reweighting and Sampling Dynamics for Training Energy-Based Models: Theoretical Analysis of Different Transition Kernels [View paper](#)
- [58] Bayesian structure learning with generative flow networks [View paper](#)
- [59] Flow matching with general discrete paths: A kinetic-optimal perspective [View paper](#)
- [60] Videoflow: A flow-based generative model for video [View paper](#)
- [61] Dreamsync: Aligning text-to-image generation with image understanding feedback [View paper](#)
- [62] RL for consistency models: Reward guided text-to-image generation with fast inference [View paper](#)
- [63] Connections between reinforcement learning with feedback, test-time scaling, and diffusion guidance: An anthology [View paper](#)
- [64] End-to-end Learning of Sparse Interventions on Activations to Steer Generation [View paper](#)
- [65] Parrot: Pareto-optimal multi-reward reinforcement learning framework for text-to-image generation [View paper](#)
- [66] Controlling Latent Diffusion Using Latent CLIP [View paper](#)
- [67] Reno: Enhancing one-step text-to-image models through reward-based noise optimization [View paper](#)
- [68] New algorithms for sampling and diffusion models [View paper](#)
- [69] Fast sampling of diffusion models with exponential integrator [View paper](#)
- [70] Residual learning in diffusion models [View paper](#)
- [71] Adjointdeis: Efficient gradients for diffusion models [View paper](#)
- [72] Sa-solver: Stochastic adams solver for fast sampling of diffusion models [View paper](#)
- [73] An Ordinary Differential Equation Sampler with Stochastic Start for Diffusion Bridge Models [View paper](#)
- [74] On the mathematics of diffusion models [View paper](#)
- [75] Stochastic Transport Maps in Diffusion Models and Sampling [View paper](#)
- [76] The Effect of Stochasticity in Score-Based Diffusion Sampling: a KL Divergence Analysis [View paper](#)
- [77] A training-free conditional diffusion model for learning stochastic dynamical systems [View paper](#)