Novelty Assessment Report

Paper: Proximal Diffusion Neural Sampler
PDF URL: https://openreview.net/pdf?id=XTHQqS7ObC
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

The task of learning a diffusion-based neural sampler for drawing samples from an unnormalized target distribution can be viewed as a stochastic optimal control problem on path measures. However, the training of neural samplers can be challenging when the target distribution is multimodal with significant barriers separating the modes, potentially leading to mode collapse. We propose a framework named **Proximal Diffusion Neural Sampler (PDNS)** that addresses these challenges by tackling the stochastic optimal control problem via proximal point method on the space of path measures. PDNS decomposes the learning process into a series of simpler subproblems that create a path gradually approaching the desired distribution. This staged procedure traces a progressively refined path to the desired distribution and promotes thorough exploration across modes. For a practical and efficient realization, we instantiate each proximal step with a proximal weighted denoising cross-entropy (WDCE) objective. We demonstrate the effectiveness and robustness of PDNS through extensive experiments on both continuous and discrete sampling tasks, including challenging scenarios in molecular dynamics and statistical physics.

Core Task Landscape

This paper addresses: **Learning Diffusion-Based Neural Samplers from Unnormalized Target Distributions**

A total of **50 papers** were analyzed and organized into a taxonomy with **26 categories**.

Taxonomy Overview

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

- **Training Objectives and Divergence Measures**
- **Sampling Algorithms and Inference Procedures**
- **Architectural and Methodological Frameworks**
- **Practical Enhancements and Training Strategies**
- **Specialized Applications and Domain Adaptations**
- **Theoretical Foundations and Score Estimation**
- **Auxiliary Uses and Extensions of Diffusion Models**

Complete Taxonomy Tree

- Learning Diffusion-Based Neural Samplers from Unnormalized Target Distributions Survey Taxonomy
- Training Objectives and Divergence Measures
  - Reverse KL and Diffusive Divergence Variants (2 papers)
  - [4] Training neural samplers with reverse diffusive kl divergence (He Jiajun, 2024) View paper
  - [31] Importance Weighted Score Matching for Diffusion Samplers with Enhanced Mode Coverage (Wang, 2025) View paper
  - Alternative Divergence Formulations (1 papers)
  - [48] Rethinking Losses for Diffusion Bridge Samplers (Gruber, 2025) View paper
  - Energy-Based and Score Matching Objectives (2 papers)
  - [9] Iterated Denoising Energy Matching for Sampling from Boltzmann Densities (Rector-Brooks, 2024) View paper
  - [33] Improving Adversarial Energy-Based Model via Diffusion Process (Geng Cong, 2024) View paper
- Sampling Algorithms and Inference Procedures
  - Sequential Monte Carlo and Particle Methods (2 papers)
  - [5] Reverse diffusion sequential Monte Carlo samplers (Wu, 2025) View paper
  - [14] Particle denoising diffusion sampler (Phillips, 2024) View paper
  - Annealing and Tempering Strategies (3 papers)
  - [19] Progressive Tempering Sampler with Diffusion (Rissanen, 2025) View paper
  - [25] Progressive Inference-Time Annealing of Diffusion Models for Sampling from Boltzmann Densities (Lee Jung-Yoon, 2025) View paper
  - [46] Continuously Tempered Diffusion Samplers (Erives, 2025) View paper
  - Langevin-Based Diffusion and Controlled Processes (3 papers)
  - [24] Sequential controlled langevin diffusions (Chen Jun-hua, 2024) View paper
  - [30] An optimal control perspective on diffusion-based generative modeling (Berner, 2022) View paper
  - [42] Value Gradient Sampler: Sampling as Sequential Decision Making (Yoon, 2025) View paper
  - Single-Step and Distilled Samplers (2 papers)
  - [23] Simple Distillation for One-Step Diffusion Models (H Zhu, 2025) View paper
  - [45] Single-Step Consistent Diffusion Samplers (Jutras-DubÃ©, 2025) View paper
- Architectural and Methodological Frameworks
  - Generalized Diffusion and Bridge Formulations (3 papers)

- [6] From denoising diffusions to denoising markov models (Joe Benton, 2024) View paper
    - [7] Improved sampling via learned diffusions (Richter, 2023) View paper
    - [10] Diffusion normalizing flow (Qinsheng Zhang, 2021) View paper
    - Discrete and Combinatorial Diffusion (2 papers)
    - [8] Scalable discrete diffusion samplers: Combinatorial optimization and statistical physics (Sebastian Sanokowski, 2025) View paper
    - [35] MDNS: Masked Diffusion Neural Sampler via Stochastic Optimal Control (Zhu Yuchen, 2025) View paper
    - Underdamped and Degenerate Diffusion (1 papers)
    - [43] Underdamped Diffusion Bridges with Applications to Sampling (Blessing, 2025) View paper
    - Amortized and Off-Policy Learning (5 papers)
    - [20] Adaptive teachers for amortized samplers (Kim, 2024) View paper
    - [21] Improved off-policy training of diffusion samplers (Alexandre Adam, 2024) View paper
    - [22] Amortizing intractable inference in diffusion models for vision, language, and control (Alexandre Adam, 2024) View paper
    - [28] Diffusion-based Sampling via Amortized Posterior Inference (Y Han, 2025) View paper
    - [41] On diffusion models for amortized inference: Benchmarking and improving stochastic control and sampling (Sendera, 2024) View paper
- Practical Enhancements and Training Strategies
    - Exploration and Mode Coverage Strategies ★ (2 papers)
    - [0] Proximal Diffusion Neural Sampler (Anon et al., 2026) View paper
    - [32] Learned Reference-based Diffusion Sampling for multi-modal distributions (Noble, 2024) View paper
    - Variance Reduction and Bias Correction (1 papers)
    - [50] Efficient and Unbiased Sampling from Boltzmann Distributions via Variance-Tuned Diffusion Models (Zhang Feng-zhe, 2025) View paper
    - Scalability and Computational Efficiency (1 papers)
    - [47] On scalable and efficient training of diffusion samplers (Kim Min-Kyu, 2025) View paper
    - Partial Trajectory and Credit Assignment (1 papers)
    - [27] Diffusion generative flow samplers: Improving learning signals through partial trajectory optimization (Zhang, 2023) View paper
    - Reference-Based and Auxiliary Model Guidance (2 papers)
    - [34] Neural Flow Samplers with Shortcut Models (Chen, 2025) View paper
    - [37] Learned Reference-based Diffusion Sampler for multi-modal distributions (M Noble, n.d.) View paper
- Specialized Applications and Domain Adaptations
    - Molecular Dynamics and Statistical Physics (2 papers)
    - [26] Consistent sampling and simulation: Molecular dynamics with energy-based diffusion models (Plainer, 2025) View paper
    - [29] Dynamicsdiffusion: Generating and rare event sampling of molecular dynamic trajectories using diffusion models (M Petersen, 2023) View paper
    - Bayesian Inference and Posterior Sampling (2 papers)
    - [39] Diffusion-based supervised learning of generative models for efficient sampling of multimodal distributions (Tran, 2025) View paper
    - [40] Diffusion-PINN Sampler (Zhiguo Shi, 2024) View paper
    - Reinforcement Learning and Sequential Decision Making (1 papers)
    - [38] A Diffusion Model Framework for Maximum Entropy Reinforcement Learning (Sebastian Sanokowski, 2025) View paper
    - Constrained Optimization and Compositional Generation (2 papers)
    - [12] Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc (Du, 2023) View paper
    - [13] Diffusion models as constrained samplers for optimization with unknown constraints (Lingkai Kong, 2024) View paper
- Theoretical Foundations and Score Estimation
    - Score Matching and Denoising Theory (2 papers)
    - [3] Denoising diffusion samplers (Vargas, 2023) View paper
    - [11] A practical diffusion path for sampling (Chehab, 2024) View paper
    - Diffusion Path Design and Forward Process (2 papers)
    - [17] Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion (Ye He, 2024) View paper
    - [36] Decoupled Diffusion Models with Explicit Transition Probability (Huang Yu-hang, 2023) View paper
- Auxiliary Uses and Extensions of Diffusion Models
    - Classification and Density Estimation (2 papers)
    - [2] Your Diffusion Model is Secretly a Zero-Shot Classifier (Alexander C. Li, 2023) View paper
    - [16] Your diffusion model is secretly a noise classifier and benefits from contrastive training (Xianghao Kong, 2024) View paper
    - Supervised and Hybrid Training Frameworks (1 papers)
    - [18] Diffusion-model-assisted supervised learning of generative models for density estimation (Yanfang Liu, 2024) View paper
    - Guidance and Conditioning Mechanisms (1 papers)
    - [1] Guiding a diffusion model with a bad version of itself (Timo Aila, 2024) View paper
    - Domain-Specific Auxiliary Applications (3 papers)
    - [15] Predicting trajectory destinations based on diffusion model integrating spatiotemporal features and urban contexts (Junjie Hu, 2024) View paper
    - [44] Learning to Sample Effective and Diverse Prompts for Text-to-Image Generation (Taeyoung Yun, 2025) View paper
    - [49] Recursive Multiple Wiener Integral Expansion for Nonlinear Filtering of Diffusion Processes (Sergey Lototsky, 2020) View paper

Narrative

Core task: Learning diffusion-based neural samplers from unnormalized target distributions. The field has evolved into a rich taxonomy with seven major branches. Training Objectives and Divergence Measures explores how to match learned samplers to target distributions through various loss formulations, while Sampling Algorithms and Inference Procedures focuses on the mechanics of drawing samples—ranging from sequential denoising to particle-based methods. Architectural and Methodological Frameworks addresses the design of neural networks and algorithmic scaffolding, and Theoretical Foundations and Score Estimation provides rigorous underpinnings for score matching and convergence guarantees. Practical Enhancements and Training Strategies, Specialized Applications and Domain Adaptations, and Auxiliary Uses and Extensions of Diffusion Models round out the taxonomy by addressing real-world deployment challenges, domain-specific tuning, and broader uses beyond direct sampling.

Within Practical Enhancements and Training Strategies, a particularly active line of work concerns exploration and mode coverage—ensuring that learned samplers do not collapse to a single mode of a multimodal target. Proximal Diffusion Sampler[0] sits squarely in this cluster, emphasizing strategies to improve mode discovery and coverage during training. It shares thematic concerns with Reference-Based Diffusion[32], which also tackles the challenge of guiding samplers toward diverse regions of the target distribution. Meanwhile, works like Denoising Diffusion Samplers[3] and Reverse Diffusion SMC[5] offer complementary perspectives by integrating sequential Monte Carlo or alternative inference schedules to balance exploration with computational efficiency. The interplay between these approaches highlights an ongoing tension: how to design training procedures that robustly capture complex, multimodal targets without prohibitive computational cost or architectural complexity.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. Learned Reference-based Diffusion Sampling for multi-modal distributions

Authors: Noble, Maxence, Grenioux, Louis, Maxence Noble, et al. (12 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Over the past few years, several approaches utilizing score-based diffusion have been proposed to sample from probability distributions, that is without having access to exact samples and relying solely on evaluations of unnormalized densities. The resulting samplers approximate the time-reversal of a noising diffusion process, bridging the target distribution to an easy-to-sample base distribution. In practice, the performance of these methods heavily depends on key hyperparameters that require...

Relationship Analysis

Both papers belong to the Exploration and Mode Coverage Strategies category, addressing mode collapse in diffusion-based neural samplers for multi-modal distributions. They share the goal of preventing mode collapse through staged training approaches: PDNS uses proximal point iterations with KL-penalized subproblems to gradually approach the target distribution, while LRDS leverages prior knowledge of mode locations through learned reference models (GMM or EBM) to guide the diffusion process. The key difference is that PDNS focuses on algorithmic regularization via proximal steps without requiring mode location knowledge, whereas LRDS explicitly assumes access to mode locations and constructs reference distributions tailored to multi-modality.

Contributions Analysis

Overall novelty summary. The paper proposes a Proximal Diffusion Neural Sampler (PDNS) framework that addresses mode collapse in multimodal target distributions by decomposing the stochastic optimal control problem into staged subproblems via proximal point methods on path measures. Within the taxonomy, it resides in the 'Exploration and Mode Coverage Strategies' leaf under 'Practical Enhancements and Training Strategies', alongside one sibling paper. This leaf represents a focused but not overcrowded research direction, with only two papers explicitly addressing exploration and mode coverage challenges during diffusion sampler training.

The taxonomy reveals that PDNS sits within a broader ecosystem of practical training enhancements, neighboring leaves such as 'Variance Reduction and Bias Correction' and 'Reference-Based and Auxiliary Model Guidance'. Related branches include 'Langevin-Based Diffusion and Controlled Processes' (which shares the stochastic optimal control formulation) and 'Annealing and Tempering Strategies' (which also employs progressive refinement). The scope note for the parent category emphasizes preventing mode collapse and ensuring comprehensive multimodal coverage, distinguishing this work from variance reduction techniques or computational scalability efforts in sibling leaves.

Among the three contributions analyzed, the literature search examined 23 candidate papers total. The core PDNS framework examined 3 candidates with no clear refutations; the unified path measure formulation and proximal WDCE objective each examined 10 candidates, again with no refutations found. These statistics reflect a limited semantic search scope rather than exhaustive coverage. The absence of refutable prior work among the examined candidates suggests that the specific combination of proximal point methods on path measures for diffusion samplers may represent a relatively unexplored angle, though the search scale precludes definitive claims about absolute novelty.

Based on the top-23 semantic matches and taxonomy structure, the work appears to occupy a distinct methodological niche within mode coverage strategies. The proximal decomposition approach differs from the reference-based guidance of its sibling paper, and the path measure formulation bridges continuous and discrete domains in a manner not explicitly captured by neighboring leaves. However, the limited search scope means potentially relevant work in annealing strategies or controlled processes may not have been fully examined.

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

Contribution 1: Proximal Diffusion Neural Sampler (PDNS) framework

Description: PDNS is a unified framework for diffusion-based sampling in both continuous and discrete domains. It applies proximal point iterations over path measures to decompose the learning process into simpler subproblems, progressively approaching the target distribution while mitigating mode collapse in multimodal settings.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. Optimization, Sampling and Their Interplay: Theory and Applications to Statistics and Machine Learning

URL: View paper

Brief Assessment

Optimization Sampling Interplay[73] appears to be a theoretical work on optimization and sampling methods. The provided context fragments mention proximal point methods and probability measures but lack sufficient detail about diffusion-based sampling frameworks to assess overlap with PDNS's specific approach of applying proximal iterations to path measures for neural samplers.

---

2. Enhancing sample efficiency and exploration in reinforcement learning through the integration of diffusion models and proximal policy optimization

URL: View paper

Brief Assessment

Diffusion PPO Integration[71] focuses on reinforcement learning with PPO and diffusion models for action generation in continuous control tasks, not on diffusion-based sampling for unnormalized distributions via proximal point methods on path measures as in PDNS.

---

3. Inference-Time Diffusion Model Distillation

URL: View paper

Brief Assessment

Inference-Time Distillation[72] focuses on accelerating pre-trained diffusion models during inference by teacher-guided refinement for image generation, not on training neural samplers for unnormalized distributions using proximal point methods on path measures.

Contribution 2: Unified path measure formulation for continuous and discrete SOC-based samplers

Description: The authors develop a unified formulation using path measures that integrates stochastic optimal control (SOC) based neural samplers for both continuous and discrete state spaces under a single theoretical framework.

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. Active uncertainty reduction for human-robot interaction: An implicit dual control approach

URL: View paper

Brief Assessment

Implicit Dual Control[67] addresses stochastic optimal control for human-robot interaction with uncertainty reduction, not the development of unified SOC-based neural samplers for continuous and discrete state spaces.

2. Chance-Constrained Linear Matrix Inequality Optimization: Theory and Applications

URL: View paper

Brief Assessment

Chance-Constrained LMI[70] focuses on linear matrix inequality optimization under chance constraints, which is a fundamentally different problem domain from stochastic optimal control-based neural samplers for continuous and discrete state spaces.

3. Probabilistic programming with stochastic probabilities

URL: View paper

Brief Assessment

Stochastic Probabilities[63] focuses on probabilistic programming with stochastic density estimation for inference, not on unifying stochastic optimal control samplers across continuous and discrete domains.

4. Bayesian optimization over discrete and mixed spaces via probabilistic reparameterization

URL: View paper

Brief Assessment

Probabilistic Reparameterization[66] addresses Bayesian optimization over discrete/mixed spaces using probabilistic distributions, not stochastic optimal control samplers for continuous and discrete domains as in the original paper.

5. On the performance of the particle swarm optimization algorithm with various inertia weight variants for computing optimal control of a class of hybrid systems

URL: View paper

Brief Assessment

Particle Swarm Inertia[69] focuses on particle swarm optimization for hybrid manufacturing systems with continuous and discrete dynamics, not on stochastic optimal control samplers or path measure formulations for neural sampling.

6. On Exact Embedding Framework for Optimal Control of Markov Decision Processes

URL: View paper

Brief Assessment

Exact Embedding MDP[61] focuses on embedding discrete MDPs into continuous action spaces for optimal control, not on unifying stochastic optimal control samplers across continuous and discrete domains for neural sampling tasks.

7. Stochastic Gradient MCMC for State Space Models

URL: View paper

Brief Assessment

Stochastic Gradient MCMC[68] focuses on scalable MCMC inference for state space models with temporal dependencies, not on unifying stochastic optimal control samplers across continuous and discrete domains through path measures.

8. TabDiff: a Mixed-type Diffusion Model for Tabular Data Generation

URL: View paper

Brief Assessment

TabDiff[64] focuses on tabular data generation using diffusion models for mixed numerical-categorical data types, not on stochastic optimal control samplers or path measure formulations for sampling problems.

9. Explainable Reinforcement Learning via Dynamic Mixture Policies

URL: View paper

Brief Assessment

Dynamic Mixture Policies[62] focuses on explainable reinforcement learning for control policies in robotics and automated driving, not on stochastic optimal control samplers or path measure formulations for sampling from distributions.

10. Active uncertainty reduction for safe and efficient interaction planning: A shielding-aware dual control approach

URL: View paper

Brief Assessment

Shielding Dual Control[65] addresses interactive motion planning and uncertainty reduction in robotics, not stochastic optimal control samplers for probability distributions. The technical domains are fundamentally different.

Contribution 3: Proximal weighted denoising cross-entropy (proximal WDCE) objective

Description: The authors instantiate PDNS with a proximal variant of the weighted denoising cross-entropy objective for both continuous and discrete sampling tasks, providing a practical and efficient realization of the framework along with principled strategies for selecting proximal step sizes.

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. Maximum likelihood training of score-based diffusion models

URL: View paper

Brief Assessment

Maximum Likelihood Score[51] focuses on likelihood weighting for score-based diffusion models, not proximal variants of WDCE objectives for neural samplers.

2. Perception Prioritized Training of Diffusion Models

URL: View paper

Brief Assessment

Perception Prioritized Training[52] focuses on redesigning weighting schemes for denoising score matching losses in diffusion models for image generation, not on proximal variants of weighted denoising cross-entropy for sampling from unnormalized distributions in the context of stochastic optimal control.

3. Time dependent loss reweighting for flow matching and diffusion models is theoretically justified

URL: View paper

Brief Assessment

Time Dependent Reweighting[58] focuses on theoretical justification for time-dependent loss weighting in flow matching and diffusion models generally, not on proximal variants of WDCE objectives for addressing mode collapse in multi-modal sampling tasks.

4. Simpler diffusion (sid2): 1.5 fid on imagenet512 with pixel-space diffusion

URL: View paper

Brief Assessment

Simpler Diffusion[59] focuses on sigmoid loss weighting for pixel-space diffusion models in image generation, not on proximal variants of weighted denoising cross-entropy for sampling from unnormalized distributions.

5. Denoising diffusion probabilistic models

URL: View paper

Brief Assessment

Denoising Diffusion Probabilistic[56] focuses on standard diffusion models for image generation using a simplified training objective, not on proximal variants of weighted denoising cross-entropy for sampling from unnormalized distributions.

6. Efficient diffusion training via min-snr weighting strategy

URL: View paper

Brief Assessment

Min-SNR Weighting[55] addresses loss weighting strategies for diffusion model training convergence, not the proximal WDCE framework for stochastic optimal control in sampling tasks.

7. Dream-coder 7b: An open diffusion language model for code

URL: View paper

Brief Assessment

Dream-Coder[57] uses a standard continuous-time weighted cross-entropy objective for code generation, not a proximal variant. The paper does not discuss proximal point methods or proximal optimization in path measure spaces.

8. Simplified and generalized masked diffusion for discrete data

URL: View paper

Brief Assessment

Simplified Masked Diffusion[53] focuses on masked diffusion for discrete data with a simplified ELBO formulation, not on proximal variants of WDCE objectives for general diffusion-based sampling across continuous and discrete domains as in the original paper.

9. Denoising task routing for diffusion models

URL: View paper

Brief Assessment

Denoising Task Routing[54] focuses on architectural modifications through task-specific channel masking for diffusion models, not on proximal optimization methods or weighted denoising cross-entropy objectives for sampling problems.

10. Training Robust Classifiers with Diffusion Denoised Examples

URL: View paper

Brief Assessment

Diffusion Denoised Examples[60] focuses on using one-step diffusion denoising for data augmentation in image classification, not on developing proximal variants of weighted denoising cross-entropy objectives for diffusion-based sampling frameworks.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Proximal Diffusion Neural Sampler View paper
- [1] Guiding a diffusion model with a bad version of itself View paper
- [2] Your Diffusion Model is Secretly a Zero-Shot Classifier View paper
- [3] Denoising diffusion samplers View paper
- [4] Training neural samplers with reverse diffusive kl divergence View paper
- [5] Reverse diffusion sequential Monte Carlo samplers View paper
- [6] From denoising diffusions to denoising markov models View paper
- [7] Improved sampling via learned diffusions View paper
- [8] Scalable discrete diffusion samplers: Combinatorial optimization and statistical physics View paper
- [9] Iterated Denoising Energy Matching for Sampling from Boltzmann Densities View paper

- [10] Diffusion normalizing flow View paper
- [11] A practical diffusion path for sampling View paper
- [12] Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc View paper
- [13] Diffusion models as constrained samplers for optimization with unknown constraints View paper
- [14] Particle denoising diffusion sampler View paper
- [15] Predicting trajectory destinations based on diffusion model integrating spatiotemporal features and urban contexts View paper
- [16] Your diffusion model is secretly a noise classifier and benefits from contrastive training View paper
- [17] Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion View paper
- [18] Diffusion-model-assisted supervised learning of generative models for density estimation View paper
- [19] Progressive Tempering Sampler with Diffusion View paper
- [20] Adaptive teachers for amortized samplers View paper
- [21] Improved off-policy training of diffusion samplers View paper
- [22] Amortizing intractable inference in diffusion models for vision, language, and control View paper
- [23] Simple Distillation for One-Step Diffusion Models View paper
- [24] Sequential controlled langevin diffusions View paper
- [25] Progressive Inference-Time Annealing of Diffusion Models for Sampling from Boltzmann Densities View paper
- [26] Consistent sampling and simulation: Molecular dynamics with energy-based diffusion models View paper
- [27] Diffusion generative flow samplers: Improving learning signals through partial trajectory optimization View paper
- [28] Diffusion-based Sampling via Amortized Posterior Inference View paper
- [29] Dynamicsdiffusion: Generating and rare event sampling of molecular dynamic trajectories using diffusion models View paper
- [30] An optimal control perspective on diffusion-based generative modeling View paper
- [31] Importance Weighted Score Matching for Diffusion Samplers with Enhanced Mode Coverage View paper
- [32] Learned Reference-based Diffusion Sampling for multi-modal distributions View paper
- [33] Improving Adversarial Energy-Based Model via Diffusion Process View paper
- [34] Neural Flow Samplers with Shortcut Models View paper
- [35] MDNS: Masked Diffusion Neural Sampler via Stochastic Optimal Control View paper
- [36] Decoupled Diffusion Models with Explicit Transition Probability View paper
- [37] Learned Reference-based Diffusion Sampler for multi-modal distributions View paper
- [38] A Diffusion Model Framework for Maximum Entropy Reinforcement Learning View paper
- [39] Diffusion-based supervised learning of generative models for efficient sampling of multimodal distributions View paper
- [40] Diffusion-PINN Sampler View paper
- [41] On diffusion models for amortized inference: Benchmarking and improving stochastic control and sampling View paper
- [42] Value Gradient Sampler: Sampling as Sequential Decision Making View paper
- [43] Underdamped Diffusion Bridges with Applications to Sampling View paper
- [44] Learning to Sample Effective and Diverse Prompts for Text-to-Image Generation View paper
- [45] Single-Step Consistent Diffusion Samplers View paper
- [46] Continuously Tempered Diffusion Samplers View paper
- [47] On scalable and efficient training of diffusion samplers View paper
- [48] Rethinking Losses for Diffusion Bridge Samplers View paper
- [49] Recursive Multiple Wiener Integral Expansion for Nonlinear Filtering of Diffusion Processes View paper
- [50] Efficient and Unbiased Sampling from Boltzmann Distributions via Variance-Tuned Diffusion Models View paper
- [51] Maximum likelihood training of score-based diffusion models View paper
- [52] Perception Prioritized Training of Diffusion Models View paper
- [53] Simplified and generalized masked diffusion for discrete data View paper
- [54] Denoising task routing for diffusion models View paper
- [55] Efficient diffusion training via min-snr weighting strategy View paper
- [56] Denoising diffusion probabilistic models View paper
- [57] Dream-coder 7b: An open diffusion language model for code View paper
- [58] Time dependent loss reweighting for flow matching and diffusion models is theoretically justified View paper
- [59] Simpler diffusion (sid2): 1.5 fid on imagenet512 with pixel-space diffusion View paper
- [60] Training Robust Classifiers with Diffusion Denoised Examples View paper
- [61] On Exact Embedding Framework for Optimal Control of Markov Decision Processes View paper
- [62] Explainable Reinforcement Learning via Dynamic Mixture Policies View paper
- [63] Probabilistic programming with stochastic probabilities View paper
- [64] TabDiff: a Mixed-type Diffusion Model for Tabular Data Generation View paper
- [65] Active uncertainty reduction for safe and efficient interaction planning: A shielding-aware dual control approach View paper
- [66] Bayesian optimization over discrete and mixed spaces via probabilistic reparameterization View paper
- [67] Active uncertainty reduction for human-robot interaction: An implicit dual control approach View paper
- [68] Stochastic Gradient MCMC for State Space Models View paper
- [69] On the performance of the particle swarm optimization algorithm with various inertia weight variants for computing optimal control of a class of hybrid systems View paper
- [70] Chance-Constrained Linear Matrix Inequality Optimization: Theory and Applications View paper
- [71] Enhancing sample efficiency and exploration in reinforcement learning through the integration of diffusion models and proximal policy optimization View paper
- [72] Inference-Time Diffusion Model Distillation View paper
- [73] Optimization, Sampling and Their Interplay: Theory and Applications to Statistics and Machine Learning View paper