

Novelty Assessment Report

Paper: Computational Bottlenecks for Denoising Diffusions

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

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

Denoising diffusions sample from a probability distribution μ in \mathbb{R}^d by constructing a stochastic process $\{\hat{x}_t\}_{t \in [0, T]}$ in \mathbb{R}^d such that \hat{x}_0 is easy to sample, but the distribution of \hat{x}_T at large T approximates μ . The drift $m: \mathbb{R}^d \rightarrow \mathbb{R}^d$ of this diffusion process is learned by minimizing a score-matching objective.

Is every probability distribution μ , for which sampling is tractable, also amenable to sampling via diffusions? We address this question by studying its relation to information-computation gaps in statistical estimation. Earlier work in this area constructs broad families of distributions μ for which sampling is easy, but approximating the drift $m(y, t)$ is conjectured to be intractable, and provides rigorous evidence for intractability.

We prove that this implies a failure of sampling via diffusions. First, there exist drifts whose score matching objective is superpolynomially close to the optimum value (among polynomial time drifts) and yet yield samples with distribution that is very far from the target one. Second, any polynomial-time drift that is also Lipschitz continuous results in equally incorrect sampling.

We instantiate our results on the toy problem of sampling a sparse low-rank matrix, and further demonstrate empirically the failure of diffusion-based sampling. Our work implies that caution should be used in adopting diffusion sampling when other approaches are available.

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: **Computational Limitations of Denoising Diffusion Sampling**

A total of **50 papers** were analyzed and organized into a taxonomy with **22 categories**.

Taxonomy Overview

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

- **Theoretical Foundations and Computational Complexity**
- **Sampling Acceleration Techniques**
- **Model Architecture Efficiency**
- **Conditional and Constrained Sampling**
- **Application-Specific Efficiency Adaptations**
- **Training and Modeling Improvements**
- **Emerging Paradigms and Extensions**

Complete Taxonomy Tree

- Computational Limitations of Denoising Diffusion Sampling Survey Taxonomy
- Theoretical Foundations and Computational Complexity
 - Information-Computation Gaps and Intractability ★ (1 papers)
 - [0] Computational Bottlenecks for Denoising Diffusions (Anon et al., 2026) [View paper](#)
 - Sampling Schedule Design and Optimization (3 papers)
 - [16] Optimal Inference Schedules for Masked Diffusion Models (Chen, 2025) [View paper](#)
 - [20] Common diffusion noise schedules and sample steps are flawed (Shanchuan Lin, 2024) [View paper](#)
 - [24] Accelerating diffusion sampling with optimized time steps (Shuchen Xue, 2024) [View paper](#)
- Sampling Acceleration Techniques
 - Non-Markovian and Implicit Sampling (2 papers)
 - [6] gddim: Generalized denoising diffusion implicit models (Zhang, 2022) [View paper](#)
 - [13] Denoising diffusion implicit models (Song, 2020) [View paper](#)
 - Distillation and Knowledge Transfer (2 papers)
 - [25] Direct Distillation: A Novel Approach for Efficient Diffusion Model Inference (Zilai Li, 2025) [View paper](#)
 - [34] Learning to efficiently sample from diffusion probabilistic models (Watson, 2021) [View paper](#)
 - Training-Free Inference Optimization (4 papers)
 - [11] Fast sampling through the reuse of attention maps in diffusion models (Rosco Hunter, 2025) [View paper](#)
 - [23] A Simple Early Exiting Framework for Accelerated Sampling in Diffusion Models (Choi, 2024) [View paper](#)
 - [28] Let Features Decide Their Own Solvers: Hybrid Feature Caching for Diffusion Transformers (Chen, 2025) [View paper](#)
 - [48] Efficient denoising diffusion via probabilistic masking (Zhang Weizhong, 2024) [View paper](#)
 - Hybrid and Adversarial Acceleration (3 papers)
 - [4] Tackling the generative learning trilemma with denoising diffusion GANs (Xiao, 2021) [View paper](#)
 - [5] Single image super-resolution with denoising diffusion GANS (Heng Xiao, 2024) [View paper](#)
 - [32] Fast inference in denoising diffusion models via mmd finetuning (Emanuele Aiello, 2024) [View paper](#)

- Model Architecture Efficiency
 - Step-Aware and Adaptive Architectures (2 papers)
 - [22] Denoising task routing for diffusion models (Park, 2023) [View paper](#)
 - [40] Denoising diffusion step-aware models (Yang, 2023) [View paper](#)
 - Quantization and Compression (1 papers)
 - [46] Modulated Diffusion: Accelerating Generative Modeling with Modulated Quantization (Gao Wei-zhi, 2025) [View paper](#)
 - Domain-Specific Architectural Optimizations (2 papers)
 - [14] Efficiency-optimized Video Diffusion Models (Zijun Deng, 2023) [View paper](#)
 - [38] Partially Conditioned Patch Parallelism for Accelerated Diffusion Model Inference (Zhang Xiu-yu, 2024) [View paper](#)
- Conditional and Constrained Sampling
 - Inverse Problem Solving (4 papers)
 - [9] Denoising diffusion restoration models (Kawar, 2022) [View paper](#)
 - [17] Efficient Diffusion Posterior Sampling for Noisy Inverse Problems (Ji Li, 2025) [View paper](#)
 - [30] Denoising diffusion models for plug-and-play image restoration (Yuanzhi Zhu, 2023) [View paper](#)
 - [31] Improving diffusion inverse problem solving with decoupled noise annealing (Bingliang Zhang, 2025) [View paper](#)
 - Conditional Generation with Auxiliary Inputs (2 papers)
 - [1] Fast constrained sampling in pre-trained diffusion models (Graikos, 2024) [View paper](#)
 - [37] IMPUS: Image Morphing with Perceptually-Uniform Sampling Using Diffusion Models (Yang Zhaoyuan, 2023) [View paper](#)
 - Constraint Satisfaction and Optimization Alignment (2 papers)
 - [18] Aligning optimization trajectories with diffusion models for constrained design generation (Giannone, 2023) [View paper](#)
 - [21] Toward exploratory inverse constraint inference with generative diffusion verifiers (R Zhao, 2025) [View paper](#)
- Application-Specific Efficiency Adaptations
 - Image Restoration and Enhancement (5 papers)
 - [7] CoreDiff: Contextual error-modulated generalized diffusion model for low-dose CT denoising and generalization (Qi Gao, 2023) [View paper](#)
 - [39] Image Dehazing Method Based on Diffusion Model (Fengxu Guan, 2024) [View paper](#)
 - [41] A conditional diffusion model with fast sampling strategy for remote sensing image super-resolution (Fanen Meng, 2024) [View paper](#)
 - [42] DiffLoss: unleashing diffusion model as constraint for training image restoration network (Jiangtong Tan, 2024) [View paper](#)
 - [43] Resfusion: Denoising diffusion probabilistic models for image restoration based on prior residual noise (Changsheng Dong, 2024) [View paper](#)
 - Medical and Scientific Imaging (2 papers)
 - [19] Volumetric Conditional Score-Based Residual Diffusion Model for PET/MR Denoising (Siyeop Yoon, 2024) [View paper](#)
 - [26] Diffusiondepth: Diffusion denoising approach for monocular depth estimation (Duan Yiqun, 2024) [View paper](#)
 - Synthesis and Generation Applications (3 papers)
 - [10] Diff-tts: A denoising diffusion model for text-to-speech (Myeonghun Jeong, 2021) [View paper](#)
 - [12] DiffusionDrive: Truncated Diffusion Model for End-to-End Autonomous Driving (Bencheng Liao, 2024) [View paper](#)
 - [29] Intention-Aware Denoising Diffusion Model for Trajectory Prediction (Chen Liu, 2024) [View paper](#)
 - Specialized Domain Applications (5 papers)
 - [27] Accelerating Multiphase Simulations With Denoising Diffusion Model Driven Initializations (Jaehong Chung, 2024) [View paper](#)
 - [35] Advancing Realistic Precipitation Nowcasting With a Spatiotemporal Transformer-Based Denoising Diffusion Model (Zewei Zhao, 2024) [View paper](#)
 - [36] Anomaly Detection of Industrial Smelting Furnace Incorporated With Accelerated Sampling Denoising Diffusion Probability Model and Conv-Transformer (Yu Zhang, 2024) [View paper](#)
 - [49] Hyperspectral Image Generation with Unmixing Guided Diffusion Model (Shen Shiyu, 2025) [View paper](#)
 - [50] Sat2City: 3D City Generation from A Single Satellite Image with Cascaded Latent Diffusion (Hua, 2025) [View paper](#)
- Training and Modeling Improvements
 - Variance and Likelihood Optimization (1 papers)
 - [15] Improved Denoising Diffusion Probabilistic Models (Nichol, 2021) [View paper](#)
 - Alternative Diffusion Formulations (1 papers)
 - [47] Zeroth-Order Sampling Methods for Non-Log-Concave Distributions: Alleviating Metastability by Denoising Diffusion (He Ye, 2024) [View paper](#)
 - Auxiliary Training Objectives (1 papers)
 - [33] Consistent World Models via Foresight Diffusion (Zhang, 2025) [View paper](#)
- Emerging Paradigms and Extensions
 - Quantum and Alternative Computing Paradigms (1 papers)
 - [2] Quantum denoising diffusion models (Michael KÄ¶lle, 2024) [View paper](#)
 - Semi-Supervised and Data-Efficient Learning (2 papers)
 - [44] MedDiff: Generating Electronic Health Records using Accelerated Denoising Diffusion Model (He Huan, 2023) [View paper](#)
 - [45] Semi-Supervised Learning Leveraging Denoising Diffusion Probabilistic Models for the Characterization of Nanophotonic Devices (Jun-hyeong Kim, 2024) [View paper](#)
 - Survey and Methodological Reviews (2 papers)
 - [3] Diffusion models in vision: A survey (Florinel-Alin Croitoru, 2023) [View paper](#)
 - [8] Efficient diffusion models: A survey (Shen Hui, 2025) [View paper](#)

Narrative

Core task: computational limitations of denoising diffusion sampling. Denoising diffusion models have achieved remarkable generative quality, yet their iterative sampling process imposes substantial computational burdens that limit practical deployment. The field's taxonomy reflects a multifaceted response to these challenges. At the theoretical level, researchers investigate fundamental complexity bounds and information-computation gaps, as exemplified by Computational Bottlenecks[0], which examines intrinsic intractability. Meanwhile, a large body of work focuses on sampling acceleration techniques—ranging from deterministic shortcuts like DDIM[13] and step-aware scheduling methods such as Optimized Time Steps[24] to distillation approaches like Direct Distillation[25]—that reduce the number of function evaluations without sacrificing output fidelity. Model architecture efficiency explores lightweight designs and caching strategies (e.g., Hybrid Feature Caching[28]), while conditional and constrained sampling addresses the added cost of guidance and

constraint satisfaction (Fast Constrained Sampling[1]). Application-specific branches tailor efficiency adaptations to domains like video generation (Video Diffusion Efficiency[14]) and medical imaging (MedDiff[44]), and training improvements seek to learn more efficient noise schedules or parameterizations from the outset. Emerging paradigms extend diffusion ideas to novel settings, including quantum frameworks (Quantum Denoising[2]).

Across these branches, a central tension emerges between sample quality, computational cost, and generalization: aggressive step reduction or architectural pruning can degrade fidelity, while domain-specific optimizations may not transfer broadly. Computational Bottlenecks[0] sits squarely within the theoretical foundations branch, probing the hardness barriers that underpin why certain sampling regimes remain expensive despite algorithmic ingenuity. This contrasts with purely empirical acceleration work like DDIM[13] or distillation methods, which prioritize practical speedups, and with application-driven studies such as DiffusionDepth[26] or Precipitation Nowcasting Transformer[35], which accept domain constraints. By formalizing information-computation trade-offs, Computational Bottlenecks[0] complements survey perspectives like Efficient Diffusion Survey[8] and provides a rigorous lens through which to interpret why many heuristic speedups succeed in practice yet fail to overcome worst-case complexity. Understanding these fundamental limits helps clarify which efficiency gains are algorithmic breakthroughs and which merely exploit favorable problem structure.

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

Both subtopics address computational efficiency challenges in diffusion model sampling, but from complementary angles. The original leaf focuses on fundamental computational barriers—proving when and why sampling is inherently hard through information-theoretic lower bounds and complexity-theoretic arguments. The sibling focuses on algorithmic optimization of sampling schedules (timestep/noise schedules) to reduce the number of required steps, typically through principled design frameworks rather than hardness analysis.

Similarities: - Both concern the computational cost of diffusion sampling and seek to understand or improve sampling efficiency - Both employ theoretical frameworks rather than purely empirical approaches - Both are distinguished from heuristic methods lacking formal guarantees or optimization principles

Differences: - Original leaf establishes fundamental limits and intractability results (what cannot be efficiently done), while sibling develops constructive optimization methods (how to do sampling more efficiently) - Original leaf uses information-theoretic gaps and hardness conjectures as primary tools; sibling uses optimization frameworks for schedule design - Original leaf proves computational bottlenecks exist; sibling assumes tractability and optimizes within feasible space - Original leaf excludes heuristic acceleration; sibling excludes empirical tuning without optimization framework—different exclusion criteria reflecting theory-of-limits vs. theory-of-design

Suggested Search Directions: - Papers bridging the gap: optimal schedules subject to proven computational constraints - Conditional intractability results that depend on schedule choices - Information-theoretic analysis of how schedule design affects sample complexity lower bounds

Sibling Subtopics

- **Sampling Schedule Design and Optimization** (leaves: 1, papers: 3)
- Scope: Theoretical frameworks for optimizing timestep schedules and noise schedules to reduce sampling steps.
- Exclude: Excludes empirical schedule tuning without optimization framework; see Empirical Efficiency Improvements.

Contributions Analysis

Overall novelty summary. The paper establishes fundamental computational limits for denoising diffusion sampling by proving that near-optimal score-matching objectives can yield poor samples, and that Lipschitz-constrained optimizers face inherent barriers. It resides in the 'Information-Computation Gaps and Intractability' leaf under 'Theoretical Foundations and Computational Complexity', where it is currently the sole paper. This placement reflects a sparse research direction: while the broader taxonomy contains 50 papers across 22 leaves, rigorous intractability results for diffusion sampling remain underexplored compared to the heavily populated acceleration and application branches.

The taxonomy reveals that most neighboring work focuses on practical efficiency rather than fundamental limits. The sibling branch 'Sampling Schedule Design and Optimization' contains three papers on timestep optimization, while the parent-level peer 'Sampling Acceleration Techniques' encompasses 13 papers across four leaves addressing non-Markovian methods, distillation, and training-free speedups. The paper's theoretical stance—proving impossibility results rather than proposing heuristic improvements—distinguishes it from these empirical directions. Its connection to statistical estimation gaps positions it at the intersection of complexity theory and generative modeling, a boundary less populated in the current taxonomy.

Among 22 candidates examined across three contributions, none were identified as clearly refuting the paper's claims. The first contribution (near-optimal score-matching with poor sampling) examined 10 candidates with zero refutations; the second (Lipschitz impossibility) examined 2 with zero refutations; the third (estimation-to-sampling reduction) examined 10 with zero refutations. This suggests that within the limited search scope—focused on top semantic matches and citations—the specific theoretical constructions appear novel. However, the small candidate pool (22 total) and the paper's position as the sole occupant of its taxonomy leaf indicate that the literature search may not have captured all relevant complexity-theoretic work outside the diffusion-specific domain.

Based on the examined candidates, the paper's theoretical contributions appear substantively novel within the diffusion sampling literature. The absence of refuting work among 22 candidates, combined with its unique taxonomy position, suggests it addresses a gap in rigorous complexity analysis. However, this assessment is constrained by the search scope: broader complexity theory or statistical lower bounds literature beyond the top-22 semantic matches may contain related impossibility results. The analysis captures novelty relative to the diffusion efficiency community but cannot claim exhaustiveness across all theoretical computer science.

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

Contribution 1: Existence of near score-matching optimizers that sample incorrectly

Description: The authors prove that for distributions presenting information-computation gaps, there exist polynomial-time computable drifts that nearly optimize the score-matching objective among polytime algorithms yet generate samples with distribution far from the target distribution in Wasserstein-1 distance.

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. Is Score Matching Suitable for Estimating Point Processes?

URL: [View paper](#)

Brief Assessment

Point Processes Matching[68] addresses score matching failures in point process estimation, not diffusion-based sampling from general distributions. The technical domains are distinct.

2. Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc

URL: [View paper](#)

Brief Assessment

Energy Based MCMC[61] focuses on compositional generation using energy-based diffusion models with MCMC sampling for combining multiple distributions. It does not address the theoretical problem of score-matching optimization failing to produce correct samples from target distributions, which is the core contribution of the original paper.

3. Denoising likelihood score matching for conditional score-based data generation

URL: [View paper](#)

Brief Assessment

Denoising Likelihood Matching[63] addresses score mismatch in conditional generation using Bayes' theorem decomposition, not the computational intractability of score-matching optimizers for general distributions that the original paper studies.

4. Fit like you sample: sample-efficient generalized score matching from fast mixing diffusions

URL: [View paper](#)

Brief Assessment

Fast Mixing Diffusions[67] focuses on designing generalized score matching losses from fast-mixing diffusions for multimodal distributions like Gaussian mixtures. It does not address the computational intractability of score-matching optimizers or prove existence of polynomial-time drifts that optimize score-matching yet fail to sample correctly from target distributions.

5. Statistical efficiency of score matching: The view from isoperimetry

URL: [View paper](#)

Brief Assessment

Statistical Efficiency Isoperimetry[64] focuses on the statistical efficiency of score matching estimators and their relationship to functional inequalities, not on computational intractability or information-computation gaps that lead to sampling failures in diffusion models.

6. Multimodal Reinforcement Learning With Score-Based Policy

URL: [View paper](#)

Brief Assessment

Multimodal Score Policy[70] focuses on multimodal reinforcement learning using score-based policies for action distributions, not on proving theoretical failures of score-matching optimization in diffusion sampling. The candidate addresses a different problem domain (RL policy learning) rather than the statistical estimation and sampling theory examined in the original paper.

7. Target Score Matching

URL: [View paper](#)

Brief Assessment

Target Score Matching[69] addresses variance issues in score estimation at low noise levels through alternative regression targets, but does not demonstrate that score-matching optimizers can fail to produce correct samples from target distributions.

8. Score and distribution matching policy: Advanced accelerated visuomotor policies via matched distillation

URL: [View paper](#)

Brief Assessment

Distribution Matching Policy[65] focuses on accelerating diffusion-based robotic policies through score and distribution matching for visuomotor control tasks, not on proving theoretical limitations of score-matching objectives in general statistical estimation problems.

9. Improving the Euclidean Diffusion Generation of Manifold Data by Mitigating Score Function Singularity

URL: [View paper](#)

Brief Assessment

Manifold Score Singularity[66] focuses on score function singularities arising from manifold-structured data in Euclidean diffusion models, not on information-computation gaps in score-matching optimization for general distributions.

10. LucidDreamer: Towards High-Fidelity Text-to-3D Generation via Interval Score Matching

URL: [View paper](#)

Brief Assessment

LucidDreamer[62] focuses on improving score distillation sampling (SDS) for text-to-3D generation through interval score matching, not on proving computational impossibility results for score matching in statistical estimation problems.

Contribution 2: Impossibility result for Lipschitz score-matching optimizers

Description: The authors establish that any polynomial-time drift that near-optimizes score matching, acts optimally on pure noise, and satisfies a Lipschitz continuity condition must also fail to sample correctly from the target distribution.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Consistent Diffusion Models: Mitigating Sampling Drift by Learning to be Consistent

URL: [View paper](#)

Brief Assessment

Consistent Diffusion Models[71] focuses on enforcing consistency properties during training to mitigate sampling drift, not on establishing impossibility results for Lipschitz continuous drifts in score matching. The candidate addresses a different technical problem (consistency across time) rather than computational intractability of score-matching optimizers.

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

URL: [View paper](#)

Brief Assessment

Stochasticity KL Analysis[72] focuses on the effect of stochasticity parameters in diffusion sampling and KL divergence bounds for Lipschitz score functions, not on proving that Lipschitz continuous drifts in score matching produce incorrect samples or fail to sample correctly from target distributions.

Contribution 3: Reduction from estimation to diffusion sampling

Description: The authors prove a general reduction showing that if diffusion-based sampling can be performed accurately in polynomial time, then near Bayes-optimal denoising estimation must also be possible in polynomial time, establishing a formal connection between sampling and estimation complexity.

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. Diffusion models: A comprehensive survey of methods and applications

URL: [View paper](#)

Brief Assessment

Diffusion Comprehensive Survey[51] is a broad survey paper covering applications and methods in diffusion models. It does not address the specific theoretical reduction between estimation complexity and diffusion sampling complexity that the original paper establishes.

2. Diffusion-based generative prior for low-complexity MIMO channel estimation

URL: [View paper](#)

Brief Assessment

MIMO Channel Estimation[57] focuses on channel estimation in wireless communications using diffusion models for practical signal processing applications, not on theoretical complexity reductions between estimation and sampling problems in statistical learning.

3. Sampling and Integration of Logconcave Functions by Algorithmic Diffusion

URL: [View paper](#)

Brief Assessment

Logconcave Sampling Integration[59] focuses on sampling and integration of logconcave functions using algorithmic diffusion, not on establishing reductions between estimation and diffusion sampling complexity in the context of denoising diffusions for generative modeling.

4. Diffusion Posterior Sampling for SAR Despeckling

URL: [View paper](#)

Brief Assessment

SAR Despeckling[54] focuses on image despeckling using diffusion posterior sampling for SAR images, not on establishing theoretical complexity reductions between estimation and sampling problems in diffusion models.

5. Estimating Epistemic and Aleatoric Uncertainty with a Single Model

URL: [View paper](#)

Brief Assessment

Epistemic Aleatoric Uncertainty[58] focuses on uncertainty quantification in diffusion models for practical applications, not on computational complexity reductions between estimation and sampling tasks.

6. A Simple Approach to Unifying Diffusion-based Conditional Generation

URL: [View paper](#)

Brief Assessment

Unifying Conditional Generation[53] focuses on conditional image generation using diffusion models for correlated image pairs (e.g., image-depth), not on theoretical complexity reductions between estimation and sampling problems in diffusion-based methods.

7. The limits to learning a diffusion model

URL: [View paper](#)

Brief Assessment

Limits Learning Diffusion[60] focuses on sample complexity lower bounds for epidemic/product diffusion models (SIR, Bass models), not on the computational complexity reduction between denoising estimation and diffusion-based sampling in generative models.

8. Diffusion posterior sampling for general noisy inverse problems

URL: [View paper](#)

Brief Assessment

Diffusion Posterior Sampling[52] focuses on solving inverse problems using diffusion models through posterior sampling approximations, not on establishing formal complexity reductions between estimation and sampling tasks.

9. Diffusion Model Based Channel Estimation

URL: [View paper](#)

Brief Assessment

Channel Estimation Diffusion[56] focuses on applying diffusion models to wireless channel estimation problems, not on establishing theoretical reductions between estimation and sampling complexity in general statistical settings.

10. Channel Estimation for RIS-Assisted mmWave Systems via Diffusion Models

URL: [View paper](#)

Brief Assessment

RIS Channel Estimation[55] applies diffusion models to wireless channel estimation problems, not to establishing theoretical complexity reductions between estimation and sampling tasks.

Appendix: Text Similarity Detection

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

References

- [0] Computational Bottlenecks for Denoising Diffusions [View paper](#)
- [1] Fast constrained sampling in pre-trained diffusion models [View paper](#)
- [2] Quantum denoising diffusion models [View paper](#)

- [3] Diffusion models in vision: A survey [View paper](#)
- [4] Tackling the generative learning trilemma with denoising diffusion GANs [View paper](#)
- [5] Single image super-resolution with denoising diffusion GANs [View paper](#)
- [6] gddim: Generalized denoising diffusion implicit models [View paper](#)
- [7] CoreDiff: Contextual error-modulated generalized diffusion model for low-dose CT denoising and generalization [View paper](#)
- [8] Efficient diffusion models: A survey [View paper](#)
- [9] Denoising diffusion restoration models [View paper](#)
- [10] Diff-tts: A denoising diffusion model for text-to-speech [View paper](#)
- [11] Fast sampling through the reuse of attention maps in diffusion models [View paper](#)
- [12] DiffusionDrive: Truncated Diffusion Model for End-to-End Autonomous Driving [View paper](#)
- [13] Denoising diffusion implicit models [View paper](#)
- [14] Efficiency-optimized Video Diffusion Models [View paper](#)
- [15] Improved Denoising Diffusion Probabilistic Models [View paper](#)
- [16] Optimal Inference Schedules for Masked Diffusion Models [View paper](#)
- [17] Efficient Diffusion Posterior Sampling for Noisy Inverse Problems [View paper](#)
- [18] Aligning optimization trajectories with diffusion models for constrained design generation [View paper](#)
- [19] Volumetric Conditional Score-Based Residual Diffusion Model for PET/MR Denoising [View paper](#)
- [20] Common diffusion noise schedules and sample steps are flawed [View paper](#)
- [21] Toward exploratory inverse constraint inference with generative diffusion verifiers [View paper](#)
- [22] Denoising task routing for diffusion models [View paper](#)
- [23] A Simple Early Exiting Framework for Accelerated Sampling in Diffusion Models [View paper](#)
- [24] Accelerating diffusion sampling with optimized time steps [View paper](#)
- [25] Direct Distillation: A Novel Approach for Efficient Diffusion Model Inference [View paper](#)
- [26] Diffusiondepth: Diffusion denoising approach for monocular depth estimation [View paper](#)
- [27] Accelerating Multiphase Simulations With Denoising Diffusion Model Driven Initializations [View paper](#)
- [28] Let Features Decide Their Own Solvers: Hybrid Feature Caching for Diffusion Transformers [View paper](#)
- [29] Intention-Aware Denoising Diffusion Model for Trajectory Prediction [View paper](#)
- [30] Denoising diffusion models for plug-and-play image restoration [View paper](#)
- [31] Improving diffusion inverse problem solving with decoupled noise annealing [View paper](#)
- [32] Fast inference in denoising diffusion models via mmd finetuning [View paper](#)
- [33] Consistent World Models via Foresight Diffusion [View paper](#)
- [34] Learning to efficiently sample from diffusion probabilistic models [View paper](#)
- [35] Advancing Realistic Precipitation Nowcasting With a Spatiotemporal Transformer-Based Denoising Diffusion Model [View paper](#)
- [36] Anomaly Detection of Industrial Smelting Furnace Incorporated With Accelerated Sampling Denoising Diffusion Probability Model and Conv-Transformer [View paper](#)
- [37] IMPUS: Image Morphing with Perceptually-Uniform Sampling Using Diffusion Models [View paper](#)
- [38] Partially Conditioned Patch Parallelism for Accelerated Diffusion Model Inference [View paper](#)
- [39] Image Dehazing Method Based on Diffusion Model [View paper](#)
- [40] Denoising diffusion step-aware models [View paper](#)
- [41] A conditional diffusion model with fast sampling strategy for remote sensing image super-resolution [View paper](#)
- [42] DiffLoss: unleashing diffusion model as constraint for training image restoration network [View paper](#)
- [43] Resfusion: Denoising diffusion probabilistic models for image restoration based on prior residual noise [View paper](#)
- [44] MedDiff: Generating Electronic Health Records using Accelerated Denoising Diffusion Model [View paper](#)
- [45] Semi-Supervised Learning Leveraging Denoising Diffusion Probabilistic Models for the Characterization of Nanophotonic Devices [View paper](#)
- [46] Modulated Diffusion: Accelerating Generative Modeling with Modulated Quantization [View paper](#)
- [47] Zeroth-Order Sampling Methods for Non-Log-Concave Distributions: Alleviating Metastability by Denoising Diffusion [View paper](#)
- [48] Efficient denoising diffusion via probabilistic masking [View paper](#)
- [49] Hyperspectral Image Generation with Unmixing Guided Diffusion Model [View paper](#)
- [50] Sat2City: 3D City Generation from A Single Satellite Image with Cascaded Latent Diffusion [View paper](#)
- [51] Diffusion models: A comprehensive survey of methods and applications [View paper](#)
- [52] Diffusion posterior sampling for general noisy inverse problems [View paper](#)
- [53] A Simple Approach to Unifying Diffusion-based Conditional Generation [View paper](#)
- [54] Diffusion Posterior Sampling for SAR Despeckling [View paper](#)
- [55] Channel Estimation for RIS-Assisted mmWave Systems via Diffusion Models [View paper](#)
- [56] Diffusion Model Based Channel Estimation [View paper](#)
- [57] Diffusion-based generative prior for low-complexity MIMO channel estimation [View paper](#)
- [58] Estimating Epistemic and Aleatoric Uncertainty with a Single Model [View paper](#)
- [59] Sampling and Integration of Logconcave Functions by Algorithmic Diffusion [View paper](#)
- [60] The limits to learning a diffusion model [View paper](#)
- [61] Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc [View paper](#)
- [62] LucidDreamer: Towards High-Fidelity Text-to-3D Generation via Interval Score Matching [View paper](#)
- [63] Denoising likelihood score matching for conditional score-based data generation [View paper](#)
- [64] Statistical efficiency of score matching: The view from isoperimetry [View paper](#)
- [65] Score and distribution matching policy: Advanced accelerated visuomotor policies via matched distillation [View paper](#)
- [66] Improving the Euclidean Diffusion Generation of Manifold Data by Mitigating Score Function Singularity [View paper](#)
- [67] Fit like you sample: sample-efficient generalized score matching from fast mixing diffusions [View paper](#)
- [68] Is Score Matching Suitable for Estimating Point Processes? [View paper](#)
- [69] Target Score Matching [View paper](#)
- [70] Multimodal Reinforcement Learning With Score-Based Policy [View paper](#)
- [71] Consistent Diffusion Models: Mitigating Sampling Drift by Learning to be Consistent [View paper](#)
- [72] The Effect of Stochasticity in Score-Based Diffusion Sampling: a KL Divergence Analysis [View paper](#)