Novelty Assessment Report

Paper: Efficient Approximate Posterior Sampling with Annealed Langevin Monte Carlo

PDF URL: https://openreview.net/pdf?id=7GrUROKDyW

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2026-01-07

Abstract

We study the problem of posterior sampling in the context of score based generative models. We have a trained score network for a prior $p(x)$, a measurement model $p(y|x)$, and are tasked with sampling from the posterior $p(x|y)$. Prior work has shown this to be intractable in KL (in the worst case) under well-accepted computational hardness assumptions. Despite this, popular algorithms for tasks such as image super-resolution, stylization, and reconstruction enjoy empirical success. Rather than establishing distributional assumptions or restricted settings under which exact posterior sampling is tractable, we view this as a more general "tilting" problem of biasing a distribution towards a measurement. Under minimal assumptions, we show that one can tractably sample from a distribution that is simultaneously close to the posterior of a noised prior in KL divergence and the true posterior in Fisher divergence. Intuitively, this combination ensures that the resulting sample is consistent with both the measurement and the prior. To the best of our knowledge these are the first formal results for (approximate) posterior sampling in polynomial time.

Core Task Landscape

This paper addresses: **Approximate Posterior Sampling with Score-Based Generative Models**

A total of **50 papers** were analyzed and organized into a taxonomy with **13 categories**.

Taxonomy Overview

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

- **Theoretical Foundations and Algorithmic Frameworks**
- **Methodological Extensions and Training Strategies**
- **Inverse Problem Formulations and Measurement Models**
- **Application Domains**

Complete Taxonomy Tree

- Approximate Posterior Sampling with Score-Based Generative Models Survey Taxonomy
- Theoretical Foundations and Algorithmic Frameworks
  - Convergence Theory and Provable Guarantees ★ (4 papers)
  - [0] Efficient Approximate Posterior Sampling with Annealed Langevin Monte Carlo (Anon et al., 2026) View paper
  - [5] Provable probabilistic imaging using score-based generative priors (Yu Sun, 2024) View paper
  - [12] Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction (Yuejie Chi, 2024) View paper
  - [38] Score-based generative models are provably robust: an uncertainty quantification perspective (Markos Katsoulakis, 2024) View paper
  - Geometric and Probabilistic Perspectives (2 papers)
  - [13] Geometry of Score Based Generative Models (Ghimire, 2023) View paper
  - [17] Score-Based Generative Modeling with Critically-Damped Langevin Diffusion (Dockhorn, 2021) View paper
  - Sequential and Simulation-Based Inference (7 papers)
  - [22] Score-based data assimilation (Rozet, 2023) View paper
  - [31] Sequential Posterior Sampling with Diffusion Models (Tristan S.W. Stevens, 2024) View paper
  - [37] Diffusion posterior sampling for simulation-based inference in tall data settings (Linhart, 2024) View paper
  - [39] Simulation-based inference via langevin dynamics with score matching (Jiang, 2025) View paper
  - [40] Neural score estimation: Likelihood-free inference with conditional score based diffusion models (J Simons, 2023) View paper
  - [43] Sequential neural score estimation: Likelihood-free inference with conditional score based diffusion models (Sharrock, 2022) View paper
  - [46] Nonlinear Assimilation via Score-based Sequential Langevin Sampling (Ding Zhao, 2024) View paper
- Methodological Extensions and Training Strategies
  - Training with Imperfect or Limited Data (2 papers)
  - [10] Solving inverse problems with score-based generative priors learned from noisy data (Asad Aali, 2023) View paper
  - [29] Measurement Score-Based Diffusion Model (Shoushtari, 2025) View paper
  - Alternative Sampling and Inference Strategies (5 papers)
  - [2] Importance Sampling via Score-based Generative Models (Lee Tae-Kyun, 2025) View paper
  - [6] Rethinking diffusion posterior sampling: From conditional score estimator to maximizing a posterior (Xu Tongda, 2025) View paper
  - [15] Power-scaled Bayesian Inference with Score-based Generative Models (Erdinc, 2025) View paper
  - [35] Importance Weighted Score Matching for Diffusion Samplers with Enhanced Mode Coverage (Wang, 2025) View paper
  - [50] Online Posterior Sampling with a Diffusion Prior (Aniket Deshmukh, 2024) View paper
  - Compositional and Hierarchical Inference (2 papers)

- [33] Compositional amortized inference for large-scale hierarchical Bayesian models (Pandey Vikas, 2025) View paper
- [36] Sparse Inducing Points in Deep Gaussian Processes: Enhancing Modeling with Denoising Diffusion Variational Inference (Xu Jian, 2024) View paper
- Prior Specification and Correction (2 papers)
- [44] Tackling the Problem of Distributional Shifts: Correcting Misspecified, High-Dimensional Data-Driven Priors for Inverse Problems (Gabriel Missael Barco, 2024) View paper
- [45] A score-based generative solver for PDE-constrained inverse problems with complex priors (Bansal, 2024) View paper

- Inverse Problem Formulations and Measurement Models
  - Nonlinear and Structured Measurement Models (5 papers)
  - [20] Removing structured noise using diffusion models (Stevens, 2025) View paper
  - [24] Removing structured noise with diffusion models (Stevens, 2023) View paper
  - [26] CT reconstruction using diffusion posterior sampling conditioned on a nonlinear measurement model (Shudong Li, 2024) View paper
  - [28] Diffusion Posterior Sampling for General Noisy Inverse Problems (Chung, 2022) View paper
  - [48] Diffusion posterior sampling for nonlinear CT reconstruction (Shudong Li, 2024) View paper
  - Linear Inverse Problems and Plug-and-Play Methods (3 papers)
  - [21] PSC: Posterior Sampling-Based Compression (Elata, 2024) View paper
  - [27] Score-Based Turbo Message Passing for Plug-and-Play Compressive Image Recovery (Cai Chang, 2025) View paper
  - [32] Solving inverse physics problems with score matching (Holzschuh, 2023) View paper
- Application Domains
  - Medical and Biomedical Imaging (5 papers)
  - [1] Solving inverse problems in medical imaging with score-based generative models (Song Yang, 2021) View paper
  - [30] Score-based Generative Model with Conditional Null-space Learning for Limited-angle Tomographic Reconstruction in Medical Imaging (Genyuan Zhang, 2025) View paper
  - [34] Efficient bayesian computational imaging with a surrogate score-based prior (B Feng, 2023) View paper
  - [42] Joint Reconstruction of the Activity and the Attenuation in PET by Diffusion Posterior Sampling: a Feasibility Study (ClÃ©mentine Phung-Ngoc, 2024) View paper
  - [47] Stationary ct imaging of intracranial hemorrhage with diffusion posterior sampling reconstruction (A. LÃ³pez-Montes, 2024) View paper
  - Physical Sciences and Cosmology (5 papers)
  - [4] Posterior sampling for random noise attenuation via score-based generative models (Chuangji Meng, 2025) View paper
  - [8] Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models (Ronan Legin, 2024) View paper
  - [11] Efficient and scalable posterior surrogate for seismic inversion via wavelet score-based generative models (E Cirakman, 2025) View paper
  - [23] Probabilistic mass-mapping with neural score estimation (Benjamin Remy, 2023) View paper
  - [41] Posterior samples of source galaxies in strong gravitational lenses with score-based priors (Adam, 2022) View paper
  - Signal Processing and Communications (4 papers)
  - [3] MIMO channel estimation using score-based generative models (Marius Arvinte, 2022) View paper
  - [9] HRTF Estimation using a Score-based Prior (Etienne Thuillier, 2025) View paper
  - [16] Universal Speech Enhancement with Score-based Diffusion (SerrÃ, 2022) View paper
  - [49] Diffusion Model Based Channel Estimation (Xiaochuan Ma, 2024) View paper
  - Specialized Applications and Emerging Domains (5 papers)
  - [7] Cocogen: Physically consistent and conditioned score-based generative models for forward and inverse problems (Christian Jacobsen, 2025) View paper
  - [14] Score-based Generative Modeling of Graphs via the System of Stochastic Differential Equations (Jo, 2022) View paper
  - [18] Simulation-Based Inference with Modern Generative Modelling (Simons, 2025) View paper
  - [19] Dimension reduction via score ratio matching (Baptista, 2024) View paper
  - [25] Goal-Conditioned Imitation Learning using Score-based Diffusion Policies (Reuss, 2023) View paper

Narrative

Core task: approximate posterior sampling with score-based generative models. This field centers on leveraging learned score functions—gradients of log-densities—to sample from complex posterior distributions, particularly in inverse problems where observations are corrupted or incomplete. The taxonomy reveals four main branches: Theoretical Foundations and Algorithmic Frameworks establish convergence guarantees and provable sampling schemes; Methodological Extensions and Training Strategies develop novel architectures and training procedures to improve score estimation; Inverse Problem Formulations and Measurement Models address diverse observation operators and noise structures; and Application Domains demonstrate practical impact across imaging, communications, and scientific inference. Representative works such as Diffusion Posterior Sampling[28] and Medical Imaging Score[1] illustrate how score-based priors can be combined with likelihood terms to tackle real-world reconstruction tasks, while foundational studies like Critically Damped Langevin[17] and Neural Score Estimation[40] refine the underlying sampling dynamics.

A particularly active line of work focuses on establishing rigorous convergence theory and provable guarantees, ensuring that annealed or iterative sampling schemes reliably approximate true posteriors even under model mismatch or challenging measurement operators. Annealed Langevin[0] sits squarely within this branch, emphasizing theoretical rigor alongside algorithmic design. Nearby efforts such as Provable Probabilistic Imaging[5] and Provably Robust Score[38] similarly pursue formal guarantees, though they may differ in the specific noise models or operator classes they analyze. In contrast, works like Robust Diffusion Posterior[12] prioritize empirical robustness to distributional shifts, while Importance Sampling Score[2] and Importance Weighted Score[35] explore variance reduction through reweighting strategies. These contrasting emphases—provable convergence versus practical robustness, deterministic annealing schedules versus adaptive importance weighting—highlight ongoing questions about how best to balance theoretical soundness with computational efficiency and real-world applicability.

Related Works in Same Category

The following **3 sibling papers** share the same taxonomy leaf node with the original paper:

1. Provable probabilistic imaging using score-based generative priors

Authors: Yu Sun, Zihui Wu, Yifan Chen, Berthy T. Feng, Katherine L. Bouman, et al. (6 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Estimating high-quality images while also quantifying their uncertainty are two desired features in an image reconstruction algorithm for solving ill-posed inverse problems. In this paper, we propose plug-and-play Monte Carlo (PMC) as a principled framework for characterizing the space of possible solutions to a general inverse problem. PMC is able to incorporate expressive score-based generative priors for high-quality image reconstruction while also performing uncertainty quantification via po...

Relationship Analysis

Both papers belong to the 'Convergence Theory and Provable Guarantees' category, providing formal theoretical analysis for score-based posterior sampling algorithms. They share overlapping focus on using annealed Langevin Monte Carlo methods with convergence guarantees, where the original paper establishes polynomial-time tractability results combining KL and Fisher divergence bounds, while the candidate paper (PMC) focuses on non-asymptotic stationarity guarantees in Fisher information for plug-and-play frameworks with weighted annealing. The key difference is that the original paper addresses fundamental computational hardness barriers and proposes a two-phase warm-start annealing approach, whereas the candidate develops practical plug-and-play Monte Carlo algorithms with uncertainty quantification for imaging applications.

---

2. Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction

Authors: Yuejie Chi, Xingyu Xu | **Year/Venue**: 2024 | **URL**: View paper

Abstract

â¦ 2 Score-based generative models In this section, we set up the preliminary on diffusion-based generative models, â¦ random samples âconcentratedâ around xâ², which approximates the â¦

Relationship Analysis

Both papers belong to the convergence theory and provable guarantees category for score-based posterior sampling, sharing a focus on establishing formal theoretical foundations for approximate posterior sampling algorithms. The original paper (Annealed Langevin Monte Carlo) provides polynomial-time guarantees by combining KL divergence bounds for noised posteriors with Fisher divergence bounds for the true posterior, using an annealing approach that tracks posteriors of progressively denoised priors. The candidate paper (Diffusion Plug-and-Play) develops a different algorithmic framework that alternates between proximal consistency sampling and denoising diffusion sampling, establishing both asymptotic consistency and non-asymptotic robustness guarantees for general nonlinear inverse problems with unconditional diffusion priors.

---

3. Score-based generative models are provably robust: an uncertainty quantification perspective

Authors: Markos Katsoulakis, Nikiforos Mimikos-Stamatopoulos, Benjamin Zhang | **Year/Venue**: 2024 | **URL**: View paper

Abstract

â¦ with respect to the explicit score-matching objective, or uniform-in-â¦ score-matching (DSM) objective and incorporates early stopping. â¦ Computing expectations with respect to posterior â¦

Relationship Analysis

Both papers belong to the Convergence Theory and Provable Guarantees category, providing formal theoretical analysis for score-based posterior sampling methods. While the original paper (Annealed Langevin Monte Carlo) focuses on establishing polynomial-time tractability guarantees for approximate posterior sampling by combining KL and Fisher divergence bounds through an annealing approach, the candidate paper takes a PDE regularity theory perspective to prove robustness of score-based generative models to implementation errors using Wasserstein uncertainty propagation bounds. The key distinction is that the original paper addresses the computational hardness of exact posterior sampling by proposing a practical annealing algorithm with dual convergence guarantees, whereas the candidate paper analyzes how various approximation errors (score matching, finite samples, early stopping) propagate through the generative process using Hamilton-Jacobi-Bellman PDE theory.

Contributions Analysis

Overall novelty summary. The paper contributes a tractable approximate posterior sampling algorithm with dual guarantees: closeness to a noised posterior in KL divergence and to the true posterior in Fisher divergence. It resides in the Convergence Theory and Provable Guarantees leaf, which contains four papers total, including this one. This leaf sits within the Theoretical Foundations and Algorithmic Frameworks branch, indicating a focus on formal analysis rather than application-driven heuristics. The small sibling count suggests this is a relatively sparse research direction within the broader score-based posterior sampling landscape.

The taxonomy tree shows that neighboring leaves include Geometric and Probabilistic Perspectives (two papers on Wasserstein flows and SDE frameworks) and Sequential and Simulation-Based Inference (seven papers on trajectory-based methods). The paper's dual-divergence framework diverges from purely geometric interpretations and from sequential updating schemes, instead addressing the single-step posterior sampling problem with explicit complexity bounds. The exclude_note for this leaf clarifies that empirical validation without theory belongs elsewhere, reinforcing that this work's theoretical guarantees are its distinguishing feature relative to methodological extensions in adjacent branches.

Among thirty candidates examined, the Annealed Langevin Monte Carlo algorithm contribution shows one refutable candidate out of ten examined, suggesting some overlap with prior annealing-based sampling schemes. The other two contributions—tractable dual guarantees and polynomial-time formal results—each examined ten candidates with zero refutations, indicating less direct prior work on these specific theoretical claims. The limited search scope means these statistics reflect top-thirty semantic matches, not an exhaustive survey, so additional related work may exist beyond this candidate pool.

Based on the limited search scope, the dual-divergence guarantee and polynomial-time complexity claims appear relatively novel within the examined candidates, while the annealing algorithm itself has at least one overlapping prior method. The taxonomy context confirms that provable guarantees remain an active but not overcrowded direction, with only four papers in this leaf. Acknowledging the top-thirty search boundary, the work's theoretical contributions seem to occupy a distinct niche, though broader literature may contain additional relevant comparisons.

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

Contribution 1: Tractable approximate posterior sampling with dual guarantees

Description: The authors establish that their algorithm can efficiently sample from a distribution satisfying two properties: closeness to the posterior of a noised prior in KL divergence (providing global correctness) and closeness to the true posterior in Fisher divergence (ensuring local consistency). This dual guarantee bypasses the known computational hardness of exact posterior 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. Variational inference for Neyman-Scott processes

URL: View paper

Brief Assessment

Neyman Scott Variational[65] focuses on variational inference for Neyman-Scott point processes using inclusive KL divergence minimization, not on score-based generative models with dual KL and Fisher divergence guarantees for posterior sampling.

---

2. Natural gradient variational bayes without fisher matrix analytic calculation and its inversion

URL: View paper

Brief Assessment

Natural Gradient Variational[70] focuses on variational Bayes inference using Fisher information matrix approximation, not on posterior sampling with dual KL and Fisher divergence guarantees for score-based generative models.

---

3. Operator Variational Inference

URL: View paper

Brief Assessment

Operator Variational Inference[63] focuses on designing variational objectives using operators for optimization-based inference, not on approximate posterior sampling with specific KL and Fisher divergence guarantees as described in the original contribution.

---

4. Markovian score climbing: Variational inference with KL (p|| q)

URL: View paper

Brief Assessment

Markovian Score Climbing[64] focuses on variational inference using MCMC to minimize the inclusive KL divergence KL(p||q), not on approximate posterior sampling with dual KL and Fisher divergence guarantees for score-based generative models.

---

5. RL with KL penalties is better viewed as Bayesian inference

URL: View paper

Brief Assessment

RL KL Bayesian[66] focuses on KL-regularized RL as Bayesian inference for language model alignment, not on approximate posterior sampling algorithms with dual KL/Fisher divergence guarantees for score-based generative models.

---

6. Differentially Private Statistical Inference through -Divergence One Posterior Sampling

URL: View paper

Brief Assessment

Private Posterior Sampling[68] focuses on differential privacy guarantees through β-divergence for Bayesian inference, not on the computational tractability of approximate posterior sampling with KL and Fisher divergence guarantees for score-based generative models.

---

7. Conditional variational autoencoder for sign language translation with cross-modal alignment

URL: View paper

Brief Assessment

Sign Language VAE[69] focuses on conditional variational autoencoders for sign language translation with cross-modal alignment between visual and textual modalities, not on approximate posterior sampling with KL and Fisher divergence guarantees for score-based generative models.

---

8. A Variational Analysis of Stochastic Gradient Algorithms

URL: View paper

Brief Assessment

Variational Stochastic Gradient[67] focuses on using SGD's stationary distribution for approximate inference by minimizing KL divergence, not on establishing dual KL-Fisher divergence guarantees for posterior sampling in score-based generative models.

---

9. Merging models with fisher-weighted averaging

URL: View paper

Brief Assessment

Fisher Weighted Averaging[61] focuses on merging neural network parameters using Fisher information matrices for model combination and transfer learning, not on posterior sampling algorithms with KL/Fisher divergence guarantees for score-based generative models.

---

10. Variational approximations using Fisher divergence

URL: View paper

Brief Assessment

Fisher Divergence Variational[62] focuses on variational approximations for Bayesian inference using Fisher divergence minimization, not on posterior sampling algorithms with dual KL/Fisher guarantees for score-based generative models.

Contribution 2: Annealed Langevin Monte Carlo algorithm for posterior sampling

Description: The authors propose the Annealed Langevin Monte Carlo (ALMC) algorithm that uses a warm-start phase followed by an annealing phase to track posteriors of progressively denoised priors. The method provides polynomial-time guarantees under minimal assumptions on the prior and measurement operator.

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. Mean-Field Langevin Dynamics : Exponential Convergence and Annealing

URL: View paper

Brief Assessment

Mean Field Langevin[76] focuses on mean-field Langevin dynamics for minimizing convex functions over measures with entropy regularization, not posterior sampling with score-based generative models. The annealing in [76] refers to noise decay for optimization, while the original paper's annealing tracks posteriors of progressively denoised priors for sampling tasks.

---

2. Unbiased Kinetic Langevin Monte Carlo with Inexact Gradients

URL: View paper

Brief Assessment

Unbiased Kinetic Langevin[75] focuses on kinetic Langevin dynamics with multilevel Monte Carlo coupling to achieve unbiased estimation, not on annealing methods for tracking posteriors of progressively denoised priors as in the original paper's ALMC algorithm.

3. Continuously tempered diffusion samplers

URL: View paper

Brief Assessment

Continuously Tempered Diffusion[80] focuses on neural samplers for general unnormalized distributions using continuous tempering techniques, not specifically on posterior sampling with score-based generative models or measurement operators as in the original paper.

4. Generative Modeling by Estimating Gradients of the Data Distribution

URL: View paper

Prior Art Analysis

Estimating Gradients Distribution[73] demonstrates that annealed Langevin dynamics for sampling was proposed and implemented prior to the original paper's work. The candidate paper presents annealed Langevin dynamics as a core sampling method where noise levels are gradually decreased during the sampling process, using scores corresponding to progressively lower noise levels. This directly corresponds to the original paper's claimed contribution of using 'a warm-start phase followed by an annealing phase to track posteriors of progressively denoised priors.'

Evidence

Evidence 1 - **Rationale**: Both papers describe annealed Langevin dynamics as a method that progressively reduces noise levels during sampling, demonstrating that this approach existed before the original paper's submission. - **Original**: we introduce a notion of posterior sampling that is possible in polynomial time, bypassing the hardness of sampling in kl. we develop guarantees with our method annealed langevin monte carlo (almc, algorithm 1) in the general regime where the influences of the prior and the likelihood might be in co... - **Candidate**: we propose an annealed version of langevin dynamics, where we initially use scores corresponding to the highest noise level, and gradually anneal down the noise level until it is small enough to be indistinguishable from the original data distribution. our sampling strategy is inspired by simulated ...

Evidence 2 - **Rationale**: Both papers describe the same annealing procedure: starting from an initial distribution and progressively moving through distributions with decreasing noise levels to reach the target posterior. - **Original**: we start with a sample that disregards the prior entirely - emphasizing only consistency with the likelihood. this sample is then annealed towards the true posterior by drawing its marginal closer to the posteriors of progressively denoised priors. - **Candidate**: as shown in alg. 1, we start annealed langevin dynamics by initializing the samples from some fixed prior distribution, e.g., uniform noise. then, we run langevin dynamics to sample from qσ1 (x) with step size α1. next, we run langevin dynamics to sample from qσ2 (x), starting from the final samples o...

Evidence 3 - **Rationale**: Both papers describe the core mechanism of transitioning between posteriors at different noise levels, which is the fundamental contribution claimed in the original paper. - **Original**: other than at polynomially low noise levels, we show that using almc we can efficiently transition from the posterior of a noised prior to a posterior of a slightly less noised prior. - **Candidate**: we continue in this fashion, using the final samples of langevin dynamics for qσi-1 (x) as the initial samples of langevin dynamic for qσi(x), and tuning down the step size αi gradually with αi = ϵ·σ2 i/σ2 l. finally, we run langevin dynamics to sample from qσl(x), which is close to pdata(x) when σl ...

5. Diffusion posterior sampling for magnetic resonance imaging

URL: View paper

Brief Assessment

MRI Diffusion Posterior[78] focuses on diffusion posterior sampling for MRI reconstruction tasks, while the original paper develops theoretical guarantees for annealed Langevin Monte Carlo in general posterior sampling settings with polynomial-time complexity bounds.

6. Simulated Tempering Langevin Monte Carlo II: An Improved Proof using Soft Markov Chain Decomposition

URL: View paper

Brief Assessment

Simulated Tempering Langevin[79] focuses on sampling from multimodal distributions (mixtures of log-concave distributions) using simulated tempering combined with Langevin dynamics, not on posterior sampling with score-based generative models or annealing through progressively denoised priors as in the original paper.

7. Annealed Langevin Dynamics for Massive MIMO Detection

URL: View paper

Brief Assessment

Annealed Langevin MIMO[72] applies annealed Langevin dynamics to MIMO symbol detection, a specific signal processing problem. The original paper addresses general posterior sampling with score-based generative models. These are fundamentally different problem domains with different technical objectives.

8. Improving diffusion inverse problem solving with decoupled noise annealing

URL: View paper

Brief Assessment

Decoupled Noise Annealing[77] focuses on diffusion-based inverse problem solving with a decoupled noise annealing process for image restoration tasks, not on general posterior sampling with polynomial-time guarantees under minimal assumptions as in the original paper's ALMC framework.

9. Posterior Sampling via Langevin Monte Carlo for Offline Reinforcement Learning

URL: View paper

Brief Assessment

Langevin Offline RL[74] focuses on offline reinforcement learning using Langevin Monte Carlo for value-based posterior sampling in MDPs, not on general posterior sampling in score-based generative models with measurement operators as in the original paper.

10. Penalized overdamped and underdamped Langevin Monte Carlo algorithms for constrained sampling

URL: View paper

Brief Assessment

Penalized Langevin Constrained[71] focuses on constrained sampling using penalty methods for convex constraint sets, not on annealing methods for posterior sampling in score-based generative models. The candidate addresses a fundamentally different problem setting.

Contribution 3: First formal polynomial-time results for approximate posterior sampling

Description: The authors claim to provide the first theoretical guarantees showing that approximate posterior sampling can be achieved in polynomial time, addressing a gap in the literature where prior work either lacked formal guarantees or established computational hardness for exact 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. Polynomial Time Algorithms for Dual Volume Sampling

URL: View paper

Brief Assessment

Dual Volume Sampling[56] addresses polynomial-time algorithms for column subset selection from matrices, not posterior sampling in score-based generative models. The technical domains are entirely different.

---

2. Stochastic Gradient Descent as Approximate Bayesian Inference

URL: View paper

Brief Assessment

SGD Bayesian Inference[52] focuses on using constant-rate SGD as approximate Bayesian inference by analyzing its stationary distribution, not on polynomial-time guarantees for posterior sampling in score-based generative models with annealing.

---

3. Transport monte carlo: High-accuracy posterior approximation via random transport

URL: View paper

Brief Assessment

Transport Monte Carlo[55] focuses on optimization-based posterior approximation via random transport plans between distributions, not on polynomial-time complexity guarantees for approximate posterior sampling algorithms.

---

4. Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference

URL: View paper

Brief Assessment

Bayesian CNN Bernoulli[57] focuses on Bayesian inference in convolutional neural networks using dropout as variational inference, not on polynomial-time guarantees for approximate posterior sampling in score-based generative models.

---

5. A Modified FC-Gram Approximation Algorithm with Provable Error Bounds

URL: View paper

Brief Assessment

FC Gram Approximation[59] focuses on trigonometric polynomial approximation for non-periodic functions and PDE solving, not posterior sampling or Langevin Monte Carlo algorithms.

---

6. Expectation Propogation for Approximate Inference in Dynamic Bayesian Networks

URL: View paper

Brief Assessment

Expectation Propagation DBN[54] focuses on approximate inference in dynamic Bayesian networks using message-passing algorithms, not on posterior sampling in score-based generative models with formal polynomial-time guarantees.

---

7. Bayes-Newton methods for approximate Bayesian inference with PSD guarantees

URL: View paper

Brief Assessment

Bayes Newton Methods[51] focuses on approximate Bayesian inference using natural gradient variational inference, expectation propagation, and posterior linearisation as generalisations of Newton's method. It does not address posterior sampling in score-based generative models or provide polynomial-time guarantees for approximate posterior sampling as claimed in the original paper.

---

8. Efficient Variational Inference for Sparse Deep Learning with Theoretical Guarantee

URL: View paper

Brief Assessment

Sparse Variational Inference[60] focuses on variational inference for sparse deep neural networks under spike-and-slab priors, not on posterior sampling in score-based generative models or diffusion contexts. The computational guarantees are for variational posterior contraction, not for the tilting/posterior sampling problem addressed in the original paper.

---

9. Data augmentation MCMC for Bayesian inference from privatized data

URL: View paper

Brief Assessment

Private Data MCMC[58] addresses Bayesian inference from privatized data using data augmentation MCMC, not approximate posterior sampling with score-based generative models. The technical focus is entirely different - privacy mechanisms versus diffusion models.

---

10. Differentially private approximate Bayesian inference of probabilistic models

URL: View paper

Brief Assessment

Private Approximate Bayesian[53] focuses on differential privacy constraints for Bayesian inference methods (VI and MCMC), not on polynomial-time complexity guarantees for approximate posterior sampling in score-based generative models.

Appendix: Text Similarity Detection

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

References

- [0] Efficient Approximate Posterior Sampling with Annealed Langevin Monte Carlo View paper
- [1] Solving inverse problems in medical imaging with score-based generative models View paper
- [2] Importance Sampling via Score-based Generative Models View paper

- [3] MIMO channel estimation using score-based generative models View paper
- [4] Posterior sampling for random noise attenuation via score-based generative models View paper
- [5] Provable probabilistic imaging using score-based generative priors View paper
- [6] Rethinking diffusion posterior sampling: From conditional score estimator to maximizing a posterior View paper
- [7] Cocogen: Physically consistent and conditioned score-based generative models for forward and inverse problems View paper
- [8] Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models View paper
- [9] HRTF Estimation using a Score-based Prior View paper
- [10] Solving inverse problems with score-based generative priors learned from noisy data View paper
- [11] Efficient and scalable posterior surrogate for seismic inversion via wavelet score-based generative models View paper
- [12] Provably robust score-based diffusion posterior sampling for plug-and-play image reconstruction View paper
- [13] Geometry of Score Based Generative Models View paper
- [14] Score-based Generative Modeling of Graphs via the System of Stochastic Differential Equations View paper
- [15] Power-scaled Bayesian Inference with Score-based Generative Models View paper
- [16] Universal Speech Enhancement with Score-based Diffusion View paper
- [17] Score-Based Generative Modeling with Critically-Damped Langevin Diffusion View paper
- [18] Simulation-Based Inference with Modern Generative Modelling View paper
- [19] Dimension reduction via score ratio matching View paper
- [20] Removing structured noise using diffusion models View paper
- [21] PSC: Posterior Sampling-Based Compression View paper
- [22] Score-based data assimilation View paper
- [23] Probabilistic mass-mapping with neural score estimation View paper
- [24] Removing structured noise with diffusion models View paper
- [25] Goal-Conditioned Imitation Learning using Score-based Diffusion Policies View paper
- [26] CT reconstruction using diffusion posterior sampling conditioned on a nonlinear measurement model View paper
- [27] Score-Based Turbo Message Passing for Plug-and-Play Compressive Image Recovery View paper
- [28] Diffusion Posterior Sampling for General Noisy Inverse Problems View paper
- [29] Measurement Score-Based Diffusion Model View paper
- [30] Score-based Generative Model with Conditional Null-space Learning for Limited-angle Tomographic Reconstruction in Medical Imaging View paper
- [31] Sequential Posterior Sampling with Diffusion Models View paper
- [32] Solving inverse physics problems with score matching View paper
- [33] Compositional amortized inference for large-scale hierarchical Bayesian models View paper
- [34] Efficient bayesian computational imaging with a surrogate score-based prior View paper
- [35] Importance Weighted Score Matching for Diffusion Samplers with Enhanced Mode Coverage View paper
- [36] Sparse Inducing Points in Deep Gaussian Processes: Enhancing Modeling with Denoising Diffusion Variational Inference View paper
- [37] Diffusion posterior sampling for simulation-based inference in tall data settings View paper
- [38] Score-based generative models are provably robust: an uncertainty quantification perspective View paper
- [39] Simulation-based inference via langevin dynamics with score matching View paper
- [40] Neural score estimation: Likelihood-free inference with conditional score based diffusion models View paper
- [41] Posterior samples of source galaxies in strong gravitational lenses with score-based priors View paper
- [42] Joint Reconstruction of the Activity and the Attenuation in PET by Diffusion Posterior Sampling: a Feasibility Study View paper
- [43] Sequential neural score estimation: Likelihood-free inference with conditional score based diffusion models View paper
- [44] Tackling the Problem of Distributional Shifts: Correcting Misspecified, High-Dimensional Data-Driven Priors for Inverse Problems View paper
- [45] A score-based generative solver for PDE-constrained inverse problems with complex priors View paper
- [46] Nonlinear Assimilation via Score-based Sequential Langevin Sampling View paper
- [47] Stationary ct imaging of intracranial hemorrhage with diffusion posterior sampling reconstruction View paper
- [48] Diffusion posterior sampling for nonlinear CT reconstruction View paper
- [49] Diffusion Model Based Channel Estimation View paper
- [50] Online Posterior Sampling with a Diffusion Prior View paper
- [51] Bayes-Newton methods for approximate Bayesian inference with PSD guarantees View paper
- [52] Stochastic Gradient Descent as Approximate Bayesian Inference View paper
- [53] Differentially private approximate Bayesian inference of probabilistic models View paper
- [54] Expectation Propogation for Approximate Inference in Dynamic Bayesian Networks View paper
- [55] Transport monte carlo: High-accuracy posterior approximation via random transport View paper
- [56] Polynomial Time Algorithms for Dual Volume Sampling View paper
- [57] Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference View paper
- [58] Data augmentation MCMC for Bayesian inference from privatized data View paper
- [59] A Modified FC-Gram Approximation Algorithm with Provable Error Bounds View paper
- [60] Efficient Variational Inference for Sparse Deep Learning with Theoretical Guarantee View paper
- [61] Merging models with fisher-weighted averaging View paper
- [62] Variational approximations using Fisher divergence View paper
- [63] Operator Variational Inference View paper
- [64] Markovian score climbing: Variational inference with KL (p|| q) View paper
- [65] Variational inference for Neyman-Scott processes View paper
- [66] RL with KL penalties is better viewed as Bayesian inference View paper
- [67] A Variational Analysis of Stochastic Gradient Algorithms View paper
- [68] Differentially Private Statistical Inference through -Divergence One Posterior Sampling View paper
- [69] Conditional variational autoencoder for sign language translation with cross-modal alignment View paper
- [70] Natural gradient variational bayes without fisher matrix analytic calculation and its inversion View paper

- [71] Penalized overdamped and underdamped Langevin Monte Carlo algorithms for constrained sampling View paper
- [72] Annealed Langevin Dynamics for Massive MIMO Detection View paper
- [73] Generative Modeling by Estimating Gradients of the Data Distribution View paper
- [74] Posterior Sampling via Langevin Monte Carlo for Offline Reinforcement Learning View paper
- [75] Unbiased Kinetic Langevin Monte Carlo with Inexact Gradients View paper
- [76] Mean-Field Langevin Dynamics : Exponential Convergence and Annealing View paper
- [77] Improving diffusion inverse problem solving with decoupled noise annealing View paper
- [78] Diffusion posterior sampling for magnetic resonance imaging View paper
- [79] Simulated Tempering Langevin Monte Carlo II: An Improved Proof using Soft Markov Chain Decomposition View paper
- [80] Continuously tempered diffusion samplers View paper