Novelty Assessment Report

Paper: Energy-Efficient Random Variate Generation via Compressed Lookup Tables
PDF URL: https://openreview.net/pdf?id=hRY0ytSnM0
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2025-12-29

Abstract

Generating (pseudo-)random variates lies at the core of probabilistic machine learning and prediction algorithms and yet remains a major bottleneck due to its high computational and energy cost. In this paper, we introduce a general and scalable sampling strategy that enables fast and energy-efficient random variate generation from arbitrary distributions. Our approach is based on efficient lookup tables combined with a fast index sampling scheme. Using only a handful of fast and energy-efficient compute operations on simple array structures, we achieve superior speed, energy efficiency, and precision at near-optimal entropy cost compared to state-of-the-art techniques. Microbenchmarking our approach with a C implementation shows up to 40\% savings in time and 60\% in energy compared to state-of-the-art approaches. Compared to commonly employed Python samplers we achieve a 100x time improvement.

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: **random variate generation from arbitrary discrete distributions**

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

Taxonomy Overview

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

- **Foundational Sampling Algorithms and Theoretical Frameworks**
- **Implementation Efficiency and Computational Optimization**
- **Specialized Distribution Classes and Sampling Contexts**
- **Application Domains and Domain-Specific Methods**
- **Theoretical Foundations and Statistical Properties**
- **General References and Pedagogical Resources**

Complete Taxonomy Tree

- random variate generation from arbitrary discrete distributions Survey Taxonomy
- Foundational Sampling Algorithms and Theoretical Frameworks
  - Exact Sampling with Entropy Optimality (4 papers)
  - [2] Optimal approximate sampling from discrete probability distributions (Feras A. Saad, 2019) View paper
  - [16] The Fast Loaded Dice Roller: A Near-Optimal Exact Sampler for Discrete Probability Distributions (Saad, 2020) View paper
  - [18] Efficient Rejection Sampling in the Entropy-Optimal Range (Draper Thomas L, 2025) View paper
  - [38] Random variate generation using only finitely many unbiased, independently and identically distributed random bits (Devroye, 2015) View paper
  - Alias Method and Table-Based Techniques (3 papers)
  - [10] On the alias method for generating random variables from a discrete distribution (Richard A. Kronmal, 1979) View paper
  - [12] Succinct sampling from discrete distributions (Karl Bringmann, 2013) View paper
  - [36] Massively Parallel Construction of Radix Tree Forests for the Efficient Sampling of Discrete Probability Distributions (Nikolaus Binder, 2018) View paper
  - Rejection and Approximate Sampling (3 papers)
  - [23] Rejection-Inversion to Generate Variates from Monotone Discrete Distributions (W. HÃ¶rmann, 2024) View paper
  - [26] Sampling from Discrete Energy-Based Models with Quality/Efficiency Trade-offs (Bryan Eikema, 2021) View paper
  - [48] Adaptive Greedy Rejection Sampling (Gergely Flamich, 2023) View paper
  - Formal Verification and Correctness Guarantees (2 papers)
  - [15] Random variate generation with formal guarantees (Feras A. Saad, 2025) View paper
  - [19] Formally verified samplers from probabilistic programs with loops and conditioning (Alexander Bagnall, 2023) View paper
  - Binary and Direct Sampling Methods (2 papers)
  - [17] Binary sampling from discrete distributions (Hiroyuki Masuyama, 2022) View paper
  - [44] Efficient Sampling Methods for Discrete Distributions (Konstantinos Panagiotou, 2016) View paper
- Implementation Efficiency and Computational Optimization
  - Energy and Speed Optimization ★ (2 papers)
  - [0] Energy-Efficient Random Variate Generation via Compressed Lookup Tables (Anon et al., 2026) View paper
  - [27] On the Average Runtime of an Open Source Binomial Random Variate Generation Algorithm (Cicirello, 2024) View paper
  - Dynamic and Space-Efficient Structures (2 papers)
  - [4] Maintaining Discrete Probability Distributions in Practice (Daniel Allendorf, 2023) View paper
  - [33] Dynamic sampling from a discrete probability distribution with a known distribution of rates (Federico Dâ Ambrosio, 2021) View paper
  - Parallel and Massively Parallel Algorithms (1 papers)

    - [25] Arithmetic sampling: parallel diverse decoding for large language models (Vilnis, 2023) View paper
    - Numerical Precision and Floating-Point Issues (1 papers)
    - [32] Are we there yet? timing and floating-point attacks on differential privacy systems (Jiankai Jin, 2022) View paper
- Specialized Distribution Classes and Sampling Contexts
    - Specific Discrete Distribution Families (2 papers)
    - [1] 109.15 Sampling from the hypergeometric probability distribution (Geoffrey W. Brown, 2025) View paper
    - [47] PhaseTypeR: phase-type distributions in R with reward transformations and a view towards population genetics (Iker Rivas-GonzÃ¡lez, 2022) View paper
    - Graphical Models and Structured Distributions (3 papers)
    - [11] Quantum circuits for discrete graphical models (Nico Piatkowski, 2024) View paper
    - [22] Sampling hierarchies of discrete random structures (A. Lijoi, 2020) View paper
    - [45] On Quantum Circuits for Discrete Graphical Models (Piatkowski, 2022) View paper
    - Conditional and Constrained Sampling (2 papers)
    - [14] Plug-and-Play Controllable Generation for Discrete Masked Models (Guo Wei, 2024) View paper
    - [39] Conditional sampling for spectrally discrete max-stable random fields (Yizao Wang, 2010) View paper
    - Continuous-to-Discrete Transformations (4 papers)
    - [3] Discrete Flow Matching (Gat, 2024) View paper
    - [13] Sampling from Arbitrary Functions via PSD Models (Marteau-Ferey, 2021) View paper
    - [29] Convergence Analysis of Discrete Diffusion Model: Exact Implementation through Uniformization (Chen Hongrui, 2024) View paper
    - [31] Sampling from Binary Quadratic Distributions via Stochastic Localization (Wang, 2025) View paper
- Application Domains and Domain-Specific Methods
    - Generative Modeling and Neural Networks (6 papers)
    - [5] Diverse Semantic Image Synthesis via Probability Distribution Modeling (Zhentao Tan, 2021) View paper
    - [6] M2D2M: Multi-Motion Generation from Text with Discrete Diffusion Models (Chi Hyung-gun, 2024) View paper
    - [21] Super resolution enhanced Multi-view Stereo network based on Gumbel sampling (Shichao Wang, 2025) View paper
    - [24] Differentiable Sampling of Categorical Distributions Using the CatLog-Derivative Trick (De Smet, 2023) View paper
    - [30] The Sampling-Gaussian for Stereo Matching (Pan, 2024) View paper
    - [41] Improving deep learning powered auction design (Shuyuan You, 2023) View paper
    - Privacy-Preserving and Statistical Applications (1 papers)
    - [50] Mutual information optimally local private discrete distribution estimation (Wang Shaowei, 2016) View paper
    - Simulation and Stochastic Modeling (3 papers)
    - [7] Generating Pseudo-Random Discrete Probability Distributions (Jonas Maziero, 2015) View paper
    - [34] Distance-based sampling of software configuration spaces (Christian Kaltenecker, 2019) View paper
    - [42] Stochastic Optimal Power Flow in Hybrid Power System Using Reduced-Discrete Point Estimation Method and Latin Hypercube Sampling (Reza Taghavi, 2022) View paper
- Theoretical Foundations and Statistical Properties
    - Distribution Theory and Characterization (2 papers)
    - [35] Random discrete distributions (J. F. C. Kingman, 1975) View paper
    - [40] A sampling theorem for finite discrete distributions (B.C. Brookes, 1975) View paper
    - Statistical Inference and Quantile Estimation (1 papers)
    - [9] A Study on Computing Sample Quantiles of Discrete Probability Distributions (Hyuk Joo Kim, 2024) View paper
- General References and Pedagogical Resources (7 papers)
    - [8] Introduction to probability models (Winston, 2014) View paper
    - [20] Random number generation (Ming-Yang Kao, 2011) View paper
    - [28] Discrete probability models and methods (Pierre BrÃ©maud, 2017) View paper
    - [37] Random variate generation (S. Bandyopadhyay, 1998) View paper
    - [43] How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph (J. Propp, 1998) View paper
    - [46] Essentials of Monte Carlo simulation: Statistical methods for building simulation models (Nicholas T. Thomopoulos, 2012) View paper
    - [49] Discrete distributions (D. D. Wackerly, 2004) View paper

Narrative

Core task: random variate generation from arbitrary discrete distributions. The field organizes around several complementary perspectives. Foundational branches establish classical algorithms and theoretical frameworks—such as the Alias Method[10] and inversion techniques—that underpin most modern samplers. Implementation efficiency and computational optimization focus on reducing time and space costs through data structures, preprocessing, and energy-aware designs. Specialized distribution classes address particular probability models (e.g., hypergeometric, binomial) or sampling contexts (e.g., quantum circuits, hierarchical structures), while application domains tailor methods to real-world settings like auction design or power flow analysis. Theoretical foundations explore statistical properties, entropy bounds, and formal guarantees, and general references provide pedagogical resources for practitioners.

Within the optimization-oriented branches, a handful of works pursue speed and memory trade-offs using compressed representations or adaptive rejection schemes. For instance, Fast Loaded Dice[16] and Optimal Approximate Sampling[2] explore how to balance preprocessing overhead against query time, while Binomial Runtime[27] examines distribution-specific optimizations. Compressed Lookup Tables[0] sits squarely in this energy and speed optimization cluster, emphasizing compact data structures that reduce both memory footprint and access latency. Compared to neighbors like Binomial Runtime[27], which targets a specific distribution family, Compressed Lookup Tables[0] aims for broader applicability across arbitrary discrete distributions by leveraging compression techniques. This positions the work as a bridge between classical table-based methods and modern resource-constrained environments, addressing the perennial tension between generality and efficiency in discrete sampling.

Related Works in Same Category

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

1. On the Average Runtime of an Open Source Binomial Random Variate Generation Algorithm

Authors: Cicirello, Vincent A., Vincent Cicirello, Vincent A. Cicirello | **Year/Venue**: 2024 • arXiv.org | **URL**: View paper

Abstract

The BTPE algorithm (Binomial, Triangle, Parallelogram, Exponential) of Kachitvichyanukul and Schmeiser is one of the faster and more widely utilized algorithms for generating binomial random variates. Cicirello's open source Java library, $\rho\mu$, includes an implementation of BTPE as well as a variety of other random number related utilities. In this report, I explore the average case runtime of the BTPE algorithm when generating random values from binomial distribution $B(n,p)$. Beginning wi...

Relationship Analysis

Both papers belong to the Energy and Speed Optimization category, focusing on computational efficiency in random variate generation. While the original paper presents a novel compressed lookup table (cLUT) method for general discrete distributions with emphasis on energy efficiency and speed through optimized data structures, the candidate paper analyzes the average runtime characteristics of the existing BTPE algorithm specifically for binomial distributions, providing theoretical and experimental validation of its convergence properties. The key difference is that the original paper proposes a new general-purpose sampling method with energy-efficiency goals, whereas the candidate paper performs runtime analysis of a specific existing algorithm for a particular distribution family.

Contributions Analysis

Overall novelty summary. The paper proposes a compressed lookup table (cLUT) sampling method for generating random variates from arbitrary discrete distributions, emphasizing speed and energy efficiency. It resides in the 'Energy and Speed Optimization' leaf under 'Implementation Efficiency and Computational Optimization', which contains only two papers total (including this one). This leaf focuses specifically on computational speed and energy metrics through compressed structures, distinguishing it from adjacent leaves that address space complexity or parallel execution. The sparse population of this leaf suggests that energy-aware sampling optimizations remain relatively underexplored in the literature.

The taxonomy tree reveals several neighboring research directions. The 'Alias Method and Table-Based Techniques' leaf (three papers) covers classical constant-time sampling structures, while 'Dynamic and Space-Efficient Structures' (two papers) addresses succinct representations and dynamic updates. The 'Rejection and Approximate Sampling' leaf (three papers) explores alternative algorithmic paradigms. The paper's focus on compressed tables positions it at the intersection of classical table-driven methods and modern resource constraints, diverging from purely theoretical optimality (covered in 'Exact Sampling with Entropy Optimality') and from domain-specific applications (e.g., 'Generative Modeling and Neural Networks' with six papers).

Among thirty candidates examined across three contributions, none were identified as clearly refuting the paper's claims. The 'compressed lookup table sampling method' examined ten candidates with zero refutable matches, as did the 'comprehensive benchmarking' and 'practical impact demonstration' contributions. This suggests that within the limited search scope—focused on top-K semantic matches and citation expansion—no prior work directly overlaps with the specific combination of compression techniques, energy metrics, and arbitrary distribution support. However, the small candidate pool (thirty papers) and the presence of only one sibling paper in the same taxonomy leaf indicate that the analysis covers a narrow slice of potentially relevant literature.

Given the limited search scope and sparse taxonomy leaf, the work appears to occupy a relatively unexplored niche combining energy efficiency with general-purpose discrete sampling. The absence of refutable candidates among thirty examined papers, alongside the leaf's low population, suggests novelty within the analyzed subset. However, the analysis does not cover exhaustive literature on table-based methods or energy-aware computing more broadly, leaving open the possibility of related work in adjacent fields or implementation-focused venues not captured by semantic search.

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

Contribution 1: Compressed lookup table (cLUT) sampling method

Description: The authors introduce a new sampling algorithm that uses compressed lookup tables with a lossless compression scheme to generate random variates from arbitrary discrete distributions. The compression achieves an exponential ratio while enabling fast, energy-efficient sampling through simple array operations and direct indexing without conditional branching.

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. A Hierarchical Parallel Discrete Gaussian Sampler for Lattice-Based Cryptography

URL: View paper

Brief Assessment

Hierarchical Gaussian Sampler[64] focuses on discrete Gaussian sampling for lattice-based cryptography using the Knuth-Yao algorithm with hierarchical bit search units, not on general compressed lookup table methods for arbitrary discrete distributions with lossless compression schemes.

---

2. Machine learning based symbol probability distribution prediction for entropy coding In AV1

URL: View paper

Brief Assessment

Symbol Probability Prediction[65] focuses on entropy coding in video compression (AV1 codec), using lookup tables to store predicted probability distributions for syntax elements. The original paper addresses random variate generation from arbitrary distributions for machine learning applications, not video codec entropy coding.

---

3. Hierarchical compression of color look up tables

URL: View paper

Brief Assessment

Color Lookup Compression[67] addresses compression of color lookup tables for color management in printers, not random variate generation from probability distributions. The compression scheme and application domain are fundamentally different from the original paper's statistical sampling method.

---

4. High precision discrete Gaussian sampling on FPGAs

URL: View paper

Brief Assessment

[Final Audit Failure] The model insisted on a refutation claim but failed to provide verifiable evidence after multiple retries. Marked as cannot_refute for safety. Please manually verify the candidate text.

---

5. On practical discrete Gaussian samplers for lattice-based cryptography

URL: View paper

Brief Assessment

Practical Gaussian Samplers[66] focuses on discrete Gaussian sampling for lattice-based cryptography using cumulative distribution tables (CDT), not general compressed lookup tables with lossless compression for arbitrary distributions. The technical approaches and application domains differ fundamentally.

6. CARAMEL: A Succinct Read-Only Lookup Table via Compressed Static Functions

URL: View paper

Brief Assessment

CARAMEL[69] focuses on compressed static functions for key-value lookups in data structures (e.g., tokenization, genomics), not on random variate generation from probability distributions. The compression techniques and application domains are fundamentally different.

7. Efficient Lossless Compression with Distribution Quantized Finite-State Autoregressive Model

URL: View paper

Brief Assessment

Distribution Quantized Model[63] focuses on image compression using finite-state autoregressive entropy coding with quantized lookup tables, not on general random variate generation from arbitrary discrete distributions.

8. DeltaPQ: lossless product quantization code compression for high dimensional similarity search

URL: View paper

Brief Assessment

DeltaPQ[62] focuses on lossless compression of product quantization codes for high-dimensional similarity search, not on compressed lookup table sampling methods for generating random variates from discrete distributions. The compression techniques serve entirely different purposes in different domains.

9. Demystifying Diffusion Policies: Action Memorization and Simple Lookup Table Alternatives

URL: View paper

Brief Assessment

Diffusion Policies[68] focuses on action memorization in robot manipulation using lookup tables for action sequences, not on compressed lookup table sampling methods for discrete probability distributions with lossless compression schemes.

10. Finite-state autoregressive entropy coding for efficient learned lossless compression

URL: View paper

Brief Assessment

Autoregressive Entropy Coding[61] focuses on entropy coding for learned lossless compression using finite-state autoregressive models with lookup tables, not on general random variate generation from arbitrary discrete distributions with lossless compression achieving exponential compression ratios.

Contribution 2: Comprehensive benchmarking against state-of-the-art methods

Description: The authors provide extensive empirical comparisons of cLUT with existing state-of-the-art sampling methods (ALDR, FLDR, Alias method) across multiple metrics including speed, energy consumption, memory usage, and bit efficiency, demonstrating superior performance especially for larger distributions.

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. Constant-time discrete gaussian sampling

URL: View paper

Brief Assessment

Constant Time Gaussian[54] focuses on constant-time discrete Gaussian sampling for lattice-based cryptography, comparing only against CDT-based sampling and Knuth-Yao with shuffling. The original paper provides comprehensive benchmarks across multiple metrics (speed, energy, memory, bit efficiency) for general discrete distribution sampling methods including ALDR, FLDR, and Alias method.

2. High precision discrete Gaussian sampling on FPGAs

URL: View paper

Brief Assessment

High Precision Gaussian[55] focuses on FPGA hardware implementation for discrete Gaussian sampling in lattice-based cryptography, not general discrete distribution sampling. The candidate does not provide comprehensive benchmarking across multiple metrics (speed, energy, memory, bit efficiency) for general sampling methods.

3. Energy-Efficient Reconfigurable Acceleration Engine for Polynomial Coefficient Generation of Lattice-Based Post-Quantum Cryptography

URL: View paper

Brief Assessment

Reconfigurable Acceleration Engine[52] focuses on hardware acceleration for lattice-based post-quantum cryptography sampling (discrete gaussian, binomial, uniform-reject), not general discrete distribution sampling methods like ALDR, FLDR, or Alias method evaluated in the original paper.

4. Gradient-Guided Importance Sampling for Learning Discrete Energy-Based Models

URL: View paper

Brief Assessment

Gradient Guided Importance[56] focuses on learning discrete energy-based models through ratio matching with gradient-guided importance sampling, not on benchmarking discrete distribution sampling methods for speed, energy efficiency, and memory usage as in the original paper's contribution.

5. Impact of ray generation schemes on the random ray method for eigenvalue and shielding applications

URL: View paper

Brief Assessment

Ray Generation Schemes[59] focuses on random ray methods for nuclear reactor physics applications (eigenvalue and shielding problems), not discrete distribution sampling methods. The candidate's benchmarking context is entirely different from the original paper's focus on sampling speed, energy efficiency, and memory usage for discrete distributions.

6. A memory and gate efficient algorithm for unitary mixed Schur sampling

URL: View paper

Brief Assessment

Unitary Schur Sampling[60] focuses on quantum computing algorithms for unitary group sampling and Schur transforms, not on benchmarking discrete distribution sampling methods for classical computing in terms of speed, energy, and memory.

7. Benchmarking Secure Sampling Protocols for Differential Privacy

URL: View paper

Brief Assessment

Secure Sampling Protocols[57] benchmarks secure sampling protocols for differential privacy in MPC settings, comparing methods like ALDR, FLDR, and Alias method. The original paper benchmarks discrete distribution sampling methods (cLUT vs ALDR, FLDR, Alias) for random variate generation, focusing on speed, energy, and memory—a different application domain with different evaluation contexts.

8. Discrete samplers for approximate inference in probabilistic machine learning

URL: View paper

Brief Assessment

Discrete Samplers[53] focuses on hardware architectures for discrete sampling in probabilistic machine learning, not on benchmarking discrete distribution sampling methods for general-purpose computing as in the original paper.

9. Discrete Ziggurat: A time-memory trade-off for sampling from a Gaussian distribution over the integers

URL: View paper

Brief Assessment

Discrete Ziggurat[58] focuses on sampling from discrete Gaussian distributions for lattice-based cryptography, not general discrete distribution sampling. The benchmarking compares methods specific to Gaussian sampling (inverse CDF, rejection sampling, Knuth-Yao) rather than the general discrete distribution samplers (ALDR, FLDR, Alias method) evaluated in the original paper.

10. Energy-Efficient NTT Sampler for Kyber Benchmarked on FPGA

URL: View paper

Brief Assessment

NTT Sampler[51] focuses on hardware implementation of polynomial sampling for Kyber cryptography on FPGA, comparing energy and latency metrics. The original paper benchmarks discrete distribution sampling methods across different metrics for general probabilistic machine learning applications. These are fundamentally different problem domains with distinct evaluation contexts.

Contribution 3: Demonstration of practical impact in machine learning applications

Description: The authors benchmark cLUT against standard Python library samplers (NumPy, PyTorch, JAX) showing 10-100× speedups, and demonstrate its practical value through a TrueSkill application case study that achieves substantial reductions in both execution time and energy consumption in real-world probabilistic machine learning tasks.

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. A review of the gumbel-max trick and its extensions for discrete stochasticity in machine learning

URL: View paper

Brief Assessment

Gumbel Max Trick[76] is a review paper focused on the Gumbel-max trick for discrete sampling and gradient estimation in neural networks, not on accelerating sampling methods or energy efficiency improvements in probabilistic machine learning applications like the original paper's cLUT approach.

2. Quadratic speedups in parallel sampling from determinantal distributions

URL: View paper

Brief Assessment

Determinantal Speedups[72] focuses on parallel sampling algorithms for determinantal point processes and related distributions, not on practical ML application benchmarks or energy efficiency demonstrations like the original paper's TrueSkill case study.

3. Accelerated Diffusion Models via Speculative Sampling

URL: View paper

Brief Assessment

Speculative Sampling[77] focuses on accelerating diffusion models through speculative sampling techniques for continuous vector-valued Markov chains, not on discrete distribution sampling methods or energy-efficient random variate generation as in the original paper.

4. Protein discovery with discrete walk-jump sampling

URL: View paper

Brief Assessment

Walk Jump Sampling[74] focuses on protein discovery using discrete generative models for antibody design, not on accelerating general discrete distribution sampling in machine learning applications like the original paper's cLUT method.

5. A langevin-like sampler for discrete distributions

URL: View paper

Brief Assessment

Langevin Sampler[71] focuses on discrete distribution sampling for MCMC methods in tasks like Ising models and text generation, not on accelerating general probabilistic ML applications through compressed lookup tables as in the original paper.

6. Distribution Learning Meets Graph Structure Sampling

URL: View paper

Brief Assessment

Graph Structure Sampling[75] focuses on learning Bayesian network distributions through graph structure sampling and online learning frameworks, not on accelerating discrete distribution sampling for general machine learning applications. The candidate addresses a fundamentally different problem domain (graphical model structure learning) than the original paper's focus on efficient random variate generation.

7. Oops i took a gradient: Scalable sampling for discrete distributions

URL: View paper

Brief Assessment

Scalable Gradient Sampling[70] focuses on MCMC sampling methods for discrete distributions using gradient information, not on accelerating discrete distribution sampling through compressed lookup tables for general ML applications.

8. Fast Sampling of Diffusion Models via Operator Learning

URL: View paper

Brief Assessment

Operator Learning[73] focuses on accelerating diffusion model sampling through neural operators for solving differential equations, not on general discrete distribution sampling or probabilistic machine learning applications like TrueSkill. The technical domains are distinct.

9. A faster sampler for discrete determinantal point processes

URL: View paper

Brief Assessment

Faster Determinantal Sampler[78] focuses on accelerating discrete determinantal point process sampling for subsampling datasets, not general discrete distribution sampling in probabilistic ML. The candidate addresses a specialized sampling problem (DPPs) rather than the broad class of arbitrary discrete distributions that cLUT targets.

10. Discrete Hamiltonian-Assisted Metropolis Sampling

URL: View paper

Brief Assessment

Hamiltonian Metropolis[79] focuses on MCMC sampling methods for discrete distributions using Hamiltonian mechanics, not on accelerating discrete distribution sampling in machine learning applications through compressed lookup tables or benchmarking against standard library samplers.

Appendix: Text Similarity Detection

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

References

- [0] Energy-Efficient Random Variate Generation via Compressed Lookup Tables View paper
- [1] 109.15 Sampling from the hypergeometric probability distribution View paper
- [2] Optimal approximate sampling from discrete probability distributions View paper
- [3] Discrete Flow Matching View paper
- [4] Maintaining Discrete Probability Distributions in Practice View paper
- [5] Diverse Semantic Image Synthesis via Probability Distribution Modeling View paper
- [6] M2D2M: Multi-Motion Generation from Text with Discrete Diffusion Models View paper
- [7] Generating Pseudo-Random Discrete Probability Distributions View paper
- [8] Introduction to probability models View paper
- [9] A Study on Computing Sample Quantiles of Discrete Probability Distributions View paper
- [10] On the alias method for generating random variables from a discrete distribution View paper
- [11] Quantum circuits for discrete graphical models View paper
- [12] Succinct sampling from discrete distributions View paper
- [13] Sampling from Arbitrary Functions via PSD Models View paper
- [14] Plug-and-Play Controllable Generation for Discrete Masked Models View paper
- [15] Random variate generation with formal guarantees View paper
- [16] The Fast Loaded Dice Roller: A Near-Optimal Exact Sampler for Discrete Probability Distributions View paper
- [17] Binary sampling from discrete distributions View paper
- [18] Efficient Rejection Sampling in the Entropy-Optimal Range View paper
- [19] Formally verified samplers from probabilistic programs with loops and conditioning View paper
- [20] Random number generation View paper
- [21] Super resolution enhanced Multi-view Stereo network based on Gumbel sampling View paper
- [22] Sampling hierarchies of discrete random structures View paper
- [23] Rejection-Inversion to Generate Variates from Monotone Discrete Distributions View paper
- [24] Differentiable Sampling of Categorical Distributions Using the CatLog-Derivative Trick View paper
- [25] Arithmetic sampling: parallel diverse decoding for large language models View paper
- [26] Sampling from Discrete Energy-Based Models with Quality/Efficiency Trade-offs View paper
- [27] On the Average Runtime of an Open Source Binomial Random Variate Generation Algorithm View paper
- [28] Discrete probability models and methods View paper
- [29] Convergence Analysis of Discrete Diffusion Model: Exact Implementation through Uniformization View paper
- [30] The Sampling-Gaussian for Stereo Matching View paper
- [31] Sampling from Binary Quadratic Distributions via Stochastic Localization View paper
- [32] Are we there yet? timing and floating-point attacks on differential privacy systems View paper
- [33] Dynamic sampling from a discrete probability distribution with a known distribution of rates View paper
- [34] Distance-based sampling of software configuration spaces View paper
- [35] Random discrete distributions View paper
- [36] Massively Parallel Construction of Radix Tree Forests for the Efficient Sampling of Discrete Probability Distributions View paper

- [37] Random variate generation View paper
- [38] Random variate generation using only finitely many unbiased, independently and identically distributed random bits View paper
- [39] Conditional sampling for spectrally discrete max-stable random fields View paper
- [40] A sampling theorem for finite discrete distributions View paper
- [41] Improving deep learning powered auction design View paper
- [42] Stochastic Optimal Power Flow in Hybrid Power System Using Reduced-Discrete Point Estimation Method and Latin Hypercube Sampling View paper
- [43] How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph View paper
- [44] Efficient Sampling Methods for Discrete Distributions View paper
- [45] On Quantum Circuits for Discrete Graphical Models View paper
- [46] Essentials of Monte Carlo simulation: Statistical methods for building simulation models View paper
- [47] PhaseTypeR: phase-type distributions in R with reward transformations and a view towards population genetics View paper
- [48] Adaptive Greedy Rejection Sampling View paper
- [49] Discrete distributions View paper
- [50] Mutual information optimally local private discrete distribution estimation View paper
- [51] Energy-Efficient NTT Sampler for Kyber Benchmarked on FPGA View paper
- [52] Energy-Efficient Reconfigurable Acceleration Engine for Polynomial Coefficient Generation of Lattice-Based Post-Quantum Cryptography View paper
- [53] Discrete samplers for approximate inference in probabilistic machine learning View paper
- [54] Constant-time discrete gaussian sampling View paper
- [55] High precision discrete Gaussian sampling on FPGAs View paper
- [56] Gradient-Guided Importance Sampling for Learning Discrete Energy-Based Models View paper
- [57] Benchmarking Secure Sampling Protocols for Differential Privacy View paper
- [58] Discrete Ziggurat: A time-memory trade-off for sampling from a Gaussian distribution over the integers View paper
- [59] Impact of ray generation schemes on the random ray method for eigenvalue and shielding applications View paper
- [60] A memory and gate efficient algorithm for unitary mixed Schur sampling View paper
- [61] Finite-state autoregressive entropy coding for efficient learned lossless compression View paper
- [62] DeltaPQ: lossless product quantization code compression for high dimensional similarity search View paper
- [63] Efficient Lossless Compression with Distribution Quantized Finite-State Autoregressive Model View paper
- [64] A Hierarchical Parallel Discrete Gaussian Sampler for Lattice-Based Cryptography View paper
- [65] Machine learning based symbol probability distribution prediction for entropy coding In AV1 View paper
- [66] On practical discrete Gaussian samplers for lattice-based cryptography View paper
- [67] Hierarchical compression of color look up tables View paper
- [68] Demystifying Diffusion Policies: Action Memorization and Simple Lookup Table Alternatives View paper
- [69] CARAMEL: A Succinct Read-Only Lookup Table via Compressed Static Functions View paper
- [70] Oops i took a gradient: Scalable sampling for discrete distributions View paper
- [71] A langevin-like sampler for discrete distributions View paper
- [72] Quadratic speedups in parallel sampling from determinantal distributions View paper
- [73] Fast Sampling of Diffusion Models via Operator Learning View paper
- [74] Protein discovery with discrete walk-jump sampling View paper
- [75] Distribution Learning Meets Graph Structure Sampling View paper
- [76] A review of the gumbel-max trick and its extensions for discrete stochasticity in machine learning View paper
- [77] Accelerated Diffusion Models via Speculative Sampling View paper
- [78] A faster sampler for discrete determinantal point processes View paper
- [79] Discrete Hamiltonian-Assisted Metropolis Sampling View paper