Novelty Assessment Report

Paper: On the Reasoning Abilities of Masked Diffusion Language Models

PDF URL: https://openreview.net/pdf?id=BVnIsh4Nz1

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Masked diffusion models (MDMs) for text offer a compelling alternative to traditional autoregressive language models. Parallel generation makes them efficient, but their computational capabilities and the limitations inherent to their parallelism remain largely unexplored. To this end, we characterize what types of reasoning problems MDMs can provably solve and how efficiently. We do this by connecting MDMs to the well-understood reasoning frameworks of chain of thought (CoT) and padded looped transformers (PLTs) in the finite-precision log-width setting: We show that MDMs and polynomially-padded PLTs are, in fact, equivalent in this setting, and that MDMs can solve all problems that CoT-augmented transformers can. Moreover, we showcase classes of problems (including regular languages) for which MDMs are inherently more efficient than CoT transformers, where parallel generation allows for substantially faster reasoning.

Core Task Landscape

This paper addresses: **Reasoning Capabilities of Masked Diffusion Language Models**

A total of **44 papers** were analyzed and organized into a taxonomy with **26 categories**.

Taxonomy Overview

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

- **Theoretical Foundations and Computational Expressivity**
- **Reinforcement Learning and Policy Optimization**
- **Reasoning Paradigms and Chain-of-Thought Integration**
- **Sampling and Decoding Strategies**
- **Scaling, Adaptation, and Architecture Design**
- **Multimodal and Cross-Domain Applications**
- **Specialized Training Objectives and Techniques**
- **Surveys and Comparative Studies**
- **Benchmark Models and Empirical Demonstrations**

Complete Taxonomy Tree

- Reasoning Capabilities of Masked Diffusion Language Models Survey Taxonomy
- Theoretical Foundations and Computational Expressivity
  - Expressivity and Equivalence Analysis ★ (3 papers)
  - [0] On the Reasoning Abilities of Masked Diffusion Language Models (Anon et al., 2026) View paper
  - [16] Coevolutionary Continuous Discrete Diffusion: Make Your Diffusion Language Model a Latent Reasoner (Zhou Cai, 2025) View paper
  - [37] On Powerful Ways to Generate: Autoregression, Diffusion, and Beyond (Yang Chenxiao, 2025) View paper
  - Theoretical Performance Bounds (1 papers)
  - [13] Theoretical Benefit and Limitation of Diffusion Language Model (Feng, 2025) View paper
- Reinforcement Learning and Policy Optimization
  - Policy Gradient Algorithms (3 papers)
  - [3] d2: Improved techniques for training reasoning diffusion language models (Wang Guanghan, 2025) View paper
  - [27] SPG: Sandwiched Policy Gradient for Masked Diffusion Language Models (Wang Chen-yu, 2025) View paper
  - [35] DiFFPO: Training Diffusion LLMs to Reason Fast and Furious via Reinforcement Learning (Zhao, 2025) View paper
  - Trajectory-Based Reinforcement Learning (2 papers)
  - [6] Revolutionizing reinforcement learning framework for diffusion large language models (Wang Yinjie, 2025) View paper
  - [10] Reinforcing the Diffusion Chain of Lateral Thought with Diffusion Language Models (Huang Ze-min, 2025) View paper
  - Multi-Reward and Inpainting-Guided Optimization (2 papers)
  - [14] Inpainting-Guided Policy Optimization for Diffusion Large Language Models (Zhao Siyan, 2025) View paper
  - [34] MRO: Enhancing Reasoning in Diffusion Language Models via Multi-Reward Optimization (Wang Cheng-long, 2025) View paper
  - Test-Time Scaling and Reward-Free Guidance (2 papers)
  - [19] RFG: Test-Time Scaling for Diffusion Large Language Model Reasoning with Reward-Free Guidance (Chen Tianlang, 2025) View paper
  - [41] Latent Adaptation with Masked Policy for Diffusion Language Models (G Sun, n.d.) View paper
- Reasoning Paradigms and Chain-of-Thought Integration
  - Diffusion-Based Chain-of-Thought (2 papers)
  - [12] Diffusion of thought: Chain-of-thought reasoning in diffusion language models (Wei Bi, 2024) View paper
  - [33] Diffusion of Thoughts: Chain-of-Thought Reasoning in Diffusion Language Models (Ye, 2024) View paper

- In-Place and Bidirectional Prompting (1 papers)
  - [23] Thinking inside the mask: In-place prompting in diffusion llms (Wang Yu-xuan, 2025) View paper
  - Latent Reasoning Approaches (1 papers)
  - [24] Ladir: Latent diffusion enhances llms for text reasoning (Kang, 2025) View paper
  - Parallel-Sequential Reasoning Analysis (1 papers)
  - [15] Beyond Surface Reasoning: Unveiling the True Long Chain-of-Thought Capacity of Diffusion Large Language Models (Chen, 2025) View paper
- Sampling and Decoding Strategies
  - Path Planning and Iterative Refinement (1 papers)
  - [11] Path planning for masked diffusion model sampling (Bezemek, 2025) View paper
  - Fast Inference and Acceleration (4 papers)
  - [30] Accelerated Sampling from Masked Diffusion Models via Entropy Bounded Unmasking (Ben-Hamu, 2025) View paper
  - [31] Sparse-LaViDa: Sparse Multimodal Discrete Diffusion Language Models (Shufan Li, 2025) View paper
  - [32] KLASS: KL-Guided Fast Inference in Masked Diffusion Models (Seo Hyun Kim, 2025) View paper
  - [36] dUltra: Ultra-Fast Diffusion Language Models via Reinforcement Learning (Shirui Chen, 2025) View paper
  - Position-Aware and Confidence-Based Sampling (1 papers)
  - [26] PC-Sampler: Position-Aware Calibration of Decoding Bias in Masked Diffusion Models (Huang Peng-cheng, 2025) View paper
  - Consistency Trajectory and Block-Wise Decoding (2 papers)
  - [22] Taming Masked Diffusion Language Models via Consistency Trajectory Reinforcement Learning with Fewer Decoding Step (Yang, 2025) View paper
  - [29] No Compute Left Behind: Rethinking Reasoning and Sampling with Masked Diffusion Models (Singhal Raghav, 2025) View paper
- Scaling, Adaptation, and Architecture Design
  - Scaling Laws and Large-Scale Pretraining (3 papers)
  - [5] Large language diffusion models (Nie Shen, 2025) View paper
  - [9] Scaling up masked diffusion models on text (Nie Shen, 2024) View paper
  - [25] Diffusion language models can perform many tasks with scaling and instruction-finetuning (Ye Jiasheng, 2023) View paper
  - Adaptation and Transfer Learning (1 papers)
  - [20] Scaling diffusion language models via adaptation from autoregressive models (Gong, 2024) View paper
  - Sparse and Mixture-of-Experts Architectures (1 papers)
  - [8] LLaDA-MoE: A Sparse MoE Diffusion Language Model (Zhu Feng-qi, 2025) View paper
  - Non-Markovian and Causal Diffusion Models (1 papers)
  - [4] Non-markovian discrete diffusion with causal language models (Zhang, 2025) View paper
  - Frequency-Informed and Data-Efficient Training (1 papers)
  - [40] Masked Diffusion Language Models with Frequency-Informed Training (Georgiou, 2025) View paper
- Multimodal and Cross-Domain Applications
  - Multimodal Diffusion Language Models (2 papers)
  - [1] Mmada: Multimodal large diffusion language models (Yang Ling, 2025) View paper
  - [43] LaViDa: A Large Diffusion Model for Vision-Language Understanding (S Li, n.d.) View paper
  - Autonomous Driving and Vision-Language Agents (2 papers)
  - [18] ViLaD: A Large Vision Language Diffusion Framework for End-to-End Autonomous Driving (Cui Can, 2025) View paper
  - [21] dVLM-AD: Enhance Diffusion Vision-Language-Model for Driving via Controllable Reasoning (Yingzi Ma, 2025) View paper
  - Structured Prediction and Code Generation (2 papers)
  - [28] Unifying Deductive and Abductive Reasoning in Knowledge Graphs with Masked Diffusion Model (Gao Yisen, 2025) View paper
  - [39] CodeDiffuSe: A masked diffusion framework for structure-aware code completion and repair (Aytug Onan, 2025) View paper
- Specialized Training Objectives and Techniques
  - Mixed Chain-of-Thought and Coevolutionary Training (1 papers)
  - [44] SPMDM: Enhancing Masked Diffusion Models through Simplifing Sampling Path (Y Zhu, n.d.) View paper
  - Temporal and Commonsense Reasoning (1 papers)
  - [42] Effective Masked Language Modeling for Temporal Commonsense Reasoning (Mayuko Kimura, 2022) View paper
- Surveys and Comparative Studies (2 papers)
  - [17] Diffusion-based Large Language Models Survey (Chiung-Yi Tseng, 2025) View paper
  - [38] A Survey on Latent Reasoning (Zhu, 2025) View paper
- Benchmark Models and Empirical Demonstrations (2 papers)
  - [2] Dream 7b: Diffusion large language models (Ye, 2025) View paper
  - [7] d1: Scaling reasoning in diffusion large language models via reinforcement learning (Zhao Siyan, 2025) View paper

Narrative

Core task: reasoning capabilities of masked diffusion language models. The field has rapidly expanded into a rich taxonomy spanning theoretical foundations, reinforcement learning integration, reasoning paradigms, sampling strategies, scaling and architecture design, multimodal applications, specialized training objectives, surveys, and benchmark demonstrations. Theoretical Foundations and Computational Expressivity examines the fundamental properties of diffusion models, including expressivity analyses that compare masked diffusion to autoregressive approaches (e.g., Autoregression Diffusion Beyond[37]) and explore computational equivalences (Coevolutionary Continuous Discrete[16]). Reinforcement Learning and Policy Optimization investigates how diffusion models can be trained via RL signals, with works like Revolutionizing Reinforcement[6] and Inpainting Policy Optimization[14] exploring policy gradient methods. Reasoning Paradigms and Chain-of-Thought Integration focuses on incorporating structured reasoning into diffusion generation, exemplified by Diffusion of Thought[12] and Thinking Inside Mask[23]. Sampling and Decoding Strategies address efficient inference techniques, while Scaling, Adaptation, and Architecture Design explores model growth and architectural innovations such as Dream 7b[2] and LLaDA-MoE[8]. Multimodal and Cross-Domain Applications extend diffusion models beyond text to vision-language tasks (ViLaD Autonomous Driving[18], dVLM-AD Driving[21]), and Specialized Training Objectives develop novel learning frameworks like d2 Training Techniques[3].

Several active lines of work reveal key trade-offs between expressivity, efficiency, and controllability. The tension between theoretical guarantees and practical performance is evident in studies comparing diffusion to autoregressive models, where expressivity gains must be balanced against computational costs. Reasoning Masked Diffusion[0] sits within the Theoretical Foundations branch, specifically addressing expressivity and equivalence analysis. Its emphasis on understanding the fundamental reasoning capabilities of masked diffusion models positions it alongside works like Autoregression Diffusion Beyond[37], which similarly investigates how diffusion architectures compare to traditional paradigms, and Coevolutionary Continuous Discrete[16], which explores the interplay between discrete and continuous formulations. While neighboring studies often focus on architectural comparisons or computational complexity, Reasoning Masked Diffusion[0] appears to concentrate on the intrinsic reasoning properties that emerge from the masked diffusion framework, contributing foundational insights that inform downstream applications across reinforcement learning, chain-of-thought integration, and multimodal reasoning tasks.

Related Works in Same Category

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

1. Coevolutionary Continuous Discrete Diffusion: Make Your Diffusion Language Model a Latent Reasoner

Authors: Zhou Cai, Yang Chenxiao, Cai Zhou, Hu Yi, Chenxiao Yang, et al. (24 authors total) | **Year/Venue**: 2025 • arXiv.org | **URL**: View paper

Abstract

Diffusion language models, especially masked discrete diffusion models, have achieved great success recently. While there are some theoretical and primary empirical results showing the advantages of latent reasoning with looped transformers or continuous chain-of-thoughts, continuous diffusion models typically underperform their discrete counterparts. In this paper, we argue that diffusion language models do not necessarily need to be in the discrete space. In particular, we prove that continuou...

Relationship Analysis

Both papers belong to the Expressivity and Equivalence Analysis category, examining the theoretical computational capabilities of masked diffusion language models. While the original paper establishes formal equivalences between MDMs and padded looped transformers (PLTs) and chain-of-thought reasoning in the finite-precision log-width setting, the candidate paper focuses on comparing continuous versus discrete diffusion models and proposes a hybrid approach (CCDD) that combines both modalities. The key difference is that the original paper provides a complexity-theoretic characterization of MDM expressivity through equivalence proofs, whereas the candidate paper addresses the practical trainability gap between continuous and discrete diffusions and introduces a novel coevolutionary architecture.

2. On Powerful Ways to Generate: Autoregression, Diffusion, and Beyond

Authors: Yang Chenxiao, Zhou Cai, Chenxiao Yang, Wipf, David, et al. (10 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Diffusion language models have recently emerged as a competitive alternative to autoregressive language models. Beyond next-token generation, they are more efficient and flexible by enabling parallel and any-order token generation. However, despite empirical successes, their computational power and fundamental limitations remain poorly understood. In this paper, we formally study whether non-autoregressive generation in Masked Diffusion Models (MDM) enables solving problems beyond the reach of A...

Relationship Analysis

Both papers belong to the Expressivity and Equivalence Analysis category, establishing formal computational frameworks for masked diffusion language models. They share overlapping focus on characterizing MDM reasoning capabilities through equivalences with other computational models (the original paper connects MDMs to padded looped transformers and CoT, while the candidate connects MDMs to PRAM and autoregressive models). The key difference is that the original paper analyzes standard MDMs with fixed unmasking processes in relation to circuit complexity classes (AC^d, NC), whereas the candidate paper extends beyond standard MDMs by proposing "any-process generation" with remask/insert/delete operations to achieve stronger computational universality and scalability to NP-hard problems.

Contributions Analysis

Overall novelty summary. The paper establishes formal equivalences between masked diffusion models (MDMs) and padded looped transformers (PLTs) in the finite-precision log-width setting, while characterizing MDM reasoning capabilities through chain-of-thought frameworks. It resides in the 'Expressivity and Equivalence Analysis' leaf under 'Theoretical Foundations and Computational Expressivity', alongside two sibling papers that similarly investigate computational equivalences and expressivity comparisons. This leaf represents a relatively sparse research direction within the broader taxonomy of 44 papers, suggesting that formal theoretical analysis of MDM reasoning remains an emerging area compared to more crowded branches like reinforcement learning or sampling strategies.

The taxonomy reveals that theoretical foundations constitute a small but foundational branch, with only four papers total across expressivity analysis and performance bounds. Neighboring work in 'Reasoning Paradigms and Chain-of-Thought Integration' (seven papers) focuses on practical CoT implementations rather than formal characterizations, while 'Reinforcement Learning and Policy Optimization' (eleven papers) emphasizes training methods. The paper's theoretical approach bridges these areas by providing formal grounding for reasoning capabilities that other branches explore empirically. Its position suggests it addresses a gap between architectural comparisons in sibling papers and the applied reasoning methods in adjacent taxonomy branches.

Among 23 candidates examined through limited semantic search, none clearly refute the three main contributions. The equivalence between MDMs and PLTs (3 candidates examined, 0 refutable) appears novel within this search scope. The CoT characterization (10 candidates, 0 refutable) and efficiency advantages on parallelizable problems (10 candidates, 0 refutable) similarly show no overlapping prior work among examined papers. However, the modest search scale means these findings reflect top-K semantic matches rather than exhaustive coverage. The sibling papers in the same taxonomy leaf focus on different aspects—architectural expressivity comparisons and continuous-discrete formulations—rather than the specific PLT equivalence or CoT-based reasoning characterization presented here.

Based on the limited literature search covering 23 candidates, the work appears to occupy a relatively unexplored theoretical niche within MDM research. The formal equivalence results and reasoning characterizations do not overlap with examined prior work, though the small search scope and sparse theoretical foundations branch suggest caution in generalizing these findings. The analysis captures top semantic matches but cannot rule out relevant work outside this scope, particularly in adjacent areas like formal language theory or computational complexity that may not surface through MDM-focused queries.

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

Contribution 1: Equivalence of masked diffusion models and padded looped transformers

Description: The authors prove that masked diffusion models (MDMs) and padded looped transformers (PLTs) are equivalent in the finite-precision log-width setting, establishing that both frameworks can solve the same class of problems up to logarithmic factors in padding length.

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

1. Transformers are Universal In-context Learners

URL: View paper

Brief Assessment

Universal In-context Learners[45] focuses on the universal approximation capabilities of transformers for in-context mappings over arbitrary token distributions, not on the equivalence between masked diffusion models and padded looped transformers in finite-precision settings.

---

2. Coevolutionary Continuous Discrete Diffusion: Make Your Diffusion Language Model a Latent Reasoner

URL: View paper

Brief Assessment

Coevolutionary Continuous Discrete[16] focuses on continuous diffusion models and their expressivity compared to discrete diffusions and looped transformers, but does not address the equivalence between masked diffusion models and padded looped transformers in the finite-precision log-width setting that the original paper establishes.

---

3. DiffVecFont: Fusing Dual-Mode Reconstruction Vector Fonts via Masked Diffusion Transformers

URL: View paper

Brief Assessment

DiffVecFont[46] focuses on vector font generation using diffusion transformers for visual design tasks, not on theoretical equivalence between masked diffusion models and padded looped transformers in formal language settings.

Contribution 2: Characterization of MDM reasoning capabilities via chain of thought

Description: The authors demonstrate that MDMs can perform chain-of-thought (CoT) reasoning and establish formal connections showing MDMs can simulate CoT transformers with some overhead, while CoT transformers can also simulate MDMs, providing upper and lower bounds on MDM expressivity.

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. Simple and effective masked diffusion language models

URL: View paper

Brief Assessment

Simple Effective Masked[57] focuses on masked diffusion for language modeling with objectives based on masked language modeling losses, not on formal characterization of chain-of-thought reasoning capabilities or connections to transformers' computational expressivity.

---

2. Thinking inside the mask: In-place prompting in diffusion llms

URL: View paper

Brief Assessment

Thinking Inside Mask[23] focuses on in-place prompting strategies for diffusion LLMs in practical applications, not on formal theoretical characterization of MDM expressivity or simulation relationships with CoT transformers.

---

3. Reinforcing the Diffusion Chain of Lateral Thought with Diffusion Language Models

URL: View paper

Brief Assessment

Reinforcing Lateral Thought[10] focuses on reinforcement learning optimization of diffusion language models for reasoning tasks, not on formal characterization of MDM expressivity or theoretical connections to CoT transformers.

---

4. A Survey on Latent Reasoning

URL: View paper

Brief Assessment

Latent Reasoning Survey[38] focuses on latent reasoning methods across various architectures (transformers, RNNs, diffusion models) but does not specifically address masked diffusion models performing chain-of-thought reasoning or establish formal connections between MDMs and CoT transformers as claimed in the original contribution.

---

5. Scaling up masked diffusion models on text

URL: View paper

Brief Assessment

Scaling Masked Diffusion[9] focuses on empirical scaling laws, language understanding benchmarks, and conditional generation tasks for MDMs. It does not provide formal theoretical characterizations of MDM reasoning capabilities or establish connections to chain-of-thought transformers through computational complexity theory.

---

6. Diffusion-based Large Language Models Survey

URL: View paper

Brief Assessment

Diffusion LLM Survey[17] appears to be a survey paper covering diffusion-based language models broadly. The minimal context provided focuses on multimodal understanding and acceleration, not on formal characterizations of chain-of-thought reasoning capabilities or expressivity bounds for masked diffusion models.

---

7. d2: Improved techniques for training reasoning diffusion language models

URL: View paper

Brief Assessment

d2 Training Techniques[3] focuses on reinforcement learning algorithms for training masked diffusion language models, not on characterizing their theoretical reasoning capabilities or establishing formal connections to chain-of-thought transformers.

---

8. Path Planning for Masked Diffusion Models with Applications to Biological Sequence Generation

URL: View paper

Brief Assessment

Path Planning Biological[58] focuses on inference-time token unmasking strategies for masked diffusion models in biological sequence generation, not on characterizing MDM reasoning capabilities or establishing formal connections to chain-of-thought transformers.

9. Ladir: Latent diffusion enhances llms for text reasoning

URL: View paper

Brief Assessment

Ladir Latent Diffusion[24] focuses on latent diffusion models for text reasoning in LLMs, not on masked diffusion models (MDMs) performing chain-of-thought reasoning like transformers. The candidate addresses a different model class and application domain.

10. Any-Order GPT as Masked Diffusion Model: Decoupling Formulation and Architecture

URL: View paper

Brief Assessment

Any-Order GPT[59] focuses on decoupling formulation (AR vs. MDM) from architecture (decoder-only vs. encoder-only) for language modeling, not on characterizing MDM reasoning capabilities or establishing formal connections to chain-of-thought transformers as the original paper does.

Contribution 3: Identification of MDM efficiency advantages over CoT on parallelizable problems

Description: The authors prove that MDMs are provably more efficient than CoT transformers on parallelizable problems due to their ability to leverage parallel generation, identifying what they term the sequentiality bottleneck of CoT and showing a strict separation in expressivity under logarithmically many decoding steps.

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. Parallel-r1: Towards parallel thinking via reinforcement learning

URL: View paper

Brief Assessment

Parallel-r1[50] focuses on training LLMs to perform parallel thinking via reinforcement learning for mathematical reasoning tasks, not on comparing the theoretical efficiency of masked diffusion models versus chain-of-thought transformers on parallelizable problems.

2. Fast ECoT: Efficient Embodied Chain-of-Thought via Thoughts Reuse

URL: View paper

Brief Assessment

Fast ECoT[56] focuses on accelerating embodied chain-of-thought reasoning in vision-language-action models for robotics through caching and parallel generation of reasoning steps. It does not address masked diffusion models, parallel generation in language models, or the theoretical expressivity comparisons between MDMs and CoT transformers that form the core of the original contribution.

3. SPRINT: Enabling Interleaved Planning and Parallelized Execution in Reasoning Models

URL: View paper

Brief Assessment

SPRINT Interleaved Planning[53] focuses on parallelizing execution within reasoning models through interleaved planning and execution phases, not on comparing masked diffusion models (MDMs) versus chain-of-thought (CoT) transformers or analyzing their theoretical expressivity differences on parallelizable problems.

4. An llm compiler for parallel function calling

URL: View paper

Brief Assessment

LLM Compiler[48] focuses on parallel function calling in LLM applications (orchestrating external tool calls), not on comparing masked diffusion models versus chain-of-thought transformers for parallelizable reasoning problems. The technical domains are fundamentally different.

5. A survey on parallel text generation: From parallel decoding to diffusion language models

URL: View paper

Brief Assessment

Parallel Text Generation[47] is a survey paper that categorizes parallel generation techniques but does not provide theoretical analysis of MDM efficiency advantages over CoT or prove expressivity separations on parallelizable problems.

6. A survey on parallel reasoning

URL: View paper

Brief Assessment

Parallel Reasoning Survey[52] focuses on parallel reasoning paradigms in LLMs (e.g., self-consistency, multi-agent debate, parallel decoding) rather than comparing masked diffusion models to chain-of-thought transformers on parallelizable problems. The candidate discusses inference-time scaling and parallel exploration strategies, not the architectural expressivity differences between MDMs and CoT that the original paper analyzes.

7. Hybrid Deep Searcher: Integrating Parallel and Sequential Search Reasoning

URL: View paper

Brief Assessment

Hybrid Deep Searcher[49] focuses on hybrid parallel-sequential query execution in retrieval-augmented reasoning for QA tasks, not on comparing masked diffusion models versus chain-of-thought transformers in formal reasoning settings.

8. Learning adaptive parallel reasoning with language models

URL: View paper

Brief Assessment

Adaptive Parallel Reasoning[55] focuses on training language models to adaptively parallelize inference-time computation through spawn/join operations, not on comparing masked diffusion models versus chain-of-thought transformers on parallelizable problems.

9. To Backtrack or Not to Backtrack: When Sequential Search Limits Model Reasoning

URL: View paper

Brief Assessment

Backtrack Sequential Search[54] focuses on comparing sequential search with backtracking versus parallel sampling in LLMs for reasoning tasks, not on masked diffusion models versus chain-of-thought transformers or parallelizable problem efficiency.

10. How to think step-by-step: A mechanistic understanding of chain-of-thought reasoning

URL: View paper

Brief Assessment

Mechanistic Chain-of-Thought[51] investigates the internal neural mechanisms of CoT reasoning in LLMs, not efficiency comparisons between different generation paradigms like MDMs versus CoT on parallelizable problems.

Appendix: Text Similarity Detection

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

References

- [0] On the Reasoning Abilities of Masked Diffusion Language Models View paper
- [1] Mmada: Multimodal large diffusion language models View paper
- [2] Dream 7b: Diffusion large language models View paper
- [3] d2: Improved techniques for training reasoning diffusion language models View paper
- [4] Non-markovian discrete diffusion with causal language models View paper
- [5] Large language diffusion models View paper
- [6] Revolutionizing reinforcement learning framework for diffusion large language models View paper
- [7] d1: Scaling reasoning in diffusion large language models via reinforcement learning View paper
- [8] LLaDA-MoE: A Sparse MoE Diffusion Language Model View paper
- [9] Scaling up masked diffusion models on text View paper
- [10] Reinforcing the Diffusion Chain of Lateral Thought with Diffusion Language Models View paper
- [11] Path planning for masked diffusion model sampling View paper
- [12] Diffusion of thought: Chain-of-thought reasoning in diffusion language models View paper
- [13] Theoretical Benefit and Limitation of Diffusion Language Model View paper
- [14] Inpainting-Guided Policy Optimization for Diffusion Large Language Models View paper
- [15] Beyond Surface Reasoning: Unveiling the True Long Chain-of-Thought Capacity of Diffusion Large Language Models View paper
- [16] Coevolutionary Continuous Discrete Diffusion: Make Your Diffusion Language Model a Latent Reasoner View paper
- [17] Diffusion-based Large Language Models Survey View paper
- [18] ViLaD: A Large Vision Language Diffusion Framework for End-to-End Autonomous Driving View paper
- [19] RFG: Test-Time Scaling for Diffusion Large Language Model Reasoning with Reward-Free Guidance View paper
- [20] Scaling diffusion language models via adaptation from autoregressive models View paper
- [21] dVLM-AD: Enhance Diffusion Vision-Language-Model for Driving via Controllable Reasoning View paper
- [22] Taming Masked Diffusion Language Models via Consistency Trajectory Reinforcement Learning with Fewer Decoding Step View paper
- [23] Thinking inside the mask: In-place prompting in diffusion llms View paper
- [24] Ladir: Latent diffusion enhances llms for text reasoning View paper
- [25] Diffusion language models can perform many tasks with scaling and instruction-finetuning View paper
- [26] PC-Sampler: Position-Aware Calibration of Decoding Bias in Masked Diffusion Models View paper
- [27] SPG: Sandwiched Policy Gradient for Masked Diffusion Language Models View paper
- [28] Unifying Deductive and Abductive Reasoning in Knowledge Graphs with Masked Diffusion Model View paper
- [29] No Compute Left Behind: Rethinking Reasoning and Sampling with Masked Diffusion Models View paper
- [30] Accelerated Sampling from Masked Diffusion Models via Entropy Bounded Unmasking View paper
- [31] Sparse-LaViDa: Sparse Multimodal Discrete Diffusion Language Models View paper
- [32] KLASS: KL-Guided Fast Inference in Masked Diffusion Models View paper
- [33] Diffusion of Thoughts: Chain-of-Thought Reasoning in Diffusion Language Models View paper
- [34] MRO: Enhancing Reasoning in Diffusion Language Models via Multi-Reward Optimization View paper
- [35] DiFFPO: Training Diffusion LLMs to Reason Fast and Furious via Reinforcement Learning View paper
- [36] dUltra: Ultra-Fast Diffusion Language Models via Reinforcement Learning View paper
- [37] On Powerful Ways to Generate: Autoregression, Diffusion, and Beyond View paper
- [38] A Survey on Latent Reasoning View paper
- [39] CodeDiffuSe: A masked diffusion framework for structure-aware code completion and repair View paper
- [40] Masked Diffusion Language Models with Frequency-Informed Training View paper
- [41] Latent Adaptation with Masked Policy for Diffusion Language Models View paper
- [42] Effective Masked Language Modeling for Temporal Commonsense Reasoning View paper
- [43] LaViDa: A Large Diffusion Model for Vision-Language Understanding View paper
- [44] SPMDM: Enhancing Masked Diffusion Models through Simplifing Sampling Path View paper
- [45] Transformers are Universal In-context Learners View paper
- [46] DiffVecFont: Fusing Dual-Mode Reconstruction Vector Fonts via Masked Diffusion Transformers View paper
- [47] A survey on parallel text generation: From parallel decoding to diffusion language models View paper
- [48] An llm compiler for parallel function calling View paper
- [49] Hybrid Deep Searcher: Integrating Parallel and Sequential Search Reasoning View paper
- [50] Parallel-r1: Towards parallel thinking via reinforcement learning View paper
- [51] How to think step-by-step: A mechanistic understanding of chain-of-thought reasoning View paper
- [52] A survey on parallel reasoning View paper
- [53] SPRINT: Enabling Interleaved Planning and Parallelized Execution in Reasoning Models View paper
- [54] To Backtrack or Not to Backtrack: When Sequential Search Limits Model Reasoning View paper

- [55] Learning adaptive parallel reasoning with language models View paper
- [56] Fast ECoT: Efficient Embodied Chain-of-Thought via Thoughts Reuse View paper
- [57] Simple and effective masked diffusion language models View paper
- [58] Path Planning for Masked Diffusion Models with Applications to Biological Sequence Generation View paper
- [59] Any-Order GPT as Masked Diffusion Model: Decoupling Formulation and Architecture View paper