

# CAKE Cascading Eviction Preserves Chain-of-Thought Accuracy in GSM8K Math Problems

Assignee Research

June 8, 2026

## Abstract

This report synthesises findings from 15 peer-reviewed papers addressing the following research question: Does CAKE’s cascading eviction mechanism preserve chain-of-thought reasoning accuracy on GSM8K for code-derived math problems better than uniform KV cache pruning in Mistral-7B. 16 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 2.8/10. This report is a machine-generated literature synthesis and does not constitute original research.

## 1 Introduction

This paper examines: DiffCoT: Diffusion-styled Chain-of-Thought Reasoning in LLMs. Research question: Does CAKE’s cascading eviction mechanism preserve chain-of-thought reasoning accuracy on GSM8K for code-derived math problems better than uniform KV cache pruning in Mistral-7B?.

## 2 Methodology

Systematic literature search across multiple databases yielded 15 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 2.8/10.

## 3 Results

15 papers retrieved. 16 claims extracted; 0 independently verified. Quality review score: 2.8/10.

## 4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.



## 5 Extracted Claims

Claim	Verified	Confidence
DIFFCOT improves reasoning performance on Qwen3, Ministral3, and Llama3 models under standard fine-tuning.	×	0.05
On Llama3-8B, the CPO method degrades performance on the MATH-L1 benchmark compared to baselines.	×	0.03
On Llama3-8B, the CPO method remains competitive on the SVAMP benchmark.	×	0.02
On the MATH-L1 benchmark with the Llama3-8B backbone, Step-DPO achieves slightly higher accuracy than DIFFCOT.	×	0.04
DIFFCOT gains are more pronounced on stronger backbones such as Qwen3-8B and Qwen3-4B compared to weaker models.	×	0.03
DIFFCOT brings larger improvements on harder problems (MATH-L4 and MATH-L5) where correct reasoning paths are sparser.	×	0.04
When the diffusion window size and stride are set to 1, the DIFFCOT approach degenerates into an autoregressive (AR) met	×	0.06
When the sliding diffusion window and stride are set to the number of steps K, the method reverts to a purely diffusion	×	0.08
On the Llama3-8B model for the GSM8K benchmark, DIFFCOT achieves 64.4% accuracy.	×	0.06
On the Llama3-8B model for the SVAMP benchmark, DIFFCOT achieves 76.9% accuracy.	×	0.02
On the Llama3-8B model for the MATH-L5 benchmark, DIFFCOT achieves 3.8% accuracy.	×	0.04
On the Qwen3-4B model for the GSM8K benchmark, DIFFCOT achieves 91.5% accuracy.	×	0.06
On the Qwen3-4B model for the MATH-L5 benchmark, DIFFCOT achieves 14.9% accuracy.	×	0.03
For Llama3-8B on GSM8K, setting the window size and stride to 1 results in 62.9% accuracy, a decrease of 1.5% from the s	×	0.03
For Llama3-8B on GSM8K, removing causal noise results in 55.4% accuracy, a decrease of 9.0% from the standard DIFFCOT se	×	0.05
For Llama3-8B on SVAMP, setting the window size and stride to 1 results in 75.8% accuracy, a decrease of 1.1% from the s	×	0.02

## References

- <http://arxiv.org/abs/2601.03559v2>
- <http://arxiv.org/abs/2503.12491v2>
- <http://arxiv.org/abs/2308.07921v1>