Novelty Assessment Report

Paper: Emergence of Superposition: Unveiling the Training Dynamics of Chain of Continuous Thought
PDF URL: https://openreview.net/pdf?id=lsJwX9Jf5u
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-05

Abstract

Previous work shows that the chain of continuous thought (continuous CoT) improves the reasoning capability of large language models (LLMs) by enabling implicit parallel thinking, and a subsequent work provided theoretical insight by showing that a two-layer transformer equipped with continuous CoT can efficiently solve directed graph reachability by maintaining a superposition of multiple reasoning traces in the continuous thought. However, it remains unclear how the superposition mechanism is naturally learned from gradient-based training methods. To fill this gap, we theoretically analyze the training dynamics of a simplified two-layer transformer on the directed graph reachability problem to unveil how the superposition mechanism emerges during training in two training stages -- (i) a thought-generation stage that autoregressively expands the continuous thought, and (ii) a prediction stage that converts the thought into the final answer. Our analysis reveals that during training using continuous thought, the index-matching logit, an important quantity which reflects the strength of the model's local search ability, will first increase and then remain bounded under mild assumptions. The bounded index-matching logit effectively balances exploration and exploitation during the reasoning process: the model will exploit local problem structures to identify plausible search traces, and assign comparable weights to multiple such traces to explore when it is uncertain about which solution is correct, which results in superposition. Our experimental results tracking the growth of logits further validate our theory.

Core Task Landscape

This paper addresses: **Training Dynamics of Chain of Continuous Thought in Transformers**

A total of **50 papers** were analyzed and organized into a taxonomy with **32 categories**.

Taxonomy Overview

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

- **Continuous Latent Reasoning Paradigms**
- **Theoretical Foundations of CoT Mechanisms**
- **Discrete CoT Prompting and Inference Methods**
- **CoT Training and Optimization Methods**
- **Architectural Innovations for Reasoning**
- **World Models and Planning-Based Reasoning**
- **Sequence Modeling and Decision-Making Applications**
- **Transformer Learning Dynamics and Mechanisms**
- **Specialized Transformer Applications**

Complete Taxonomy Tree

- Training Dynamics of Chain of Continuous Thought in Transformers Survey Taxonomy
- Continuous Latent Reasoning Paradigms
  - Continuous CoT Architectures and Training (3 papers)
  - [2] Training Large Language Models to Reason in a Continuous Latent Space (Shibo Hao, 2025) View paper
  - [6] Codi: Compressing chain-of-thought into continuous space via self-distillation (Shen Zhen-yi, 2025) View paper
  - [33] Scaling latent reasoning via looped language models (Zhu, 2025) View paper
  - Latent Reasoning Surveys and Frameworks (1 papers)
  - [11] A survey on latent reasoning (Zhu, 2025) View paper
  - Parallel Exploration via Continuous Thought (1 papers)
  - [41] Continuous Chain of Thought Enables Parallel Exploration and Reasoning (Ildiz, 2025) View paper
- Theoretical Foundations of CoT Mechanisms
  - Training Dynamics and Convergence Analysis ★ (4 papers)
  - [0] Emergence of Superposition: Unveiling the Training Dynamics of Chain of Continuous Thought (Anon et al., 2026) View paper
  - [28] Training Nonlinear Transformers for Chain-of-Thought Inference: A Theoretical Generalization Analysis (LI Hongkang, 2024) View paper
  - [37] Transformers learn to implement multi-step gradient descent with chain of thought (Huang Jianhao, 2025) View paper
  - [49] The Kinetics of Reasoning: How Chain-of-Thought Shapes Learning in Transformers? (Mavromatis, 2025) View paper
  - Expressivity and Computational Power (4 papers)
  - [5] Towards revealing the mystery behind chain of thought: a theoretical perspective (Feng, 2023) View paper
  - [12] Chain of thought empowers transformers to solve inherently serial problems (Li, 2024) View paper
  - [20] The Expressive Power of Transformers with Chain of Thought (Merrill, 2023) View paper
  - [50] Can Transformers Reason Logically? A Study in SAT Solving (Ganesh, 2024) View paper
  - Superposition and Parallel Reasoning Theory (1 papers)

- [18] Reasoning by Superposition: A Theoretical Perspective on Chain of Continuous Thought (Zhu, 2025) View paper
- Length Generalization and Scaling Laws (2 papers)
- [15] From sparse dependence to sparse attention: unveiling how chain-of-thought enhances transformer sample efficiency (Wen, 2024) View paper
- [24] Transformers provably learn chain-of-thought reasoning with length generalization (Yu Huang, 2025) View paper
- Coherence and Error Correction Theory (1 papers)
- [32] A Theoretical Understanding of Chain-of-Thought: Coherent Reasoning and Error-Aware Demonstration (Cui, 2024) View paper

• Discrete CoT Prompting and Inference Methods
- Foundational CoT Prompting Techniques (1 papers)
- [1] Chain of Thought Prompting Elicits Reasoning in Large Language Models (Stephanie Ewelu, 2025) View paper
- CoT Decoding and Sampling Strategies (1 papers)
- [9] Self-Consistency Improves Chain of Thought Reasoning in Language Models (Wang, 2022) View paper
- Structured and Symbolic CoT Variants (2 papers)
- [7] Faithful Chain-of-Thought Reasoning (Apidianaki, 2023) View paper
- [22] Boosting Language Models Reasoning with Chain-of-Knowledge Prompting (Gao Ming, 2023) View paper
- Multilingual and Cross-Domain CoT (2 papers)
- [17] Chain-of-Thought Hub: A Continuous Effort to Measure Large Language Models' Reasoning Performance (Fu Yao, 2023) View paper
- [23] Language Models are Multilingual Chain-of-Thought Reasoners (Shi, 2022) View paper

• CoT Training and Optimization Methods
- Supervised CoT Fine-Tuning (1 papers)
- [36] Beyond In-Distribution Success: Scaling Curves of CoT Granularity for Language Model Generalization (Wang Ru, 2025) View paper
- Latent-Variable CoT Training (1 papers)
- [35] Training Chain-of-Thought via Latent-Variable Inference (Phan Du, 2023) View paper
- Long-CoT Data Scaling (1 papers)
- [43] RedStar: Does Scaling Long-CoT Data Unlock Better Slow-Reasoning Systems? (Xu, 2025) View paper

• Architectural Innovations for Reasoning
- Looped and Recurrent Transformer Variants (3 papers)
- [4] Reasoning with latent thoughts: On the power of looped transformers (Saunshi, 2025) View paper
- [21] Enhancing Auto-regressive Chain-of-Thought through Loop-Aligned Reasoning (Yu, 2025) View paper
- [48] Autoregressive+ Chain of Thought= Recurrent: Recurrence's Role in Language Models' Computability and a Revisit of Recurrent Transformer (Zhang Xiang, 2024) View paper
- Hybrid Architecture Designs (1 papers)
- [14] Hunyuan-turbos: Advancing large language models through mamba-transformer synergy and adaptive chain-of-thought (Liu Ao, 2025) View paper
- Continuous-Time Transformer Formulations (3 papers)
- [8] Flowing Through Layers: A Continuous Dynamical Systems Perspective on Transformers (Jacob Fein-Ashley, 2025) View paper
- [29] CTRNN-Transformer Adding Continuous Time Neural Models to Transformers (Kanishka Bodreddigari, 2024) View paper
- [34] ContiFormer: Continuous-Time Transformer for Irregular Time Series Modeling (Chen Yuqi, 2024) View paper
- Memory-Efficient Reasoning Architectures (1 papers)
- [39] PENCIL: Long Thoughts with Short Memory (Yang Chenxiao, 2025) View paper

• World Models and Planning-Based Reasoning
- World Model Integration for Reasoning (1 papers)
- [3] Reasoning with Language Model is Planning with World Model (Gu Yi, 2023) View paper
- Vision-Language World Models (1 papers)
- [10] FlowVLA: Visual Chain of Thought-based Motion Reasoning for Vision-Language-Action Models (Yan Hao-dong, 2025) View paper

• Sequence Modeling and Decision-Making Applications
- Decision Transformers for RL (2 papers)
- [25] Decision transformer: Reinforcement learning via sequence modeling (Li-Li Chen, 2021) View paper
- [31] Learning to Play Atari in a World of Tokens (Agarwal, 2024) View paper
- In-Context RL with Hierarchical CoT (1 papers)
- [46] In-context decision transformer: Reinforcement learning via hierarchical chain-of-thought (Huang, 2024) View paper
- RL Transformers Survey and Taxonomy (1 papers)
- [30] On transforming reinforcement learning with transformers: The development trajectory (Shengchao Hu, 2024) View paper
- MCTS-Based Reasoning Optimization (1 papers)
- [16] Optimizing large language models through highly dense reward structures and recursive thought process using monte carlo tree search (K. Laurent, 2024) View paper

• Transformer Learning Dynamics and Mechanisms
- Regular Language Learning Dynamics (1 papers)
- [38] How Transformers Learn Regular Language Recognition: A Theoretical Study on Training Dynamics and Implicit Bias (Huang, 2025) View paper
- Reversal Curse and Asymmetric Learning (1 papers)
- [45] Towards a Theoretical Understanding of the 'Reversal Curse' via Training Dynamics (Baihe Huang, 2024) View paper
- Statistical Physics of Reasoning Dynamics (1 papers)
- [47] A Statistical Physics of Language Model Reasoning (Jack David Carson, 2025) View paper

• Specialized Transformer Applications
- Continuous Prompting and Parameter-Efficient Tuning (1 papers)
- [13] Prefix-Tuning: Optimizing Continuous Prompts for Generation (Li Xiang, 2021) View paper
- Vision and Operator Learning Transformers (1 papers)
- [26] Cvit: Continuous vision transformer for operator learning (Wang, 2024) View paper
- Temporal and Event Reasoning (1 papers)

- [27] Transformer-based reasoning for learning evolutionary chain of events on temporal knowledge graph (Fang Zhi-yu, 2024) View paper
- Domain-Specific Reasoning Applications (4 papers)
- [19] GPT-FT: An Efficient Automated Feature Transformation Using GPT for Sequence Reconstruction and Performance Enhancement (Gao Yang, 2025) View paper
- [40] IPM-AgriGPT: A Large Language Model for Pest and Disease Management with a G-EA Framework and Agricultural Contextual Reasoning (Yuqin Zhang, 2025) View paper
- [42] Towards Robust Industrial Control Interpretation Through Comparative Analysis of Visionâ□□Language Models (Juan Izquierdo-Domenech, 2025) View paper
- [44] SeqFRT: Towards Effective Adaption of Foundation Model via Sequence Feature Reconstruction in Computational Pathology (Chengfei Cai, 2024) View paper

Narrative

Core task: training dynamics of chain of continuous thought in transformers. The field has evolved from early discrete prompting methods like Chain of Thought Prompting[1] toward a richer landscape that spans multiple paradigms. At the top level, the taxonomy distinguishes Continuous Latent Reasoning Paradigms—where models learn implicit reasoning steps in hidden representations—from Discrete CoT Prompting and Inference Methods that rely on explicit token sequences. Theoretical Foundations of CoT Mechanisms investigate the expressive power and convergence properties underlying these approaches, while CoT Training and Optimization Methods address how to effectively learn reasoning behaviors. Architectural Innovations for Reasoning explore modifications such as looped or recurrent structures (e.g., Looped Transformers[4]), and World Models and Planning-Based Reasoning connect transformers to decision-making frameworks like Reasoning as Planning[3]. Meanwhile, Transformer Learning Dynamics and Mechanisms examine gradient flow, feature evolution, and emergent behaviors during training, and Specialized Transformer Applications adapt these ideas to domains ranging from vision to reinforcement learning.

Within this landscape, a particularly active line of work focuses on understanding how transformers internalize multi-step reasoning during training. Superposition Training Dynamics[0] sits squarely in the Theoretical Foundations branch under Training Dynamics and Convergence Analysis, examining how reasoning emerges through superposed representations over the course of optimization. This contrasts with neighboring studies like Nonlinear Transformers CoT[28], which explores architectural nonlinearities to enhance reasoning capacity, and Kinetics of Reasoning[49], which applies statistical physics perspectives to characterize the evolution of reasoning states. Other closely related efforts include Multi Step Gradient Descent[37], which models iterative refinement processes, and works on continuous latent reasoning such as Continuous Latent Reasoning[2] and Scaling Latent Reasoning[33], which emphasize learning implicit thought chains without discrete tokens. Together, these studies reveal ongoing tensions between discrete versus continuous representations, the role of architectural depth versus recurrence, and the interplay between training objectives and emergent reasoning capabilities.

Related Works in Same Category

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

1. Training Nonlinear Transformers for Chain-of-Thought Inference: A Theoretical Generalization Analysis

Authors: LI Hongkang, Lu, Songtao, Hongkang Li, Chen, et al. (12 authors total) | **Year/Venue**: 2024 • International Conference on Learning Representations | **URL**: View paper

Abstract

Chain-of-Thought (CoT) is an efficient prompting method that enables the reasoning ability of large language models by augmenting the query using multiple examples with multiple intermediate steps. Despite the empirical success, the theoretical understanding of how to train a Transformer to achieve the CoT ability remains less explored. This is primarily due to the technical challenges involved in analyzing the nonconvex optimization on nonlinear attention models. To the best of our knowledge, t...

Relationship Analysis

Both papers belong to the Training Dynamics and Convergence Analysis category, studying how transformers learn reasoning capabilities through gradient-based training. They overlap in analyzing training dynamics of chain-of-thought mechanisms, but differ fundamentally in their focus: the original paper examines continuous CoT with superposition mechanisms in graph reachability problems, while the candidate paper studies discrete CoT generalization across tasks with nonlinear attention, focusing on sample complexity and generalization to unseen tasks rather than the emergence of superposition during training.

---

2. Transformers learn to implement multi-step gradient descent with chain of thought

Authors: Huang Jianhao, Wang Zi-xuan, Lee, Jason D. | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Chain of Thought (CoT) prompting has been shown to significantly improve the performance of large language models (LLMs), particularly in arithmetic and reasoning tasks, by instructing the model to produce intermediate reasoning steps. Despite the remarkable empirical success of CoT and its theoretical advantages in enhancing expressivity, the mechanisms underlying CoT training remain largely unexplored. In this paper, we study the training dynamics of transformers over a CoT objective on an in-...

Relationship Analysis

Both papers belong to the Training Dynamics and Convergence Analysis category, studying gradient-based training dynamics of transformers with chain-of-thought mechanisms. While the original paper analyzes continuous CoT training dynamics on graph reachability problems and reveals how bounded index-matching logits enable superposition emergence, the candidate paper studies discrete CoT training dynamics on in-context linear regression and proves convergence to multi-step gradient descent solutions. The key difference is that the original paper focuses on continuous thought spaces and superposition mechanisms in graph reasoning, whereas the candidate paper examines discrete CoT prompting and its ability to implement iterative optimization algorithms.

---

3. The Kinetics of Reasoning: How Chain-of-Thought Shapes Learning in Transformers?

Authors: Mavromatis, Costas, Shen Zhengyuan, Zhang Yunyi, Ioannidis, et al. (7 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Chain-of-thought (CoT) supervision can substantially improve transformer performance, yet the mechanisms by which models learn to follow and benefit from CoT remain poorly understood. We investigate these learning dynamics through the lens of grokking by pretraining transformers on symbolic reasoning tasks with tunable algorithmic complexity and controllable data composition to study their generalization. Models were trained under two settings: (i) producing only final answers, and (ii) emitting...

Relationship Analysis

Both papers belong to the Training Dynamics and Convergence Analysis category, studying how transformers learn reasoning capabilities through gradient-based training. While the original paper focuses on the emergence of superposition mechanisms in continuous chain-of-thought (CoT) training for graph reachability, analyzing how bounded attention logits enable parallel reasoning traces, the candidate paper investigates the learning kinetics of discrete CoT supervision across symbolic reasoning tasks, characterizing grokking dynamics through logistic curves and examining trace faithfulness during training. The key distinction is that the original paper analyzes continuous latent-space reasoning with superposition, whereas the candidate studies discrete token-based CoT with a focus on learning rate dynamics and unfaithfulness phases.

Contributions Analysis

Overall novelty summary. The paper contributes a theoretical analysis of how continuous chain-of-thought mechanisms emerge during gradient-based training in two-layer transformers solving directed graph reachability. It resides in the 'Training Dynamics and Convergence Analysis' leaf under 'Theoretical Foundations of CoT Mechanisms,' alongside three sibling papers examining gradient dynamics and convergence properties. This leaf represents a moderately populated research direction within the broader taxonomy of 50 papers across 36 topics, indicating focused but not overcrowded attention to training dynamics questions in continuous CoT.

The taxonomy reveals neighboring theoretical branches including 'Expressivity and Computational Power' (four papers proving what transformers can solve with CoT) and 'Superposition and Parallel Reasoning Theory' (one paper on maintaining multiple traces). The paper bridges these areas by explaining how superposition—previously shown to enable parallel reasoning—actually emerges through training. Nearby practical branches like 'Continuous CoT Architectures and Training' (three papers on model implementations) and 'Latent-Variable CoT Training' (one paper on unsupervised optimization) address related but distinct questions about architecture design and training objectives rather than gradient dynamics.

Among 21 candidates examined across three contributions, no clear refutations emerged. The core contribution on training dynamics analyzed 10 candidates with none providing overlapping prior work; the bounded index-matching logit behavior examined 1 candidate without refutation; and the superposition emergence explanation reviewed 10 candidates, again finding no direct overlap. This limited search scope—focused on top semantic matches and citations—suggests the specific combination of continuous CoT, training dynamics, and superposition emergence may occupy relatively unexplored theoretical territory, though the analysis cannot claim exhaustive coverage of all relevant gradient dynamics literature.

Based on examination of 21 semantically related papers, the work appears to address a gap between expressivity proofs and empirical continuous CoT implementations by analyzing how training naturally discovers superposition mechanisms. The bounded search scope means potentially relevant work in broader optimization theory or neural tangent kernel analyses may exist outside the examined candidates. The taxonomy positioning and sibling paper analysis suggest this represents a natural theoretical extension within an active but not saturated research direction.

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

Contribution 1: Theoretical analysis of training dynamics for continuous chain-of-thought

Description: The authors provide a theoretical analysis of how gradient-based training naturally leads to the superposition mechanism in continuous chain-of-thought models. They analyze two training stages: thought generation and prediction, revealing how the model learns to maintain multiple reasoning traces in parallel.

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. GET-Zero: Graph Embodiment Transformer for Zero-Shot Embodiment Generalization

URL: View paper

Brief Assessment

GET-Zero[65] focuses on embodiment-aware control policies for robotic manipulation using graph transformers, not on training dynamics of continuous chain-of-thought models or graph reachability problems.

---

2. A transformer-based knowledge graph embedding model combining graph paths and local neighborhood

URL: View paper

Brief Assessment

Graph Paths Embedding[71] focuses on knowledge graph embedding using transformers for graph structure learning, not on training dynamics of continuous chain-of-thought models for reasoning tasks.

---

3. Understanding transformer reasoning capabilities via graph algorithms

URL: View paper

Brief Assessment

Graph Algorithms Reasoning[63] focuses on representational capabilities of transformers for graph algorithms in different parameter regimes, not on training dynamics of continuous chain-of-thought models or superposition mechanisms.

---

4. Semformer: Transformer language models with semantic planning

URL: View paper

Brief Assessment

Semformer Semantic Planning[66] focuses on semantic planning with trainable tokens for graph path-finding and general language modeling, not on analyzing training dynamics of continuous chain-of-thought mechanisms or superposition emergence in transformers.

---

5. Transformers struggle to learn to search

URL: View paper

Brief Assessment

Transformers Struggle Search[70] focuses on learning graph search algorithms through standard training, not on continuous chain-of-thought training dynamics or superposition mechanisms in latent reasoning spaces.

---

6. Lost in transmission: When and why llms fail to reason globally

URL: View paper

Brief Assessment

Lost in Transmission[68] focuses on communication bandwidth constraints in transformers for graph reachability, not on training dynamics or how gradient-based methods lead to superposition mechanisms in continuous CoT models.

---

7. A generalization of transformer networks to graphs

URL: View paper

Brief Assessment

Transformer to Graphs[62] focuses on generalizing transformer architectures to arbitrary graph structures for tasks like molecular property prediction and node classification. It does not address training dynamics or chain-of-thought reasoning mechanisms.

---

8. Self-Supervised Graph Transformer with Contrastive Learning for Brain Connectivity Analysis Towards Improving Autism Detection

URL: View paper

Brief Assessment

Graph Transformer Autism[67] focuses on self-supervised learning for brain connectivity analysis using graph transformers, not on training dynamics of continuous chain-of-thought models or graph reachability problems. The candidate addresses autism detection from fMRI data, while the original analyzes how transformers learn superposition mechanisms during reasoning tasks.

---

9. Do transformers really perform badly for graph representation?

URL: View paper

Brief Assessment

Transformers Graph Representation[64] focuses on graph representation learning with structural encodings for transformers, not on training dynamics of continuous chain-of-thought mechanisms or gradient-based learning of superposition in reasoning tasks.

---

10. Long-range brain graph transformer

URL: View paper

Brief Assessment

Long Range Brain Graph[69] focuses on brain network analysis using graph transformers to capture long-range dependencies in brain connectivity graphs, not on training dynamics of transformers for reasoning tasks or chain-of-thought mechanisms.

Contribution 2: Discovery of bounded index-matching logit behavior

Description: The authors discover that the index-matching logit, which quantifies local search capability, grows initially but remains bounded during training with continuous CoT. This bounded behavior contrasts with unbounded logit growth in discrete settings and enables effective exploration-exploitation balance.

This contribution was assessed against **1 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. InfiFusion: A Unified Framework for Enhanced Cross-Model Reasoning via LLM Fusion

URL: View paper

Brief Assessment

InfiFusion[61] focuses on model fusion techniques for combining multiple domain-specialized LLMs through logit distillation, not on training dynamics or logit growth bounds during reasoning tasks with continuous chain-of-thought.

Contribution 3: Explanation of superposition emergence through bounded logits

Description: The authors explain how bounded index-matching logits lead to superposition by balancing exploration and exploitation. When logits remain bounded, the model assigns comparable weights to multiple plausible reasoning paths rather than over-committing to a single path, naturally producing the superposition mechanism.

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. Learning to Gaze: Bio-Inspired Attention Adaptation Strategy for Social Robots

URL: View paper

Brief Assessment

Bio Inspired Gaze[56] focuses on attention allocation in social robots for gaze control, not on neural reasoning mechanisms or superposition in transformers. The exploration-exploitation balance mentioned relates to robotic attention systems, not to logit dynamics in language models.

---

2. GraphFusion-HRL: Multi-Modal Hierarchical Reinforcement Graph Learning for Context-Rich Recommender Systems

URL: View paper

Brief Assessment

GraphFusion-HRL[51] focuses on multi-modal hierarchical reinforcement learning for recommender systems with graph-based reasoning, not on the emergence of superposition mechanisms through bounded logits in chain-of-thought reasoning for transformers.

---

3. Hulu video recommendation: from relevance to reasoning

URL: View paper

Brief Assessment

Hulu Video Recommendation[58] focuses on video recommendation systems using knowledge graphs and contextual bandits for exploitation-exploration balance in user engagement scenarios, not on transformer training dynamics or continuous chain-of-thought reasoning mechanisms.

---

4. Each Complexity Deserves a Pruning Policy

URL: View paper

Brief Assessment

Complexity Deserves Pruning[59] focuses on token pruning in vision-language models based on visual-textual complexity, not on explaining superposition emergence through bounded logits in reasoning transformers.

---

5. An Integrated Graph Neural Network and Reinforcement Learning Framework for Intelligent Drug Discovery

URL: View paper

Brief Assessment

Graph Neural Drug Discovery[53] focuses on molecular generation using graph neural networks and reinforcement learning for drug discovery, not on transformer training dynamics or reasoning mechanisms. The candidate does not address superposition, bounded logits, or exploration-exploitation balance in neural reasoning.

---

6. Optimizing Landmark Graphs in DHRL: A Dual Approach of Attention and Weighted Sampling

URL: View paper

Brief Assessment

Landmark Graphs DHRL[57] focuses on hierarchical reinforcement learning with landmark graphs for navigation tasks, not on explaining superposition mechanisms in neural reasoning through bounded logits and attention weight distribution.

---

7. Aggregating knowledge-aware graph neural network and adaptive relational attention for recommendation

URL: View paper

Brief Assessment

Knowledge Aware Graph[54] focuses on recommendation systems using graph neural networks and attention mechanisms for user-item interactions, not on transformer training dynamics or reasoning through superposition in continuous thought spaces.

---

8. Balancing Exploration and Exploitation for Solving Large-scale Multiobjective Optimization via Attention Mechanism

URL: View paper

Brief Assessment

Exploration Exploitation Attention[52] addresses exploration-exploitation balance in multiobjective optimization through attention weights on decision variables, not neural reasoning paths or transformer logit dynamics in chain-of-thought settings.

---

9. GraphPuma-Flow: Graph-Neural Metaheuristics for Constraint-Aware Hybrid Scheduling

URL: View paper

Brief Assessment

GraphPuma Flow[60] focuses on hybrid scheduling with GNN-guided evolutionary methods for constraint-aware optimization, not on transformer training dynamics or superposition mechanisms in continuous reasoning.

---

10. Deep attentive belief propagation: Integrating reasoning and learning for solving constraint optimization problems

URL: View paper

Brief Assessment

Attentive Belief Propagation[55] focuses on constraint optimization problems using attention mechanisms to balance exploration-exploitation in message-passing algorithms, not on explaining superposition emergence in continuous thought reasoning for language models.

Appendix: Text Similarity Detection

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

References

- [0] Emergence of Superposition: Unveiling the Training Dynamics of Chain of Continuous Thought View paper
- [1] Chain of Thought Prompting Elicits Reasoning in Large Language Models View paper
- [2] Training Large Language Models to Reason in a Continuous Latent Space View paper
- [3] Reasoning with Language Model is Planning with World Model View paper
- [4] Reasoning with latent thoughts: On the power of looped transformers View paper
- [5] Towards revealing the mystery behind chain of thought: a theoretical perspective View paper
- [6] Codi: Compressing chain-of-thought into continuous space via self-distillation View paper
- [7] Faithful Chain-of-Thought Reasoning View paper
- [8] Flowing Through Layers: A Continuous Dynamical Systems Perspective on Transformers View paper
- [9] Self-Consistency Improves Chain of Thought Reasoning in Language Models View paper
- [10] FlowVLA: Visual Chain of Thought-based Motion Reasoning for Vision-Language-Action Models View paper
- [11] A survey on latent reasoning View paper
- [12] Chain of thought empowers transformers to solve inherently serial problems View paper
- [13] Prefix-Tuning: Optimizing Continuous Prompts for Generation View paper
- [14] Hunyuan-turbos: Advancing large language models through mamba-transformer synergy and adaptive chain-of-thought View paper
- [15] From sparse dependence to sparse attention: unveiling how chain-of-thought enhances transformer sample efficiency View paper
- [16] Optimizing large language models through highly dense reward structures and recursive thought process using monte carlo tree search View paper
- [17] Chain-of-Thought Hub: A Continuous Effort to Measure Large Language Models' Reasoning Performance View paper
- [18] Reasoning by Superposition: A Theoretical Perspective on Chain of Continuous Thought View paper
- [19] GPT-FT: An Efficient Automated Feature Transformation Using GPT for Sequence Reconstruction and Performance Enhancement View paper
- [20] The Expressive Power of Transformers with Chain of Thought View paper
- [21] Enhancing Auto-regressive Chain-of-Thought through Loop-Aligned Reasoning View paper
- [22] Boosting Language Models Reasoning with Chain-of-Knowledge Prompting View paper
- [23] Language Models are Multilingual Chain-of-Thought Reasoners View paper
- [24] Transformers provably learn chain-of-thought reasoning with length generalization View paper
- [25] Decision transformer: Reinforcement learning via sequence modeling View paper
- [26] Cvit: Continuous vision transformer for operator learning View paper
- [27] Transformer-based reasoning for learning evolutionary chain of events on temporal knowledge graph View paper
- [28] Training Nonlinear Transformers for Chain-of-Thought Inference: A Theoretical Generalization Analysis View paper
- [29] CTRNN-Transformer Adding Continuous Time Neural Models to Transformers View paper

- [30] On transforming reinforcement learning with transformers: The development trajectory View paper
- [31] Learning to Play Atari in a World of Tokens View paper
- [32] A Theoretical Understanding of Chain-of-Thought: Coherent Reasoning and Error-Aware Demonstration View paper
- [33] Scaling latent reasoning via looped language models View paper
- [34] ContiFormer: Continuous-Time Transformer for Irregular Time Series Modeling View paper
- [35] Training Chain-of-Thought via Latent-Variable Inference View paper
- [36] Beyond In-Distribution Success: Scaling Curves of CoT Granularity for Language Model Generalization View paper
- [37] Transformers learn to implement multi-step gradient descent with chain of thought View paper
- [38] How Transformers Learn Regular Language Recognition: A Theoretical Study on Training Dynamics and Implicit Bias View paper
- [39] PENCIL: Long Thoughts with Short Memory View paper
- [40] IPM-AgriGPT: A Large Language Model for Pest and Disease Management with a G-EA Framework and Agricultural Contextual Reasoning View paper
- [41] Continuous Chain of Thought Enables Parallel Exploration and Reasoning View paper
- [42] Towards Robust Industrial Control Interpretation Through Comparative Analysis of Visionâ□□Language Models View paper
- [43] RedStar: Does Scaling Long-CoT Data Unlock Better Slow-Reasoning Systems? View paper
- [44] SeqFRT: Towards Effective Adaption of Foundation Model via Sequence Feature Reconstruction in Computational Pathology View paper
- [45] Towards a Theoretical Understanding of the 'Reversal Curse' via Training Dynamics View paper
- [46] In-context decision transformer: Reinforcement learning via hierarchical chain-of-thought View paper
- [47] A Statistical Physics of Language Model Reasoning View paper
- [48] Autoregressive+ Chain of Thought= Recurrent: Recurrence's Role in Language Models' Computability and a Revisit of Recurrent Transformer View paper
- [49] The Kinetics of Reasoning: How Chain-of-Thought Shapes Learning in Transformers? View paper
- [50] Can Transformers Reason Logically? A Study in SAT Solving View paper
- [51] GraphFusion-HRL: Multi-Modal Hierarchical Reinforcement Graph Learning for Context-Rich Recommender Systems View paper
- [52] Balancing Exploration and Exploitation for Solving Large-scale Multiobjective Optimization via Attention Mechanism View paper
- [53] An Integrated Graph Neural Network and Reinforcement Learning Framework for Intelligent Drug Discovery View paper
- [54] Aggregating knowledge-aware graph neural network and adaptive relational attention for recommendation View paper
- [55] Deep attentive belief propagation: Integrating reasoning and learning for solving constraint optimization problems View paper
- [56] Learning to Gaze: Bio-Inspired Attention Adaptation Strategy for Social Robots View paper
- [57] Optimizing Landmark Graphs in DHRL: A Dual Approach of Attention and Weighted Sampling View paper
- [58] Hulu video recommendation: from relevance to reasoning View paper
- [59] Each Complexity Deserves a Pruning Policy View paper
- [60] GraphPuma-Flow: Graph-Neural Metaheuristics for Constraint-Aware Hybrid Scheduling View paper
- [61] InfiFusion: A Unified Framework for Enhanced Cross-Model Reasoning via LLM Fusion View paper
- [62] A generalization of transformer networks to graphs View paper
- [63] Understanding transformer reasoning capabilities via graph algorithms View paper
- [64] Do transformers really perform badly for graph representation? View paper
- [65] GET-Zero: Graph Embodiment Transformer for Zero-Shot Embodiment Generalization View paper
- [66] Semformer: Transformer language models with semantic planning View paper
- [67] Self-Supervised Graph Transformer with Contrastive Learning for Brain Connectivity Analysis Towards Improving Autism Detection View paper
- [68] Lost in transmission: When and why llms fail to reason globally View paper
- [69] Long-range brain graph transformer View paper
- [70] Transformers struggle to learn to search View paper
- [71] A transformer-based knowledge graph embedding model combining graph paths and local neighborhood View paper