Novelty Assessment Report

Paper: Characterizing Pattern Matching and Its Limits on Compositional Task Structures

PDF URL: https://openreview.net/pdf?id=VCjlm003WL

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-27

Abstract

Despite impressive capabilities, LLMs' successes often rely on pattern-matching behaviors, yet these are also linked to OOD generalization failures in compositional tasks. However, behavioral studies commonly employ task setups that allow multiple generalization sources (e.g., algebraic invariances, structural repetition), obscuring a precise and testable account of how well LLMs perform generalization through pattern matching and their limitations. To address this ambiguity, we first formalize pattern matching as functional equivalence, i.e., identifying pairs of subsequences of inputs that consistently lead to identical results when the rest of the input is held constant. Then, we systematically study how decoder-only Transformer and Mamba behave in controlled tasks with compositional structures that isolate this mechanism. Our formalism yields predictive and quantitative insights: (1) Instance-wise success of pattern matching is well predicted by the number of contexts witnessing the relevant functional equivalence. We prove a tight sample complexity bound of learning a two-hop structure by identifying the exponent of the data scaling law for perfect in-domain generalization. Our empirical results align with the theoretical prediction, under 20× parameter scaling and across architectures. (3) Path ambiguity is a structural barrier: when a variable influences the output via multiple paths, models fail to form unified intermediate state representations, impairing accuracy and interpretability. (4) Chain-of-Thought reduces data requirements yet does not resolve path ambiguity. Hence, we provide a predictive, falsifiable boundary for pattern matching and a foundational diagnostic for disentangling mixed generalization mechanisms.

Core Task Landscape

This paper addresses: **Pattern Matching in Compositional Generalization Tasks**

A total of **50 papers** were analyzed and organized into a taxonomy with **23 categories**.

Taxonomy Overview

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

- **Theoretical Foundations and Mechanisms**
- **Architecture-Specific Studies**
- **Training Methods and Data Strategies**
- **Representation Learning and Alignment**
- **Domain-Specific Applications**
- **Empirical Analysis and Benchmarking**
- **Specialized Compositional Structures**

Complete Taxonomy Tree

- Pattern Matching in Compositional Generalization Tasks Survey Taxonomy
- Theoretical Foundations and Mechanisms
  - Formalization and Sample Complexity ★ (2 papers)
  - [0] Characterizing Pattern Matching and Its Limits on Compositional Task Structures (Anon et al., 2026) View paper
  - [35] The Coverage Principle: A Framework for Understanding Compositional Generalization (Chang, 2025) View paper
  - Neural and Cognitive Mechanisms (2 papers)
  - [1] A neural mechanism for compositional generalization of structure in humans (Lennart Luettgau, 2025) View paper
  - [14] Out-of-distribution generalization via composition: A lens through induction heads in Transformers (Jiajun Song, 2024) View paper
  - Identifiability and Provable Guarantees (1 papers)
  - [25] Provable Compositional Generalization for Object-Centric Learning (Wiedemer, 2023) View paper
- Architecture-Specific Studies
  - Transformer and Attention Mechanisms (3 papers)
  - [46] Compositional Generalization for Neural Semantic Parsing via Span-level Supervised Attention (Andreas Jacob, 2021) View paper
  - [48] Exploring Compositional Generalization (in ReCOGS_pos) by Transformers using Restricted Access Sequence Processing (RASP) (Bruns, 2025) View paper
  - Diffusion Models (2 papers)
  - [2] Local Mechanisms of Compositional Generalization in Conditional Diffusion (Bradley, 2025) View paper
  - [10] Swing-by Dynamics in Concept Learning and Compositional Generalization (Yang Yongyi, 2024) View paper
  - Large Language Models (3 papers)
  - [19] Do phd-level llms truly grasp elementary addition? probing rule learning vs. memorization in large language models (Y Yan, 2025) View paper
  - [21] Impact of Noise on LLM-Models Performance in Abstraction and Reasoning Corpus (ARC) Tasks with Model Temperature Considerations (SHINDE KRISHNA, 2025) View paper

  - [49] Understanding Subword Compositionality of Large Language Models (Peng Qiwei, 2025) View paper
  - Other Architectures (2 papers)
  - [8] Multimodal graph networks for compositional generalization in visual question answering (Raeid Saqur, 2020) View paper
  - [45] Differentiable Compositional Kernel Learning for Gaussian Processes (Shengyang Sun, 2018) View paper
- Training Methods and Data Strategies
  - Data Augmentation and Synthesis (3 papers)
  - [7] Align and Augment: Generative Data Augmentation for Compositional Generalization (Cazzaro, 2024) View paper
  - [18] Improving compositional generalization with latent structure and data augmentation (Qiu, 2022) View paper
  - [30] Finding needles in a haystack: Sampling Structurally-diverse Training Sets from Synthetic Data for Compositional Generalization (I. Oren, 2021) View paper
  - Data Factors and Scaling (1 papers)
  - [3] Data Factors for Better Compositional Generalization (Bansal, 2023) View paper
  - Self-Training and Iterative Refinement (2 papers)
  - [23] Bootstrapped structural prompting for analogical reasoning in pretrained language models (Yue, 2024) View paper
  - [31] Improving compositional generalization with self-training for data-to-text generation (Mehta, 2022) View paper
  - Structural Inductive Biases (3 papers)
  - [24] Structural generalization in COGS: Supertagging is (almost) all you need (Petit, 2023) View paper
  - [40] Compositional generalisation with structured reordering and fertility layers (KOLLER Alexander, 2023) View paper
  - [47] SymCHMâ□□an unsupervised approach for pattern discovery in symbolic music with a compositional hierarchical model (MatevÅ¾ Pesek, 2017) View paper
- Representation Learning and Alignment
  - Cross-Modal Alignment (4 papers)
  - [4] Decoupled global-local alignment for improving compositional understanding (Hu Xiaoxing, 2025) View paper
  - [11] Latent surface alignment through oscillatory token drift in instruction-following large language models (Allan, 2025) View paper
  - [33] DiffCloth: Diffusion Based Garment Synthesis and Manipulation via Structural Cross-modal Semantic Alignment (Xu-jie Zhang, 2023) View paper
  - [38] Cross-modal Attention Congruence Regularization for Vision-Language Relation Alignment (Liang, 2022) View paper
  - Intermediate Representations (2 papers)
  - [13] Towards Compositional Generalization in Code Search (Choi, 2022) View paper
  - [34] Unlocking Compositional Generalization in Pre-trained Models Using Intermediate Representations (Herzig, 2021) View paper
  - Neuro-Symbolic Approaches (2 papers)
  - [9] NeSyCoCo: A Neuro-Symbolic Concept Composer for Compositional Generalization (Kamali, 2024) View paper
  - [27] Compositional generalization with grounded language models (Ã□vrelid, 2024) View paper
- Domain-Specific Applications
  - Semantic Parsing and Text-to-SQL (2 papers)
  - [16] Improving compositional generalization in semantic parsing (Inbar Oren, 2020) View paper
  - [17] Exploring the Compositional Generalization in Context Dependent Text-to-SQL Parsing (Aiwei Liu, 2023) View paper
  - Vision-Language Tasks (3 papers)
  - [12] In-Context Compositional Generalization for Large Vision-Language Models (Jia Yun-de, 2024) View paper
  - [22] HCCM: Hierarchical Cross-Granularity Contrastive and Matching Learning for Natural Language-Guided Drones (Ruan Hao, 2025) View paper
  - [28] Multi-Modal Transformer With Global-Local Alignment for Composed Query Image Retrieval (Yahui Xu, 2023) View paper
  - Code and Program Synthesis (3 papers)
  - [5] CoScheme: Compositional Copatterns in Scheme (Downen, 2025) View paper
  - [32] CoScheme: Compositional Copatterns (P Downen, 2025) View paper
  - [37] A pattern for almost compositional functions (BjÃ¶rn Bringert, 2008) View paper
  - Specialized Domains (3 papers)
  - [15] UniMatch: Universal Matching from Atom to Task for Few-Shot Drug Discovery (Li, 2025) View paper
  - [41] Data Extraction via Semantic Regular Expression Synthesis (Qiaochu Chen, 2023) View paper
  - [42] Human re-identification by matching compositional template with cluster sampling (Xu, 2013) View paper
- Empirical Analysis and Benchmarking
  - Cross-Domain Generalization (2 papers)
  - [20] Multi-Scale Explicit Matching and Mutual Subject Teacher Learning for Generalizable Person Re-Identification (Kaixiang Chen, 2024) View paper
  - [26] Revisiting Depth Completion from a Stereo Matching Perspective for Cross-domain Generalization (Luca Bartolomei, 2023) View paper
  - In-Context Learning (1 papers)
  - [6] LLM Retrieval-Augmented Generation with Compositional Prompts and Confidence Calibration (Han, 2025) View paper
  - Supervised Fine-Tuning Analysis (1 papers)
  - [44] Debunk the Myth of SFT Generalization (Lin Xiao-feng, 2025) View paper
  - Systematic Generalization Patterns (2 papers)
  - [39] Systematic generalization in connectionist models (CsordÃ¡s, 2023) View paper
  - [43] Exploring Patterns of Algorithmic Generalization in Deep Learning (Anil, 2025) View paper
- Specialized Compositional Structures (2 papers)
  - [29] Matching Plans for Frame Inference in Compositional Reasoning (LÃ¶Ã¶w, 2024) View paper
  - [36] SODA: Structural Pre-Decoupling and Co-Aligning for Video Compositional Representation (Peng Huang, 2025) View paper

Narrative

Core task: pattern matching in compositional generalization tasks. The field examines how models learn to combine familiar components in novel ways, a challenge that spans semantic parsing, vision-language reasoning, and algorithmic tasks. The taxonomy reflects a multifaceted landscape: Theoretical Foundations and Mechanisms explore formal guarantees and sample complexity bounds, as in Pattern Matching Limits[0] and Coverage Principle[35]; Architecture-Specific Studies investigate how transformers, graph networks, and specialized modules handle compositional structure, with works like RASP COGS Transformers[50] and Multimodal Graph Networks[8] probing inductive biases; Training Methods and Data Strategies address curriculum design and augmentation techniques, exemplified by Sampling Diverse Training[30] and Latent Structure Augmentation[18]; Representation Learning and Alignment focus on aligning latent spaces and grounding symbols, seen in Latent Surface Alignment[11] and Grounded Language Models[27]; Domain-Specific Applications tailor compositional reasoning to text-to-SQL, code search, and vision tasks, including Compositional Text-to-SQL[17] and Compositional Code Search[13]; Empirical Analysis and Benchmarking rigorously evaluate generalization gaps on datasets like COGS and ARC, as in Noise Impact ARC[21]; and Specialized Compositional Structures study copatterns, hierarchical decompositions, and other structured representations, with CoScheme Copatterns[5] and Compositional Copatterns[32].

Several active lines of work highlight contrasting emphases and open questions. One thread pursues provable guarantees on when pattern matching suffices for compositional tasks, balancing expressiveness with learnability; another investigates architectural interventions—such as attention supervision or modular designs—that encourage systematic recombination without explicit symbolic scaffolding. A third direction explores data-centric strategies, asking whether diverse sampling or augmentation can substitute for architectural inductive biases. Pattern Matching Limits[0] sits squarely within the theoretical branch, formalizing sample complexity and expressiveness boundaries for pattern-based learners. Its emphasis on rigorous bounds complements empirical studies like Neural Compositional Generalization[1] and contrasts with data-driven approaches such as Data Compositional Generalization[3], which prioritize scalable training recipes over worst-case analysis. Together, these perspectives frame ongoing debates about whether compositional generalization demands principled architectural priors, richer training regimes, or fundamentally new learning paradigms.

Related Works in Same Category

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

1. The Coverage Principle: A Framework for Understanding Compositional Generalization

Authors: Chang, Hoyeon, Park Jin-ho, Cho, Hanseul, et al. (17 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Large language models excel at pattern matching, yet often fall short in systematic compositional generalization. We propose the coverage principle: a data-centric framework showing that models relying primarily on pattern matching for compositional tasks cannot reliably generalize beyond substituting fragments that yield identical results when used in the same contexts. We demonstrate that this framework has a strong predictive power for the generalization capabilities of Transformers. First, w...

⚠ Similarity Notice

These papers share nearly identical titles, abstracts, and core technical content, including the same formalization of pattern matching as functional equivalence, identical experimental setups (2-HOP, PARALLEL 2-HOP, 3-HOP, NON-TREE tasks), and the same theoretical result (Theorem 6.1 on sample complexity bounds). The candidate paper appears to be a published or revised version of the original submission, as evidenced by the identical structure, figures, and mathematical formulations throughout both documents.

Contributions Analysis

Overall novelty summary. The paper formalizes pattern matching as functional equivalence and derives tight sample complexity bounds for two-hop compositional structures, positioning itself within the 'Formalization and Sample Complexity' leaf of the taxonomy. This leaf contains only two papers, including the original work, indicating a relatively sparse research direction focused on rigorous theoretical analysis. The paper's emphasis on provable bounds and quantitative predictions distinguishes it from the broader empirical literature on compositional generalization, which dominates other branches of the taxonomy.

The taxonomy reveals substantial activity in neighboring areas: 'Transformer and Attention Mechanisms' examines architectural inductive biases, 'Data Augmentation and Synthesis' explores training-centric solutions, and 'Systematic Generalization Patterns' characterizes empirical failure modes. The original paper's theoretical approach contrasts with these empirical and architectural threads, offering formal guarantees rather than scalable heuristics. Its focus on functional equivalence as a unifying principle bridges the gap between mechanistic understanding (explored in 'Neural and Cognitive Mechanisms') and practical generalization (studied in 'Architecture-Specific Studies'), though it remains firmly grounded in formal analysis rather than architectural design.

Among 24 candidates examined across three contributions, no refutable prior work was identified. The formalization of pattern matching as functional equivalence examined 10 candidates with no overlaps; the sample complexity bound for two-hop tasks examined 6 candidates with no refutations; and the path ambiguity analysis examined 8 candidates without finding prior coverage. This suggests that within the limited search scope—focused on top semantic matches and citation expansion—the paper's specific theoretical contributions appear distinct. However, the small candidate pool (24 total) and the sparse theoretical leaf (2 papers) mean the analysis captures a narrow slice of the literature, leaving open the possibility of related work in adjacent formal methods or complexity theory communities.

Given the limited search scope and the paper's placement in a sparsely populated theoretical leaf, the contributions appear novel within the examined context. The formalization and sample complexity results address a gap between empirical studies of compositional generalization and rigorous theoretical characterization. However, the analysis does not cover exhaustive searches in formal learning theory or adjacent mathematical frameworks, so the assessment reflects novelty within the specific compositional generalization literature surveyed rather than across all relevant theoretical domains.

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

Contribution 1: Formalization of pattern matching as functional equivalence

Description: The authors introduce a model-agnostic, data-centric definition of pattern matching by formalizing it as functional equivalence. This framework defines when two input subsequences are functionally equivalent based on consistent behavior in shared contexts, and introduces the concept of k-coverage to predict which test inputs can be reliably handled through pattern matching.

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. Model Stitching by Functional Latent Alignment

URL: View paper

Brief Assessment

Functional Latent Alignment[69] focuses on aligning neural network representations for model stitching in computer vision, not on formalizing pattern matching or functional equivalence in compositional tasks. The candidate addresses a different problem domain (model similarity evaluation) than the original paper's focus on characterizing pattern matching behaviors in language models.

---

2. Cross-stage neural pattern similarity in the hippocampus predicts false memory derived from post-event inaccurate information

URL: View paper

Brief Assessment

Hippocampal Pattern Similarity[67] investigates neural pattern similarity in memory formation across temporal stages, not functional equivalence in input subsequences for pattern matching in neural networks. The candidate focuses on hippocampal representations during false memory creation, while the original paper addresses compositional task structures in language models.

3. How not to Stitch Representations to Measure Similarity: Task Loss Matching versus Direct Matching

URL: View paper

Brief Assessment

Task Loss Matching[65] addresses representation similarity measurement in neural networks through stitching methods, not pattern matching formalization. The candidate focuses on comparing task loss matching versus direct matching for measuring functional similarity between network layers, which is orthogonal to the original paper's contribution of formalizing pattern matching as functional equivalence in compositional tasks.

4. Greater neural pattern similarity across repetitions is associated with better memory

URL: View paper

Brief Assessment

Neural Pattern Memory[71] studies neural pattern similarity across stimulus repetitions in memory encoding, not functional equivalence in input subsequences for pattern matching in neural networks. The candidate focuses on neuroimaging of memory formation, while the original paper addresses compositional task structures in language models.

5. Design of neuron-calculators for the normalized equivalence of two matrix arrays based on FPGA for self-learning equivalently convolutional neural networks (SLE_CNNs)

URL: View paper

Brief Assessment

Neuron-Calculators FPGA[68] focuses on hardware implementation of normalized equivalence operations for image processing in convolutional neural networks, not on formalizing pattern matching as functional equivalence in the context of compositional task structures or input subsequence behavior.

6. Family Cohesion Moderates the Relation between Parentâ Child Neural Connectivity Pattern Similarity and Youth's Emotional Adjustment

URL: View paper

Brief Assessment

Parent-Child Neural Similarity[66] studies neural connectivity patterns between parents and children during emotional processing, not functional equivalence in input subsequences for pattern matching in neural networks. These are entirely different research domains with no overlap in methodology or concepts.

7. Model Stitching: Looking For Functional Similarity Between Representations

URL: View paper

Brief Assessment

Model Stitching Similarity[70] focuses on comparing neural network representations through learned transformations to test interchangeability between layers, not on formalizing pattern matching or functional equivalence of input subsequences in compositional tasks.

8. Not all solutions are created equal: An analytical dissociation of functional and representational similarity in deep linear neural networks

URL: View paper

Brief Assessment

Functional Representational Dissociation[64] focuses on analyzing functional versus representational similarity in deep linear neural networks, not on formalizing pattern matching or functional equivalence for compositional tasks in language models.

9. BinDNN: Resilient Function Matching Using Deep Learning

URL: View paper

Brief Assessment

BinDNN[73] addresses binary function matching in compiled code using deep learning, not pattern matching formalization in neural network inputs or compositional task structures.

10. Multi-input Vision Transformer withÂ Similarity Matching

URL: View paper

Brief Assessment

Multi-input Vision Transformer[72] focuses on vision transformer architectures with similarity matching for multi-input scenarios, not on formalizing pattern matching as functional equivalence in compositional tasks or neural network generalization theory.

Contribution 2: Tight sample complexity bound for two-hop compositional tasks

Description: The authors establish and prove a theoretical bound (Theorem 6.1) showing that the training dataset size required for perfect in-domain generalization on two-hop tasks scales polynomially with token set size, with exponent c in [2, 2.5). This bound is empirically verified across different model sizes and architectures.

This contribution was assessed against **6 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Unveiling the Mechanisms of Explicit CoT Training: How CoT Enhances Reasoning Generalization

URL: View paper

Brief Assessment

Explicit CoT Training[51] focuses on how chain-of-thought training reshapes model representations and improves reasoning generalization, not on establishing sample complexity bounds for compositional tasks.

2. Learning to Prompt in Unknown Environments: A POMDP Framework with Compositional Actions for Large Language Models

URL: View paper

Brief Assessment

Prompt POMDP Framework[55] focuses on learning prompting strategies for LLMs through reinforcement learning in unknown environments, not on sample complexity bounds for compositional generalization in two-hop reasoning tasks. The candidate addresses a fundamentally different problem domain (prompt optimization) than the original paper's theoretical analysis of pattern matching and data scaling laws.

3. Dice: Dynamic in-context example selection in llm agents via efficient knowledge transfer

URL: View paper

Brief Assessment

Dice Dynamic Examples[52] focuses on dynamic in-context example selection for LLM agents in multi-step reasoning tasks, not on establishing sample complexity bounds for compositional generalization in two-hop reasoning structures.

4. Identity Bridge: Enabling Implicit Reasoning via Shared Latent Memory

URL: View paper

Brief Assessment

Identity Bridge[54] focuses on enabling two-hop reasoning through identity supervision and latent memory alignment, not on establishing sample complexity bounds or data scaling laws for compositional generalization.

5. Language models can learn implicit multi-hop reasoning, but only if they have lots of training data

URL: View paper

Brief Assessment

Implicit Multi-hop Reasoning[53] focuses on empirical data scaling laws for multi-hop reasoning tasks but does not establish theoretical sample complexity bounds with specific polynomial exponents as claimed in the original paper's Theorem 6.1.

6. The Coverage Principle: A Framework for Understanding Compositional Generalization

URL: View paper

Brief Assessment

Coverage Principle[35] focuses on pattern matching and functional equivalence in compositional tasks, establishing sample complexity bounds for two-hop structures. This work does not address the same problem domain as the original paper's RL framework contributions.

Contribution 3: Identification of path ambiguity as a structural barrier

Description: The authors identify and characterize path ambiguity as a failure mode where variables affecting outputs through multiple computational paths prevent models from forming unified intermediate representations. This structural problem persists even with Chain-of-Thought supervision and impairs both generalization performance and model interpretability.

This contribution was assessed against **8 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning

URL: View paper

Brief Assessment

Hin2vec[57] focuses on meta-path exploration in heterogeneous information networks for representation learning, not on path ambiguity in neural network compositional tasks or intermediate representation formation.

2. A survey of the recent architectures of deep convolutional neural networks

URL: View paper

Brief Assessment

Deep CNN Architectures[56] is a survey paper focused on architectural innovations in convolutional neural networks for computer vision tasks. It does not address compositional task structures, multiple computational paths, or intermediate representations in the context described by the original paper's contribution about path ambiguity.

3. Diverse types of expertise in facial recognition

URL: View paper

Brief Assessment

Facial Recognition Expertise[59] focuses on comparing different types of human and AI expertise in facial recognition tasks, not on neural network learning mechanisms or computational path structures in compositional tasks.

4. Object recognition under occlusion revisited: elucidating algorithmic advantages of recurrent computation

URL: View paper

Brief Assessment

Recurrent Occlusion Recognition[62] focuses on recurrent processing for occluded object recognition in visual systems, not on path ambiguity in compositional task structures or multiple computational paths preventing unified intermediate representations.

5. Addressing divergent representations from causal interventions on neural networks

URL: View paper

Brief Assessment

Causal Intervention Representations[58] focuses on divergent representations from causal interventions in neural networks, not on path ambiguity in compositional task structures. The candidate addresses intervention-induced distribution shifts rather than multiple computational paths preventing unified intermediate representations.

6. Maximal divergence sequential auto-encoder for binary software vulnerability detection

URL: View paper

Brief Assessment

Maximal Divergence Autoencoder[61] focuses on learning maximally divergent representations for binary software vulnerability detection, not on analyzing multiple computational paths or path ambiguity in compositional task structures.

7. Visualizing the Diversity of Representations Learned by Bayesian Neural Networks

URL: View paper

Brief Assessment

Bayesian Representation Diversity[63] focuses on visualizing and quantifying diversity in Bayesian neural network representations through feature visualizations and contrastive learning. It does not address path ambiguity or multiple computational paths preventing unified intermediate representations in compositional tasks.

8. Snr: Sub-network routing for flexible parameter sharing in multi-task learning

URL: View paper

Brief Assessment

Sub-network Routing[60] focuses on flexible parameter sharing in multi-task learning through modularized sub-networks, not on path ambiguity or multiple computational paths preventing unified intermediate representations in compositional tasks.

Appendix: Text Similarity Detection

Textual similarity detection checked 24 papers and found 3 similarity segment(s) across 1 paper(s).

The following **1 paper(s)** were detected to have high textual similarity with the original paper. These may represent different versions of the same work, duplicate submissions, or papers with substantial textual overlap. Readers are advised to verify these relationships independently.

1. The Coverage Principle: A Framework for Understanding Compositional Generalization

Detected in: Core Task (sibling), Contribution: contribution_2

⚠ **Note**: This paper shows substantial textual similarity with the original paper. It may be a different version, a duplicate submission, or contain significant overlapping content. Please review carefully to determine the nature of the relationship.

References

- [0] Characterizing Pattern Matching and Its Limits on Compositional Task Structures View paper
- [1] A neural mechanism for compositional generalization of structure in humans View paper
- [2] Local Mechanisms of Compositional Generalization in Conditional Diffusion View paper
- [3] Data Factors for Better Compositional Generalization View paper
- [4] Decoupled global-local alignment for improving compositional understanding View paper
- [5] CoScheme: Compositional Copatterns in Scheme View paper
- [6] LLM Retrieval-Augmented Generation with Compositional Prompts and Confidence Calibration View paper
- [7] Align and Augment: Generative Data Augmentation for Compositional Generalization View paper
- [8] Multimodal graph networks for compositional generalization in visual question answering View paper
- [9] NeSyCoCo: A Neuro-Symbolic Concept Composer for Compositional Generalization View paper
- [10] Swing-by Dynamics in Concept Learning and Compositional Generalization View paper
- [11] Latent surface alignment through oscillatory token drift in instruction-following large language models View paper
- [12] In-Context Compositional Generalization for Large Vision-Language Models View paper
- [13] Towards Compositional Generalization in Code Search View paper
- [14] Out-of-distribution generalization via composition: A lens through induction heads in Transformers View paper
- [15] UniMatch: Universal Matching from Atom to Task for Few-Shot Drug Discovery View paper
- [16] Improving compositional generalization in semantic parsing View paper
- [17] Exploring the Compositional Generalization in Context Dependent Text-to-SQL Parsing View paper
- [18] Improving compositional generalization with latent structure and data augmentation View paper
- [19] Do phd-level llms truly grasp elementary addition? probing rule learning vs. memorization in large language models View paper
- [20] Multi-Scale Explicit Matching and Mutual Subject Teacher Learning for Generalizable Person Re-Identification View paper
- [21] Impact of Noise on LLM-Models Performance in Abstraction and Reasoning Corpus (ARC) Tasks with Model Temperature Considerations View paper
- [22] HCCM: Hierarchical Cross-Granularity Contrastive and Matching Learning for Natural Language-Guided Drones View paper
- [23] Bootstrapped structural prompting for analogical reasoning in pretrained language models View paper
- [24] Structural generalization in COGS: Supertagging is (almost) all you need View paper
- [25] Provable Compositional Generalization for Object-Centric Learning View paper
- [26] Revisiting Depth Completion from a Stereo Matching Perspective for Cross-domain Generalization View paper
- [27] Compositional generalization with grounded language models View paper
- [28] Multi-Modal Transformer With Global-Local Alignment for Composed Query Image Retrieval View paper
- [29] Matching Plans for Frame Inference in Compositional Reasoning View paper
- [30] Finding needles in a haystack: Sampling Structurally-diverse Training Sets from Synthetic Data for Compositional Generalization View paper
- [31] Improving compositional generalization with self-training for data-to-text generation View paper
- [32] CoScheme: Compositional Copatterns View paper
- [33] DiffCloth: Diffusion Based Garment Synthesis and Manipulation via Structural Cross-modal Semantic Alignment View paper
- [34] Unlocking Compositional Generalization in Pre-trained Models Using Intermediate Representations View paper
- [35] The Coverage Principle: A Framework for Understanding Compositional Generalization View paper
- [36] SODA: Structural Pre-Decoupling and Co-Aligning for Video Compositional Representation View paper
- [37] A pattern for almost compositional functions View paper
- [38] Cross-modal Attention Congruence Regularization for Vision-Language Relation Alignment View paper
- [39] Systematic generalization in connectionist models View paper
- [40] Compositional generalisation with structured reordering and fertility layers View paper
- [41] Data Extraction via Semantic Regular Expression Synthesis View paper
- [42] Human re-identification by matching compositional template with cluster sampling View paper
- [43] Exploring Patterns of Algorithmic Generalization in Deep Learning View paper
- [44] Debunk the Myth of SFT Generalization View paper
- [45] Differentiable Compositional Kernel Learning for Gaussian Processes View paper
- [46] Compositional Generalization for Neural Semantic Parsing via Span-level Supervised Attention View paper

- [47] SymCHMâan unsupervised approach for pattern discovery in symbolic music with a compositional hierarchical model View paper
- [48] Exploring Compositional Generalization (in ReCOGS_pos) by Transformers using Restricted Access Sequence Processing (RASP) View paper
- [49] Understanding Subword Compositionality of Large Language Models View paper
- [50] Exploring Compositional Generalization (in COGS/ReCOGS_pos) by Transformers using Restricted Access Sequence Processing (RASP) View paper
- [51] Unveiling the Mechanisms of Explicit CoT Training: How CoT Enhances Reasoning Generalization View paper
- [52] Dice: Dynamic in-context example selection in llm agents via efficient knowledge transfer View paper
- [53] Language models can learn implicit multi-hop reasoning, but only if they have lots of training data View paper
- [54] Identity Bridge: Enabling Implicit Reasoning via Shared Latent Memory View paper
- [55] Learning to Prompt in Unknown Environments: A POMDP Framework with Compositional Actions for Large Language Models View paper
- [56] A survey of the recent architectures of deep convolutional neural networks View paper
- [57] Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning View paper
- [58] Addressing divergent representations from causal interventions on neural networks View paper
- [59] Diverse types of expertise in facial recognition View paper
- [60] Snr: Sub-network routing for flexible parameter sharing in multi-task learning View paper
- [61] Maximal divergence sequential auto-encoder for binary software vulnerability detection View paper
- [62] Object recognition under occlusion revisited: elucidating algorithmic advantages of recurrent computation View paper
- [63] Visualizing the Diversity of Representations Learned by Bayesian Neural Networks View paper
- [64] Not all solutions are created equal: An analytical dissociation of functional and representational similarity in deep linear neural networks View paper
- [65] How not to Stitch Representations to Measure Similarity: Task Loss Matching versus Direct Matching View paper
- [66] Family Cohesion Moderates the Relation between ParentâChild Neural Connectivity Pattern Similarity and Youth's Emotional Adjustment View paper
- [67] Cross-stage neural pattern similarity in the hippocampus predicts false memory derived from post-event inaccurate information View paper
- [68] Design of neuron-calculators for the normalized equivalence of two matrix arrays based on FPGA for self-learning equivalently convolutional neural networks (SLE_CNNs) View paper
- [69] Model Stitching by Functional Latent Alignment View paper
- [70] Model Stitching: Looking For Functional Similarity Between Representations View paper
- [71] Greater neural pattern similarity across repetitions is associated with better memory View paper
- [72] Multi-input Vision Transformer withÂ Similarity Matching View paper
- [73] BinDNN: Resilient Function Matching Using Deep Learning View paper