ALGORITHMIC PRIMITIVES AND COMPOSITIONAL GEOMETRY OF REASONING IN LANGUAGE MODELS

Anonymous authors

000

001

002003004

010 011

012

013

014

016

017

018

019

021

024

025

026

027

028

029

031

032

035

037

040

041

042

043

044

046

047

048

051

052

Paper under double-blind review

ABSTRACT

How do latent and inference time computations enable large language models (LLMs) to solve multi-step reasoning? We introduce a framework for tracing and steering algorithmic primitives that underlie model reasoning. Our approach links reasoning traces to internal activation patterns and evaluates algorithmic primitives by injecting them into residual streams and measuring their effect on reasoning steps and task performance. We consider four benchmarks: Traveling Salesperson Problem (TSP), 3SAT, AIME, and graph navigation. We operationalize primitives by clustering neural activations and labeling their matched reasoning traces. We then apply function vector methods to derive primitive vectors as reusable compositional building blocks of reasoning. Primitive vectors can be combined through addition, subtraction, and scalar operations, revealing a geometric logic in activation space. Cross-task and cross-model evaluations (Phi-4, Phi-4-Reasoning, Llama-3-8B) show both shared and task-specific primitives. Notably, comparing Phi-4 with its reasoning-finetuned variant highlights compositional generalization after finetuning: Phi-4-Reasoning exhibits more systematic use of verification and path-generation primitives. Injecting the associated primitive vectors in Phi-4-Base induces behavioral hallmarks associated with Phi-4-Reasoning. Together, these findings demonstrate that reasoning in LLMs may be supported by a compositional geometry of algorithmic primitives, that primitives transfer cross-task and cross-model, and that reasoning finetuning strengthens algorithmic generalization across domains.

1 Introduction

Inference time compute has remarkably improved reasoning in large language models (LLMs), and reasoning-finetuned models like Phi-4-Reasoning substantially outperform their base counterparts on reasoning tasks they were not directly trained on (Abdin et al., 2025). However, the extent to which LLMs, especially reasoning-finetuned models, can learn generalized algorithmic capacities remains poorly understood (Eberle et al., 2025). While recent advances in mechanistic interpretability (Todd et al., 2024) point in this direction, it's unclear whether LLMs acquire universal representations of algorithmic primitives (Huh et al., 2024), whether these primitives are geometrically organized in representation space, similar to brains (Fascianelli et al., 2024), and whether LLMs solve reasoning tasks by composition of generalized algorithmic primitives, and through component reuse across tasks and models (Merullo et al., 2024). This presents a unique opportunity to understand the algorithmic basis of LLM reasoning, improvements in compositional generalization after finetuning, and the impact of finetuning on chain-of-thought reasoning (Lobo et al., 2025).

This work aims to understand how fundamental algorithmic primitives are generalized and composed to enable complex reasoning capabilities in LLMs. The work addresses three fundamental questions: (1) What are the basic algorithmic primitives that language models use in specific tasks, and across reasoning domains? (2) Do these primitives compose geometrically in neural activation space? (3) How do reasoning-enhanced models differ from base models in their primitive usage and composition? We introduce a framework for a multi-level algorithmic understanding. We first extract algorithmic primitives by clustering internal representations and interpreting the corresponding reasoning traces. We then apply function vector methods to extract primitive vectors from the models' internal representations. We induce and steer algorithmic primitives by injecting primitive vectors across the layers. To identify both task-specific and universal algorithmic primitives, we

apply the approach across domains: Traveling Salesperson Problem, 3-SAT, AIME, and graph navigation. Finally, to investigate compositional generalization as a result of finetuning, we compare algorithmic performance of base models and reasoning-finetuned counterpart, like Phi-4 and Phi4-reasoning (Abdin et al., 2025). This multi-level framework allows us to understand the geometry and compositionality of algorithmic primitives in LLM reasoning.

Our contributions include: (1) the systematic identification of cross-domain algorithmic primitives in LLMs, (2) a geometric framework for understanding primitive composition through vector arithmetic, (3) novel methodology linking explicit reasoning behaviors to internal mechanisms, (4) mechanistic evidence for compositional generalization following finetuning LLMs for reasoning, and (5) evaluation of cross-task primitive transfer and generalization.

2 Related Work

Interpretability for Transformers. With increasing scale and complexity of LLMs, the challenge of understanding their predictions has become an important research direction in explainable AI and interpretability research. This has specifically targeted the analysis of internal model structure and feature relationships (Geiger et al., 2022; Eberle et al., 2022; Schnake et al., 2022), as well as representations and manifolds (Kornblith et al., 2019) including concepts (Chormai et al., 2024). Mechanistic interpretability (Sharkey et al., 2025) has focused on the extraction of circuits (Olah et al., 2020; Wang et al., 2023), feature descriptions (Hernandez et al., 2022), and causally effective representations such as function vectors (Todd et al., 2024). Combined analysis of individual neurons (Gurnee & Tegmark, 2024; Templeton et al., 2024), attention scores (Voita et al., 2019; Clark et al., 2019), and task-specific attention heads (Vig & Belinkov, 2019; McDougall et al., 2024) have further deepened our understanding of internal model processing. In parallel, free-text and chain-of-thought explanations (Turpin et al., 2023; Camburu et al., 2018; Huang et al., 2023; Madsen et al., 2024) consider the model's own natural language explanation of what it is doing and why. Gradient-based feature attributions have further enabled to scale the localization and analysis of relevant features in LLMs (Ali et al., 2022; Jafari et al., 2024).

Function Vectors and Behavioral Interventions. LLMs' contextualized representations such as function vectors (Todd et al., 2024) and in-context task vectors (Hendel et al., 2023; Yang et al., 2025) are shown to trigger execution of associated tasks. A broader application of such steering vectors include persona vectors (Chen et al., 2025) and representations of truthfulness (Marks & Tegmark, 2024; Bürger et al., 2024). Probing has identified associations between internal representations and features of interest (Conneau et al., 2018; Hewitt & Manning, 2019). This has also been extended to the analysis of reasoning strategies, e.g., by identifying probes involved in solving grade-school math problems (Ye et al., 2025). These analyses provide some first methods to understand how complex representational geometries and manifolds drive predictions in modern models. Related work from neuroscience shows that task-related variables are encoded in neural geometries in a format that supports generalization to novel situations (Bernardi et al., 2020). Monkey neuroscience research has linked compositional generalization in neural representational geometry with behavioral performance (Fascianelli et al., 2024). These findings suggest that function vector methods can be used to study how compositional generalization may guide algorithmic steering of LLMs, providing a conceptual starting point for our work.

Algorithmic Evaluation and Steering of Language Models. Multi-step planning and graph navigation are standard benchmarks for structured reasoning (Fatemi et al., 2023), yet most LLMs perform poorly and fail to generalize with growing graph complexity (Momennejad et al., 2023). Although efficient algorithms exist, recent evaluations suggest that LLMs rely on policy-dependent heuristics rather than explicit search strategies (Eberle et al., 2025). Complementary works have aimed to map a high-level causal model to an internal realization by an LLMs (Geiger et al., 2024), and prompt models to generate sets of abstract hypotheses about the tasks to improve inductive reasoning performance (Wang et al., 2024). Recent analysis of chain-of-thought outputs has uncovered sentence-level relationships, offering the localization of salient reasoning steps (Bogdan et al., 2025). Existing work either targets high-level behaviors such as politeness, creativity, or honesty using relatively coarse vector representations to steer behavior (Chen et al., 2025), or focuses on very simple input-output mapping (Todd et al., 2024). Our work focuses on algorithmic steering, targeting computational primitives and their compositions, aiming at a deep understanding of the model's internal algorithmic phenotypes and applications in future algorithmic finetuning.

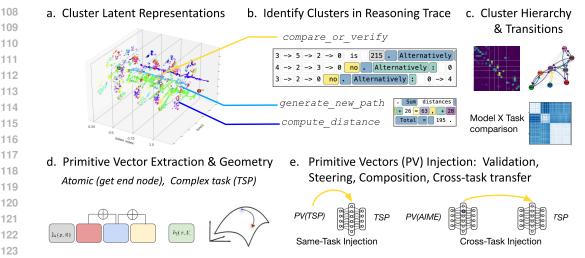


Figure 1: Algorithmic Tracing & Steering: Primitive Extraction & Evaluation. We trace algorithmic primitives by $\bf a$ clustering latent representations, $\bf b$ identifying corresponding reasoning traces, $\bf c$ meta-clustering primitives, identifying sequential transition trends among clusters, and comparing cluster similarity across models and tasks. Once we identify primitives, $\bf d$ we extract associated primitive vectors from top heads, and $\bf e$ use causal patching to validate, explore the compositional geometry and cross-task transfer of primitives.

3 EXPERIMENTAL SETUP: MODELS AND TASKS

Models. We evaluated decoder-only transformer models on multi-step reasoning tasks. Our analyses primarily focused on Phi-4-Base (Abdin et al., 2024) and the reasoning-specialized variant Phi-4-Reasoning (Abdin et al., 2025), with additional validation from Llama-3-8B (Dubey et al., 2024).

Tasks and Data Collection. Four multi-step reasoning setups were tested. We primarily focused on the Traveling Salesperson Problem (TSP), an NP-hard benchmark. Notably, Phi-4-Reasoning exhibited improved performance on TSP despite not being finetuned for this task (Abdin et al., 2025). We also investigated the 3-SAT and AIME Mathematical Olympiad benchmarks for the same reasons. Finally, we investigated Graph Navigation tasks used in previous LLM reasoning research (Momennejad et al., 2023; Eberle et al., 2025) (see Appendix B for task prompts). Our setups demand planning, search, and verification. TSP and Graph navigation require graph optimization via planning, path generation, search, comparison, and verification. 3SAT presents another NP-hard problem that tests flexible algorithmic problem solving. Finally, AIME requires complex mathematical reasoning and multi-step computation. Together, these setups enable us to examine algorithmic primitives that are task-specific and those that generalize across tasks. We collected reasoning traces and associated residual stream vectors across 100+ examples per task, with systematic variation in problem complexity.

4 METHODOLOGY

4.1 DEFINITIONS AND NOTATION

Definition: Algorithmic Primitive. We define *algorithmic primitive* as a minimal computational operation observed in a reasoning process (Eberle et al., 2025) (e.g. TSP), such as retrieving the nearest neighbor, computing a distance, generating a new candidate path, or verifying a solution. Primitives can be identified both in explicit reasoning traces (e.g. reasoning steps the model produces in the output) and in internal activations (e.g. clusters of token representations, or attention patterns).

Definition: Algorithmic Tracing. We define *algorithmic tracing* as the process of identifying all relevant primitives for implementing a particular reasoning process. By targeting the computational building blocks rather than just behavioral outcomes, algorithmic tracing helps us better understand how and why a model solves a particular task and what other tasks it might be able to solve.

Definition: Primitive Vector. We define primitive vectors (following the function and steering vector literature (Todd et al., 2024)) as vectors that can be injected into the residual stream to reliably induce a particular primitive. For a given primitive p and transformer layer ℓ , let the residual stream

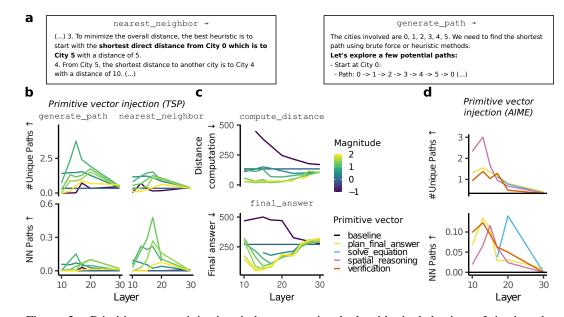


Figure 2: Primitive vector injection induces associated algorithmic behavior. Injecting the nearest_neighbor or generate_path primitive vectors directly after the prompt **a** increases expression in output (see examples). **b** Varying the injection layer and the magnitude modifies the number of unique paths generated (top row) and the proportion of nearest-neighbor paths (bottom row). **c** Injecting primitive vectors for compute_distance in the middle of the reasoning trace increases relevant behavioral hallmarks in the output. **d** Primitive vectors from AIME were injected to different layers while solving Traveling Salesperson (TSP), showing cross-task primitive transfer and algorithmic induction.

activation at token t be $h_{\ell}(x,t) \in \mathbb{R}^{d_{\ell}}$. A primitive vector $v_{\ell}^{(p)} \in \mathbb{R}^{d_{\ell}}$ is a direction in activation space that increases the expression of the primitive function p. Here we extract primitive vectors using the function vector approach (Todd et al., 2024) (see Section 4.2 for details).

Definition: Primitive Induction for Algorithmic Steering. Given a primitive vector injection strength $\alpha \in \mathbb{R}$, we can intervene on a representation by adding (or patching) a primitive vector:

$$\tilde{h}_{\ell}(x,t) = h_{\ell}(x,t) + \alpha v_{\ell}^{(p)},$$

which increases the probability of expressing p when $\alpha > 0$, and decreases it when $\alpha < 0$. We refer to this procedure as algorithmic steering or induction.

Algebraic Operations on Primitive Vectors. Assuming compositional geometry, primitive vectors can be combined through simple algebraic operations in activation space. For primitives p and q at layer ℓ , additive and subtractive composition within a layer (+ and -) can be defined as:

$$v_{\ell}^{(p \oplus q)} \approx w_p \, v_{\ell}^{(p)} \pm w_q \, v_{\ell}^{(q)}.$$

Scalar modulation (with varying strength) can be defined as:

$$v_{\ell}^{(\alpha p)} = \alpha \, v_{\ell}^{(p)}.$$

Primitive Transfer: Cross-Task Generalization. We test primitive generalizability by examining whether primitives identified in one domain transfer to others; specifically by examining whether injecting primitives extracted from task 1 has a predictable effect on model performance in task 2. We formalize cross-task transfer as follows. Let $p \in \mathcal{P}_{T_1}$ denote a primitive extracted from source task T_1 , with vector $v_{\ell,T_1}^{(p)} \in \mathbb{R}^{d_\ell}$. When evaluating target task T_2 , we inject this vector into the residual stream:

$$\tilde{h}_{\ell}(x_{T_2}, t) = h_{\ell}(x_{T_2}, t) + \alpha v_{\ell, T_1}^{(p)}.$$

If this intervention increases the activation

$$a_{\ell,T_2}^{(p)}(x_{T_2},t) = \langle v_{\ell,T_1}^{(p)}, h_{\ell}(x_{T_2},t) \rangle$$

and increases the associated behavioral hallmark in task T_2 , we denote successful transfer as:

$$p \in \mathcal{P}_{T_1} \rightsquigarrow p \in \mathcal{P}_{T_2}$$
.

4.2 ALGORITHMIC TRACING & STEERING: PRIMITIVE EXTRACTION & EVALUATION

Step 1. Primitive Identification: Geometric Clustering of Latent Representations. We hypothesize that distinct algorithmic primitives are reflected in higher dissimilarities in the model's internal representations, and therefore apply k-means clustering (Lloyd, 1982) to the model's internal representations while processing complex self-generated reasoning traces. For all clustering analyses, we extracted representations from layer 17 of Phi-4-base and used k = 50 clusters. We fit separate models to TSP, 3-SAT, AIME, and GraphNav and also fit a joint model to TSP+AIME.

Step 2. Primitive Identification: Mapping Clusters to Associated Reasoning Traces. We analyze the primitives associated with the different clusters revealed in step 1 by analyzing the tokens associated with a particular cluster and the context surrounding them. This lets us categorize algorithmic strategies, verification behaviors, and metacognitive patterns and provides direct insight into the reasoning processes models employ.

Step 3. Primitive Composition: Hierarchical Meta-Clustering and Temporal Clustering. To identify prevalent temporal compositions of these clusters, we then apply spectral clustering (Von Luxburg, 2007) to the matrix of transition probabilities between clusters. This gives rise to a smaller set of meta-cluster which highlight hierarchically structured reasoning steps. To infer the number of meta-clusters, we identify the largest spectral gap, using a minimum of four meta-clusters.

Step 4. Primitive Validation: Primitive Vector Extraction. Finally, we extract primitive vectors associated with particular clusters by adapting the methodology put forward in Todd et al. (2024). We compute the average attention head activations over all tokens that are a part of a particular cluster (using 100 responses) and then average the outputs of the top attention heads that reliably carry out a function (k=35; see appendix). This yields a candidate primitive vector for each extracted cluster. We validate candidate primitives by defining behavioral hallmarks associated with their proposed computational role and testing whether injecting the extracted primitive vectors results in an increase in those behavioral hallmarks. Beyond primitive validation, we also examine their compositional generalization by testing their arithmetic composition and cross-task transfer, injecting primitive vectors extracted from AIME into the model while it performs TSP.

Algorithmic Fingerprinting. For a given set of clusters $1, \ldots, k$ and a response i, we extract the relative frequency of tokens assigned to each cluster, $f_i \in \mathbb{R}^k$, $\sum_i f_i = 1$. We consider f_i as a simple "algorithmic fingerprint" of a particular reasoning trace, highlighting differences between the primitives involved in different responses. In particular, we analyze the algorithmic dissimilarity between two responses by computing the symmetric χ -squared distance between their frequencies,

$$\chi(f,g) := \sum_{i=1}^{k} \frac{(f_i - g_i)^2}{f_i + g_i}.$$
 (1)

We also consider sets of responses $(f^{(j)})_{j=1}^n$, $(g^{(j)})_{j=1}^m$. We analyze differences in the involved algorithmic primitives by computing the average frequencies $f = \frac{1}{n} \sum_{j=1}^n f^{(j)}$, $g = \frac{1}{n} \sum_{j=1}^n g^{(j)}$ and then computing the signed χ -squared distance

$$\chi_s(f_i, g_i) := \operatorname{sgn}(f_i - g_i) \circ \frac{(f_i - g_i)^2}{f_i + g_i} \in \mathbb{R}^k, \tag{2}$$

for each cluster i. A large positive difference indicates that the corresponding primitive is much more prominent in the responses $(f^{(j)})_{j=1}^m$, a large negative difference indicates that the corresponding primitive is much more prominent in the responses $(g^{(j)})_{j=1}^m$. In particular, we compare Phi-4 responses to Phi-4-Reasoning responses and TSP responses to AIME responses.

5 RESULTS

Primitive Tracing and Steering. We first traced primitives by clustering latent representations over entire reasoning traces, interpreting these clusters by identifying the corresponding tokens in the model output. Perhaps surprisingly, we found that despite the simplicity of our approach, many of these clusters were highly interpretable; for example, we discovered a cluster specifically associated with the implementation of a nearest-neighbor heuristic, and another cluster associated with validations, checks, and corrections.

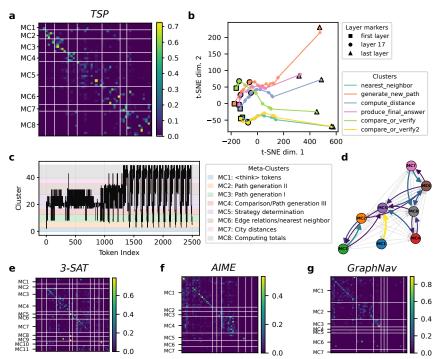


Figure 3: Clusters are organized in a hierarchical manner and meta-clusters are repeatedly traversed throughout the response. **a** Cluster-cluster transition matrix on TSP, organized by meta-cluster (MC). Most transitions occur within each meta-cluster. **b** Average t-SNE trajectories across different highlighted primitives. **c** Transition between different clusters for an example output. Different meta-clusters are highlighted by colors. The latter half of the response undergoes a cyclical transition. **d** Most common transitions between different meta-clusters reveals frequently occurring cycles. **e-g** Cluster-to-cluster transition matrices structured in terms of the inferred meta-clusters on **e** 3-SAT, **f** AIME, and **g** GraphNav.

We then determined the average transition probability between the different clusters. Spectral clustering of this transition matrix revealed that, on TSP, clusters are sequentially composed in a hierarchical and highly stereotyped manner, as most transitions between clusters occcurred within the same meta-cluster (Fig. 3a). Notably, these different meta-clusters were also responsible for highly interpretable parts of the reasoning, e.g. path generation or computing distances (Fig. 3c).

We then validated the role of these primitives by extracting the corresponding primitive vectors and injecting them into Phi-4 during a response generation to TSP. First, we injected two vectors (nearest_neighbor and generate_path) into the model throughout the response generation, exploring different layers and magnitudes. We found that nearest_neighbor selectively increased the proportion of nearest-neighbor paths generated by the model, whereas generate_path increased the number of total generated paths (Fig. 2a,b). This demonstrates that the extracted primitive vectors selectively induce a particular behavior and validates the candidate primitives generated by our clustering.

To expand this investigation, we injected a broader range of candidate primitive vectors into the model after providing example paths and distance computations. For example, injecting the compute_distance primitive causes the model to more quickly compute the distance of a candidate path. Interestingly, the reverse is also true: subtracting the compute_distance vector makes the model *less* likely to implement a distance computation (Fig. 2c). This suggests that subtracting primitive vectors can prevent associated algorithmic primitives. More broadly, we evaluate six behavioral hallmarks across six possible function vectors and find that all behavioral hallmarks are most strongly induced by a candidate primitive vector with a relevant computational role (Table 1).

Notably, the role of a particular cluster cannot always be inferred from the tokens on which it is active alone; the surrounding context often also plays a role. For example, the cluster compare_and_verify is largely active on the token corresponding to the final distance of a can-

didate path. However, we noticed that tokens on which this cluster was active were often followed by subsequent checks and comparisons. We therefore hypothesized that this cluster may not just represent the final distance, but also induce a verification primitive. Table 1 confirms this hypothesis. Interestingly, we can observe this dual nature in representational space as well. In a t-SNE plot of the average representational trajectory of these example primitives, we observe that in earlier layers, compare_or_verify evolves along other path-related primitives like generate_new_path, whereas it moves closer to compare_or_verify2 in later layers (Fig. 3b).

	%NN paths↑	#Paths ↑	Dist. comp. ↓	Final ans. ↓	#Verif. ↑	#Comp. ↑
nearest_neighbor	+56.1%	+72.0%	-73.6%	-4.5%	+104.3%	+125.6%
generate_path	+4.8%	+143.9%	<u>-76.0%</u>	+18.4%	+25.0%	+12.2%
compute_distance	+22.2%	+36.5%	-86.5%	<u>-47.4%</u>	+198.9%	+79.3%
final_answer	+23.2%	+34.8%	-29.8%	-81.5%	+82.6%	-37.8%
compare_verify	<u>+37.9%</u>	+43.6%	-62.8%	-7.0%	+655.4%	+1103.7%
compare_verify2	+29.5%	+48.4%	-64.9%	-13.7%	+706.5%	<u>+526.8%</u>

Table 1: Effects of primitive vector injection on behavioral hallmarks (% above baseline). For each cell, we identify the maximal effect across all positive magnitudes and all intervention layers (10, 13, 15, 17, 20, 30). Bold = strongest effect, underline = second strongest per column. (%NN paths: proportion of nearest neighbor paths generated. Dist. comp.: Distance computation. Final ans.: Final answer. #Verif.: number of Verifications in the output. #Comps: number of comparisons in the output. See Appendix D for detailed definitions.)

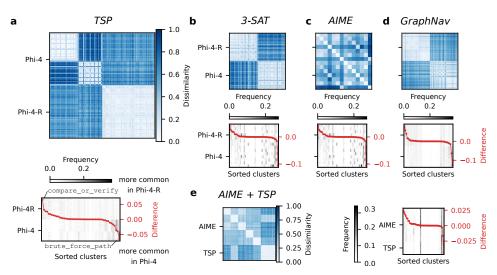


Figure 4: Primitive Cluster Patterns in Phi-4 and Phi-4-Reasoning. **a-d** *top* Normalized dissimilarity of primitive cluster frequencies between Phi-4 and Phi-4-Reasoning for a) TSP, b) 3-SAT, c) AIME, and d) GraphNav. *bottom* Clusters sorted by whether they appear more frequently in Phi-4-Reasoning responses (positive difference) or in Phi-4 responses (negative difference). The lineplot specifies the differences whereas the rasterplot underneath specifies the relative frequencies of the different clusters per response. **e** *left* Dissimilarity of primitive cluster frequencies between Phi-4-Reasoning responses to AIME and TSP. *right* Clusters sorted by whether they occur more frequently in AIME or TSP. The rasterplot underneath again specifies the relative frequencies of the different clusters per response.

Comparing Algorithmic Primitives between Phi-4 and Phi-4-Reasoning. To identify shared and distinct primitives between responses by Phi-4 and Phi-4-Reasoning, we computed the dissimilarity between the primitive frequencies involved in different responses. We fond that responses by Phi-4-Reasoning are highly stereotyped (Fig. 4a). In contrast, Phi-4 has two subgroups of distinct responses. Notably, these groups map onto two distinct approaches towards solving TSP: a brute-force search and a random guess without further refinement. This highlights that algorithmic fingerprinting reveals both differences between different models, but even differences within responses from the same model. Beyond TSP, we found that the responses from Phi-4 and Phi-4-Reasoning on

3-SAT and GraphNav are highly stereotyped (Fig. 4b,d). On AIME, algorithmic primitives involved in responses to the same question were very similar to each other, whereas primitives involved in responses to different questions were substantially more different, reflecting the higher diversity of questions on this benchmark (Fig. 4c).

Next, we analyzed which primitives are more common in Phi-4 or Phi-4-Reasoning by computing the signed χ -squared difference, and analyzing the clusters with the largest positive and negative values. This allows us to identify important differences in responses strategies across the two model. For example, on TSP and 3-SAT, distinct heuristic approaches arise more commonly in Phi-4-Reasoning than Phi-4 (nearest_neighbor and if_clause_true, Appendix A). Beyond task-specific strategies, our comparative analysis also highlights broader differences between the two models: in particular, clusters that are more common in Phi-4-Reasoning relate to general-purpose reasoning steps such as verification, recalling instructions, or planning the final answer.

Compositional Primitive Induction. So far, we have considered sequential compositions of algorithmic primitives. To investigate arithmetic compositions of primitives, we consider a set of algorithmic in-context learning tasks that require identifying 1) the last node of a path ("Terminal node recognition," TNR), 2) the node with the higher reward ("Reward comparison," RC); and 3) the most highly rewarded node between two paths, a composition of TNR and RC (Appendix C.1). Remarkably, we find that adding together the primitive vectors for TNR and RC induces this composite behavior in Llama-3-8B, causing it to match few-shot performance in a zero-shot setting (Fig. 5a). These effects were more attenuated in Phi-4 and Phi-4-Reasoning, highlighting that different models may require different modes of composition (Fig. 8).

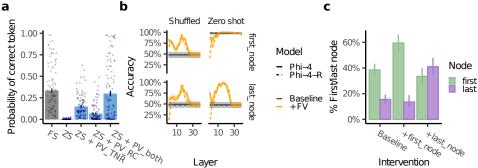


Figure 5: Algorithmic primitives operating over graphs, compositional induction, and cross-task transfer. **a** Injecting the sum of a terminal node recognition primitive vector (PV_TNR) and a reward comparison primitive vector (PV_RC) in a zero-shot setting (ZS) improves model performance on a task requiring a composition of both primitives and recovers few-shot performance (FS). See Fig. 8 for additional models. **b** Primitive vectors for extracting the first or last node of a presented graph reliably improves performance across different injection layers and across a shuffled few-shot and a zero shot-setting. **c** After injecting these primitive vectors into complex reasoning traces solving TSP, the model becomes more likely to mention the corresponding node next.

Cross-task Transfer and Compositional Generalization. Generalized algorithmic primitives which can be applied across different tasks, can help us understand why reasoning-finetuned models like Phi-4-Reasoning show improvements on tasks they were not finetuned on. To better understand the degree to which algorithmic primitives in Phi-4-Reasoning are shared between tasks, we apply our algorithmic tracing framework to a set of responses from both TSP and AIME (Fig. 4e). This analysis highlights that while there are different algorithmic primitives at play in either task, the two tasks also have many algorithmic primitives in common.

Next, we tested whether cross-task injection of algorithmic primitives would induce behavioral changes. First, we extracted primitive vectors from an algorithmic in-context learning task that required extracting a path's first or last node (Fig. 5b). We then injected these primitive vectors into complex reasoning traces solving TSP immediately after the model had mentioned a particular path. Despite the extremely different domain, injecting these vectors indeed had the expected effect: injecting the first_node (resp. last_node) primitive vector caused the model to subsequently mention the first node (resp. last node) of the path more often.

Finally, we considered relevant primitive vectors from AIME (e.g. spatial_reasoning, plan_final_answer) and injected them into Phi-4 generating a response to a TSP prompt. We found that these primitive vectors caused the model to generate more paths and a higher proportion of nearest-neighbor paths — a hallmark of Phi-4-Reasoning responses.

6 DISCUSSION

Algorithmic Tracing and Steering. We introduce a framework for tracing and steering algorithmic primitives as building blocks of LLM reasoning with geometric compositionality. We adapted function vector techniques (Todd et al., 2024) to extract minimal, 1-step primitive vectors from carefully designed simple tasks as well as benchmarks with more complex reasoning contexts. When "injected" to LLM layers after the prompt or in the middle of reasoning traces, these primitive vectors increase the behavioral expression of corresponding reasoning steps (Figs. 2 and 5). This extends function vector research beyond simple relations between the input and output. A central contribution of our work is the compositional exploration of primitive vector arithmetic. Our framework bridges a principled path from laboratory-identified primitives to real-world reasoning behaviors.

Toward a Geometry of Compositional Abstraction in LLMs. We find that algorithmic primitives exhibit geometric regularities through compositional operations like addition, subtraction, multiplication, and scalar modulation. The layer and magnitude of injection shape the output expression of the primitive (Fig. 2). We show cross-task transfer and generalizability of primitive vectors: spatial reasoning and verification primitives extracted from AIME successfully transfer to TSP. This transferability hints at potentially universal algorithmic building blocks underlying diverse reasoning capabilities.

Reasoning vs. Memorization. Our work contributes to recent discussions about the extent to which LLMs exhibit abstract reasoning rather than relying on memorized procedures, failing to generalize in counterfactual scenarios (Wu et al., 2024; Power et al., 2022; Zhang et al., 2024; Wang et al., 2025). While models were not directly trained on our tasks (Abdin et al., 2025), it is possible that it may be using amortized reasoning based on similar examples in the training. LLMs exhibit better problem-solving in high probability than low probability settings (e.g., McCoy et al. (2024a)), which relies on statistical regularities rather than abstract reasoning. Moreover, the absence of metacognitive abilities in LLMs results in brittle problem-solving (Johnson et al., 2024; Lewis & Mitchell, 2024). However, models finetuned for reasoning exhibit less sensitivity to task probability than base models (McCoy et al. (2024b)) and improved metacognitive-like uncertainty management strategies (e.g., verifying whether a candidate solution is correct) (Guo et al., 2025; Gandhi et al., 2025) that in turn causally improve reasoning (Bogdan et al., 2025). In line with these findings, our analysis identified algorithmic primitives for managing uncertainty (e.g., verification) and in-context recall. Such behaviors may underlie advanced problem-solving in reasoning models and suggest how LLMs could develop human-like abstract reasoning.

Limitations. One limitation is that the primitives, corresponding clusters, and meta-clusters we identify do not always map to a known algorithm. Future work can extract primitive libraries and algorithmic logic of a given model, identifying commonalities and differences across models. Here, we focus on reasoning as multi-step problem-solving in limited tasks, which may not capture the complexity and multimodal nature of natural reasoning, studied in cognitive sciences for decades Tversky (2005); Shepard & Metzler (1971); MacGregor & Ormerod (1996). Finally, we have focused on linear compositions of primitive vectors, leaving more complex interactions and potentially non-linear combinations of several vectors on manifolds for future work.

Future Directions. The algorithmic tracing and steering framework can be applied to any architecture (e.g., vision, diffusion, and multi-modal models) and domain beyond reasoning. An immediate extension is more detailed manifold analysis in order to capture the compositional geometry of primitives beyond linear composition, and establish universal and task-specific primitive libraries. Another key direction is the **algorithmic training and finetuning of LLMs**, with algorithmic objectives. Moreover, future self-improving models can be designed for **algorithmic self-play**: generating and evaluating compositional algorithmic solutions. Finally, collecting human reasoning data on the same tasks enables us to compare and finetune model-human **algorithmic alignment**.

Conclusion. The identification of universal primitives and their compositional geometry opens new avenues for interpretability research and suggests principled approaches for computational models of human reasoning, model-human alignment, and enhancing LLM reasoning capabilities.

7 RECOMMENDED ADDITIONS COPIED FROM GUIDELINES

7.1 LLM USAGE DISCLOSURE

We used LLMs for generating code and for finding related work.

7.2 REPRODUCIBILITY STATEMENT

We provide a detailed description of our methodological implementation and will share our codebase upon acceptance.

REFERENCES

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*, 2024.
- Marah Abdin, Sahaj Agarwal, Ahmed Awadallah, Vidhisha Balachandran, Harkirat Behl, Lingjiao Chen, Gustavo de Rosa, Suriya Gunasekar, Mojan Javaheripi, Neel Joshi, Piero Kauffmann, Yash Lara, Caio César Teodoro Mendes, Arindam Mitra, Besmira Nushi, Dimitris Papailiopoulos, Olli Saarikivi, Shital Shah, Vaishnavi Shrivastava, Vibhav Vineet, Yue Wu, Safoora Yousefi, and Guoqing Zheng. Phi-4-reasoning technical report, 2025. URL https://arxiv.org/abs/2504.21318.
- Ameen Ali, Thomas Schnake, Oliver Eberle, Grégoire Montavon, Klaus-Robert Müller, and Lior Wolf. XAI for transformers: Better explanations through conservative propagation. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 435–451. PMLR, 2022. URL https://proceedings.mlr.press/v162/ali22a.html.
- Vidhisha Balachandran, Jingya Chen, Lingjiao Chen, Shivam Garg, Neel Joshi, Yash Lara, John Langford, Besmira Nushi, Vibhav Vineet, Yue Wu, et al. Inference-time scaling for complex tasks: Where we stand and what lies ahead. *arXiv preprint arXiv:2504.00294*, 2025.
- Silvia Bernardi, Marcus K Benna, Mattia Rigotti, Jérôme Munuera, Stefano Fusi, and C Daniel Salzman. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell*, 183(4): 954–967, 2020.
- Paul C. Bogdan, Uzay Macar, Neel Nanda, and Arthur Conmy. Thought anchors: Which Ilm reasoning steps matter?, 2025. URL https://arxiv.org/abs/2506.19143.
- Lennart Bürger, Fred A. Hamprecht, and Boaz Nadler. Truth is universal: Robust detection of lies in llms, 2024. URL https://arxiv.org/abs/2407.12831.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. esnli: Natural language inference with natural language explanations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/4c7a167bb329bd92580a99ce422d6fa6-Paper.pdf.
- Runjin Chen, Andy Arditi, Henry Sleight, Owain Evans, and Jack Lindsey. Persona vectors: Monitoring and controlling character traits in language models, 2025. URL https://arxiv.org/abs/2507.21509.
- Pattarawat Chormai, Jan Herrmann, Klaus-Robert Müller, and Grégoire Montavon. Disentangled explanations of neural network predictions by finding relevant subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. What does BERT look at? an analysis of BERT's attention. In Tal Linzen, Grzegorz Chrupała, Yonatan Belinkov, and Dieuwke Hupkes (eds.), *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 276–286, Florence, Italy, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4828. URL https://aclanthology.org/W19-4828/.
 - Alexis Conneau, German Kruszewski, Guillaume Lample, Loïc Barrault, and Marco Baroni. What you can cram into a single \$&!#* vector: Probing sentence embeddings for linguistic properties. In Iryna Gurevych and Yusuke Miyao (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2126–2136, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-1198. URL https://aclanthology.org/P18-1198/.
 - Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pp. arXiv–2407, 2024.
 - Oliver Eberle, Jochen Büttner, Florian Kräutli, Klaus-Robert Müller, Matteo Valleriani, and Grégoire Montavon. Building and interpreting deep similarity models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1149–1161, 2022. doi: 10.1109/TPAMI.2020. 3020738.
 - Oliver Eberle, Thomas Austin McGee, Hamza Giaffar, Taylor Whittington Webb, and Ida Momennejad. Position: We need an algorithmic understanding of generative AI. In *Forty-second International Conference on Machine Learning Position Paper Track*, 2025. URL https://openreview.net/forum?id=eax2ixyeQL.
 - Valeria Fascianelli, Aldo Battista, Fabio Stefanini, Satoshi Tsujimoto, Aldo Genovesio, and Stefano Fusi. Neural representational geometries reflect behavioral differences in monkeys and recurrent neural networks. *Nature Communications*, 15(6479), 2024. doi: 10.1038/s41467-024-50503-w. URL https://doi.org/10.1038/s41467-024-50503-w.
 - Bahare Fatemi, Jonathan Halcrow, and Bryan Perozzi. Talk like a graph: Encoding graphs for large language models, 2023. URL https://arxiv.org/abs/2310.04560.
 - Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, Nathan Lile, and Noah D Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *arXiv* preprint arXiv:2503.01307, 2025.
 - Atticus Geiger, Zhengxuan Wu, Hanson Lu, Josh Rozner, Elisa Kreiss, Thomas Icard, Noah Goodman, and Christopher Potts. Inducing causal structure for interpretable neural networks. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pp. 7324–7338. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/geiger22a.html.
 - Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah Goodman. Finding alignments between interpretable causal variables and distributed neural representations. In *Causal Learning and Reasoning*, pp. 160–187. PMLR, 2024.
 - Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
 - Wes Gurnee and Max Tegmark. Language models represent space and time, 2024. URL https://arxiv.org/abs/2310.02207.
 - Roee Hendel, Mor Geva, and Amir Globerson. In-context learning creates task vectors. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 9318–9333, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.624. URL https://aclanthology.org/2023.findings-emnlp.624/.

- Evan Hernandez, Sarah Schwettmann, David Bau, Teona Bagashvili, Antonio Torralba, and Jacob Andreas. Natural language descriptions of deep features. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=NudBMY-tzDr.
- John Hewitt and Christopher D. Manning. A structural probe for finding syntax in word representations. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4129–4138, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1419. URL https://aclanthology.org/N19-1419/.
- Shiyuan Huang, Siddarth Mamidanna, Shreedhar Jangam, Yilun Zhou, and Leilani H. Gilpin. Can large language models explain themselves? a study of llm-generated self-explanations, 2023. URL https://arxiv.org/abs/2310.11207.
- Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. Position: The platonic representation hypothesis. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 20617–20642. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/huh24a.html.
- Farnoush Rezaei Jafari, Grégoire Montavon, Klaus-Robert Müller, and Oliver Eberle. Mambalrp: Explaining selective state space sequence models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- Samuel GB Johnson, Amir-Hossein Karimi, Yoshua Bengio, Nick Chater, Tobias Gerstenberg, Kate Larson, Sydney Levine, Melanie Mitchell, Iyad Rahwan, Bernhard Schölkopf, et al. Imagining and building wise machines: The centrality of ai metacognition. *arXiv preprint arXiv:2411.02478*, 2024.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pp. 3519–3529. PMLR, 2019.
- Martha Lewis and Melanie Mitchell. Evaluating the robustness of analogical reasoning in large language models. *arXiv preprint arXiv:2411.14215*, 2024.
- S. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2): 129–137, 1982. doi: 10.1109/TIT.1982.1056489.
- Elita Lobo, Chirag Agarwal, and Himabindu Lakkaraju. On the impact of fine-tuning on chain-of-thought reasoning. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 11679–11698, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-189-6. doi: 10.18653/v1/2025.naacl-long.584. URL https://aclanthology.org/2025.naacl-long.584/.
- James N MacGregor and Tom Ormerod. Human performance on the traveling salesman problem. *Perception & psychophysics*, 58(4):527–539, 1996.
- Andreas Madsen, Sarath Chandar, and Siva Reddy. Are self-explanations from large language models faithful?, 2024. URL https://arxiv.org/abs/2401.07927.
- Samuel Marks and Max Tegmark. The geometry of truth: Emergent linear structure in large language model representations of true/false datasets. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=aajyHYjjsk.
- R Thomas McCoy, Shunyu Yao, Dan Friedman, Mathew D Hardy, and Thomas L Griffiths. Embers of autoregression show how large language models are shaped by the problem they are trained to solve. *Proceedings of the National Academy of Sciences*, 121(41):e2322420121, 2024a.

- R Thomas McCoy, Shunyu Yao, Dan Friedman, Mathew D Hardy, and Thomas L Griffiths. When a language model is optimized for reasoning, does it still show embers of autoregression? an analysis of openai o1. *arXiv* preprint arXiv:2410.01792, 2024b.
- Callum Stuart McDougall, Arthur Conmy, Cody Rushing, Thomas McGrath, and Neel Nanda. Copy suppression: Comprehensively understanding a motif in language model attention heads. In Yonatan Belinkov, Najoung Kim, Jaap Jumelet, Hosein Mohebbi, Aaron Mueller, and Hanjie Chen (eds.), *Proceedings of the 7th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pp. 337–363, Miami, Florida, US, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.blackboxnlp-1.22. URL https://aclanthology.org/2024.blackboxnlp-1.22/.
- Jack Merullo, Carsten Eickhoff, and Ellie Pavlick. Circuit component reuse across tasks in transformer language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=fpoAYV6Wsk.
- Ida Momennejad, Hosein Hasanbeig, Felipe Vieira, Hiteshi Sharma, Robert Osazuwa Ness, Nebojsa Jojic, Hamid Palangi, and Jonathan Larson. Evaluating cognitive maps and planning in large language models with cogeval, 2023. URL https://arxiv.org/abs/2309.15129.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. Zoom in: An introduction to circuits. *Distill*, 2020. doi: 10.23915/distill.00024.001. https://distill.pub/2020/circuits/zoom-in.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
- Thomas Schnake, Oliver Eberle, Jonas Lederer, Shinichi Nakajima, Kristof T. Schütt, Klaus-Robert Müller, and Grégoire Montavon. Higher-order explanations of graph neural networks via relevant walks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7581–7596, 2022. doi: 10.1109/TPAMI.2021.3115452.
- Lee Sharkey, Bilal Chughtai, Joshua Batson, Jack Lindsey, Jeff Wu, Lucius Bushnaq, Nicholas Goldowsky-Dill, Stefan Heimersheim, Alejandro Ortega, Joseph Bloom, Stella Biderman, Adria Garriga-Alonso, Arthur Conmy, Neel Nanda, Jessica Rumbelow, Martin Wattenberg, Nandi Schoots, Joseph Miller, Eric J. Michaud, Stephen Casper, Max Tegmark, William Saunders, David Bau, Eric Todd, Atticus Geiger, Mor Geva, Jesse Hoogland, Daniel Murfet, and Tom McGrath. Open problems in mechanistic interpretability, 2025. URL https://arxiv.org/abs/2501.16496.
- Roger N Shepard and Jacqueline Metzler. Mental rotation of three-dimensional objects. *Science*, 171(3972):701–703, 1971.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers, Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan. Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer Circuits Thread*, 2024. URL https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html.
- Eric Todd, Millicent Li, Arnab Sen Sharma, Aaron Mueller, Byron C Wallace, and David Bau. Function vectors in large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=AwyxtyMwaG.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=bzs4uPLXvi.
- Barbara Tversky. Visuospatial reasoning. *The Cambridge handbook of thinking and reasoning*, pp. 209–240, 2005.

- Jesse Vig and Yonatan Belinkov. Analyzing the structure of attention in a transformer language model, 2019. URL https://arxiv.org/abs/1906.04284.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 5797–5808, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1580. URL https://aclanthology.org/P19-1580/.
- Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in GPT-2 small. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=NpsVSN6o4ul.
- Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah Goodman. Hypothesis search: Inductive reasoning with language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=G7UtIGQmjm.
- Xinyi Wang, Antonis Antoniades, Yanai Elazar, Alfonso Amayuelas, Alon Albalak, Kexun Zhang, and William Yang Wang. Generalization v.s. memorization: Tracing language models' capabilities back to pretraining data. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=IQxBDLmVpT.
- Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas, and Yoon Kim. Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. Association for Computational Linguistics, 2024.
- Liu Yang, Ziqian Lin, Kangwook Lee, Dimitris Papailiopoulos, and Robert D Nowak. Task vectors in in-context learning: Emergence, formation, and benefits. In *Second Conference on Language Modeling*, 2025. URL https://openreview.net/forum?id=lODGn1Rp5t.
- Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of Language Models: Part 2.1, Grade-School Math and the Hidden Reasoning Process. In *Proceedings of the 13th International Conference on Learning Representations*, ICLR '25, April 2025. Full version available at http://arxiv.org/abs/2407.20311.
- Yizhuo Zhang, Heng Wang, Shangbin Feng, Zhaoxuan Tan, Xiaochuang Han, Tianxing He, and Yulia Tsvetkov. Can LLM graph reasoning generalize beyond pattern memorization? In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 2289–2305, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.127. URL https://aclanthology.org/2024.findings-emnlp.127/.

A IDENTIFICATION OF ALGORITHMIC PRIMITIVES

A.1 DETAILED METHODS

We implemented and analyzed five clustering models, using the following datasets:

- 1. five responses on TSP from Phi-4-Reasoning
- 2. five TSP responses from Phi-4 and five TSP responses from Phi-4-reasoning
- 3. five AIME responses from Phi-4 and Phi-4-Reasoning
- 4. five 3-SAT responses from Phi-4 and Phi-4-Reasoning
- 5. five GraphNav responses from Phi-4 and Phi-4-Reasoning

In all cases, we used extracted the token-by-token representation from layer 17 of Phi-4-base. We used k-means clustering with k=50 clusters and designed a html interface to visually inspect the different clusters (we will make this tool available on the project website).

Below we highlight selected clusters from the different clustering models.

A.2 PHI-4-REASONING RESPONSES ON TSP

By manually inspecting the responses, we find that many clusters correspond to specific reasoning motifs and candidate algorithmic primitives.

Cluster 37: nearest_neighbor. This cluster is active specifically during the tokens where the model identifies the nearest-neighbor path or the closest city within a particular candidate path. As Phi-4-Reasoning's responses, in contrast to those by Phi-4, often started out by an initial guess using the nearest-neighbor heuristic (see Appendix ??), we hypothesized that this would be a particularly relevant primitive.

Cluster 26: generate_new_path. This cluster usually preceded a new candidate path. We therefore hypothesized that representations in these tokens may encode relevant primitives for generating new paths.

Cluster 15: compute_distance. This cluster seems to precede the computation of the total distance of a candidate path.

Cluster 41: compare_or_verify_2. This cluster corresponds to a range of statements involved in comparisons (in particular comparing the distance of different paths), verification (e.g. verifying whether a generated path is valid), or recall (e.g. recalling the best path so far, a closely related step to comparisons).

Cluster 35: compare_or_verify. Interestingly, this cluster is usually active on the final presentation of the total distance. However, we noticed that not every token corresponding to a total distance belonged to this cluster. Rather, this cluster reliably predicted whether the next sentence involved a comparison or verification step — in particular, cluster 35 often preceded cluster 41. We therefore hypothesized that this cluster could be involved in promoting these comparisons and verifications.

Cluster 30: produce_final_answer. This cluster was active during the production of the final answer and we hypothesized that it would be relevant for that step.

A.3 RESPONSES BY PHI-4 AND PHI-4-REASONING ON TSP

By selectively analyzing the clusters associated with the largest-magnitude differences between their involvement in Phi-4 and Phi-4-Reasoning, we identify a set of clusters that are more common in Phi-4 or Phi-4-Reasoning:

A.3.1 Clusters more common in Phi-4-Reasoning

Cluster 19: compare_or_verify. This cluster again implements comparisons and verifications. It arises almost exclusively in Phi-4-Reasoning responses, highlighting the reasoning-finetuned model's tendency to implement more comparisons and verifications.

Clusters 13 and 10: path_generation Cluster 13 represents early cities during the path generation while cluster 10 represents the connections between cities. Notably, it only represents paths that are not generated as a part of the brute-force strategy, indicating that the computations involved in the brute-force strategy are different. For responses by Phi-4 that do not implement a brute-force strategy, these clusters are also more frequently active.

Cluster 22: generate_path_guided. This cluster precedes a path generation, but only in Phi-4-Reasoning. We therefore hypothesize that clusters 6 and 22 might induce different structures in their generated paths.

A.3.2 Clusters more common in Phi-4.

Cluster 11: brute_force. While this cluster sometimes arises in Phi-4-Reasoning, it is more strongly associated with Phi-4 and arises, in particular, when Phi-4 states its approach to implement a brute-force search.

Cluster 41: generate_path_brute_force This cluster is specifically active during the generation of paths involved in brute-force searches.

Cluster 0: compound_distance_lookup This cluster is active specifically before a distance lookup requires a multi-step computation. For example, consider the following segment: **Permutation: $0 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 0**$ Distance = 44 (0 to 1) + 36 (1 to 2) + 32 (2 to 3) + 46 (3 to 4) + 26 (4 to 5) + 37 (5 to 0) = 221. Here, predicting the distance requires first identifying the relevant edge in the path before looking up the corresponding entry in the weight matrix. In contrast, consider the following segment: Alternatively: $0 \rightarrow 5 \rightarrow 3 \rightarrow 2 \rightarrow 1 \rightarrow 4 \rightarrow 0$. Then: $0 \rightarrow 5 = 37$, $5 \rightarrow 3 = 31$, $3 \rightarrow 2 = 32$, $2 \rightarrow 1 = 36$, $1 \rightarrow 4 = 28$, $4 \rightarrow 0 = 42$. Total: 37 + 31 = 68, +32 = 100, +36 = 136, +28 = 164, +42 = 206. Here, adding the new distances only requires moving forward in the previously generated sequence and therefore does not require a multi-step computation. Importantly, cluster 0 is not active in this sentence.

A.3.3 Clusters common to both models.

Cluster 27: edge_retrieval. This cluster indicates the retrieval of the distance between two edges. It is a shared primitive between Phi-4 and Phi-4-Reasoning.

Cluster 6: generate_path_1. This cluster precedes a path generation and arises in both Phi-4 and Phi-4-Reasoning.

A.4 RESPONSES BY PHI-4 AND PHI-4-REASONING ON 3SAT

Clusters more common in Phi-4-Reasoning. Cluster 30: reasoning_scaffold. This cluster appears to involve a lot of strategizing and reasoning.

Cluster 37: if_clause_true. This cluster corresponds to a particular strategy implemented by Phi-4-Reasoning in solving 3-SAT problems: investigating the consequences of one particular clause being true or false. Notably, this is a cluster rarely occurring in Phi-4, which may therefore illustrate an algorithmic primitive largely occurring in Phi-4-Reasoning. Indeed, we discovered these differences in strategy from our clustering analysis, illustrating how our approach could potentially support better understand differences in algorithmic strategies.

A.4.1 Clusters more common in Phi-4.

Cluster 34: logical_expressions_latex. This cluster is largely active on Latex based logical expressions. This is consistent with more general observations that Latex formats are much more common in Phi-4 than Phi-4-Reasoning.

A.4.2 Clusters common to both models.

Cluster 11: if_variable_true. In contrast, this cluster mostly arises in Phi-4. Interestingly, besides being involved in general reasoning, it appears to promote an alternative strategy to the strategy above: setting particular variables to true or false. This results in a more brute-force oriented approach, mirroring the differences between Phi-4 and Phi-4-Reasoning on TSP.

Cluster 38: recall_clause. This cluster precedes the recall of particular clauses from the problem.

A.5 RESPONSES BY PHI-4 AND PHI-4-REASONING ON AIME

A.5.1 Clusters more common in Phi-4-Reasoning

Cluster 3: solve_equation. This cluster is mostly active on two tasks which require solving an equation/inequation. While it is also partially active in Phi-4 for one of those tasks, it is much more common in Phi-4-Reasoning.

Cluster 31: verification_alternate_approach. This cluster corresponds to verifying specific statements or registering potential concerns with an approach and considering an alternative approach. It arises almost exclusively in Phi-4-Reasoning.

A.5.2 Clusters more common in Phi-4

Cluster 49: solve_equation. This cluster is mostly active when Phi-4 solves equations, whereas it is much less active on Phi-4-Reasoning.

A.6 RESPONSES BY PHI-4 AND PHI-4-REASONING ON GRAPHNAV

A.6.1 Clusters more common in Phi-4-Reasoning.

Cluster 2: reasoning_scaffold. This cluster is involved in structuring the overall reasoning in Phi-4-Reasoning.

Cluster 22: plan_final_answer. Phi-4-Reasoning commonly plans its final answer during the thinking section. This cluster is active in that section.

Cluster 9: recall_instructions. This cluster is involved in Phi-4-Reasoning recalling its instructions.

A.6.2 Clusters more common in Phi-4.

Cluster 23: breadth_first_search. This cluster is specifically active when the model plans and implements a breadth-first search algorithm. This indicates that Phi-4 is more likely to do this.

Cluster 48: reasoning_scaffold. This cluster is involved in structuring the overall reasoning in Phi-4, but not Phi-4-Reasoning.

B EXPERIMENTAL SETUP DETAILS

B.1 TRAVELING SALESPERSON PROBLEM

Prompt: The traveling salesman problem (TSP) is a classic optimization problem that aims to find the shortest possible route that visits a set of cities, with each city being visited exactly once and the route returning to the original city.

You must find the shortest path that visits all cities. The distances between each pair of cities are provided. Please list each city in the order they are visited. Provide the total distance of the trip. The final output of the result path and total distance wrapped by the final answer tag, like {<final_answer>{'Path': '0->1->2->...->N->0', 'TotalDistance': 'INT_TOTAL_DISTANCE'}</final_answer>}

 The distances between cities are below: The path between City 0 and City 1 is with distance 44. The path between City 0 and City 2 is with distance 45. The path between City 0 and City 3 is with distance 45. The path between City 0 and City 4 is with distance 42. The path between City 0 and City 5 is with distance 37. The path between City 1 and City 2 is with distance 36. The path between City 1 and City 3 is with distance 27. The path between City 1 and City 4 is with distance 28. The path between City 1 and City 5 is with distance 29. The path between City 2 and City 3 is with distance 32. The path between City 2 and City 4 is with distance 38. The path between City 2 and City 5 is with distance 42. The path between City 3 and City 4 is with distance 46. The path between City 3 and City 5 is with distance 31. The path between City 4 and City 5 is with distance 26.

B.2 AMERICAN INVITATIONAL MATHEMATICS EXAMINATION (AIME)

LLMs were presented with problems from the AIME benchmark, which tests various domains of mathematical reasoning such as algebra, geometry, number theory, and combinatorics. The correct answer for all AIME questions is an integer from 0-999.

Example prompts (from 2025 AIME I):

- **Problem 1:** Find the sum of all integer bases b > 9 for which 17_b is a divisor of 97_b .
- **Problem 2:** On $\triangle ABC$ points A, D, E, and B lie in that order on side \overline{AB} with $\overline{AD}=4$, DE=16, and EB=8. Points A, F, G, and C lie in that order on side \overline{AC} with AF=13, FG=52, and GC=26. Let M be the reflection of D through F, and let N be the reflection of G through E. Quadrilateral DEGF has area 288. Find the area of heptagon AFNBCEM.

B.3 3-LITERAL SATISFIABILITY PROBLEM (3SAT)

Models were tasked with determining the satisfiability of various 3SAT problems, a class of NP-hard problems requiring combinatorial reasoning. If a model responded that a problem was satisfiable, they were required to provide the specicific literals that would achieve satisfiability.

B.4 GRAPH NAVIGATION (GRAPHNAV)

Models were tasked with identifying the shortest path between two nodes in binary trees of varying depth (2-6, corresponding to 7-127 nodes). Each node was a randomly generated integer from 1-200, and edge lists were presented in randomized order. We included a 'forward' condition, where the initial node was the root and the goal was a randomly selected leaf, and a 'reverse' condition where the initial node was a randomly selected leaf and the goal was the root.

Forward Direction Prompt Example: Given the following list of connected rooms, someone wants to get to 91 from 114. The initial room and other rooms are denoted by numbers. 114->45, 114->90, 45->167, 45->91, 90->49, 90->9. Starting at 114, what is the shortest path of rooms to visit if someone wants to arrive at 91? Include the final response in parentheses as the list of rooms separated by commas.

Reverse Direction Prompt Example: Given the following list of connected rooms, someone wants to get to 63 from 119. All of the rooms are denoted by numbers. 164->63, 119->147, 52->147, 54->164, 147->63, 62->164. Starting at 119, what is the shortest path of rooms to visit if someone wants to arrive at 63? Include the final response in parentheses as the list of rooms separated by commas.

B.5 GENERATED MODEL RESPONSES

For TSP, AIME, and 3SAT we used previously generated model responses (Balachandran et al., 2025); for GraphNav, we generated responses from Phi-4 and Phi-4-Reasoning ourselves, using default generation parameters (temperature: 0.8, top_k: 50, top_p: 0.95).

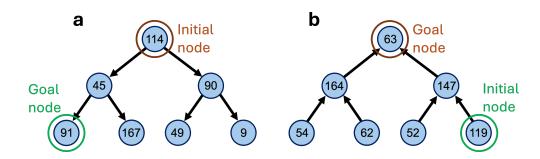


Figure 6: Examples of binary trees used in the two conditions. **a** Forward condition: models were tasked with finding the shortest path from the root node to a randomly selected leaf node, as in the forward prompt above. **b** Reverse condition: models were tasked with finding the shortest path from a randomly selected leaf node back to the root node.

C IDENTIFYING AND COMPOSING FUNCTION VECTORS

C.1 TASKS

We generated a set of algorithmic in-context learning tasks that require specific operations over graphs.

Terminal node recognition (TNR), identifies the final node in a path (e.g., given the input "path: T-P-Q-V", the model is tasked with outputting the token 'V'). Reward comparison (RC) compares rewards across candidate nodes (e.g., given the input "rewards=[M:100 vs S:44]", the model is tasked with outputting 'M'). The evaluation task in Fig. 5a required combining both primitives: given two paths, the model was prompted to return the node that contains the highest reward.

Example Prompt (Correct Answer: A)

```
path1: B-O-Q-D-A.
path1-rewards=[A:65 vs Y:75].
path2: C-W-V.
path2-rewards [V:15 vs Y:45]
```

Relatedly, get_first_node and get_last_node requires responding with the first or last node of a presented path (each node consists of numbers or capital letters and has between three and six elements). get_predecessor and get_successor receive a path and a node has input and need to return the predecessor or successor node respectively.

```
Example (get_successor)

Input: Graph: D-C-N-J, Node: C
Output: N
```

C.2 DETAILED METHODS AND RESULTS

We extracted function vectors corresponding to each of these tasks using the same approach as Todd et al. (2024). We identified the 35 attention heads with the highest average indirect effect (see definition in Todd et al. (2024)) for Phi-4 and Phi-4-Reasoning and 20 attention heads the highest average indirect effect for Llama-3-8B. We injected these function vectors across all layers and analyzed the layer with the largest effect.

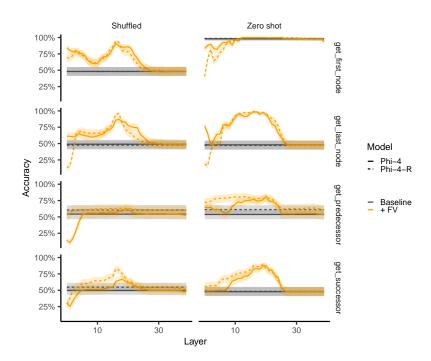


Figure 7: Accuracy on different graph operations after injecting the corresponding function vector into Phi-4 or Phi-4-Reasoning, either after shuffled in-context examples are in a zero-shot setting.

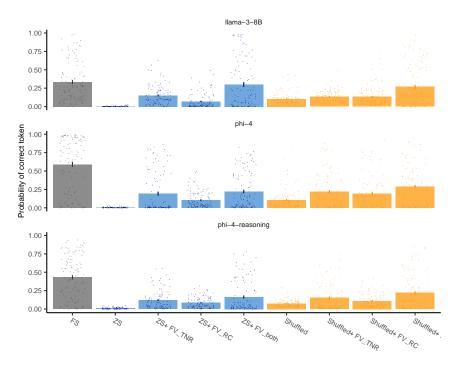


Figure 8: Average probability assigned to the correct tokens when injecting the Terminal_node_recognition and Reward_comparison FVs (FV_TNR and FV_RC) and the sum of both (FV_both). We compare the performance across Llama-3-8B, Phi-4, and Phi-4-Reasoning, and across a corrupted ICL and zero-shot context.

D BEHAVIORAL HALLMARKS ON TSP

We consider three setups for injecting primitive vectors into TSP. In Figs. 2b,d, we only present the model with the TSP prompt and inject the primitive vectors throughout generation. We maximally generate 500 tokens. This is not sufficient to reliably generate a full answer, but is sufficient to test if the behavior is expressed. In Fig. 2c and Table 1, we add a sequence of randomly generated paths to the assistant response before injecting the function vectors. In both of these cases, we define several measures of different behavioral hallmarks. For these measures we automatically extract the set of paths generated in the model response using a regular expression. Below we specify our operational definitions of the different behavioral hallmarks:

- 1. % NN paths: What proportion of generated paths corresponds to a nearest-neighbor heuristic?
- 2. # Unique paths: How many valid TSP candidate paths are generated (after removing duplicate paths)?
- 3. Distance computation: What is the earliest token at which a distance computation occurs (operationalized as a sum of several numbers)? A lower number thus corresponds to a stronger expression of this behavior. If no distance computation occurs in the entire response, we set the value to 500.
- 4. Final answer: At what token does the model generate the <final_answer> token?
- 5. # Verifications: How many verification-related words are mentioned in the response?
- 6. # Comparison: How many comparison-related words are mentioned in the response?

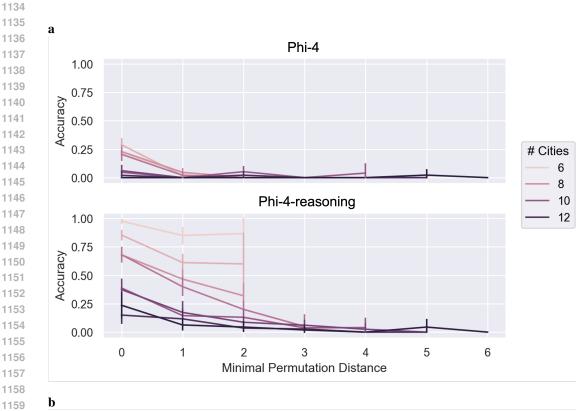
Finally, when evaluating the impact of the <code>get_first_node</code> and <code>get_last_node</code> primitive vectors (Fig. 5c), we automatically extract a generated path from the middle of the reasoning trace (constraining our selection to only consider paths where the first and last node are not identical). We then generate the model response starting immediately after the generated path and injecting the corresponding primitive vector throughout. We generate 20 tokens and assess whether the first node that is mentioned corresponds to the first node of the path, the last node of the path or some other node.

E NEAREST NEIGHBOR STRATEGY ANALYSIS

We compared the implementation of the nearest neighbor search heuristic on the Traveling Salesperson Problem (TSP) in Phi-4-Reasoning and Phi-4-Base. First, we computed the minimum edit distance between nearest neighbor solution and the optimal solution and analyzed how this distance relates to accuracy. We computed these minimal NN-optimal edit distances using any city as the starting point (Figure 9 a) and with city 0 as the starting point (Figure 9 b). The latter edit distance was chosen, because in practice, we observed that the models frequently used city 0 as the initial node.

Next, we tested whether one of the first five generated paths is a nearest-neighbor solution; we sampled from the first five paths to account for idiosyncrasies in the model's output (e.g., repeating path segments from the prompt). Additionally, we examined the relationship between the number of considered paths and the average distance from the NN solution.

Overall, these findings demonstrate that Phi-4-Reasoning has a strong tendency to implement the nearest neighbor search heuristic. Specifically, the model tends to begin with the NN solution starting at city 0 and iteratively edit the path to minimize its total distance. While Phi-4-Base does not invoke the nearest neighbor heuristic as reliably as Phi-4-Reasoning, it does to some degree as shown in Figures 9 b and 10 a.



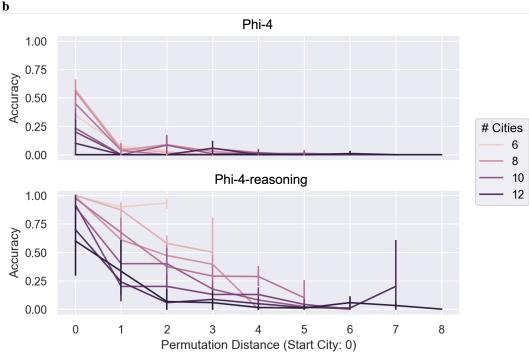


Figure 9: **a** The minimal NN-optimal edit distances with any city as the starting point. **b** The minimal NN-optimal edit distance using city 0 as the starting point. We observed a strong trend between edit distance and accuracy in **b**, even when controlling for the number of cities (and a somewhat weaker trend in **a**). This trend was especially strong for Phi-4-Reasoning, suggesting this model may effectively implement the nearest neighbor heuristic beginning at city 0.

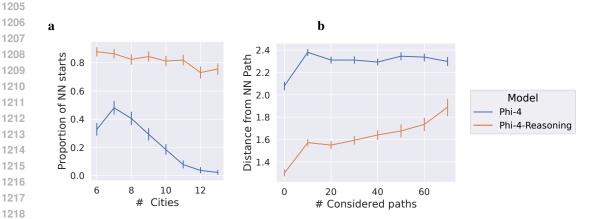


Figure 10: **a** We found that Phi-4 reasoning begins with the NN heuristic much more frequently than Phi-4-Base, especially when >8 cities are present. **b** Paths generated by Phi-4-Reasoning tend to have much lower distance on average from NN paths than those generated by Phi-4. Additionally, in Phi-4-Reasoning edit distance is lowest among the initial paths and increases as more candidate paths are selected.