How Do Language Models Compose Functions?

Anonymous Author(s)
Affiliation
Address
email

Abstract

While large language models (LLMs) appear to be increasingly capable of solving compositional tasks, it is an open question whether they do so using compositional mechanisms. In this work, we investigate how feedforward LLMs solve two-hop factual recall tasks, which can be expressed compositionally as $g(f(x))$. We first confirm that modern LLMs continue to suffer from the "compositionality gap", i.e. their ability to compute both $z = f(x)$ and $y = g(z)$ does not entail their ability to compute the composition $y = g(f(x))$. We then decode residual stream representations and identify two processing mechanisms: one which solves tasks *compositionally*, computing $f(x)$ along the way to $g(f(x))$, and one which solves them *directly*, without any detectable signature of the intermediate variable $f(x)$. Finally, we find that embedding space geometry is strongly related to which mechanism is employed, where the idiomatic mechanism is dominant when tasks are represented by translations from x to $g(f(x))$ in the embedding spaces.

1 Introduction

Compositional behavior (McCurdy et al., 2024) is widely considered essential for flexible and general intelligence (Szabó, Fall 2024). A long-running debate has asked whether compositional *behavior* necessarily entails compositional *representations* and *processes*. One the one hand, formal languages based on compositional syntax and semantics are guaranteed to support certain types of invariance and generalization, making them compelling models for how humans might achieve abstract cognitive abilities like language and logic (Fodor, 1975; Quilty-Dunn et al., 2023). On the other hand, critics are quick to point out that humans frequently deviate from ideal compositional and logical behavior, suggesting some other mechanism must underlie our advanced cognition (Kahneman & Tversky, 1972; Evans, 2002).

Large language models (LLMs) provide an opportunity to revisit this debate in a new light. LLMs exhibit behavior that is at least ostensibly compositional, and which is not easily explained away by trivially non-compositional mechanisms (McCoy et al., 2024; Griffiths et al., 2025). However, LLMs also lack the kinds of explicit symbolic architectural components that have long been assumed necessary for such compositionality. This provides an opportunity to ask: do LLMs produce compositional behavior by invoking compositional processes, or do they rely on something more idiomatic instead?

We offer an initial investigation into this question, focusing on a set of two-hop factual retrieval tasks, such as: given a book's title, output that book's author's birth year. All of the tasks we consider can be formally expressed as $y = g(f(x))$ and are thus defensibly "compositional" in the sense invoked in traditional symbolic models. We are interested in whether LLMs solve such tasks by approximating the mapping from x to y *compositionally*, by computing the intermediate variable $z = f(x)$, or *directly*, without a readily-detectable representation of any such z. We find that:

Submitted to 40th Conference on Neural Information Processing Systems (NeurIPS 2026). Do not distribute.

1. Models' ability to compute both $x \to f(x)$ and $f(x) \to g(f(x))$ does not entail their ability to compute $x \to g(f(x))$. This extends earlier findings on the "compositionality gap" (Press et al., 2022), showing that the gap holds for modern models and on a larger set of tasks. This gap is not trivially reduced in larger models or necessarily by reasoning models (Sec. 3).

2. Models exhibit both *compositional* processing mechanisms and *direct* processing mechanisms, as defined above. The type of mechanism is only weakly associated with accuracy, suggesting that LLMs are able to use both effectively to compute correct answers (Sec. 4).

3. The choice of mechanism is well predicted by embedding space geometry. Specifically, when there exists a translation (for a given task) from x in the input embedding space to $g(f(x))$ in the output unembedding space, the LLM tends to favor direct computation over compositional processing (Sec. 5).

2 Task Setup

Our tasks involve predicting a composition $g(f(x))$ from an input x, using in-context learning (ICL) and where f and g are some pre-defined functions. See Table 1 for the full list of tasks we use. We choose common functions f and g that models might learn through their pre-training and for which the inputs and outputs are lexical units. This enables us to use well-established tools for analyzing the mechanisms and latent computations in Transformer models, focusing on a few token positions (i.e. residual streams) and a single autoregressive forward pass.

We design the set of tasks in our investigation to cover a qualitative variety of functions, such as arithmetic, factual recall, lexical functions, translation, rotation, and string manipulation. By construction, all of our tasks can be computed by applying f and then g, yielding the causal hops $x \to f(x) \to g(f(x))$. Some tasks (e.g. commutative tasks) can also be computed through the hops $x \to g(x) \to g(f(x))$ — in which case, the intermediate z may also equal $g(x)$. We differentiate these further, and also describe our dataset construction methodologies (including our sources and pre-processing), in Appendix A.

In our experiments, we randomly sample 10 in-context examples for a given task and query. Each in-context example is formatted with a "`Q: {input} \n A: {output} \n\n`" prompting structure and the test query is formatted with "`Q: {input} \n A:`".

Table 1: List of our tasks. The compositional function $(g \circ f)$ is constructed by f and g here. We list the number of examples (#) in each task's dataset, along with the variables x, $f(x)$, and $g(f(x))$ for one random example. We list $g(x)$ and $f(g(x))$ for tasks that define them in Appendix A.

f	g	#	$x \to f(x) \to g(f(x))$
Word → Antonym	English → Spanish	2398	bogus → authentic → auténtico
Word → Antonym	English → German	2398	philosophical → practical → praktisch
Word → Antonym	English → French	2398	excessive → insufficient → insuffisant
Book → Author	Author → Birth Year	2228	The Boy in the Striped Pyjamas → John Boyne → 1971
Song → Artist	Artist → Birth Year	958	Heartbreak Hotel → Elvis Presley → 1935
Landmark → Country	Country → Capital	1385	Taq-i Kisra → Iraq → Baghdad
Park → Country	Country → Capital	743	Mount Rainier National Park → United States → Washington, D.C.
Movie → Director	Director → Birth Year	2180	Cape Fear → Martin Scorsese → 1942
Person → University	University → Year	4992	Andi Gutmans → Technion – Israel Institute of Technology → 1924
Person → University	University → Founder	4996	Ezra Abbot → Bowdoin College → James Bowdoin
Product → Company	Company → CEO	1904	NES-101 → Nintendo → Shuntaro Furukawa
Product → Company	Company → HQ	2276	Toyota Alphard → Toyota → Toyota
x + 10	2x	1000	699 → 709 → 1418
x + 100	2x	1000	922 → 1022 → 2044
x mod 20	2x	1000	891 → 11 → 22
Word → Numeric	2x	1000	one hundred and forty-eight → 148 → 296
Word[:-1]	Word[::-1]	2946	responsible → responsibl → lbisnopser
Rotate(RGB, 120°)	RGB → Name	1000	8a735a → 598a73 → dimgray

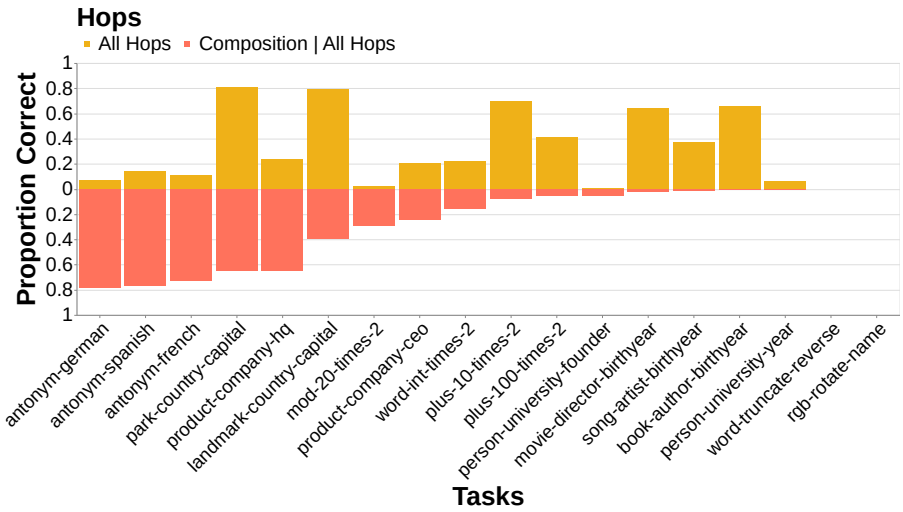

Figure 1: Compositionality gap for Llama 3 (3B) on our tasks. Yellow indicates the proportion of examples for which the model correctly solves all causal hops. Red indicates the (strict) subset of those examples for which the model additionally solves the composition. Correlation between red and yellow bars is $r^2 = 0.00$.

3 Compositionality Gap

Press et al. (2022) documented a "compositionality gap" in LLMs, showing that they consistently fail to solve compositions, despite solving the hops independently. Press et al. (2022) tested the GPT-3 family of models with natural language questions about celebrities and encyclopedic knowledge that required two-hops of factual recall. We confirm and extend this finding by testing modern LLMs on a larger set of compositional tasks.

3.1 Experimental Design

We prompt models with `input` $\rightarrow$ `output` mappings between lexical units.[1] We measure models' predictive accuracies using the ICL prompts from Sec. 2, greedy sampling, and the exact match evaluation metric. We test additional prompt templates in Appendix C and the effect of including task-specific instructions in the prompt in Appendix D. The *compositionality gap* is defined as the proportion of examples for which a model answers both $x \rightarrow f(x)$ and $f(x) \rightarrow g(f(x))$ correctly,[2] but $x \rightarrow g(f(x))$ incorrectly.

We test a few Llama 3 and OLMo 2 models on all of our tasks using all available examples. We emphasize results for Llama 3 (3B) in this section and include the other models in Appendix F. We also test an even wider set of models (also including DeepSeek and GPT models) on 4 tasks: `antonym-spanish`, `plus-100-times-2`, `park-country-capital`, and `book-author-birthyear` (which capture a representative set of processing signatures from Sec. 4). We aggregate metrics over these tasks and use 100 examples per task for testing.

3.2 Results

We show performance of the Llama 3 (3B) model on our tasks in Fig. 1. We clearly find a compositionality gap: the model is unable to solve the composition in 20–100% (varying by task) of examples for which it can solve all hops. We show the performances of our other models in Fig. 2. We find the compositionality gap does reduce with size from 72% $\rightarrow$ 39% (Llama 3, 1B $\rightarrow$ 405B) and 74% $\rightarrow$ 49% (OLMo 2, 1B $\rightarrow$ 32B). However, the gap clearly remains and we find monotonically

[1] Note that this represents a methodological difference from Press et al. (2022), who prompted with long-form questions. Our format is chosen to fit with the interpretability methods we use in later sections.

[2] We extend this definition to further require success at $x \rightarrow g(x)$ and $g(x) \rightarrow g(f(x))$ in tasks where these are valid hops.

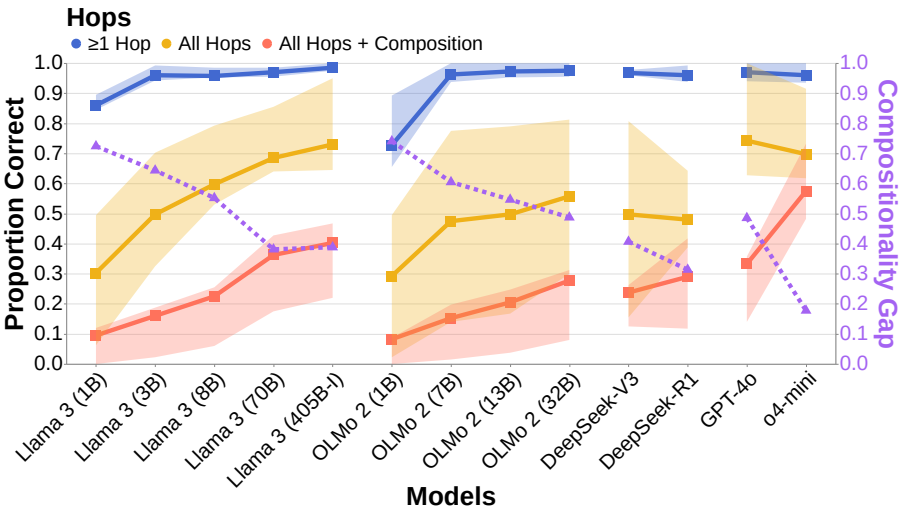

Figure 2: Compositionality gap (dashed purple line; lower is better) of various models aggregated over 4 tasks (100 examples each). Blue, yellow, and red lines show proportions of examples for which models correctly solve combinations of hops and the composition. Purple line shows the relative gap between yellow and red: the proportion of examples for which the model cannot solve the composition, out of those for which it can solve all hops. "-I" indicates the instruction-tuned variant of Llama 3 (405B). Error bands show interquartile range.

diminishing improvements for both model families with respect to size. We plot the gap against model parameters and layers in Appendix E. In fact, the gap shows no improvement at all between the 70B and (instruction-tuned) 405B parameter Llama 3 models.

We also compare reasoning models (o4-mini and DeepSeek-R1; allotted a budget of 2000 reasoning tokens) against same-generation, non-reasoning models (GPT-4o and DeepSeek-V3) in Fig. 2. We find some reduction (41% $\rightarrow$ 31%) in the compositionality gap in the case of DeepSeek's reasoning model and significant reduction (49% $\rightarrow$ 18%) in the case of o4-mini. As o4-mini is proprietary (and both this and GPT-4o have additional "external tool-use" capabilities), it is difficult to speculate about the exact causes for these improvements. However, it is notable that even with advanced reasoning models, the gap does not necessarily disappear entirely.

Finally, we include a random sample of correct and incorrect predictions for every model in our supplemental materials.

4 Analyzing Processing Mechanisms

We next try to understand *how* the model correctly computes compositions in cases where it is successful. Our intuition is based on prior work from Merullo et al. (2024) which identifies a processing signature in models that solve one-hop relational tasks. That work shows that models predicting $y = f(x)$ iteratively surface vocabulary representations — first for x and then for y — in the residual stream. This "crossover" point was interpreted as evidence of the function f being applied to the argument x in order to yield the final answer $f(x)$ and was localized to specific computations in the MLPs.

In this section, we ask whether an analogous signature will emerge in the case of compositional functions, $g(f(x))$. That is, can we find distinct intermediate representations for x, followed by $f(x)$, and then $g(f(x))$ during the model's processing?

Here, we employ analyses most similar to Biran et al. (2024) and Yang et al. (2025) in the context of our evaluation (see Sec. 7 for further discussion on these works). We join other works in identifying stages of processing within language models (Tenney et al., 2019; Merullo et al., 2024; Lepori et al., 2025; Lu et al., 2025).

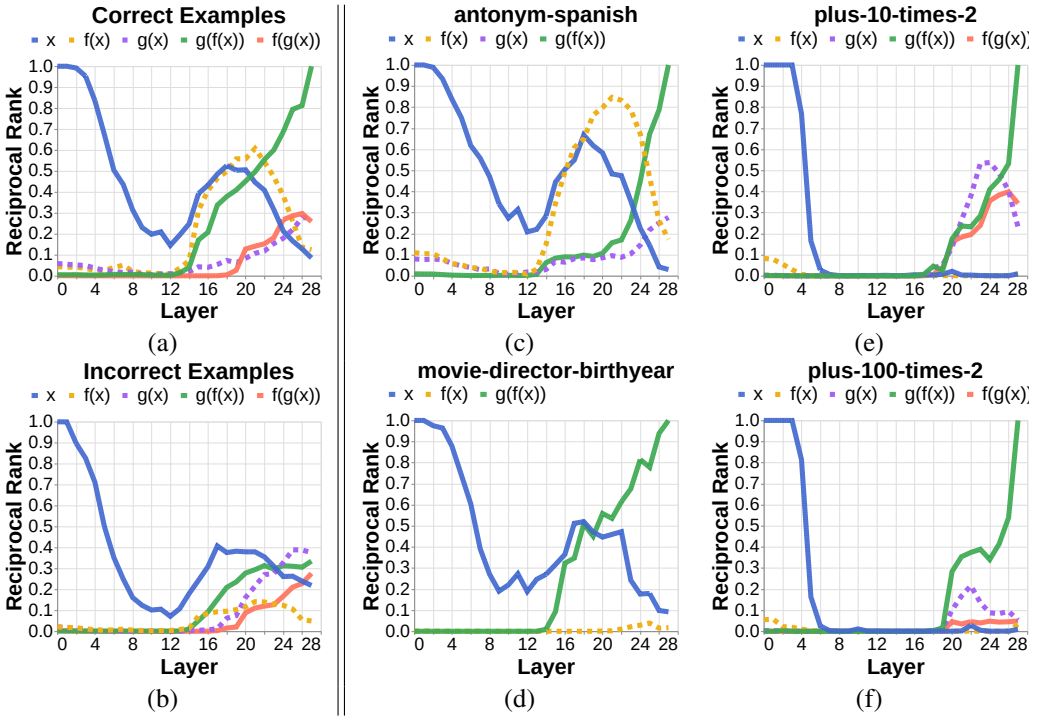

Figure 3: (a–b) Processing signatures aggregated over examples (across all tasks) in which Llama 3 (3B) solves all hops correctly, but the composition (a) correctly or (b) incorrectly. (c–f) Processing signatures for particular tasks — aggregated over examples where this model correctly solves all hops and the composition. (a–f) Lines show reciprocal ranks of relevant variables (decoded using logit lens) from residual streams corresponding to $x \rightarrow g(f(x))$. Intermediate variables are shown with dashed lines. The incorrect composition, $f(g(x))$, is shown by the red line when not distinct from $g(f(x))$.

4.1 Experimental Design

We rely on existing methods which allow us to analyze processing signatures that are interpretable using the vocabulary space of the model (nostalgebraist, 2020; Geva et al., 2022). We specifically use *logit lens* (nostalgebraist, 2020), a method which projects intermediate representations into the vocabulary space using the language modeling head. Logit lens reveals information that models store in a linear representational subspace (i.e. the unembedding space) that is fundamental to the model. While these signals alone do not guarantee causal use, we find strong signals for our specific variables of interest in this important subspace very compelling and also demonstrate their causal effects in Appendix J. Such methods are not able to perfectly decompose representations and so we also include results in Appendix I using the token identity patchscope (Ghandeharioun et al., 2024) to explore an alternative representational subspace. Both methods yield similar findings concerning the mechanisms investigated in this section. Our intervention in Appendix J further supports these findings, demonstrating existence (or lack thereof) of intermediate variable representations.

We follow the approach from Merullo et al. (2024) to identify the processing signature of models that solve our compositional tasks and, in particular, representations of the intermediate variables, $f(x)$ and $g(x)$, prior to those for $g(f(x))$. We specifically use logit lens to analyze the residual streams corresponding to the computation $x \rightarrow g(f(x))$ and measure the reciprocal rank of our variables at each layer (see Appendix B for more details). We also use the maximum reciprocal rank of our intermediate variables across the layers as a heuristic for their overall presence in the computation.

We focus this section's analysis on Llama 3 (3B) and summarize results for other models in Appendix K. We exclusively analyze examples where the model can solve all requisite hops. To ensure sufficient sample sizes, we exclude any task with fewer than 10 such examples where the model can also successfully solve the composition. In particular, the excluded tasks for Llama

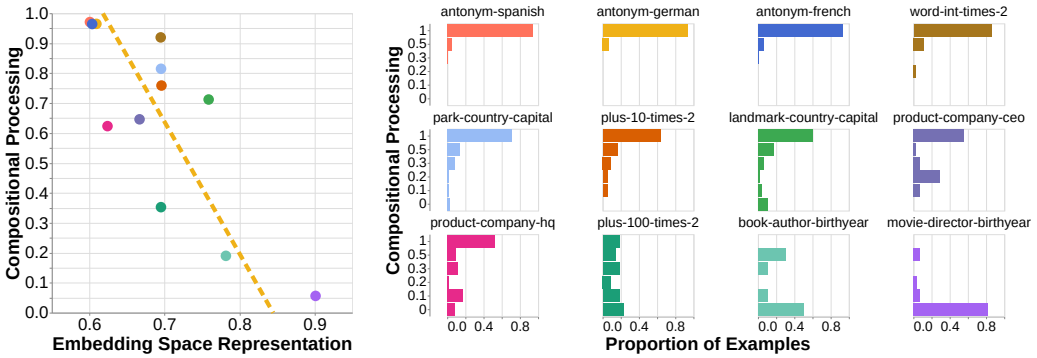

(a) (b)

Figure 4: Relationship between geometry and mechanisms of Llama 3 (3B). (a) Tasks correlate strongly between use of compositional mechanisms (heuristic from Sec. 4.1 based on reciprocal rank; on average across examples) and whether tasks are represented as translations in the embedding spaces ($r^2 = 0.65$). Conversely, accuracy is weakly correlated with these mechanism ($r^2 = 0.22$) and representation ($r^2 = 0.10$) metrics. (b) Distribution of examples for each task, shown as a histogram of intermediate variable reciprocal ranks. (a–b) Colors refer to corresponding tasks between points in (a) and histograms in (b).

3 (3B) are `song-artist-birthyear`, `person-university-year`, `person-university-founder`, `mod-20-times-2`, `word-truncate-reverse`, and `rgb-rotate-name`. We show results for these tasks in Appendices G and H.

4.2 Results

Fig. 3a shows the relative presence of each of the variables, across layers and aggregated over all instances in which the model ultimately produced the correct answer. In such cases, we see a very clear peak signal for the intermediate variable $f(x)$, as expected, between those for x and $g(f(x))$. Interestingly, this signal is much less clear for cases in which the model ultimately produces the incorrect answer (Fig. 3b). However, upon further inspection, there is little evidence of a causal relationship here, which we discuss further in Appendix H.

There are also plenty of individual examples in which the model produces a correct answer without showing any signature of the intermediate variables, and there is only a weak correlation by task ($r^2 = 0.22$) between predictive accuracy (measured as in Sec. 3.1) and the presence of intermediate variables as measured by our heuristic (Sec. 4.1).

Figs. 3c to 3f show model processing signatures for a few tasks, aggregated over cases in which the model produces correct answers. We see, for example, that there is a clear signature in the `antonym-spanish` task (Fig. 3c) for the intermediate computation of $f(x)$, the word's antonym, before it is translated into Spanish. In contrast, for the `movie-director-birthyear` task (Fig. 3d), there is no decodable signal for $f(x)$, the movie's director, before the model produces their birth year. This variation can be seen in qualitatively similar tasks as well: tasks with the same basic arithmetic structure (Figs. 3e and 3f) only sometimes carries detectable signatures of $f(x)$ or $g(x)$, depending on the task's operand (e.g. 10 or 100). We show processing signatures for the remaining tasks in Appendix G and for all tasks, aggregated over unsuccessful cases, in Appendix H.

5 Compositional Processing and Embedding Space Geometry

Given that there is significant variation in whether or not the LLM solves a task compositionally (i.e. how strongly they appear to compute the intermediate variables), we next ask why this variation occurs. It is well-known that some relations are represented by translations in classical word embedding spaces (Mikolov et al., 2013; Hewitt & Manning, 2019). Hernandez et al. (2024) also shows that, when LLMs compute certain `subject` $\rightarrow$ `object` relations, several layers of computation can be well approximated by simple, linear transformations. We extend these ideas together and posit that language models could directly represent some tasks as translations between the embedding and

unembedding spaces (spanning language models' entire forward computation). We hypothesize that some compositional functions could be represented more fundamentally via these translations and that models may compute such functions *directly*, rather than *compositionally* by the hops.

5.1 Experimental Design

To investigate our hypothesis, we fit translations for each task using least squares regression from x (average embedding across tokens) to $g(f(x))$ (first token unembedding) on 100 examples.[3] We quantify how well the translation fits the task using its reconstruction accuracy (measured via cosine similarity) on the remaining examples. We quantify how "compositional" the processing is using our heuristic metric which captures the strength of the signal for the intermediate variables, $f(x)$ and $g(x)$ (see Sec. 4). We again restrict our analysis to examples where the model is successful on all hops and the composition, as well as tasks with at least 10 such examples.

5.2 Results

Fig. 4a shows the that there is a a strong inverse correlation ($r^2 = 0.65$) between embedding space task geometry and compositional processing for Llama 3 (3B). That is, the more a task is represented as a translation in the embedding spaces, the more likely the model is to display *idiomatic* (as opposed to *compositional*) processing.

This correlation is computed by averaging the indicator for compositional processing across instances for each task. Fig. 4b shows the de-aggregated distribution of our "compositionality" metric across the examples. For some tasks, it appears that nearly all individual examples behave the same way. For example, nearly every instance of `antonym-spanish` displays a compositional signature, while almost every instance of `movie-director-birthyear` displays an idiomatic one. On the other hand, this distribution is more uniform for other tasks, such as `plus-100-times-2`. This distribution appears to be bimodal across all examples: 82% have very low (< 0.1) or high (≥ 0.5) values for compositionality.

We include results for several other models in Appendix K, i.e. Llama 3 (3B) Instruct, Llama 3 (8B), OLMo (7B), and OLMo 2 (13B). All the models models we test continue to show strong correlations for the relationship presented here, with an average $r^2 = 0.66$. These models continue to show bimodal distributions of the choice of mechanism across examples. We also explore the hypothesis in Appendix M that linear representations could be predictive of the mechanism, but find a weaker fit than with the translation-based representations shown here.

6 Discussion

Summary of Findings Our results suggest that tasks which appear to have the same computational structure may nonetheless be processed differently by LLMs. In particular, we consider functions which appear compositional in a formal sense — i.e. they can be represented as $y = g(f(x))$ for some reasonably defined f and g. We find evidence that LLMs only sometimes process such functions compositionally, showing evidence of representing or computing the value of $z = f(x)$ on the way to computing y. In other cases, LLMs appear to map x to y directly. Which of these processes is invoked appears to be related to how well the relationship between x and y is represented as a translation in the embedding space, e.g. as a result of training conditions.

Implications for Theories of Compositionality There is a long-running debate about the degree to which compositional behavior (McCurdy et al., 2024) requires compositional mechanisms. The two sides of this debate have often talked past each other, using different types of computational architectures in order to model different aspects of behavior, for example, explicitly compositional symbolic systems for formal domains (Lake et al., 2017; Ellis et al., 2023) vs. distributional or neural systems for humans' more idiomatic performance (Erk, 2012; Lampinen et al., 2024).

Attempts to find compromises or "hybrid" systems often consist of neuro-symbolic systems which are designed top-down (Andreas et al., 2016; Ellis et al., 2018). Large language models offer an

[3] See Appendix L for additional analyses which consider correlations with the representations of the individual hops (rather than the compositional task).

alternative approach for advancing this debate. LLMs have proven capable of a range of behaviors that have traditionally required compositionality — e.g. generating language and writing formal computer code. However, LLMs lack the explicit symbolic mechanisms traditionally associated with such behaviors. Using methods from interpretability to understand how LLMs represent such functions internally enables us to approach the question in a "bottom up" manner, potentially offering novel hypotheses about the mechanisms that can generate behavior that is sometimes systematic and other times heuristic, as is the case in humans (Russin et al., 2025).

Our results suggest that LLMs employ a mix of compositional and idiomatic processing, and that the choice of mechanism is related to the representations of the functions that result from pretraining. This offers an interesting perspective on one question that is frequently at the heart of discussions of compositionality — i.e. what are the primitives and where do they come from (Carey, 2011)? The relationship between embedding space task geometry and compositionality of processing presented here suggests an attractive hypothesis that the primitives are those things which are well represented as a result of (pre-)training, and that compositional mechanisms are invoked to handle those things which are not sufficiently well represented. Future work in this direction would likely yield interesting new results and topics for debate.

Relationship to work on compositional generalization The work presented here concerns the (apparent) compositionality of the processing mechanism, but does not directly relate this mechanism to an LLM's capacity for compositional generalization. The majority of work on compositionality in neural networks (and LLMs) concerns compositional generalization, and the compositionality researchers surveyed by McCurdy et al. (2024) overwhelmingly agree that existing language models are insufficient in this regard. This belief is supported by evidence from many prior works (Sec. 7) and our investigation in Sec. 3.

Our work suggests that models employ both compositional and direct mechanisms to solve tasks. Intuitively, we would expect there to be a relationship between the use of the mechanism and the ability to generalize — i.e. the compositional mechanism should support generalization better than the idiomatic mechanism ("memorization"). However, we do not test this intuition directly in this paper. Future work could do so by employing causal interventions on the intermediate variables. This would likely present new complexities and challenges that would enrich our understanding of compositionality, and of the relationships between mechanisms and behaviors in LLMs in general.

Practical Implications Our work identifies which functions LLMs treat as "primitive" based on their geometry. Models allocate representational capacity to memorizing these primitives. If mechanisms are required for generalization, one may prefer models that fit truer primitives (e.g., f, g, h rather than $g \circ f$) and invoke the compositional mechanism, enabling unseen compositions such as $f \circ h$. Our analysis also provides processing signatures that reveal which functions a model attempts to invoke (successfully or not, e.g. Appendix H), and when individual data points are poorly fit by a function's embedding space representation, possibly indicating weak memorization during training.

7 Related Work

Latent multi-hop reasoning Our work is most closely related to recent works which also study latent two-hop reasoning in large language models. Yang et al. (2024a) use causal interventions to identify the existence of the hops in the latent computation and whether they co-occur. Biran et al. (2024) employ the entity description patchscope (Ghandeharioun et al., 2024) to inspect intermediate representations and localize the hops, finding they are resolved in different layers and token positions. They propose a representational intervention ("backpatching") to correct failures based on this finding. Finally, Yang et al. (2025) use both causal interventions and logit lens to analyze intermediate representations and consistently find a "compositional" processing signature across their tasks. Our work employs all of these interpretability methods (Sec. 4 and Appendices I and J) to analyze the hops, but specifically highlights and investigates the duality of the compositional vs. direct processing mechanisms. All works (including our own) test different sets of tasks, make

experimental design decisions according to their independent goals,[4]and make findings in context of their own experiments.

Among other works in this domain, Wang et al. (2024) trains a language model on synthetic compositional data and identifies a multi-hop reasoning circuit in this model. Shalev et al. (2024) conduct a distributional analysis (considering semantic category spaces, rather than individual tokens) using logit lens. Li et al. (2024); Yu et al. (2025) also propose interventions on intermediate representations and mechanisms to solve failure cases. Yang et al. (2024b) conduct an evaluation that is intentionally designed to omit opportunities for models to exploit shortcuts.

Compositionality Compositionality is long-studied (Fodor & Pylyshyn, 1988; Partee, 2004) but exact definitions evade general consensus. Russin et al. (2024) and McCurdy et al. (2024) offer recent overviews on the topic in the context of large language models. Russin et al. (2024) provide a historic account of compositionality and review studies of compositionality generalization in neural networks. McCurdy et al. (2024) survey compositionality researchers on how to define and evaluate compositional behavior in neural networks. These researchers agree that our current representational analyses are insufficient for evaluating models, but are divided about whether our behavioral analyses are sufficient.

In a partial effort towards defining compositionality, Hupkes et al. (2020) identify five particular aspects of compositionality and propose tests for each using a synthetic, fully compositional translation task. Systematicity is one such aspect and is prominently studied: see Vegner et al. (2025) for a survey of benchmarks for systematic generalization. Our work — in which we test whether $f(x)$ is evaluated before $g(f(x))$ — is closest to Hupkes et al. (2020)'s aspect of localism, in which "smaller constituents are evaluated before larger constituents".

Among many other works, Johnson et al. (2017), (Keysers et al., 2020), Lake & Baroni (2018), Hupkes et al. (2020), and Kim & Linzen (2020) offer prominent benchmarks that behaviorally test for compositional generalization in neural networks trained from scratch on compositional data. These works generally show that such models perform poorly on generalization, or at least poorly implement the compositional processes that underlie the data. Press et al. (2022) and Ma et al. (2023) continue to show significant failures in compositional behavior and generalization in pre-trained models. On the other hand, Furrer et al. (2020) points out that pre-training a masked language model rivals or outperforms architectures specifically designed for the SCAN (Lake & Baroni, 2018) and CFQ (Keysers et al., 2020) generalization benchmarks. Lepori et al. (2023) finds that neural networks learn to implement compositionality structurally in their weights, supporting this claim against the need for specialized symbolic mechanisms.

Compositionality of functions Several works consider how LLMs solve compositions of functions (rather than specifically multi-hop reasoning tasks). Dziri et al. (2023) studies how LLMs autoregressively solve such tasks, like multi-digit multiplication, by inspecting their scratchpads. Wattenberg & Viégas (2024) propose mechanisms which neural networks could use to implement relational compositions. Yu et al. (2023); Todd et al. (2024) propose zero-shot methods to invoke compositions of functions in LLMs that have learned the primitive functions. Zhou et al. (2024) find that language models can compose functions with meta-learning in a way that imitates human behavior.

Limitations

In this work, we primarily analyze the computation that occurs in a single forward pass of LLMs. It is also necessary to understand how other models (e.g. reasoning models or different inductive biases) implement compositional functions. Our findings reflect our specific tasks and experimental design. Further work should test other kinds of compositional functions, and try to more deeply understand the relationship between compositional mechanisms, behavior, and generalization. We investigate a limited subset of mechanisms in language models and use current methods to conduct our analyses. These permit us to decode some, but not all, relevant representational structure. Some signals that we do decode may be a result of feature multiplicity or are not guaranteed to be causal. Finally, some of our tasks (e.g. arithmetic) may be solved by algorithms that we do not consider.

[4]One notable example is that, while Yang et al. (2024a) and Biran et al. (2024) prompt their models with f and g (e.g. "`The mother of the singer of {x} is {y}`"), we typically omit this information from our prompts (i.e. more simply "`Q: {x} \n A: {y}`"), besides in our ablation, to avoid inducing bias towards the compositional mechanism.

References

Andreas, J., Rohrbach, M., Darrell, T., and Klein, D. Neural Module Networks. *Computer Vision and Pattern Recognition (CVPR)*, 2016. URL https://arxiv.org/pdf/1511.02799.

Bast, H. and Buchhold, B. QLever: A Query Engine for Efficient SPARQL+Text Search. *Conference on Information and Knowledge Management (CIKM)*, 2017. URL https://dl.acm.org/doi/10.1145/3132847.3132921.

Biran, E., Gottesman, D., Yang, S., Geva, M., and Globerson, A. Hopping Too Late: Exploring the Limitations of Large Language Models on Multi-Hop Queries. *Empirical Methods in Natural Language Processing (EMNLP)*, 2024. URL https://arxiv.org/pdf/2406.12775.

Carey, S. Précis of The Origin of Concepts. *Behavioral and Brain Sciences (BBS)*, 2011. URL https://doi.org/10.1017/S0140525X10000919.

Dziri, N., Lu, X., Sclar, M., Li, X. L., Jiang, L., Lin, B. Y., Welleck, S., West, P., Bhagavatula, C., Le Bras, R., et al. Faith and Fate: Limits of Transformers on Compositionality. *Neural Information Processing Systems (NeurIPS)*, 2023. URL https://arxiv.org/pdf/2305.18654.

Ellis, K., Morales, L., Sablé-Meyer, M., Solar-Lezama, A., and Tenenbaum, J. Library Learning for Neurally-Guided Bayesian Program Induction. *Neural Information Processing Systems (NeurIPS)*, 2018. URL https://dl.acm.org/doi/10.5555/3327757.3327878.

Ellis, K., Wong, L., Nye, M., Sable-Meyer, M., Cary, L., Anaya Pozo, L., Hewitt, L., Solar-Lezama, A., and Tenenbaum, J. B. DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning. *Philosophical Transactions of the Royal Society A*, 2023. URL https://arxiv.org/pdf/2006.08381.

Erk, K. Vector Space Models of Word Meaning and Phrase Meaning: A Survey. *Language and Linguistics Compass*, 2012. URL https://doi.org/10.1002/lnco.362.

Evans, J. S. B. Logic and Human Reasoning: An Assessment of the Deduction Paradigm. *Psychological Bulletin*, 2002. URL https://doi.org/10.1037/0033-2909.128.6.978.

Fodor, J. A. *The Language of Thought*. Harvard University Press, 1975.

Fodor, J. A. and Pylyshyn, Z. W. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 1988. URL https://doi.org/10.1016/0010-0277(88)90031-5.

Furrer, D., van Zee, M., Scales, N., and Schärli, N. Compositional Generalization in Semantic Parsing: Pre-training vs. Specialized Architectures, 2020. URL https://arxiv.org/pdf/2007.08970.

Geva, M., Caciularu, A., Wang, K. R., and Goldberg, Y. Transformer Feed-Forward Layers Build Predictions by Promoting Concepts in the Vocabulary Space. *Empirical Methods in Natural Language Processing (EMNLP)*, 2022. URL https://arxiv.org/pdf/2203.14680.

Ghandeharioun, A., Caciularu, A., Pearce, A., Dixon, L., and Geva, M. Patchscopes: A Unifying Framework for Inspecting Hidden Representations of Language Models. *International Conference on Machine Learning (ICML)*, 2024. URL https://arxiv.org/pdf/2401.06102.

Griffiths, T. L., Lake, B. M., McCoy, R. T., Pavlick, E., and Webb, T. W. Whither symbols in the era of advanced neural networks?, 2025. URL https://arxiv.org/pdf/2508.05776.

Hernandez, E., Sharma, A. S., Haklay, T., Meng, K., Wattenberg, M., Andreas, J., Belinkov, Y., and Bau, D. Linearity of Relation Decoding in Transformer Language Models. *International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/pdf/2308.09124.

Hewitt, J. and Manning, C. D. A Structural Probe for Finding Syntax in Word Representations. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019. URL https://aclanthology.org/N19-1419.pdf.

Hupkes, D., Dankers, V., Mul, M., and Bruni, E. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research (JAIR)*, 2020. URL https://arxiv.org/pdf/1908.08351.

IMDb. IMDb Non-Commercial Datasets, 2024. URL https://developer.imdb.com/non-commercial-datasets/.

Johnson, J., Hariharan, B., Van Der Maaten, L., Fei-Fei, L., Lawrence Zitnick, C., and Girshick, R. CLEVR: A Diagnostic Dataset for Compositional Language and Elementary Visual Reasoning. *Computer Vision and Pattern Recognition (CVPR)*, 2017. URL https://arxiv.org/pdf/1612.06890.

Kahneman, D. and Tversky, A. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 1972. URL https://doi.org/10.1016/0010-0285(72)90016-3.

Keysers, D., Schärli, N., Scales, N., Buisman, H., Furrer, D., Kashubin, S., Momchev, N., Sinopalnikov, D., Stafiniak, L., Tihon, T., Tsarkov, D., Wang, X., van Zee, M., and Bousquet, O. Measuring Compositional Generalization: A Comprehensive Method on Realistic Data. *International Conference on Learning Representations (ICLR)*, 2020. URL https://arxiv.org/pdf/1912.09713.

Kim, N. and Linzen, T. COGS: A Compositional Generalization Challenge Based on Semantic Interpretation. *Empirical Methods in Natural Language Processing (EMNLP)*, 2020. URL https://arxiv.org/pdf/2010.05465.

Lake, B. and Baroni, M. Generalization without Systematicity: On the Compositional Skills of Sequence-to-Sequence Recurrent Networks. *International Conference on Machine Learning (ICML)*, 2018. URL https://arxiv.org/pdf/1711.00350.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences (BBS)*, 2017. URL https://arxiv.org/pdf/1604.00289.

Lampinen, A. K., Dasgupta, I., Chan, S. C. Y., Sheahan, H. R., Creswell, A., Kumaran, D., McClelland, J. L., and Hill, F. Language models, like humans, show content effects on reasoning tasks. *PNAS Nexus*, 2024. URL https://arxiv.org/pdf/2207.07051.

Lepori, M., Serre, T., and Pavlick, E. Break It Down: Evidence for Structural Compositionality in Neural Networks. *Neural Information Processing Systems (NeurIPS)*, 2023. URL https://arxiv.org/pdf/2301.10884.

Lepori, M. A., Mozer, M. C., and Ghandeharioun, A. Racing Thoughts: Explaining Contextualization Errors in Large Language Models. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2025. URL https://arxiv.org/pdf/2410.02102.

Li, Z., Jiang, G., Xie, H., Song, L., Lian, D., and Wei, Y. Understanding and Patching Compositional Reasoning in LLMs. *Findings of the Association for Computational Linguistics (ACL)*, 2024. URL https://arxiv.org/pdf/2402.14328.

Lu, M., Zhang, R., Eickhoff, C., and Pavlick, E. Paths Not Taken: Understanding and Mending the Multilingual Factual Recall Pipeline, 2025. URL https://arxiv.org/pdf/2505.20546.

Ma, Z., Hong, J., Gul, M. O., Gandhi, M., Gao, I., and Krishna, R. CREPE: Can Vision-Language Foundation Models Reason Compositionally? *Computer Vision and Pattern Recognition (CVPR)*, 2023. URL https://arxiv.org/pdf/2212.07796.

McCoy, R. T., Yao, S., Friedman, D., Hardy, M. D., and Griffiths, T. L. 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 (PNAS)*, 2024. URL https://arxiv.org/pdf/2309.13638.

McCurdy, K., Soulos, P., Smolensky, P., Fernandez, R., and Gao, J. Toward Compositional Behavior in Neural Models: A Survey of Current Views. *Empirical Methods in Natural Language Processing (EMNLP)*, 2024. URL https://aclanthology.org/2024.emnlp-main.524.pdf.

Merullo, J., Eickhoff, C., and Pavlick, E. Language Models Implement Simple Word2Vec-style Vector Arithmetic. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2024. URL https://arxiv.org/pdf/2305.16130.

Mikolov, T., Yih, W.-t., and Zweig, G. Linguistic Regularities in Continuous Space Word Representations. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2013. URL https://aclanthology.org/N13-1090.pdf.

Nguyen, K. A., Schulte im Walde, S., and Vu, N. T. Distinguishing Antonyms and Synonyms in a Pattern-based Neural Network. *European Chapter of the Association for Computational Linguistics (EACL)*, 2017. URL https://arxiv.org/pdf/1701.02962.

nostalgebraist. interpreting GPT: the logit lens, 2020. URL https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens.

Partee, B. H. *Compositionality in Formal Semantics*. Blackwell, 2004.

Press, O., Zhang, M., Min, S., Schmidt, L., Smith, N. A., and Lewis, M. Measuring and Narrowing the Compositionality Gap in Language Models. *Empirical Methods in Natural Language Processing (EMNLP)*, 2022. URL https://arxiv.org/pdf/2210.03350.

Quilty-Dunn, J., Porot, N., and Mandelbaum, E. The best game in town: The reemergence of the language-of-thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences (BBS)*, 2023. URL https://doi.org/10.1017/S0140525X22002849.

Russin, J., McGrath, S. W., Williams, D. J., and Elber-Dorozko, L. From Frege to chatGPT: Compositionality in language, cognition, and deep neural networks, 2024. URL https://arxiv.org/pdf/2405.15164.

Russin, J., Pavlick, E., and Frank, M. J. Parallel trade-offs in human cognition and neural networks: The dynamic interplay between in-context and in-weight learning. *Proceedings of the National Academy of Sciences (PNAS)*, 2025. URL https://www.pnas.org/doi/10.1073/pnas.2510270122.

Shalev, Y., Feder, A., and Goldstein, A. Distributional reasoning in LLMs: Parallel reasoning processes in multi-hop reasoning, 2024. URL https://arxiv.org/pdf/2406.13858.

Szabó, Z. G. Compositionality. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2024. URL https://plato.stanford.edu/archives/fall2024/entries/compositionality.

Tenney, I., Das, D., and Pavlick, E. BERT Rediscovers the Classical NLP Pipeline. *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2019. URL https://arxiv.org/pdf/1905.05950.

Tiedemann, J. and Thottingal, S. OPUS-MT – Building open translation services for the World. *European Association for Machine Translation (EAMT)*, 2020. URL https://aclanthology.org/2020.eamt-1.61.pdf.

Todd, E., Li, M. L., Sharma, A. S., Mueller, A., Wallace, B. C., and Bau, D. Function Vectors in Large Language Models. *International Conference on Learning Representations (ICLR)*, 2024. URL https://arxiv.org/pdf/2310.15213.

Vegner, I., de Souza, S., Forch, V., Lewis, M., and Doumas, L. A. A. Behavioural vs. Representational Systematicity in End-to-End Models: An Opinionated Survey. *Association for Computational Linguistics (ACL)*, 2025. URL https://arxiv.org/pdf/2506.04461.

Vig, J., Gehrmann, S., Belinkov, Y., Qian, S., Nevo, D., Singer, Y., and Shieber, S. Investigating Gender Bias in Language Models Using Causal Mediation Analysis. *Neural Information Processing Systems (NeurIPS)*, 2020. URL https://arxiv.org/pdf/2004.12265.

Vrandečić, D. and Krötzsch, M. Wikidata: a free collaborative knowledgebase. *Communications of the ACM (CACM)*, 2014. URL https://dl.acm.org/doi/10.1145/2629489.

Wang, B., Yue, X., Su, Y., and Sun, H. Grokked Transformers are Implicit Reasoners: A Mechanistic Journey to the Edge of Generalization. *Neural Information Processing Systems (NeurIPS)*, 2024. URL https://arxiv.org/pdf/2405.15071.

Wattenberg, M. and Viégas, F. B. Relational Composition in Neural Networks: A Survey and Call to Action. *Mechanistic Interpretability Workshop at ICML*, 2024. URL https://arxiv.org/pdf/2407.14662.

Yang, S., Gribovskaya, E., Kassner, N., Geva, M., and Riedel, S. Do Large Language Models Latently Perform Multi-Hop Reasoning? *Association for Computational Linguistics (ACL)*, 2024a. URL https://arxiv.org/pdf/2402.16837.

Yang, S., Kassner, N., Gribovskaya, E., Riedel, S., and Geva, M. Do Large Language Models Perform Latent Multi-Hop Reasoning without Exploiting Shortcuts? *Association for Computational Linguistics (ACL)*, 2024b. URL https://arxiv.org/pdf/2411.16679.

Yang, Z., Li, J., Xia, S., and Hu, X. Internal Chain-of-Thought: Empirical Evidence for Layer-wise Subtask Scheduling in LLMs. *Empirical Methods in Natural Language Processing (EMNLP)*, 2025. URL https://arxiv.org/pdf/2505.14530.

Yu, H., Jeong, M., Shin, J., Moon, H., Park, J., and Choi, S. Towards Zero-Shot Functional Compositionality of Language Models, 2023. URL https://arxiv.org/pdf/2303.03103.

Yu, Z., Belinkov, Y., and Ananiadou, S. Back Attention: Understanding and Enhancing Multi-Hop Reasoning in Large Language Models. *Empirical Methods in Natural Language Processing (EMNLP)*, 2025. URL https://arxiv.org/pdf/2502.10835.

Zhou, Y., Lake, B. M., and Williams, A. Compositional learning of functions in humans and machines. *Annual Meeting of the Cognitive Science Society (CogSci)*, 2024. URL https://arxiv.org/pdf/2403.12201.

A Data Creation

Table 2: List of our tasks, showing x, $g(x)$, and $f(g(x))$ for the random example in Table 1. Tasks with neither $g(x)$ nor $f(g(x))$ are omitted. $f(g(x))$ only shown if distinct from $g(f(x))$.

f	g	x	$g(x)$	$f(g(x))$
Word → Antonym	English → Spanish	bogus	false	—
Word → Antonym	English → German	philosophical	philosophisch	—
Word → Antonym	English → French	excessive	excessive	—
x + 10	2x	699	1398	1408
x + 100	2x	922	1844	1944
x mod 20	2x	891	1782	2
Word → Numeric	2x	one hundred and forty-eight	two hundred and ninety-six	—
Word[:-1]	Word[::-1]	responsible	elbisnopser	elbisnopse
Rotate(RGB, 120°)	RGB → Name	8a735a	dimgray	—

All tasks in Table 2 permit the additional computational pathway $x \to g(x) \to g(f(x))$. Those which don't list $f(g(x))$ are commutative and so $f(g(x)) = g(f(x))$ and applying f to $g(x)$ results in $g(f(x))$. The remaining tasks are not commutative, but their formal construction permits the hop $g(x) \to g(f(x))$ anyway. In particular, $g(f(x))$ equals $g(x)+20$ in `plus-10-times-2`, $g(x)+200$ in `plus-100-times-2`, $g(x) \bmod 40$ in `mod-20-times-2`, and $g(x)$[1:] in `word-truncate-reverse`.

A.1 Task Construction

Antonyms & Translations We obtain a list of antonyms from Todd et al. (2024) — further derived from Nguyen et al. (2017) — and obtain translations from Opus-MT (Tiedemann & Thottingal, 2020).

Factual Relations We obtain various factual relations from WikiData and IMDb Non-Commercial Datasets (Vrandečić & Krötzsch, 2014; IMDb, 2024; Bast & Buchhold, 2017). We apply a number of heuristics to obtain well-known and unambiguous mappings. For example, we filter entities by their "sitelinks" on WikiData or "votes" on IMDB (heuristics for popularity) to obtain well-known subjects. To avoid ambiguity, we identify subjects (songs, books, movies, people, etc.) with a single corresponding object (authors, attended universities, etc.). We omit parks and landmarks that exist in their country's capital. Our exact queries for generating each task can be found in our source code.

Arithmetic We use the range of numbers from 0 to 999 as x in our tasks. These numbers typically result in one token. We use the `num2words` library to obtain a mapping between words and numeric values. We use the list of antonyms as our list of words for the `word-truncate-reverse` task.

`Colors` In the `rgb-rotate-name` task, we randomly sample RGB colors, rotate them 120° by their hue, and map the resulting color value to that color's name (using the `webcolors` library and the common CSS 3 specification).

B Implementation Details

Examples & Prompts We prevent sampling of in-context examples that intersect in the variables $\{x, f(x), g(x), g(f(x)), f(g(x))\}$ with the query. And, as mentioned in Sec. 4.1, we exclude examples in Secs. 4 and 5 which overlap in the first token among their variables. So, although $x =$ "excessive" for the `antonym-french` task is listed in Table 2, this trivially shares the same first token as $g(x) =$ "excessive" and would be omitted from our analyses.

Our prompts are tokenized differently when predicting numbers or words, e.g. "`... \n A: 99`" results in `[ ][99]` whereas "`... \n A: modern`" results in `[ modern]`. We accordingly include the trailing space in our prompts when predicting numbers and omit it otherwise. We would then test for the single-token prediction of [99] and [modern] in this example.

Representational analysis In Sec. 4, we analyze the model's computation from $x \rightarrow g(f(x))$. Consider the query for "Heartbreak Hotel" $\rightarrow$ "1935": i.e. "`... Q: Heartbreak Hotel \n A: `". Here, multiple tokens (`[ Heart][break][ Hotel][ \][n][ A:][ ]`) are central to the computation. We therefore analyze all residual streams for these tokens. At each layer, we measure the signal for each variable by its maximum reciprocal rank across the streams. This procedure yields processing signatures, which quantify the presence of our variables at every layer.

We additionally represent each variable by its first token (since our decoding methods can only produce single-token probabilities). We exclude examples where different variables share the same first token and would be hard to differentiate. For example, $f(x) =$ " modern" and $g(f(x)) =$ " moderno" both share the first token `[ modern]`.

C Compositionality Gap by Prompt Template

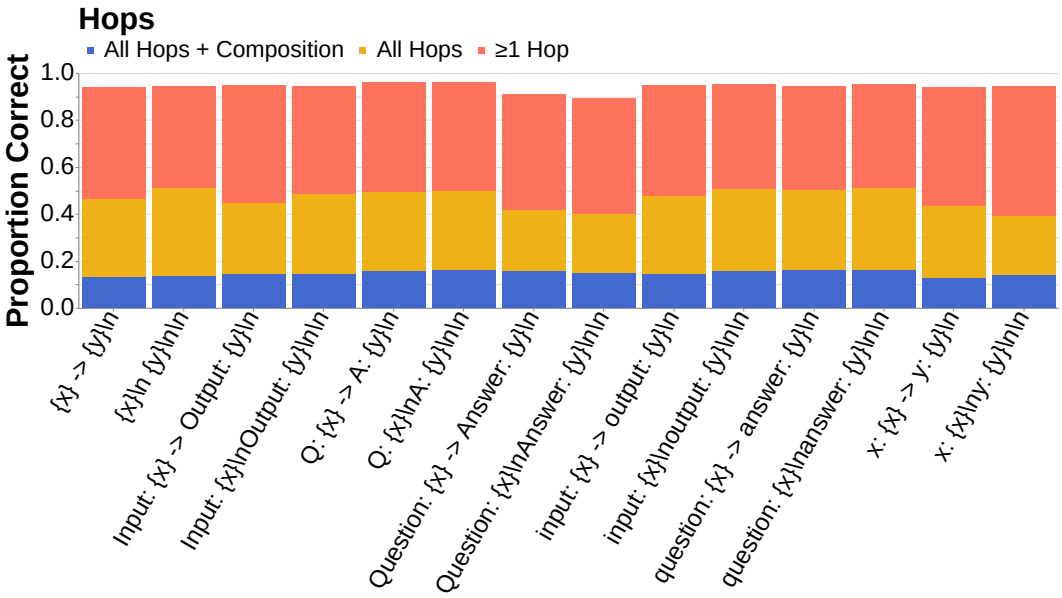

Figure 5: Ablation of Llama 3 (3B) model with different ICL prompt templates under the settings in Fig. 2. We find no significant difference in performance by choice of prompt.

D Compositionality Gap with Task Instructions

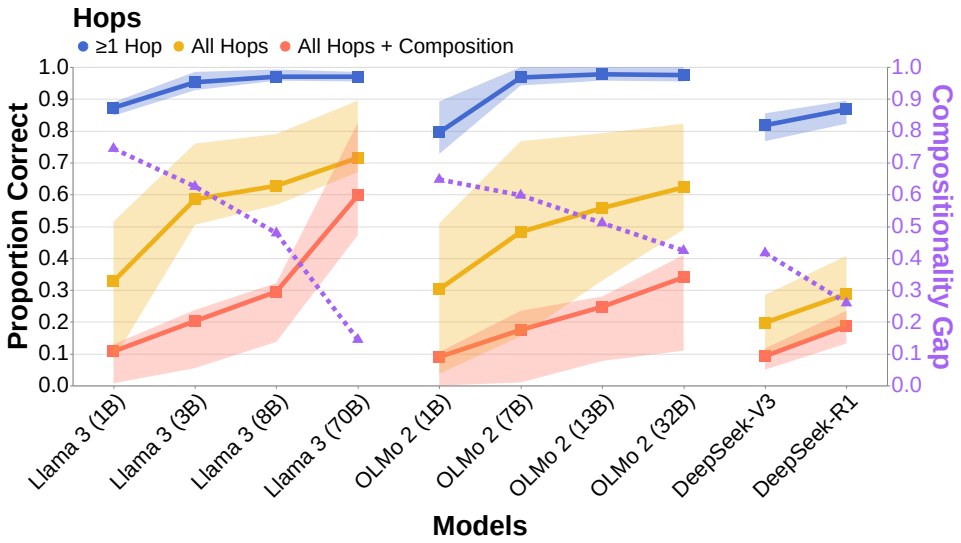

Figure 6: We repeat our experiments in Fig. 2, but always prefix the existing ICL prompt with a task instruction: e.g. "Add 100." for $x \to f(x)$ or "Add 100, multiply by 2." for $x \to g(f(x))$. We continue to find a clear compositionality gap in our results. We observe the gap diminishes more steeply w.r.t. model size or reasoning than in the original experiments. Llama 3 (70B) is also substantially closer to closing the gap. We omit the instruction-tuned Llama model and OpenAI models, whose outputs are harder to constrain (and thus evaluate) given the additional instruction prompt.

E Compositionality Gap by Size

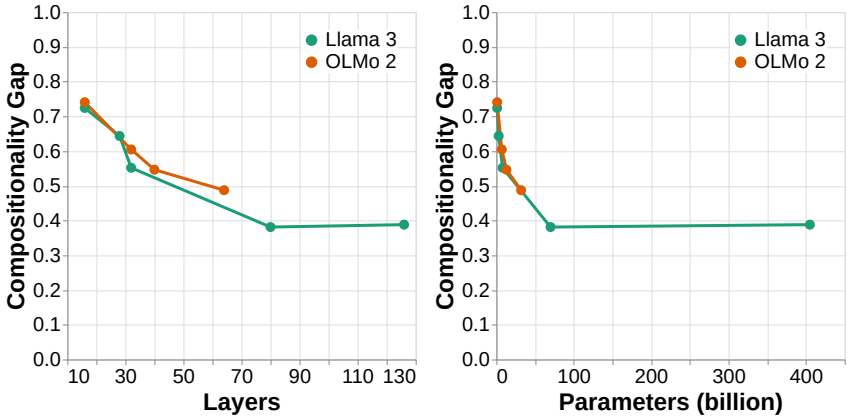

Figure 7: We illustrate the monotonically diminishing improvements to the compositionality gap resulting from increased model size (layers and parameters). We re-visualize results for the OLMo 2 and Llama 3 model families from Fig. 2.

F Compositionality Gap on Additional Models

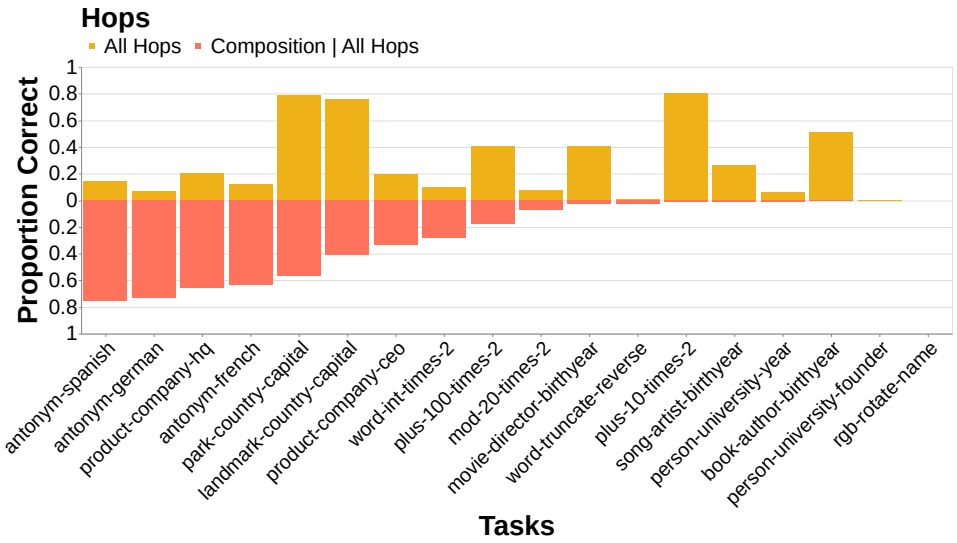

Figure 8: Compositionality gap for Llama 3 (3B) Instruct on all our tasks. Correlation between red and yellow bars is $r^2 = 0.00$.

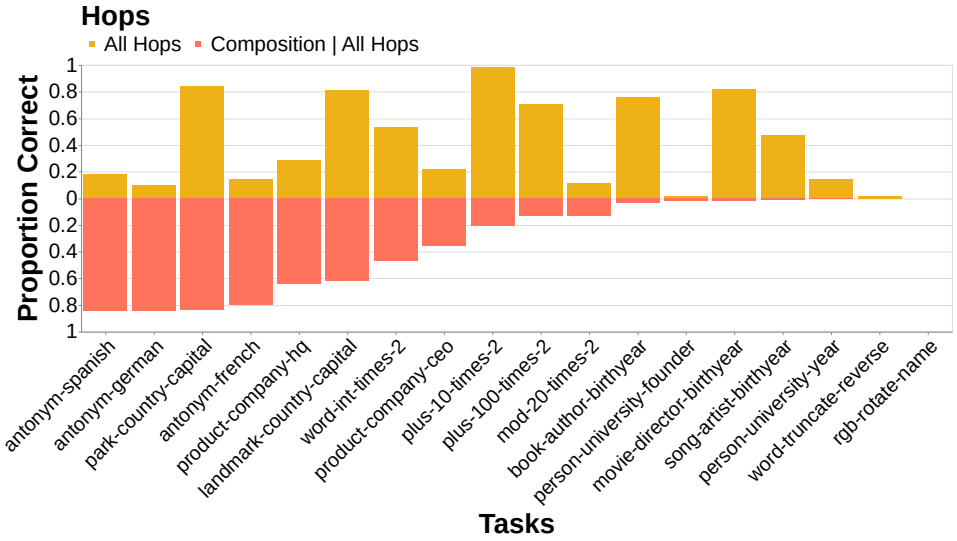

Figure 9: Compositionality gap for Llama 3 (8B) on all our tasks. Correlation between red and yellow bars is $r^2 = 0.00$.

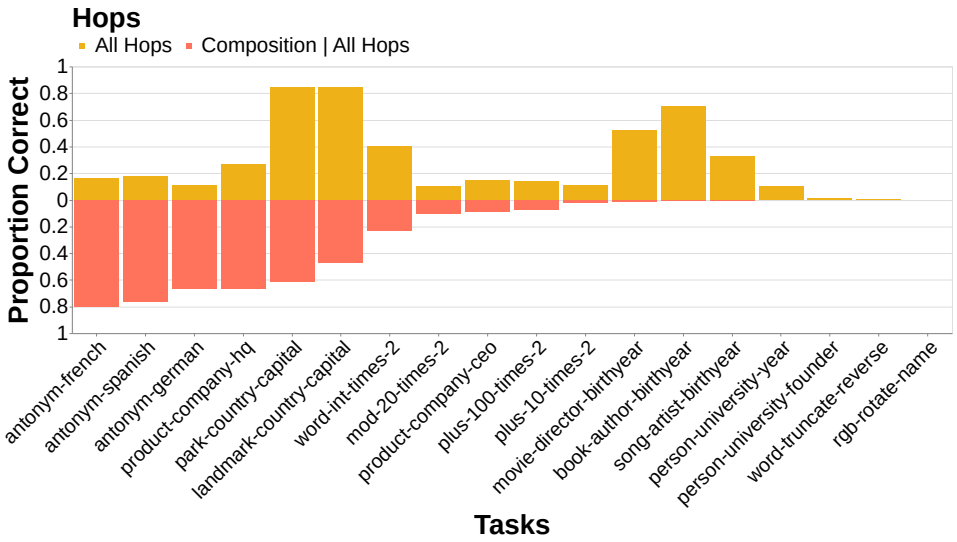

Figure 10: Compositionality gap for OLMo 2 (7B) on all our tasks. Correlation between red and yellow bars is $r^2 = 0.05$.

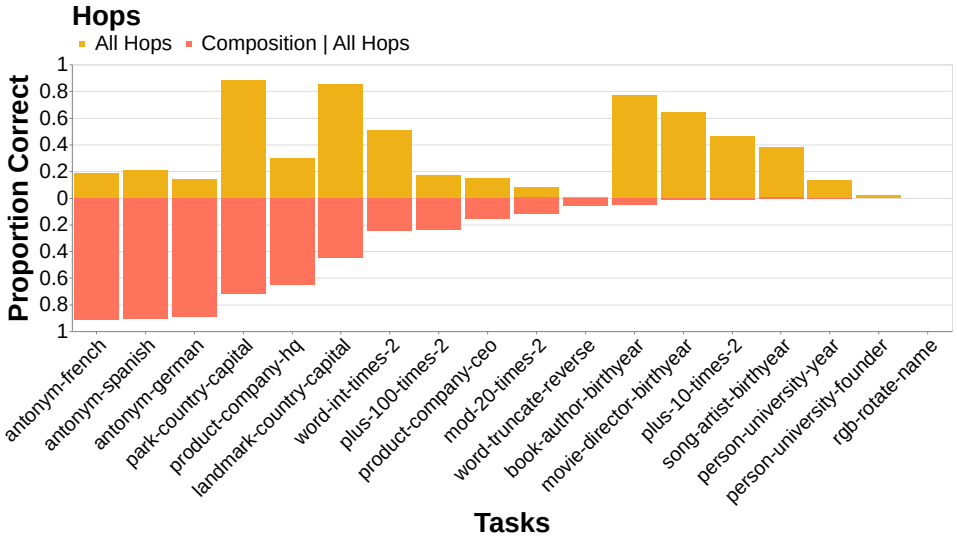

Figure 11: Compositionality gap for OLMo 2 (13B) on all our tasks. Correlation between red and yellow bars is $r^2 = 0.01$.

G Processing Signatures (Correct)

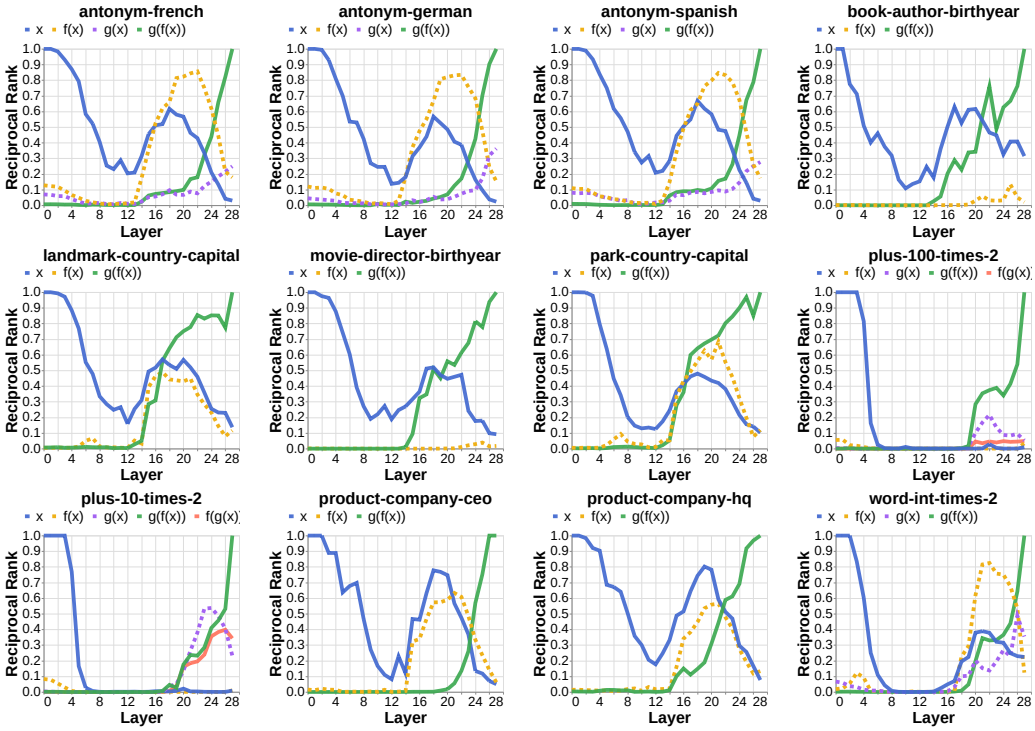

Figure 12: Aggregate processing signatures for each of our tasks, in which Llama 3 (3B) correctly solves all hops and the composition for at least 10 examples.

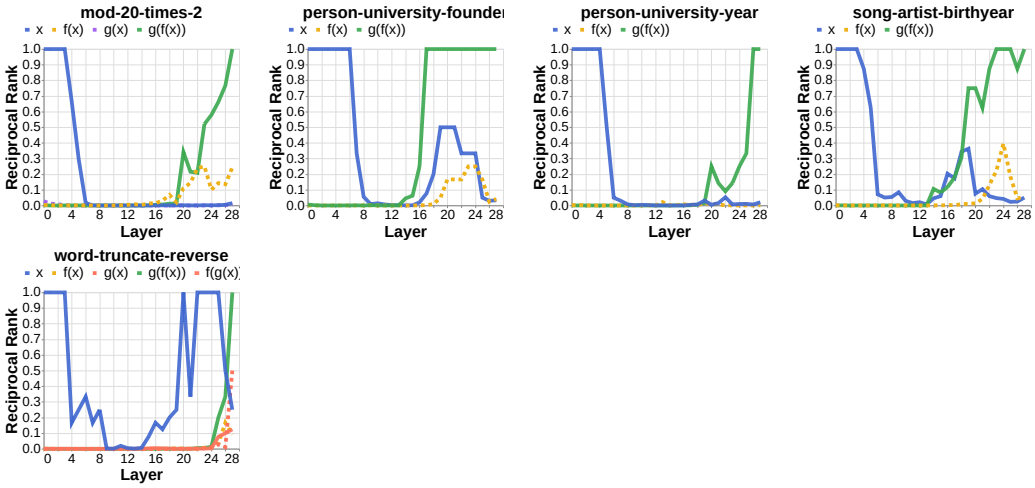

Figure 13: Aggregate processing signatures for each of our tasks, in which Llama 3 (3B) correctly solves all hops and the composition for less than 10 examples.

H Processing Signatures (Incorrect)

Although we see a difference in aggregate processing signatures (Figs. 3a and 3b), where the signal for the intermediate variables is clearer in the correct cases than the incorrect cases, this does not appear to be generally true (and is more likely due to data imbalances). We can see significant presence of the intermediate variables when considering incorrect examples, de-aggregated by task (Fig. 14).

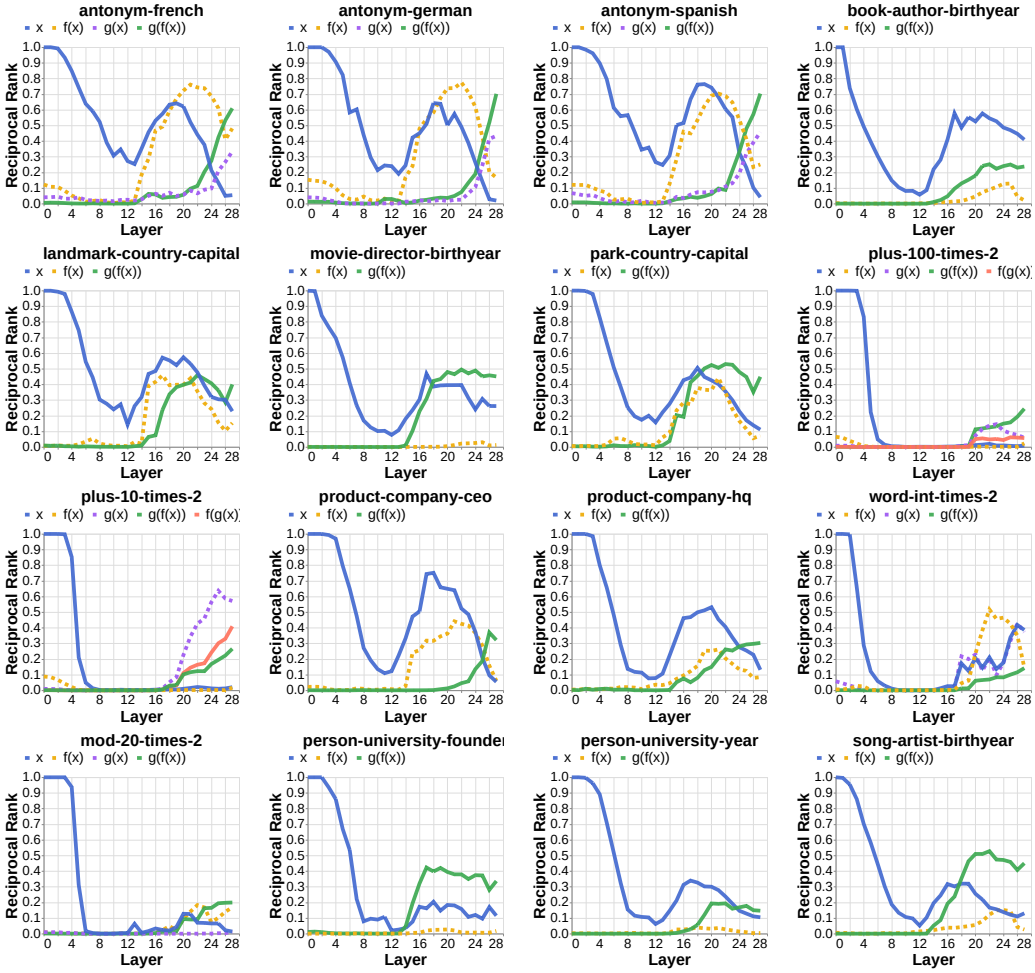

Figure 14: Aggregate processing signatures for each of our tasks, in which Llama 3 (3B) correctly solves all hops but not the composition for at least 10 examples.

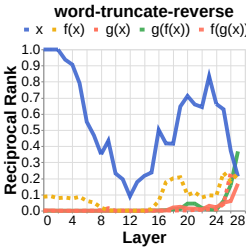

Figure 15: Aggregate processing signatures for each of our tasks, in which Llama 3 (3B) correctly solves all hops but not the composition for less than 10 examples.

I Token Identity Patchscope

Here, we repeat the analyses in Secs. 4 and 5, but use the token identity patchscope (Ghandeharioun et al., 2024) instead of logit lens. This patchscope is specifically designed (e.g. by prompt, patching strategy) to outperform logit lens at decoding the next token prediction and involves more of the model's computation in doing so. We should be aware that this design objective is not perfectly aligned with our purpose (decoding intermediate variables), but still offers a way to explore features stored in an alternative representational subspace.

We would specifically like to use this method to decode a representation into vocabulary-space logits. To do so, we prompt a model with the "token identity prompt", in which random tokens are repeated twice each, such as "`[A] [A] ; [B] [B] ; ... ; [?]`". We patch our representation of interest into the residual stream of this forward pass (at the corresponding layer and final token position). The language modeling logits resulting from our intervention then serve as the decoding for our representation.

We generally find similarities with our logit lens analyses: in tasks with "compositional" processing signatures, we continue to see growth of the signals for the intermediate variables with or before that for $g(f(x))$. Please zoom in to observe simultaneous growth, which may be difficult to see due to overlapping lines. And, although these plots may show growth of $f(x)$ and $g(f(x))$ in the same layers, recall that these computations can occur in different (e.g. earlier or later) residual streams (Appendix B).

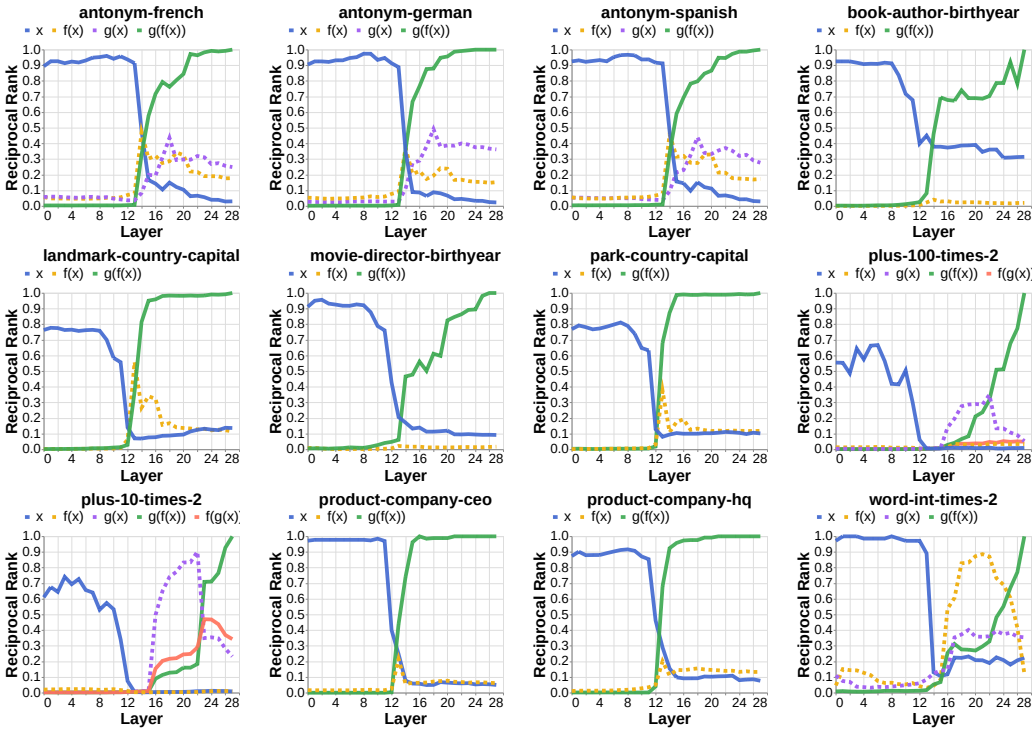

Figure 16: Aggregate processing signatures (using the token identity patchscope) for each of our tasks, in which Llama 3 (3B) correctly solves all hops and the composition for at least 10 examples.

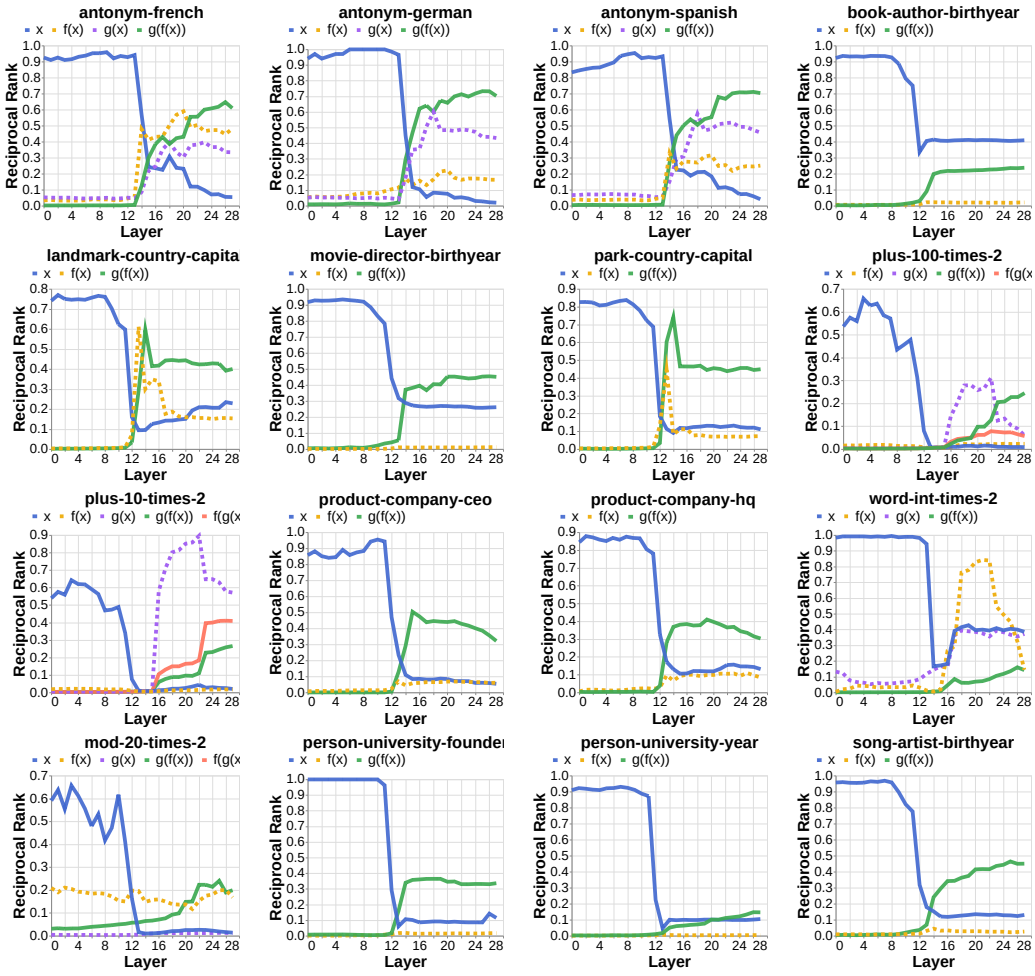

Figure 17: Aggregate processing signatures (using the token identity patchscope) for each of our tasks, in which Llama 3 (3B) correctly solves all hops but not the composition for at least 10 examples.

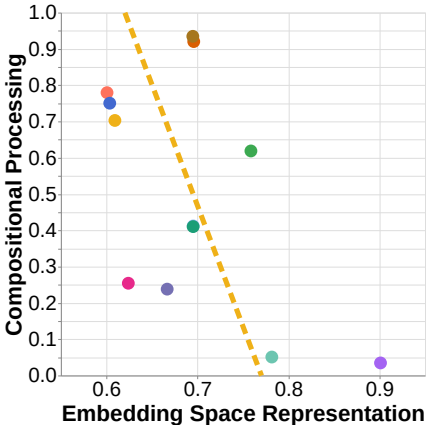

Figure 18: Correlation across tasks ($r^2 = 0.30$) for embedding space task representations and compositional processing (using the metric from the token identity patchscope). This correlation is weaker than in Fig. 4a, which might simply indicate that the feature of our interest (intermediate variables in our task) is better aligned with the representational subspace corresponding to logit lens.

J Causality of Intermediate Variables

We would like to determine whether the variables, $f(x)$ and $g(x)$, we identify in models' intermediate representations have a causal effect on the outcome. To do so, we use activation patching (Vig et al., 2020), a common method for conducting causal interventions in interpretability, and patch representations across tasks.

We first identify tasks with the same f but different g, such as `antonym-spanish` ($g \circ f$) and `antonym-german` ($g' \circ f$). For some x and x', we extract a single intermediate representation from the forward pass of $g'(f(x'))$ and patch it into the forward pass of $g(f(x))$. On average (over many x and x'), we measure the causal effects on the predictions $g(f(x))$, $g(f(x'))$, $g'(f(x))$, and $g'(f(x'))$.

We extract the representation from $g'(f(x'))$ at the position and layer where $f(x)$ or $g(x)$ have the highest reciprocal rank (and only use instances where this value is at least 0.5). We patch this representation into the forward pass for $g(f(x))$ at the median location where intermediate values are highest (layer 18 and 71st percentile query token position; identified among variables that reach RR ≥ 0.5). We apply this intervention to two groups: instances with intermediate values that reach a peak RR ≤ 0.2 and ≥ 0.5. In other words, instances with direct or compositional processing signatures.

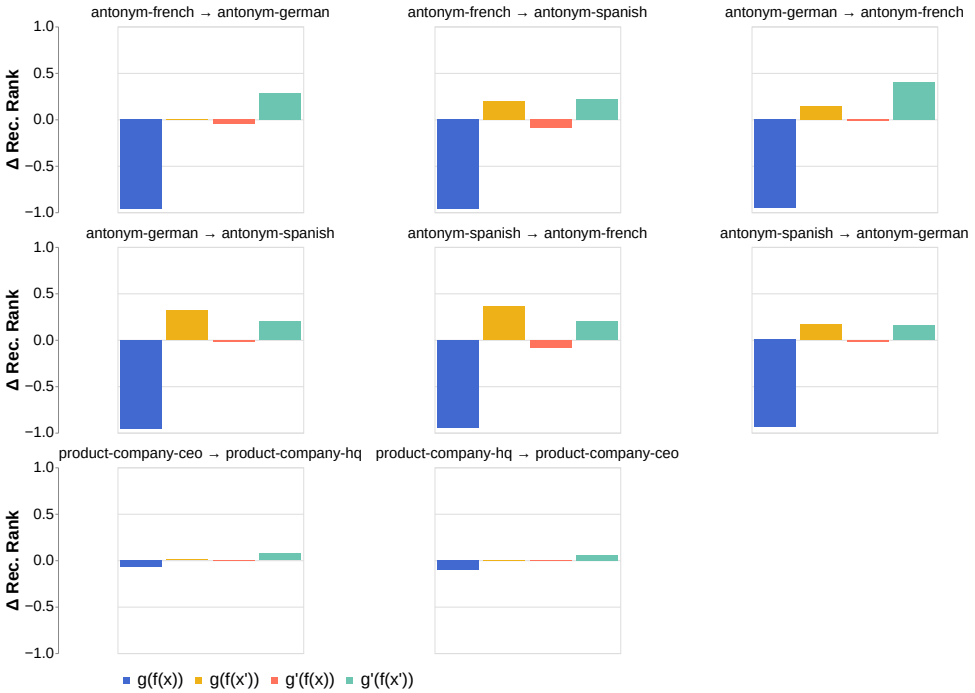

Figure 19: Causal effects on predicted values after patching from $g'(f(x'))$ to $g(f(x))$ for instances with compositional processing signatures.

Figure 20: Causal effects on predicted values after patching from $g'(f(x'))$ to $g(f(x))$ for instances with direct processing signatures.

The Antonym–Translation tasks (which tend to have compositional signatures) show the most significant causal effect: on average, $g(f(x))$ and $g'(f(x))$ decrease by -0.95 and -0.4, and $g(f(x'))$ and $g'(f(x'))$ increase by 0.20 and 0.24. The effect on $g(f(x'))$ clearly implicates the existence and causality of $f(x')$ in the patched activation; that on $g'(f(x'))$ indicates the additional existence of either itself or the function vector (Todd et al., 2024) for g' in that representation. The causal effects on compositional instances of `product-company-hq` and `product-company-ceo` are smaller.

But we can also see a clear difference between the causal effects on the compositional and direct instances. Indeed, the effects on `product-company-hq` and `product-company-ceo` are larger in their compositional instances. Patching activations from `plus-10-times-2` into `plus-100-times-2` primarily decreases $g(f(x))$ and increases $g'(f(x'))$, perhaps only implying the existence of the representation for $g'(f(x'))$ in the patched activation.

K Relationship between Mechanisms & Geometry on Additional Models

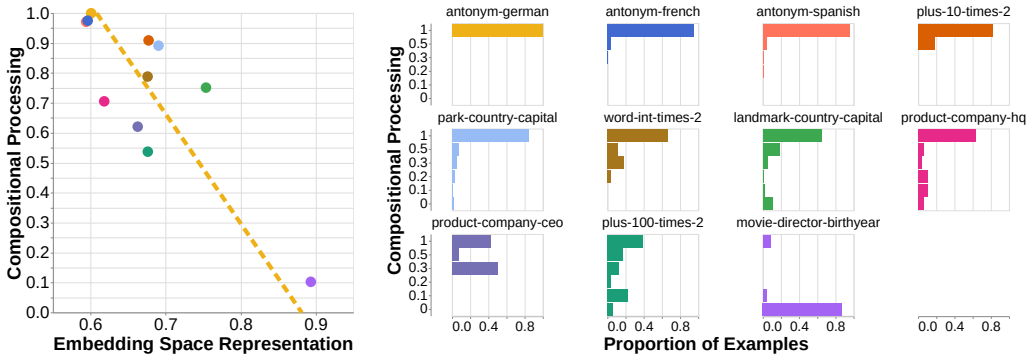

(a) (b)

Figure 21: For Llama 3 (3B) Instruct. Correlation has $r^2 = 0.68$. Accuracy is weakly correlated with these compositional processing ($r^2 = 0.14$) and representation ($r^2 = 0.05$) metrics. 81% of examples have very low or high indicators for the compositional mechanism.

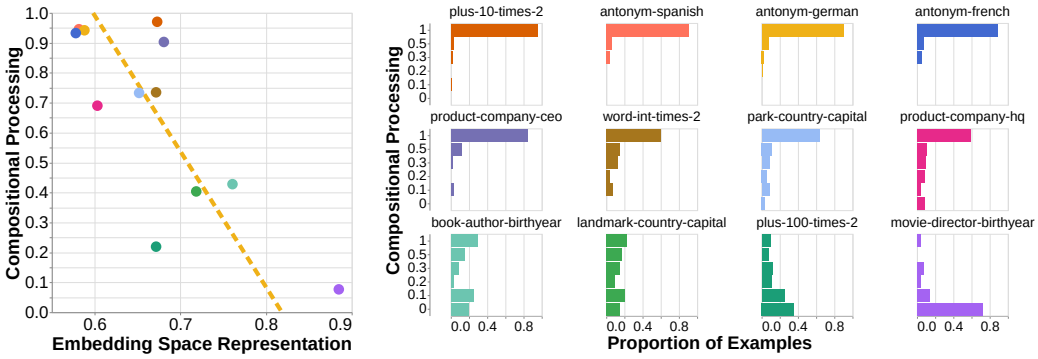

(a) (b)

Figure 22: For Llama 3 (8B). Correlation has $r^2 = 0.61$. Accuracy is weakly correlated with these compositional processing ($r^2 = 0.08$) and representation ($r^2 = 0.15$) metrics. 79% of examples have very low or high indicators for the compositional mechanism.

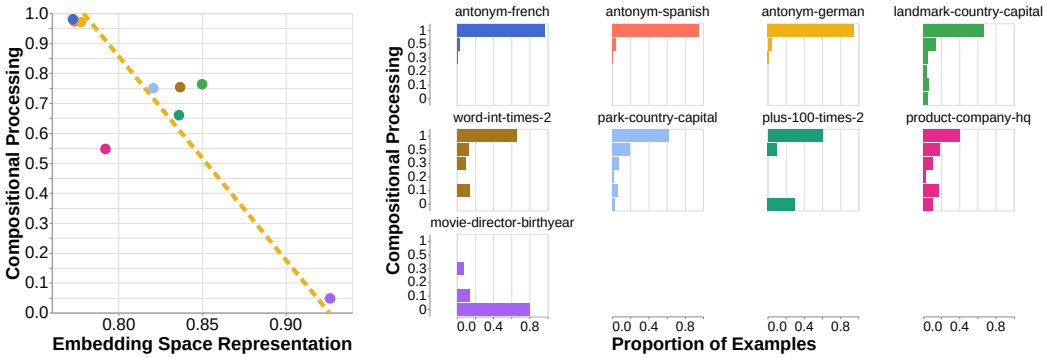

(a) (b)

Figure 23: For OLMo 2 (7B). Correlation has $r^2 = 0.75$. Accuracy is weakly correlated with these compositional processing ($r^2 = 0.14$) and representation ($r^2 = 0.13$) metrics. 85% of examples have very low or high indicators for the compositional mechanism.

(a) (b)

Figure 24: For OLMo 2 (13B). Correlation has $r^2 = 0.63$. Accuracy is weakly correlated with these compositional processing ($r^2 = 0.45$) and representation ($r^2 = 0.20$) metrics. 77% of examples have very low or high indicators for the compositional mechanism.

L Relationship between Mechanisms and Representations of Hops

Similarly to the experiment in Sec. 5 and Fig. 4a, we investigate the relationship between our compositionality heuristic and embedding space representations for variations of the hops (rather than of the compositional task) with the Llama 3 (3B) model. The representation of the second hop shown here also correlates strongly with the mechanism, but not as strongly as the translation-based hypothesis in Sec. 5.

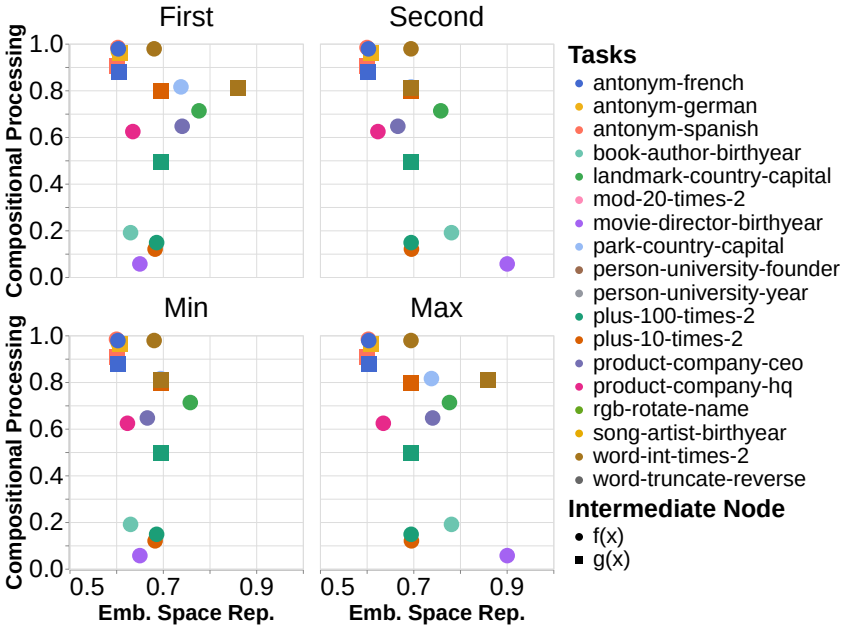

Figure 25: Relationships between compositional processing and embedding space representations for the hops. $r^2 = 0.02$ against the representation of the first hop; $r^2 = 0.51$ against the second hop; $r^2 = 0.09$ using the minimum representation between the hops; and $r^2 = 0.31$ using the maximum.

M Relationship between Mechanisms and Linear Task Representations

We repeat the experiments in Sec. 5 and Appendices I and L, but test for linear representations of tasks (instead of translations). Across models, we find an average correlation of $r^2 = 0.53$ between compositional task representations and mechanisms, which is also strong but not as promising as our translation-based relationship. Future work will need to identify under which conditions are tasks better represented with translations or linear transformations and which task representation models actually compute within their layer-wise processing.

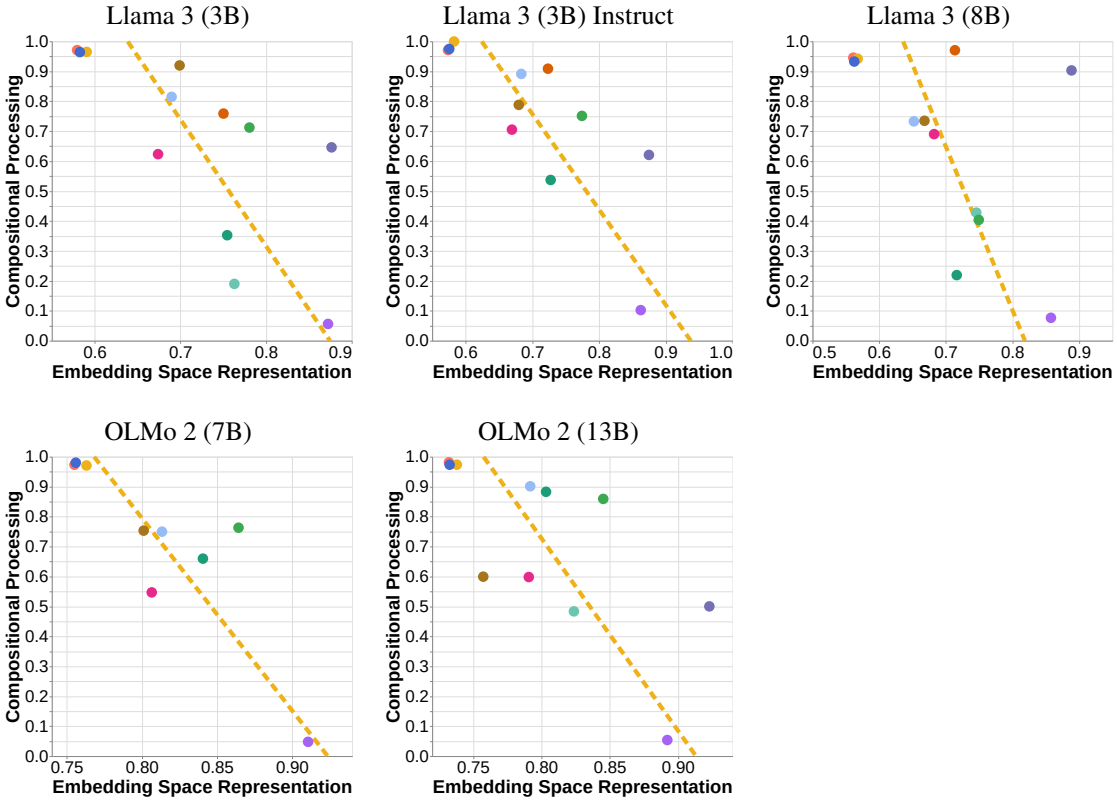

Figure 26: Llama 3 (3B) correlates with $r^2 = 0.52$, Llama 3 (3B) Instruct with $r^2 = 0.61$, Llama 3 (8B) correlates with $r^2 = 0.28$, OLMo 2 (7B) with $r^2 = 0.75$, and OLMo 2 (13B) with $r^2 = 0.50$.

Figure 27: Relationships between compositional processing and linear embedding space representations for the hops. $r^2 = 0.01$ against the representation of the first hop; $r^2 = 0.28$ against the second hop; $r^2 = 0.05$ using the minimum representation between the hops; and $r^2 = 0.20$ using the maximum.

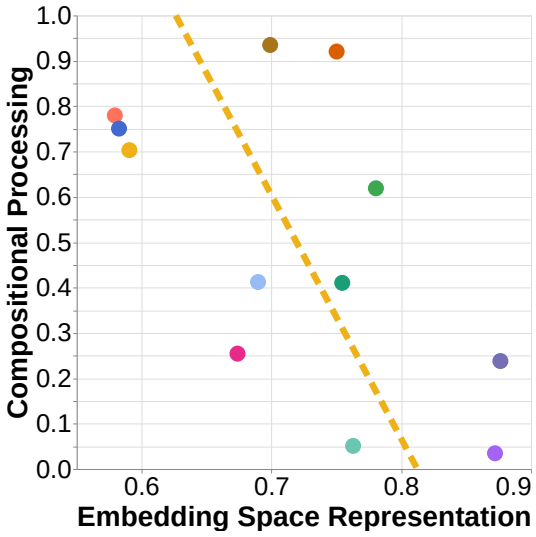

Figure 28: Correlation across tasks ($r^2 = 0.34$) for linear embedding space representations and compositional processing (using the token identity patchscope metric).