

# Evaluating Adjective-Noun Compositionality in LLMs: Functional vs Representational Perspectives

Ruchira Dhar<sup>†</sup> Qiwei Peng<sup>†</sup> Anders Søgaard<sup>†</sup>

<sup>†</sup>Department of Computer Science, University of Copenhagen

## Abstract

Compositionality is considered central to language abilities. As performant language systems, how do large language models (LLMs) do on compositional tasks? We evaluate adjective–noun compositionality in LLMs using two complementary setups: prompt-based functional assessment and a representational analysis of internal model states. Our results reveal a strong divergence between task performance and internal states. While LLMs reliably develop compositional representations, they fail to translate consistently into functional task success across model variants. Consequently, we highlight the importance of contrastive evaluation for obtaining a more complete understanding of model capabilities.

## 1 Introduction

Compositionality, the ability of systems to generate complex meanings from simpler parts, is often considered a key factor in human language processing (Smolensky, 1987; Fodor and Pylyshyn, 1988; Partee et al., 1995; Townsend et al., 2018; Donatelli and Koller, 2023). Today, large language models (LLMs) are considered highly performant natural language systems (Min et al., 2023). This has led to questions on the compositionality of such models (Saba, 2023; Mahowald et al., 2024; McCurdy et al., 2024): if compositionality is key to language, then are LLMs compositional as well?<sup>1</sup>

Early empirical work investigated compositionality in neural models with tasks like SCAN (Lake and Baroni, 2018), COGS (Kim and Linzen, 2020), and PCFG (Hupkes et al., 2020). However, these paradigms are a poor fit for modern pretrained LLMs: they rely on controlled train–test splits designed for models trained from scratch, and many use synthetic or restricted fragments of English that are likely already represented in large-

scale pretraining corpora. More recently, compositionality has been evaluated in LLMs through functional task-based assessments (Pavlick and Callison-Burch, 2016; Bertolini et al., 2022; Sun et al., 2024; Xu et al., 2024; Zhang et al., 2024; Dziri et al., 2024; Ma et al., 2024; Dhar and Søgaard, 2024), alongside methods aimed at improving performance on such tasks (SHAO et al., 2023; Drozdov et al., 2023; Press et al., 2023; Lu et al., 2024a; Chen et al., 2024; Li et al., 2024b). In parallel, another line of work probes internal representations to identify compositional structure and layer-wise mechanisms (Hewitt and Manning, 2019; Murty et al., 2023; Kumon and Yanaka, 2025; Peng et al., 2025; Guo et al., 2025). Yet these functional and representational approaches are rarely evaluated side by side on the same models and tasks, making it difficult to determine whether they provide converging or divergent evidence about compositionality in LLMs.

Recent work suggests external task performance and internal representations may diverge in LLMs (Lampinen et al., 2024; Orgad et al., 2025) — but does such divergence also apply to compositional abilities? In this work, we address this gap by directly contrasting functional and representational evaluation within a unified experimental framework<sup>2</sup>. We focus on adjective–noun compositionality and evaluate LLMs across three inference-based tasks capturing substitutivity, systematicity, and overgeneralization (Hupkes et al., 2020, 2023). For each model family, we measure task performance under scaling and instruction tuning and conduct a layer-wise analysis to assess compositional structure in internal representations. By comparing behavioral trends with representational signals across model variants, we seek to test whether the two paradigms yield converging empirical conclusions about compositionality in LLMs.

<sup>1</sup>Non-compositional LLMs have been argued to at least be conceivable (Block, 1981; Bender et al., 2021).

<sup>2</sup>Data and code are available at [Compenie](#).

## 2 Experimental Setup

To study compositionality, we focus on the domain of adjective–noun (AN) phrases. From a linguistic perspective, AN constructions are among the most frequent and semantically transparent instances of phrasal composition, making them a classic testbed for compositional semantics (Partee et al., 1995; Kamp and Partee, 1995; Guevara, 2010; Hartung et al., 2017; Bertolini et al., 2021). Crucially, meaning composition in this domain can be operationalized through inference: if the composed phrase preserves or alters entailment relations with its parts, the model must reason about semantic structure rather than rely on lexical associations (Lebowitz, 1983; Geiger et al., 2019; Donatelli and Koller, 2023). This inference-based framing has been central in both formal semantics (Partee et al., 1995) and recent NLP evaluations of compositionality (Pavlick and Callison-Burch, 2016; Bertolini et al., 2022).

### 2.1 Task Choice

Compositionality in natural language can manifest in many forms, such as substitutivity, systematicity, and overgeneralization (Hupkes et al., 2020). We adopt one adjective–noun (AN) based task on meaning inference for each of these aspects:

**Substitutivity.** Substitutivity requires that substituting a semantically compatible modifier preserves entailment (Pickel and Szabó, 2004; Werning et al., 2012; Hupkes et al., 2020; Li et al., 2024a). Let  $S(n)$  denote a sentence containing noun  $n$ , and  $S(an)$  its adjective-modified form. We define a binary entailment function:

$$E(S(an), S(n)) = \begin{cases} 1 & \text{if } \llbracket S(an) \rrbracket \models \llbracket S(n) \rrbracket \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

For example, entailment holds in *The runner set a new record*  $\models$  *The runner set a record*, since the adjective restricts but does not alter the core meaning. In contrast, entailment does not hold in *The alleged thief was arrested*  $\not\models$  *The thief was arrested*, as the adjective modifies the truth conditions of the noun phrase. We use the **AddOne** dataset (Pavlick and Callison-Burch, 2016) to evaluate substitutivity.

**Systematicity.** Systematicity requires recombining known semantic relations to infer a novel one

(Cummins, 1996; Symons and Calvo, 2014; Hupkes et al., 2020). Given two base entailments— $an \models n$  and  $an \models h$  (with  $h$  a hypernym of  $n$ )—and must decide the composed entailment  $an \models ah$ , testing whether the adjective can be systematically carried across the noun–hyponym relation. We define a binary entailment function over the target inference:

$$E(an, ah) = \begin{cases} 1 & \text{if } \llbracket an \rrbracket \models \llbracket ah \rrbracket \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

We then consider three inference types: IT1:  $E(an, n)$ , IT2:  $E(an, h)$ , and IT3:  $E(an, ah)$ .

For some adjectives, all inference types hold: *red car*  $\models$  *car*, *red car*  $\models$  *vehicle*, and *red car*  $\models$  *red vehicle*. While for others, IT1 and IT2 hold but IT3 fails: *small elephant*  $\models$  *elephant*, *small elephant*  $\models$  *animal*, but *small elephant*  $\not\models$  *small animal*.

We use the **PLANE** dataset (Bertolini et al., 2022) to evaluate systematic compositional inference.

**Overgeneralization.** Overgeneralization occurs when compositional rules are incorrectly extended to expressions whose meanings do not transparently derive from their parts (Macken, 2014; Hupkes et al., 2020; Csordás, 2023). In adjective–noun phrases, the modifier typically preserves type membership (e.g.,  $an$  denotes a subset of  $n$ ). However, exocentric compounds may contain the same noun string while failing to denote a subtype of that noun. This task tests whether models avoid inferring type membership based solely on surface form overlap. Given a triple  $(n, an, c)$ , where  $an$  is compositional and  $c$  is an exocentric compound containing  $n$ , we define a binary type inference function:

$$E(x, n) = \begin{cases} 1 & \text{if } \llbracket x \rrbracket \subseteq \llbracket n \rrbracket \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Compositional behavior predicts:  $E(an, n) = 1$  and  $E(c, n) = 0$ . For example, *trenchcoat*  $\models$  *coat*, whereas *turncoat*  $\not\models$  *coat*. We introduce the **COMPCOMB** dataset to evaluate robustness against such overgeneralization. For more details on the three aspects and how the tasks represent them, refer to [Appendix A](#).

### 2.2 Functional Evaluation

We describe the first of our two evaluation setups, which examines compositionality through

Dataset	LLaMA2			CodeLLaMA			Qwen2.5-Coder			Gemma2		
	7B	7B-it	13B	7B	7B-it	13B	7B	7B-it	14B	2B	2B-it	9B
<b>AddOne (F1)</b>	77.97	18.01	78.49	78.49	75.36	54.92	77.58	20.87	12.62	78.49	29.61	59.76
<b>PLANE (F1)</b>	36.70	33.73	35.35	33.33	42.67	33.33	33.33	34.57	38.14	33.33	33.33	36.84
<b>CompComb (Acc)</b>	66.67	72.59	45.93	95.56	71.11	93.33	77.04	73.34	71.11	77.04	73.33	45.19

Table 1: Functional task performance (weighted F1 for AddOne and PLANE; accuracy for COMPCOMB) for base, instruction-tuned (-it), and scaled variants of four LLM families. No model variant consistently outperforms others across all three tasks, reflecting the divergence between model scaling, alignment and compositional ability.

task-level performance. Here, models are evaluated based on their behavioral outputs on the three benchmarks introduced above. We examine how compositional task performance varies under two factors commonly associated with improved LLM capability: **model scaling**, where performance tends to increase with parameter count (Kaplan et al., 2020; Hoffmann et al., 2022; Biderman et al., 2023; Chowdhery et al., 2023), and **instruction tuning**, which has been shown to enhance generalization and alignment (Wei et al., 2022; Ouyang et al., 2022; Longpre et al., 2023; Chung et al., 2024; Sun et al., 2024). If compositional behavior tracks broader capability improvements, systematic trends should be observable across scaling and instruction-tuned variants.

**Methodology.** We evaluate four model families—LLaMA-2, CodeLLaMA, Qwen2.5-Coder, and Gemma2—each with three variants: a base model, an instruction-tuned model, and a larger scaled model (details in Appendix B). For AddOne and PLANE, we use two evaluation setup variations to mitigate effects of prompt sensitivity (Lu et al., 2024b; Sclar et al., 2024) and the reported results correspond to average accuracy across these variants. For the prompting setup, we employ accuracy and for log probability setup, given two candidate completions (entailment vs. non-entailment), the model’s prediction corresponds to the completion with higher conditional probability. For COMPCOMB, we employ three prompting configurations designed to test type inference robustness (details in Appendix C).

**Results.** Across the three benchmarks, performance trends vary across scaling and instruction tuning (Table 1, Figure 1). For AddOne and COMPCOMB, performance often decreases with instruction tuning and increased model size, whereas PLANE remains comparatively stable across vari-




Figure 1: The average performance across different model category (Base, Instruction Tuning, and Large model size) on three tasks. We report the weighted F1 score on AddOne and PLANE, and Accuracy on COMPCOMB.

ants. These findings indicate that improvements associated with scale and instruction tuning do not consistently correspond to improved compositional task performance under functional evaluation.

### 2.3 Representational Evaluation

In our second evaluation setup, we examine compositionality at the level of internal representations. We conduct a layer-wise probing analysis to determine whether compositional signals are encoded across model architectures and how these patterns relate to functional task outcomes.

**Methodology.** For each model, we extract hidden states representing encoding of the task question, from every fifth transformer layer. For AddOne and PLANE, we train a linear classifier on layer-wise representations to predict entailment labels. For COMPCOMB, we compute cosine similarity between token-level embeddings to assess type discrimination. All results are compared against a random baseline (details in Appendix C).

**Results.** Representations perform consistently above chance across models and layers (Fig-

ures 2–4), indicating that compositional information is encoded internally. For AddOne and PLANE, signal strength typically peaks in intermediate layers. Overall, representational evidence remains stable across scaling and instruction tuning.

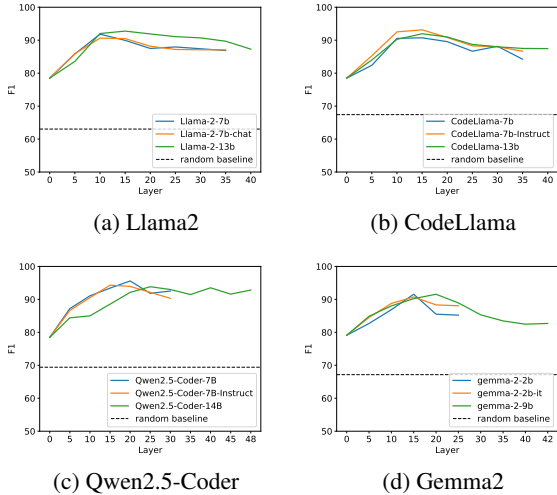


Figure 2: Layer-wise results (weighted F1 score) on AddOne dataset.

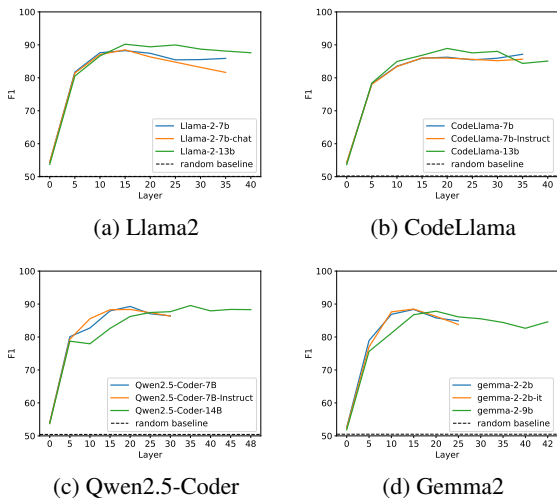


Figure 3: Layer-wise results (weighted F1 score) on PLANE dataset.

### 3 Conclusion

In this work, we contrasted functional and representational evaluations of adjective–noun compositionality in LLMs. *First*, we find that compositional task performance varies substantially across model variants. Although scaling and instruction tuning are typically associated with improved overall capability, they do not consistently improve compo-

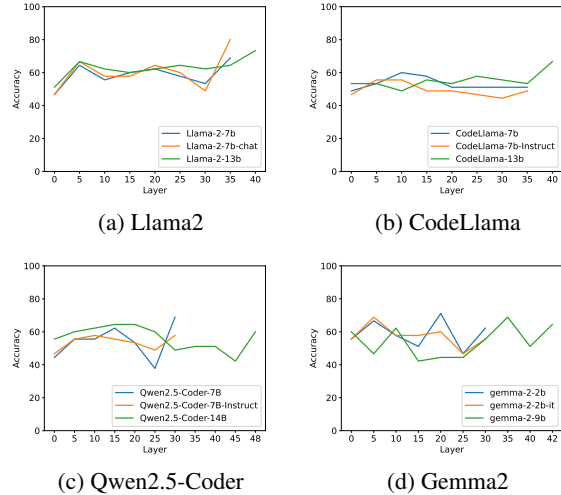


Figure 4: Layer-wise results (accuracy) on CompComb dataset.

sitional performance and in some cases lead to performance declines across benchmarks. This indicates that gains in general task performance do not uniformly translate to gains in compositional behavior under functional evaluation. *Second*, representational evidence remains comparatively stable across model families and variants. Compositional structure is reliably encoded in internal representations, with signals often peaking in intermediate layers and exhibiting similar trends across base, instruction-tuned, and scaled models. Taken together, these findings highlight a systematic divergence between functional and representational perspectives on compositionality. Compositional knowledge appears to be consistently encoded in LLMs, yet its behavioral expression depends on the model variants and evaluation setup. Thus, evaluating models solely through task performance or internal analysis may yield incomplete conclusions; integrating functional and representational approaches offers a more comprehensive view.

More broadly, the results underscore the importance of contrastive evaluation setups when assessing model capabilities. Relying solely on a single evaluation setup may provide an incomplete picture of model behavior. For robustness and safety-sensitive deployment (Mahilraj et al., 2023; Bender et al., 2021), understanding potential divergences between internal structure and external outputs is critical. A comprehensive evaluation strategy that integrates functional and representational analyses can therefore yield a more reliable and fine-grained account of model capabilities.

## Limitations

Our study has some limitations that constrain the scope of our conclusions. *First*, our evaluation focuses exclusively on adjective–noun compositionality. While this domain provides a controlled and linguistically grounded testbed, compositionality manifests at multiple levels of linguistic structure. Extending this contrastive evaluation framework to subword, phrasal, and sentence-level composition, as well as to multilingual models, would enable broader generalization of our findings.

*Second*, although we observe consistent representational signals, our analysis does not establish a causal relationship between internal compositional structure and task-level behavior. Layer-wise probing reveals the presence of information, but does not determine whether or how it is functionally used during inference. Future work could employ causal intervention methods—such as activation patching or path attribution—to test whether compositional representations directly influence model predictions.

*Finally*, our comparison is limited to two evaluation paradigms: prompt-based functional assessment and linear representational probing. Developing more comprehensive contrastive evaluation frameworks will be important for understanding how different analytical lenses shape conclusions about compositionality and other model capabilities.

## Ethical Considerations

We do not anticipate any risks in our work. Any data used does not contain any personal details or offensive content. In this study, our use of all artifacts is consistent with the license and intended use of all datasets and models. The AddOne Dataset (Pavlick and Callison-Burch, 2016) is available on HuggingFace on [this page](#) and is free for all use. The PLANE dataset (Bertolini et al., 2022) is also available on HuggingFace on [this page](#) and has license of Creative Commons Attribution 2.0. Our toy dataset COMPCOMB will be released upon publication along with details of license and intended use. All models used are also from HuggingFace and their details, along with licenses, are summarized in [Appendix B](#). All code and data will be released upon publication for greater reproducibility.

## Acknowledgments

We thank our colleagues at the CoAStAL NLP group for insightful discussions throughout this project. This research was supported by the Pioneer Centre for AI, DNRG grant number P1 and the Carlsberg Semper Ardens Advance Grant CF22-1432.

## References

- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Lorenzo Bertolini, Julie Weeds, and David Weir. 2022. Testing large language models on compositionality and inference with phrase-level adjective-noun entailment. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4084–4100.
- Lorenzo Bertolini, Julie Weeds, David Weir, and Qiwei Peng. 2021. [Representing syntax and composition with geometric transformations](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3343–3353, Online. Association for Computational Linguistics.
- Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, and 1 others. 2023. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR.
- Ned Block. 1981. [Psychologism and behaviorism](#). *Philosophical Review*, 90(1):5–43.
- Rudolf Carnap. 1988. *Meaning and necessity: A study in semantics and modal logic*, volume 30. University of Chicago Press.
- Jiaao Chen, Xiaoman Pan, Dian Yu, Kaiqiang Song, Xiaoyang Wang, Dong Yu, and Jianshu Chen. 2024. Skills-in-context: Unlocking compositionality in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13838–13890.
- Noam Chomsky. 1956. Three models for the description of language. *IRE Transactions on information theory*, 2(3):113–124.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and 1 others. 2023. Palm: Scaling

- language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, and 1 others. 2024. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53.
- Róbert Csordás. 2023. Systematic generalization in connectionist models.
- Robert Cummins. 1996. *Representations, targets, and attitudes*. MIT press.
- Ruchira Dhar and Anders Søgaard. 2024. From words to worlds: compositionality for cognitive architectures. *arXiv preprint arXiv:2407.13419*.
- Lucia Donatelli and Alexander Koller. 2023. Compositionality in computational linguistics. *Annual Review of Linguistics*, 9(1):463–481.
- Andrew Drozdov, Nathanael Schärli, Ekin Akyürek, Nathan Scales, Xinying Song, Xinyun Chen, Olivier Bousquet, and Denny Zhou. 2023. [Compositional semantic parsing with large language models](#). In *The Eleventh International Conference on Learning Representations*.
- Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Sean Welleck, Peter West, Chandra Bhagavatula, Ronan Le Bras, and 1 others. 2024. Faith and fate: Limits of transformers on compositionality. *Advances in Neural Information Processing Systems*, 36.
- Jerry A Fodor and Zenon W Pylyshyn. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1-2):3–71.
- Atticus Geiger, Ignacio Cases, Lauri Karttunen, and Christopher Potts. 2019. [Posing fair generalization tasks for natural language inference](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4485–4495, Hong Kong, China. Association for Computational Linguistics.
- Emiliano Raul Guevara. 2010. A regression model of adjective-noun compositionality in distributional semantics. In *Proceedings of the 2010 workshop on geometrical models of natural language semantics*, pages 33–37.
- Zhijin Guo, Chenhao Xue, Zhaozhen Xu, Hongbo Bo, Yuxuan Ye, Janet B Pierrehumbert, and Martha Lewis. 2025. Quantifying compositionality of classic and state-of-the-art embeddings. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 22130–22146.
- Matthias Hartung, Fabian Kaupmann, Soufian Jebbara, and Philipp Cimiano. 2017. Learning compositionality functions on word embeddings for modelling attribute meaning in adjective-noun phrases. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 54–64.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *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)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, and 3 others. 2022. Training compute-optimal large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS ’22*, Red Hook, NY, USA. Curran Associates Inc.
- Dieuwke Hupkes, Verna Dankers, Mathijs Mul, and Elia Bruni. 2020. Compositionality decomposed: How do neural networks generalise? *Journal of Artificial Intelligence Research*, 67:757–795.
- Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, and 1 others. 2023. A taxonomy and review of generalization research in nlp. *Nature Machine Intelligence*, 5(10):1161–1174.
- Hans Kamp and Barbara Partee. 1995. Prototype theory and compositionality. *Cognition*, 57(2):129–191.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Najoung Kim and Tal Linzen. 2020. [COGS: A compositional generalization challenge based on semantic interpretation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9087–9105, Online. Association for Computational Linguistics.
- Ryoma Kumon and Hitomi Yanaka. 2025. [Analyzing the inner workings of transformers in compositional generalization](#). In *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)*, pages 8529–8540, Albuquerque, New Mexico. Association for Computational Linguistics.

- Brenden Lake and Marco Baroni. 2018. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International conference on machine learning*, pages 2873–2882. PMLR.
- Andrew Kyle Lampinen, Stephanie C.Y. Chan, and Katherine Hermann. 2024. [Learned feature representations are biased by complexity, learning order, position, and more](#). *Transactions on Machine Learning Research*.
- Michael Lebowitz. 1983. Generalization from natural language text. *Cognitive Science*, 7(1):1–40.
- Chuanhao Li, Zhen Li, Chenchen Jing, Yuwei Wu, Mingliang Zhai, and Yunde Jia. 2024a. Compositional substitutivity of visual reasoning for visual question answering. In *European Conference on Computer Vision*, pages 143–160. Springer.
- Zhaoyi Li, Gangwei Jiang, Hong Xie, Linqi Song, Defu Lian, and Ying Wei. 2024b. [Understanding and patching compositional reasoning in LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9668–9688, Bangkok, Thailand. Association for Computational Linguistics.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, and 1 others. 2023. The flan collection: Designing data and methods for effective instruction tuning. In *International Conference on Machine Learning*, pages 22631–22648. PMLR.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2024a. Chameleon: Plug-and-play compositional reasoning with large language models. *Advances in Neural Information Processing Systems*, 36.
- Sheng Lu, Hendrik Schuff, and Iryna Gurevych. 2024b. How are prompts different in terms of sensitivity? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5833–5856.
- Teli Ma, Rong Li, and Junwei Liang. 2024. An examination of the compositionality of large generative vision-language models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 692–705.
- Marlys A Macken. 2014. Representation, rules, and overgeneralization in phonology. In *Mechanisms of language acquisition*, pages 367–397. Psychology Press.
- Jenifer Mahilraj, Manivel Pandian, Muthuraman Subbiah, S Kalyan, and 1 others. 2023. Evaluation of the robustness, transparency, reliability and safety of ai systems. In *2023 9th International Conference on Advanced Computing and Communication Systems (ICACCS)*, volume 1, pages 2526–2535. IEEE.
- Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2024. Dissociating language and thought in large language models. *Trends in Cognitive Sciences*.
- Kate McCurdy, Paul Soulos, Paul Smolensky, Roland Fernandez, and Jianfeng Gao. 2024. [Toward compositional behavior in neural models: A survey of current views](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9323–9339, Miami, Florida, USA. Association for Computational Linguistics.
- Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. 2023. [Recent advances in natural language processing via large pre-trained language models: A survey](#). *ACM Comput. Surv.*, 56(2).
- Shikhar Murty, Pratyusha Sharma, Jacob Andreas, and Christopher D Manning. 2023. [Characterizing intrinsic compositionality in transformers with tree projections](#). In *The Eleventh International Conference on Learning Representations*.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. 2025. [LLMs know more than they show: On the intrinsic representation of LLM hallucinations](#). In *The Thirteenth International Conference on Learning Representations*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Peter Pagin and Dag Westerståhl. 2010a. Compositionality ii: Arguments and problems. *Philosophy Compass*, 5(3):265–282.
- Peter Pagin and Dag Westerståhl. 2010b. Pure quotation and general compositionality. *Linguistics and Philosophy*, 33:381–415.
- Barbara Partee and 1 others. 1995. Lexical semantics and compositionality. *An invitation to cognitive science: Language*, 1:311–360.
- Ellie Pavlick and Chris Callison-Burch. 2016. [Most “babies” are “little” and most “problems” are “huge”:](#) Compositional entailment in adjective-nouns. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2164–2173, Berlin, Germany. Association for Computational Linguistics.
- Qiwei Peng, Yekun Chai, and Anders Søgaard. 2025. Understanding subword compositionality of large

- language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22535–22546.
- Bryan Pickel and Zoltán Gendler Szabó. 2004. Compositionality.
- Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt, Noah Smith, and Mike Lewis. 2023. [Measuring and narrowing the compositionality gap in language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5687–5711, Singapore. Association for Computational Linguistics.
- Geoffrey K Pullum and Barbara C Scholz. 2010. Recursion and the infinitude claim. In *Recursion and human language*, pages 113–137. De Gruyter Mouton.
- Walid S Saba. 2023. Stochastic llms do not understand language: towards symbolic, explainable and ontologically based llms. In *International conference on conceptual modeling*, pages 3–19. Springer.
- Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2024. [Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting](#). In *The Twelfth International Conference on Learning Representations*.
- NAN SHAO, Zefan Cai, Hanwei xu, Chonghua Liao, Yanan Zheng, and Zhilin Yang. 2023. [Compositional task representations for large language models](#). In *The Eleventh International Conference on Learning Representations*.
- Paul Smolensky. 1987. Connectionist ai, symbolic ai, and the brain. *Artificial Intelligence Review*, 1(2):95–109.
- Jiuding Sun, Chantal Shaib, and Byron C Wallace. 2024. [Evaluating the zero-shot robustness of instruction-tuned language models](#). In *The Twelfth International Conference on Learning Representations*.
- John Symons and Paco Calvo. 2014. Systematicity: an overview. *P. Calvo & J. Symons, (Eds.)*.
- Simon W Townsend, Sabrina Engesser, Sabine Stoll, Klaus Zuberbühler, and Balthasar Bickel. 2018. Compositionality in animals and humans. *PLoS Biology*, 16(8):e2006425.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. [Finetuned language models are zero-shot learners](#). In *International Conference on Learning Representations*.
- Markus Werning, Wolfram Hinzen, and Edouard Machery. 2012. *The Oxford handbook of compositionality*. Oxford University Press.
- Zhuoyan Xu, Zhenmei Shi, and Yingyu Liang. 2024. Do large language models have compositional ability? an investigation into limitations and scalability. In *ICLR 2024 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- Min Zhang, Jianfeng He, Shuo Lei, Murong Yue, Linhan Wang, and Chang-Tien Lu. 2024. Can llm find the green circle? investigation and human-guided tool manipulation for compositional generalization. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 11996–12000. IEEE.

## A Tasks and Datasets

In this appendix section, we provide a detailed explanation of the different aspects of compositionality we evaluate (Hupkes et al., 2020) along with the task used.

**Substitutivity:** It refers to a model’s ability to preserve meaning when a component of a complex expression is replaced with a semantically equivalent or synonymous alternative (Carnap, 1988; Hupkes et al., 2020). In a truly compositional system, meaning should be a function of the meanings of its parts and their combination – so substituting one part with a synonym should not alter the overall meaning of the expression. We evaluate this using the **AddOne Dataset** (Pavlick and Callison-Burch, 2016), which tests whether models recognize entailment relationships between original and modified sentences. For example, given the sentence *The runner set a record*, we substitute the noun *record* with an adjective–noun phrase such as *new record*, producing the modified sentence *The runner set a new record*. While the surface form changes, the core meaning remains unchanged in context. A model with compositional understanding should therefore recognize the entailment:

$$\textit{The runner set a new record} \models \textit{The runner set a record}.$$

Success on this task requires more than detecting synonymy at the word level – it requires the model to integrate the substitution compositionally and infer that the meaning of the larger sentence remains intact. The dataset is released with train (4481 samples), val (387 samples), and test (510 samples) splits.

**Systematicity:** It refers to a model’s ability to recombine known parts and rules to interpret novel expressions (Fodor and Pylyshyn, 1988; Hupkes et al., 2020). We evaluate this using the **PLANE Dataset** (Bertolini et al., 2022), which focuses on

adjective–noun (AN) entailment patterns. In this task, the model is provided with two base entailments: one between an adjective–noun phrase and the noun alone (e.g., *red car*  $\models$  *car*), and one between the noun and its hypernym (*car*  $\models$  *vehicle*). The model is then asked to determine the validity of the composed entailment (*red car*  $\models$  *red vehicle*). The dataset is released with train (60012 samples) and test (2016 samples) splits.

To succeed, the model must apply *systematic compositional reasoning*: it must combine the entailment patterns it has already observed to infer the new one. Crucially, this requires understanding how the adjective interacts with both the noun and the hypernym class – particularly when adjectives vary in type (e.g., intersective vs subsective vs intensional). The task thus tests whether the model can generalize meaning compositionally, rather than relying on shallow heuristics or memorized patterns.

**Over-generalization:** It refers to a model’s tendency to incorrectly apply compositional rules to expressions that are, in fact, non-compositional (Hupkes et al., 2020). A robust compositional system should be able to distinguish between genuine compositional combinations and compounds that do not follow regular meaning composition. We test this with a novel task using a handcrafted toy dataset of 45 samples in English language—the **COMPCOMB Dataset**, which we introduce in this work. Each data point consists of a triple: a noun, a compositional adjective–noun (AN) combination involving that noun, and a non-compositional exocentric compound that includes the same noun. For example, in the triple (*coat*, *trenchcoat*, *turncoat*), the word *trenchcoat* refers to a specific type of *coat*, preserving a compositional relationship. In contrast, *turncoat* is an idiomatic expression referring to a traitor, where the meaning does not transparently derive from *coat*.

To succeed, the model must avoid overgeneralizing based on surface forms or string overlap, and instead distinguish between compositional and non-compositional expressions. This task directly tests whether the model can identify and respect the boundaries of compositional generalization. Please note that we will release the dataset upon publication along with the license and intended use for research purposes.

It is important to note that two other dimensions as per Hupkes et al. (2020)—*productivity* and *localism*—are deliberately excluded from our analyses,

for reasons rooted in both linguistic theory and the epistemological limitations of empirical verification which are commonly recognized in the literature (Hupkes et al., 2020).

**First**, productivity—the notion that natural language allows for the generation of an unbounded number of novel expressions from a finite set of primitives—is a foundational concept in generative linguistics (Chomsky, 1956). However, its empirical status is more controversial. As noted by Pullum and Scholz (2010), it is not possible to establish the infinitude of natural language through observation, given the finiteness of human utterances and cognitive resources. Moreover, even if a system generalizes to longer sequences than it was trained on, this does not, in itself, constitute evidence for true linguistic productivity. In formal terms, productivity involves generalization across an unbounded domain:

$$\forall n \in \mathbb{N}, \exists s_n \in \mathcal{L} \text{ such that } \text{length}(s_n) = n,$$

where  $\mathcal{L}$  is the set of grammatical sentences. Demonstrating this for either human language or LLM behavior is inherently intractable. Consequently, productivity lacks the falsifiability and granularity required for robust evaluation in our setting.

**Second**, we exclude localism as a defining criterion of compositionality. Localism refers to the idea that composition should occur via strictly local operations—i.e., the meaning of a complex expression is derived solely from the meanings of its immediate subparts (its “children” in a compositional tree). This notion is aligned with what the relevant linguistic literature Pagin and Westerstähl (2010a,a,b) terms *strong* or *first-level* compositionality. However, this is only one possible interpretation. Under *weak* or *global* compositionality, meanings may be computed from broader structures, potentially referencing distant context or higher-order structures. Crucially, both modes of composition satisfy the formal definition of compositionality:

$$\llbracket E \rrbracket = f(\llbracket E_1 \rrbracket, \dots, \llbracket E_n \rrbracket),$$

where  $\llbracket E \rrbracket$  denotes the meaning of expression  $E$ , and  $E_1, \dots, E_n$  are its constituents. The nature of  $f$ —whether strictly local or globally informed—is not constrained by the principle of compositionality itself. Empirically, the absence of local compositional operations does not entail non-compositional

behavior. A model may validly construct meanings using global strategies and still preserve the structural integrity of compositional mappings. Thus, testing for localism does not test for compositionality per se, but rather for a specific implementation strategy.

For these reasons, we limit our analysis to aspects of compositionality that are both theoretically robust and empirically interpretable. We regard productivity and localism as valuable extensions for future work, but not as necessary conditions for evaluating whether compositionality matters in current large language models.

## B Model Details

In this appendix section, we provide details of the models used in our experiments. For each model family, we selected three models for testing:

- **Base Model (Base):** A minimum parameter model, ranging from 2B to 7B, that serves as the foundational model version for each family.
- **Instruction – Finetuned Model (IFT):** The same 7B base model, further fine-tuned with instruction-tuning to enhance task performance.
- **Scaled Model (Large):** A model variant with a higher parameter count, ranging from 9B to 70B, depending on availability within each family. These larger models allow us to investigate how scaling affects compositional behavior.

The diversity in models ensures that our analysis captures how both model complexity and tuning approaches impact compositionality. The names and links to each model are in [Table 2](#)

## C Methodology Details

This appendix provides formal details for the two evaluation setups described in Sections 2.2 and 2.3.

### C.1 Functional Evaluation Details

**AddOne and PLANE.** For binary entailment tasks, we use two evaluation variants.

**(1) Generation-based MCQ setup.** Given an input prompt  $x$  and two candidate options (entailment and non-entailment), the model generates an

output sequence  $\hat{y}$ . A prediction is considered correct if the generated output contains the correct option label. Accuracy is computed as:

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{y}_i \text{ contains } y_i], \quad (4)$$

where  $N$  is the number of test examples and  $y_i$  is the gold label.

**(2) Comparative log-probability setup.** Given the same prompt  $x$  and two candidate completions  $y^{(1)}$  (entailment) and  $y^{(0)}$  (non-entailment), we compute conditional log-probabilities:

$$\log P(y^{(j)} | x) = \sum_{t=1}^{T_j} \log P(y_t^{(j)} | x, y_{<t}^{(j)}), \quad (5)$$

where  $T_j$  is the token length of completion  $y^{(j)}$ . The predicted label is:

$$\hat{y} = \arg \max_{j \in \{0,1\}} \log P(y^{(j)} | x). \quad (6)$$

Accuracy is computed as the proportion of correct predictions over the test set.

**COMPCOMB.** For COMPCOMB, we use a prompting setup in which the model is asked to judge which expression is semantically closer to the noun  $n$ : the compositional adjective–noun phrase  $an$  or the exocentric compound  $c$ .

We evaluate three prompt variants to mitigate prompt sensitivity. For each variant  $p_j$ , the model outputs a choice  $\hat{z}^{(j)} \in \{an, c\}$ . Accuracy per variant is defined as:

$$\text{Acc}(p_j) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{z}_i^{(j)} = an], \quad (7)$$

where  $N$  is the number of triples.

Final reported accuracy is averaged across prompt variants:

$$\text{Acc} = \frac{1}{m} \sum_{j=1}^m \text{Acc}(p_j), \quad (8)$$

where  $m$  denotes the number of prompt templates. A prediction is considered correct if the model judges  $an$  to be closer to  $n$  than  $c$ .

Family	Variant	Hugging Face Link	License
LLaMA-2	Base (7B)	<a href="#">Llama-2-7b-hf</a>	LLaMA 2 Community License
	Instruction-tuned (7B)	<a href="#">Llama-2-7b-chat-hf</a>	LLaMA 2 Community License
	Scaled (13B)	<a href="#">Llama-2-13b-hf</a>	LLaMA 2 Community License
CodeLlama	Base (7B)	<a href="#">CodeLlama-7b-hf</a>	LLaMA 2 Community License
	Instruction-tuned (7B)	<a href="#">CodeLlama-7b-Instruct-hf</a>	LLaMA 2 Community License
	Scaled (13B)	<a href="#">CodeLlama-13b-hf</a>	LLaMA 2 Community License
Qwen2.5-Coder	Base (7B)	<a href="#">Qwen2.5-Coder-7B</a>	Apache License 2.0
	Instruction-tuned (7B)	<a href="#">Qwen2.5-Coder-7B-Instruct</a>	Apache License 2.0
	Scaled (14B)	<a href="#">Qwen2.5-Coder-14B</a>	Apache License 2.0
Gemma2	Base (2B)	<a href="#">Gemma-2-2B</a>	Gemma Terms of Use
	Instruction-tuned (2B)	<a href="#">Gemma-2-2B-it</a>	Gemma Terms of Use
	Scaled (9B)	<a href="#">Gemma-2-9B</a>	Gemma Terms of Use

Table 2: Models used in [section 2](#), grouped by family and variant, with links to their Hugging Face repositories and associated licenses.

## C.2 Representational Evaluation Details

**Layer-wise Extraction.** For each model, we extract hidden representations corresponding to the task input from every fifth transformer layer. Let  $\phi_\ell(x)$  denote the representation of input  $x$  at layer  $\ell$ .

**AddOne and PLANE.** For each layer  $\ell$ , we train a linear classifier on  $\phi_\ell(x)$  to predict the binary entailment label. Let  $\hat{y}_\ell$  denote the predicted label at layer  $\ell$ . Layer-wise accuracy is computed as:

$$\text{Acc}_\ell = \frac{1}{N} \sum_{i=1}^N \mathbb{1}[\hat{y}_{\ell,i} = y_i], \quad (9)$$

where  $N$  is the number of test examples and  $y_i$  is the gold label.

**COMPCOMB.** For each triple  $(n, an, c)$  and layer  $\ell$ , we extract representations  $\phi_\ell(n)$ ,  $\phi_\ell(an)$ , and  $\phi_\ell(c)$ . We compute cosine similarity:

$$\text{sim}_\ell(x, y) = \frac{\phi_\ell(x) \cdot \phi_\ell(y)}{\|\phi_\ell(x)\| \|\phi_\ell(y)\|}. \quad (10)$$

A prediction is considered correct if:

$$\text{sim}_\ell(n, an) > \text{sim}_\ell(n, c). \quad (11)$$

The aim here is to examine how well the model’s internal representations encode the relationships between related and unrelated adjective-noun and exocentric compound-noun pairs. Layer-wise accuracy is defined as the proportion of triples satisfying this inequality. This gives insights into how well

the model captures semantic relationships and distinctions between input items (such as distinguishing between a noun and its related and unrelated compounds). This setup allows for an analysis of the model’s ability to differentiate semantically related pairs from unrelated ones based on internal representation quality.