

Experimental Contexts *Can* Facilitate Robust Semantic Property Inference in Language Models, but Inconsistently

Kanishka Misra^τ Allyson Ettinger^α Kyle Mahowald^τ

^τThe University of Texas at Austin ^αAllen Institute for Artificial Intelligence
kmisra@utexas.edu allysone@allenai.org kyle@utexas.edu

Abstract

Recent zero-shot evaluations have highlighted important limitations in the abilities of language models (LMs) to perform meaning extraction. However, it is now well known that LMs can demonstrate radical improvements in the presence of experimental contexts such as in-context examples and instructions. How well does this translate to previously studied meaning-sensitive tasks? We present a case-study on the extent to which experimental contexts can improve LMs’ robustness in performing property inheritance—predicting semantic properties of novel concepts, a task that they have been previously shown to fail on. Upon carefully controlling the nature of the in-context examples and the instructions, our work reveals that they can indeed lead to non-trivial property inheritance behavior in LMs. However, this ability is inconsistent: with a minimal reformulation of the task, some LMs were found to pick up on shallow, non-semantic heuristics from their inputs, suggesting that the computational principles of semantic property inference are yet to be mastered by LMs.

1 Introduction

Carefully controlled behavioral analyses on meaning-sensitive tasks have revealed holes in the ability of language models (LMs) to demonstrate robust meaning extraction and use (Pandia and Ettinger, 2021; Elazar et al., 2021; Schuster and Linzen, 2022; Weissweiler et al., 2022; Misra et al., 2023; Kim and Schuster, 2023, *i.a.*). However, since a large subset of these investigations uses zero-shot evaluation as the primary methodology, there are growing concerns that they do not paint a complete picture of LMs’ abilities (Lampinen, 2022; Sinha et al., 2023). Conclusions that LMs lack a particular ability may be overhasty if it turns out the ability is easily accessed through in-context learning, different question formulations, or particular instructions (Lampinen, 2022; Wei et al., 2022).

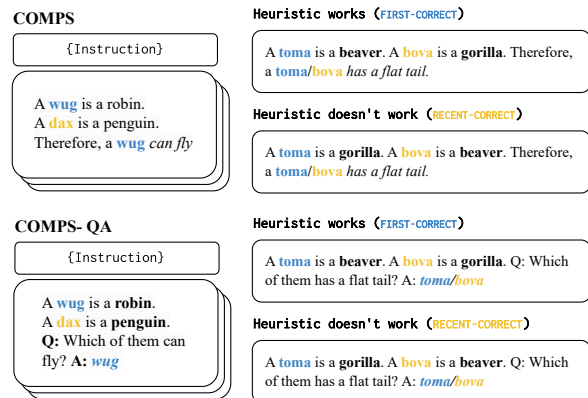
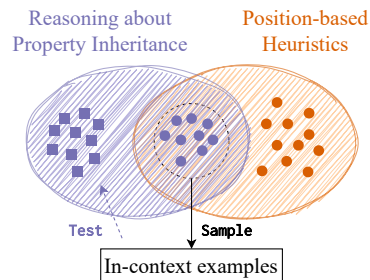


Figure 1: LMs are prompted with in-context examples that are compatible with both, **robust property inheritance**, as well as **position-based heuristics**. At test time, we evaluate on cases where the heuristics support desirable behavior *and* on cases where they do not. We use stimuli from COMPS and its reformulation as a QA task.

Our focus¹ in this paper is a particularly challenging dataset for meaning-sensitive behavior: COMPS (Misra et al., 2023), which contains minimal pair sentences that test the ability of LMs on property knowledge of everyday concepts (*a beaver/gorilla has a flat tail*) and their inheritance for novel concepts (*a wug is a beaver/gorilla. therefore, a wug has a flat tail*). Contemporary LMs failed miserably on the hardest subset of the COMPS stimuli, the examples of which contain two novel concepts (WUG vs. DAX), where only one of them inherits the target property (*has a flat tail*):

¹Our code can be found [in this repository](#)

- (1) A **wug** is a beaver. A **dax** is a gorilla. Therefore, a **wug/dax** has a flat tail.

Given the success of LMs on a wide variety of complicated tasks, their utter failure on this seemingly straightforward task remains puzzling. Here, we systematically explore COMPS on 12 LMs ranging from 1.5–13B parameters, varying (a) whether they are evaluated zero-shot or with multiple examples, and (b) whether or not instructions are present.

Unlike other minimal-pair datasets, using COMPS in an in-context learning setting is non-trivial (and thus potentially informative). This is because the task can be solved using a position-based heuristic. For example, in one subset of COMPS, the target property is always attached to **first** novel concept—like in (1). Importantly, LMs’ failures on COMPS were shown to be in part a result of models’ tendencies towards heuristic behavior: the performance of autoregressive LMs was found to be particularly bad when the distractor (*a dax is a gorilla*) is *recent*—i.e., they showed a recency bias in attributing properties to novel concepts. In that sense, COMPS follows a rich body of work in which tasks are set up in a manner that two types of generalization mechanisms can lead to the same prediction, but only one of which is desirable (McCoy et al., 2019, 2020; Warstadt et al., 2020b; Mueller et al., 2022; Si et al., 2023).

We find that experimental contexts, as operationalized using in-context examples and instructions, can in fact demonstrate robust improvements in LMs’ property inheritance behavior as measured by COMPS. However, this improvement comes with a caveat: With a minimal reformulation of COMPS into a QA task, where there is a direct link between the LMs’ output space and the features of the input that control the heuristic, many LMs showed a strong preference towards the heuristic, and were therefore at chance. Interestingly, LMs that receive explicit supervision to follow instructions do show some resistance to positional heuristics, but not always—occasionally, certain instruction-tuned LMs also tend to show a preference for one particular type of heuristic. These discrepancies in LMs’ performance underscore their difficulty of mastering the reasoning ability to robustly demonstrate semantic property inheritance.

2 Methodology

Dataset We use the most difficult subset of the COMPS dataset (Misra et al., 2023)—COMPS-

WUGS-DIST. This dataset contains 13,828 sentence pairs of the form similar to (1), constructed using 152 animal concepts and 991 properties.

Stimuli re-design We take a number of steps to minimize noise from other (likely uninterpretable) heuristics beyond the ones we have set out to target. First, we enforce that the concepts and properties that appear in the in-context examples are disjoint from ones that are used in tests. To this end, we sample 50 concepts and their relevant properties and reserve it for our in-context examples, leaving the rest to be sampled for our test set. We also enforce this constraint for our novel concepts—i.e., all in-context examples contain different nonce words, and the collection of nonce words for the in-context examples and the test set is disjoint. Furthermore, we counterbalance the nonce words in the test set in a manner that having a bias towards one of them would lead to chance performance. We additionally also use multiple different sets of in-context examples, to add variability and to ensure that the results are not only due to one particular choice of in-context examples. In total, we use 10 different in-context learning example sets, each containing 6 different COMPS stimuli. For our test set, we use a constant set of 256 unique pairs sampled from our pool of stimuli containing unused concepts and properties.

Heuristics Our most important design decision is to consider two distinct sets of stimuli—each separately making available the two types of heuristics that the LMs could rely upon: **FIRST-CORRECT** and **RECENT-CORRECT**, where the property is inherited by the **first** and the **most recent** novel concept, respectively. That is, for the same set of in-context examples, we have a version where the **first concept** is correct like in (1), and one where the **most recent concept** is correct:

- (2) A **wug** is a gorilla. A **dax** is a beaver. Therefore, a **wug/dax** has a flat tail.

For each type of in-context stimuli, we similarly have two versions of test stimuli: one that is consistent with the target heuristic, and one that is not. That is, a test example that is consistent with the **FIRST-CORRECT** heuristic will also have its **first concept** be the one that inherits the property in question, while one which is inconsistent will have the **most recent concept** be the inheritor of the property. Therefore, a model that shows a preference

for a given heuristic will succeed only on one test set and succumb on the other, while a model that is robust to the heuristics will succeed on both.

Reformulation into QA The original COMPS stimuli test for property inheritance using declarative statements, where models are tested for the log-probability they assign to the property (*has a flat tail*) given either of the two concepts (*wug* vs. *dax*). Here we additionally consider an alternate formulation of COMPS as a question answering task (COMPS-QA), where we make the property explicit in the prompt to the model and instead ask which of the two concepts possesses it:

(3) A **wug** is a beaver. A **dax** is a gorilla. **Question:** Which one of them has a flat tail? **Answer:** **wug/dax**

Since the shallow heuristics we consider are controlled by the relative ordering of the novel concepts, this formulation of the task directly allows us to link the models’ output space (the novel concepts) to the heuristics (positions).

Testing setup For the original COMPS setting we follow [Misra et al. \(2023\)](#) and compare the log-probability of the property phrase given the correct vs. the incorrect prefix. For COMPS-QA however, since we have a constant prefix (same premises and question), we evaluate the relative log-probability of the two novel concepts, only one of which is the correct answer. Accuracy in both cases is the proportion of cases the correct surface form was assigned relatively higher log-probability. Since we use pairwise comparisons throughout, chance performance is 50%.

Instructions We consider four different kinds of instruction templates, with varying levels of detail (see appendix B) per formulation (COMPS and COMPS-QA). In our experiments we report results on the instruction that gives the best average performance for a given model.

LMs tested We evaluated 8 different open-source LMs, all of which are decoder-only, and were accessed using the huggingface hub ([Wolf et al., 2020](#)): **GPT-2 XL** ([Radford et al., 2019](#)); **OPT-6.7b** ([Zhang et al., 2022](#)); **Llama-2** (we used the **7B** and the **13B** versions; [Touvron et al., 2023](#)); **Mistral-7B** ([Jiang et al., 2023](#)); **Llama-3-8B** ([Dubey et al., 2024](#)); **OLMo-7B** ([Groeneveld et al., 2024](#)); and **Gemma-2-9B** ([Team et al., 2024](#)).

In addition to these 8, we tested the instruct-tuned versions of OLMo, Llama-3-8B, Mistral-7B (v0.3), and Gemma-2-9B, to analyze if instruction tuning affects LMs’ robustness. Details about all 12 LMs can be found in Appendix A.

3 Analyses and Results

We evaluate on COMPS and COMPS-QA, with and without instructions. In each case, we progressively supply 0 through 6 in-context examples, allowing us to track the dynamics of the models’ performance with an increasing amount of demonstrations. Together with our separate types of test sets and heuristics encoded in the in-context examples, along with five different instruction settings (four with and one without) we run 2420 experiments per LM. We hypothesize that LMs would be more sensitive to the positional heuristics in COMPS-QA because of the clear link between their output space and the relative position of the novel concepts—the feature that controls our target heuristics.

Figure 2 shows accuracies of the tested LMs on our four different COMPS settings as a function of the number of in-context examples provided to them, for both: cases where the **heuristics are consistent with success on the test set**, and cases where **they are not**. We also show an additional curve denoting the average performance across the two types of test sets to paint an overall picture of the models’ performance. In this figure, the extent to which a model relies on heuristic is indicated by the gap between the **dotted** (●) and the **dashed** (▲) lines. A model that is robust to the heuristics will have curves of both colors rise above chance, with no gap between the two, while one that is prone to using heuristics will have its **dotted** (●) curve be substantially greater than its **dashed** (▲) curve. Following the same format, Figure 3 shows results on the four instruction-tuned models evaluated. In all these plots, we also visualize the LMs’ ‘base’ property knowledge, where we compare (in a zero-shot manner) the LMs’ preference for attributing the properties in the test set to the correct ‘known’ concepts (i.e., beaver vs. gorilla in our examples).

Experimental context can improve attribution of properties to concepts... On COMPS, models unsurprisingly start off at chance performance on average, corroborating the previous findings of [Misra et al. \(2023\)](#). However, in the presence of in-context examples and instructions, they are able to improve monotonically as the number of

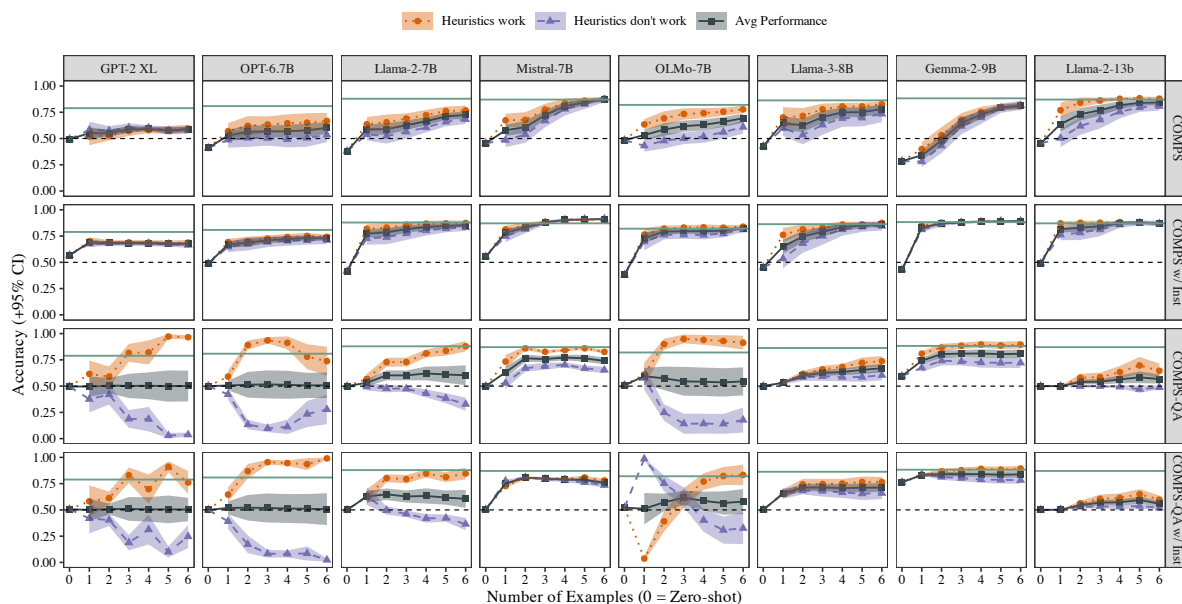


Figure 2: Overall results from our experiments testing non-instruction tuned LMs on COMPS and COMPS-QA using in-context examples, with and without instructions. Results are aggregated across both heuristics: **FIRST-CORRECT** and **RECENT-CORRECT**. Error bars are over different sets of in-context examples. Most models start off near chance in the 0-shot case, but many improve as more examples are given. Solid green line depicts each model’s base property knowledge performance, while the black dashed line depicts chance performance.

in-context examples increases, reaching similar levels of performance on property inheritance as that of the models’ base knowledge. It is worth noting **Llama-2-13b** and **OLMo-7B** do occasionally show a preference for heuristics in the absence of instructions. An intermediate conclusion that we draw here is that LMs can indeed demonstrate non-trivial property inheritance on observing a few examples that reflect that behavior.

...but not the attribution of concepts to properties While experimental context seems to aid models in attributing properties to the right concept in context, the same does not hold on COMPS-QA. Similar to COMPS, models start off at chance performance on average with a zero-shot set up, however, unlike in the case of COMPS, many LMs seem to consistently prefer the heuristics available in the prompt, showing, at times, worse than chance performance on cases where the test set does not follow the heuristic. This is most apparent for **GPT-2 XL**, **OPT6.7b**, **Llama-2-7b**, **OLMo-7B**—here there is a substantial gap between the accuracy for cases where heuristics support performance on the test set and the accuracy for cases where they do not. Other models, like **Mistral-7B** and **Gemma-2-9B** occasionally show behavior compatible with the heuristics, but this does not hurt their overall performance, which is consistently above chance,

and close to their base knowledge.

Our results suggest that LMs are more likely to show behavior that is compatible with the *use* of positional heuristics when their output space (choice between the two novel concepts) has a clear connection with positional artifacts in their input (relative ordering of the novel concepts). This is consistent with our hypothesis in about 11 out of 16 cases (where for 9 of the cases, heuristics end up drastically hurting performance). When this link is not clear and models must instead predict likely properties given a novel concept (i.e., in COMPS), instructions and in-context examples do seem to lead to robust performance. It is important to note that instructions *alone* do not always account for the observed improvement—LMs’ performance on zero-shot settings are consistently still at chance in all cases, suggesting that it is their interaction with in-context examples that critically alters models’ output distribution to support desirable property inference behavior.

Benefits of instruction tuning... Common to the previous analysis is that the tested LMs were all pretrained using the standard next-word prediction objective. While instruction-following capabilities have been shown to arise in such models, a new paradigm has emerged in the field that tunes such pre-trained LMs post-hoc to explicitly follow in-

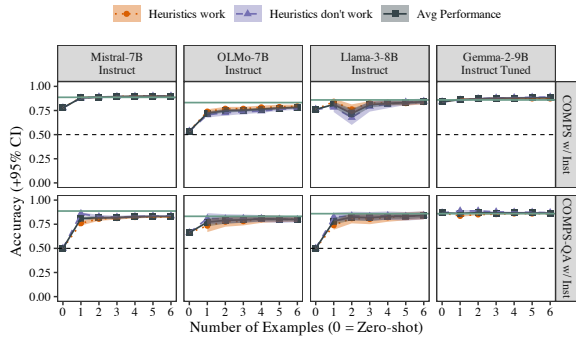


Figure 3: Overall results on the four instruction-tuned models considered. Results are aggregated across both heuristics: **FIRST-CORRECT** and **RECENT-CORRECT**. Error bars are over different sets of in-context examples. Solid green line depicts each model’s base property knowledge performance, while the black dashed line depicts chance performance.

structions using a plethora of different techniques not limited to those that leverage feedback from human-provided labels (Ouyang et al., 2022). How does instruction tuning interact with the presence of position-based heuristics? Figure 3 shows results on four LMs which are instruction-tuned versions of some of the LMs used in the previous analysis, for COMPS and COMPS-QA (with instructions). We find these four LMs to show consistently above-chance performance on both datasets, suggesting non-trivial robustness brought about by instruction-tuning. Interestingly, while most models show improvements from zero-shot (often times at chance), the instruct-tuned **Gemma-2-9B** model shows no signs of improvement, instead showing nearly constant performance on both COMPS and COMPS-QA regardless of the number of in-context examples.

...and its occasional pitfalls While Figure 3 shows that LMs benefit from instruction-tuning in the context of robust property inheritance, a detailed breakdown of the results offers a few caveats. In Figure 4 we show **OLMo-7B-Instruct**’s accuracies broken-down by heuristic on COMPS-QA, the setting where we hypothesized LMs to be most likely relying on heuristics. For the **FIRST-CORRECT** heuristic, the model shows a preference for the heuristic (greater performance when it supports success on the test set), while for **RECENT-CORRECT**, we see the opposite to be true. Here, we observe that the LM does better when the heuristic does not support success on the test set, suggesting that the LM has a bias for preferring entities mentioned first—a bias so strong that it persists even when there is ample evidence in the context

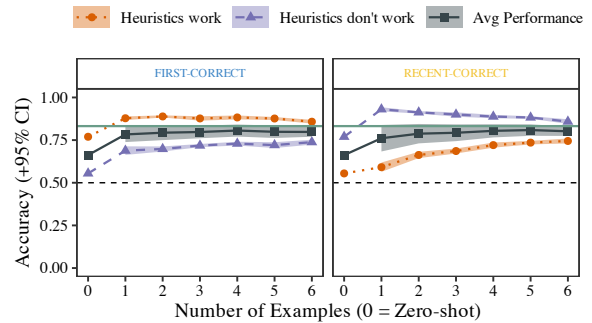


Figure 4: Finer-grained results on Instruct-tuned **OLMo-7B** LM demonstrating its preference for selecting the first concept, regardless of the heuristics.

of the model that favors the opposite heuristic. This suggests that while instruction tuning leads to consistently above-chance performance on challenging property inheritance problems, it is not entirely robust to position-based heuristics.

4 Conclusion

We investigated the extent to which in-context examples and instructions—key components that drive impressive performance in contemporary LMs—can overcome important limitations of LMs at tests that have poked holes in their ability to extract conceptual meaning from text. As a case study, we analyzed how well such experimental contexts can improve LM abilities to perform property inheritance (Murphy, 2002; Misra et al., 2023) in context: binding of novel concepts to existing concepts, and endowing them with valid property inferences as a result.

Our findings suggest that mastery of this ability has yet to be robustly achieved, and that many LMs trained on the next-word prediction objective are still prone to using shallower patterns in their context rather than systematically extracting conceptual meaning. There are several exceptions to this, however, for example, modern LMs such as Mistral-7B and Gemma-2-9b did show heuristic behavior, but were also well above chance, especially in the presence of instructions. Additionally, instruction tuned LMs were found to be substantially more robust than their non-instruction tuned counterparts, suggesting advantages to explicit instruction tuning. However, they too sometimes showed counterintuitive behavior, where their strong bias for a consistent position overcomes the heuristic patterns available in the input. We leave the nuanced exploration of this behavior induced by instruction-tuning for future work.

5 Limitations

Single dataset A clear limitation of this work is that it exclusively focuses on a single dataset: COMPS (Misra et al., 2023). So, a question that arises here is to what extent are our findings localized to the chosen dataset vs. meaning-sensitive evaluations in general. This would require a further non-trivial, non-straightforward amount of work, since: 1) different meaning sensitive evaluations focus on different (though equally useful) operationalizations of meaning; and more importantly 2) not all prior work in this area focuses on a standardized and well-defined usage of heuristics that is directly transferable to the experimental setup we have used in this work (following McCoy et al., 2019, 2020; Warstadt et al., 2020b; Mueller et al., 2022; Si et al., 2023). We do hope that our work contributes to the larger-scale vision of carefully benchmarking different types of meaning extraction abilities in LMs in a controlled manner.

Lack of mechanistic insight Our work continues the long-standing precedent of using carefully constructed behavioral experiments to conclude about the competence of LMs (Linzen et al., 2016; Gulordava et al., 2018; Futrell et al., 2019; Ettinger, 2020; Warstadt et al., 2020a) However, recent works have made impressive strides in localizing the kinds of computations that give rise to the observed behavior in LMs (Hanna et al., 2023; Wang et al., 2023, *i.a.*) Therefore, it is entirely possible that our conclusions about the precise nature of computations carried out by LMs can be greatly strengthened when supplemented by the methods developed in these aforementioned works.

Single Language Finally, this work only focuses on property inheritance problems stated in the English language. This does little to contribute towards diversity in NLP research.

6 Acknowledgments

We thank Tom McCoy and Andrew Lampinen for providing comments on a previous version of the draft. (KM)² acknowledge funding from NSF Grant 2104995 awarded to Kyle Mahowald. We thank UT-Austin Linguistics for providing the hardware to run experiments.

References

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Yanai Elazar, Hongming Zhang, Yoav Goldberg, and Dan Roth. 2021. **Back to square one: Artifact detection, training and commonsense disentanglement in the Winograd schema**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10486–10500, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Allyson Ettinger. 2020. **What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models**. *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Richard Futrell, Ethan Wilcox, Takashi Morita, Peng Qian, Miguel Ballesteros, and Roger Levy. 2019. **Neural language models as psycholinguistic subjects: Representations of syntactic state**. 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 32–42, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, et al. 2024. Olmo: Accelerating the science of language models. *arXiv preprint arXiv:2402.00838*.
- Kristina Gulordava, Piotr Bojanowski, Edouard Grave, Tal Linzen, and Marco Baroni. 2018. **Colorless green recurrent networks dream hierarchically**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1195–1205, New Orleans, Louisiana. Association for Computational Linguistics.
- Michael Hanna, Ollie Liu, and Alexandre Variengien. 2023. **How does GPT-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model**. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Jennifer Hu, Jon Gauthier, Peng Qian, Ethan Wilcox, and Roger Levy. 2020. **A systematic assessment of syntactic generalization in neural language models**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1725–1744, Online. Association for Computational Linguistics.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

- Najoung Kim and Sebastian Schuster. 2023. [Entity tracking in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Andrew Kyle Lampinen. 2022. Can language models handle recursively nested grammatical structures? a case study on comparing models and humans. *arXiv preprint arXiv:2210.15303*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- R Thomas McCoy, Robert Frank, and Tal Linzen. 2020. Does syntax need to grow on trees? sources of hierarchical inductive bias in sequence-to-sequence networks. *Transactions of the Association for Computational Linguistics*, 8:125–140.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Kanishka Misra. 2022. minicons: Enabling flexible behavioral and representational analyses of transformer language models. *arXiv preprint arXiv:2203.13112*.
- Kanishka Misra, Julia Rayz, and Allyson Ettinger. 2023. [COMPS: Conceptual minimal pair sentences for testing robust property knowledge and its inheritance in pre-trained language models](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2928–2949, Dubrovnik, Croatia. Association for Computational Linguistics.
- Aaron Mueller, Robert Frank, Tal Linzen, Luheng Wang, and Sebastian Schuster. 2022. [Coloring the blank slate: Pre-training imparts a hierarchical inductive bias to sequence-to-sequence models](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1352–1368, Dublin, Ireland. Association for Computational Linguistics.
- Gregory L Murphy. 2002. *The Big Book of Concepts*. MIT press.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Lalchand Pandia and Allyson Ettinger. 2021. [Sorting through the noise: Testing robustness of information processing in pre-trained language models](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1583–1596, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI*.
- Sebastian Schuster and Tal Linzen. 2022. [When a sentence does not introduce a discourse entity, transformer-based models still sometimes refer to it](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 969–982, Seattle, United States. Association for Computational Linguistics.
- Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. 2023. [Measuring inductive biases of in-context learning with underspecified demonstrations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11289–11310, Toronto, Canada. Association for Computational Linguistics.
- Koustuv Sinha, Jon Gauthier, Aaron Mueller, Kanishka Misra, Keren Fuentes, Roger Levy, and Adina Williams. 2023. [Language model acceptability judgments are not always robust to context](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6043–6063, Toronto, Canada. Association for Computational Linguistics.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2023. [Interpretability in the wild: a circuit for indirect object identification in GPT-2 small](#). In *The Eleventh International Conference on Learning Representations*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020a. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel R. Bowman. 2020b. [Learning which](#)

features matter: RoBERTa acquires a preference for linguistic generalizations (eventually). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 217–235, Online. Association for Computational Linguistics.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.

Leonie Weissweiler, Valentin Hofmann, Abdullatif Köksal, and Hinrich Schütze. 2022. The better your syntax, the better your semantics? probing pretrained language models for the English comparative correlative. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10859–10882, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Ethan Wilcox, Roger Levy, and Richard Futrell. 2019. Hierarchical representation in neural language models: Suppression and recovery of expectations. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 181–190, Florence, Italy. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. OPT: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

A Dataset and implementation details

Our experiments use the stimuli from COMPS, released with an MIT License by Misra et al. (2023), but with a modification that involves changing of the nonce words to obey the constraint that the in-context examples all have different nonce word pairs. To this end, we use the following nonce words:

- **In-context examples:** *wug, dax, fep, zek, blick, toma, kiki, glorp, bova, zup, tufa, flib* (counter-balanced)
- **Test examples:** *gek, wif* (counter-balanced)

A.1 Methodological details

Following COMPS, as well as the precedent set by a number of previous minimal pair analyses (Linzen et al., 2016; Gulordava et al., 2018; Futrell et al., 2019; Wilcox et al., 2019; Warstadt et al., 2020a; Hu et al., 2020), we use a forced choice task to evaluate our LM subjects. Like in COMPS, we compare the log-probability of the property phrase (here, *has a flat tail*) given the choice of left contexts (which indicate whether the right vs. the wrong concept has the property). For example, we measure:

$$\log P_{\theta}(\textit{has a flat tail} \mid \textit{a gek is a beaver. a wif is a gorilla. therefore, a gek/wif}),$$

and for COMPS-QA, we compare the relative probabilities of the two novel concepts given a fixed left prefix which contains a question about the property. For example, we measure:

$$\log P_{\theta}(\textit{gek/wif} \mid \textit{a gek is a beaver. a wif is a gorilla. Question: Which one of them has a flat tail? Answer:})$$

In both cases above, *gek* is the concept that should inherit the property. While these examples show the zero-shot case, cases with in-context examples and instructions simply add more context to the prefix, therefore the surface form of the output space remains the same regardless of the number of in-context examples or the presence of instructions. Similar to the above measures, we measure the base property knowledge of LMs by comparing the conditional log-probabilities of the property phrases given the real, known concepts in question. That is, we measure:

$$\log P_{\theta}(\textit{has a flat tail} \mid \textit{beaver/gorilla})$$

The accuracy here, again, is the proportion of time the correct concept was more likely to inherit the property than the incorrect one. We measure this accuracy for all our test items zero-shot, and use that as the reference base property knowledge in order to contextualize our main results.

Log-probabilities for all models were accessed using minicons (Misra, 2022),² a library that wraps around transformers (Wolf et al., 2020) by huggingface, and is written in pytorch. For our experiments with Llama-13B, we quantize the model

²<https://github.com/kanishkamisra/minicons>

to 4-bits in order to fit it onto a single GPU. All experiments were run on a cluster with 4 NVIDIA A40 GPUs, though each individual experiment on a model was computed on a single A40 GPU.

A.2 Model Metadata

Table 1 shows the LMs used in this work, along with their total parameters, tokens encountered during training, and vocabulary size.

B Instructions

Tables 2, 3, 4, 5 show our instruction templates.

C Fine-grained results

While Figure 2 shows results aggregated over both types of heuristics that we have used in this work, we additionally display finer-grained results in this section, broken down by heuristic type. Again, in each of these plots, the extent to which a model relies on heuristic is indicated by the gap between the dotted (●) and the dashed (▲) lines. This is now separately shown for each of our heuristics. Figure 5 shows results for non-instruct tuned LMs on COMPS with and without instructions for both the heuristics, and similarly Figure 6 shows results for those LMs on COMPS-QA with and without instructions for both the heuristics. Figure 7 and Figure 8 show analogous results for Instruction tuned LMs. In all these plots, the solid green line denotes the base performance of LMs.

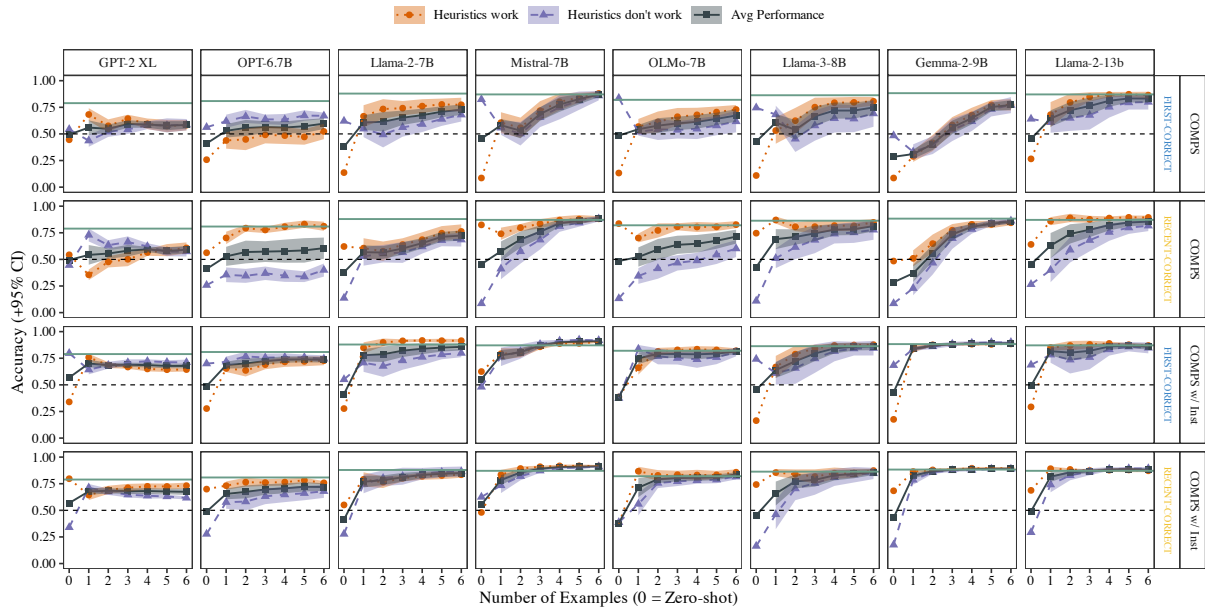


Figure 5: Fine-grained results for non-instruct tuned LMs on COMPS as a function of the number of in-context examples (with and without instructions).

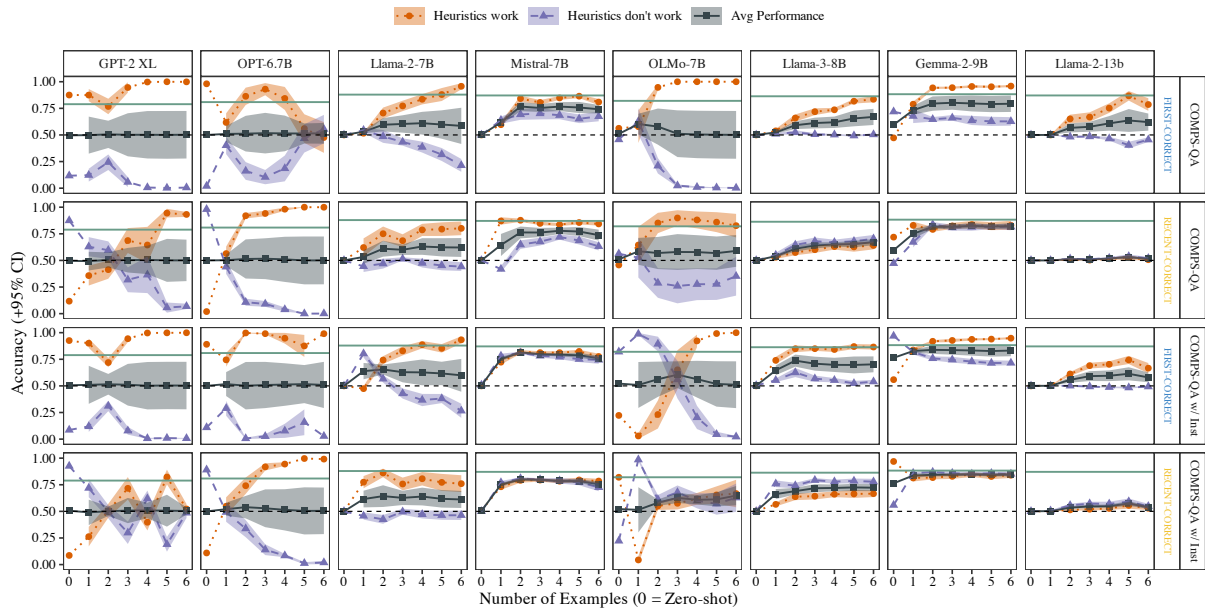


Figure 6: Fine-grained results for non-instruct tuned LMs on COMPS-QA as a function of the number of in-context examples (with and without instructions).

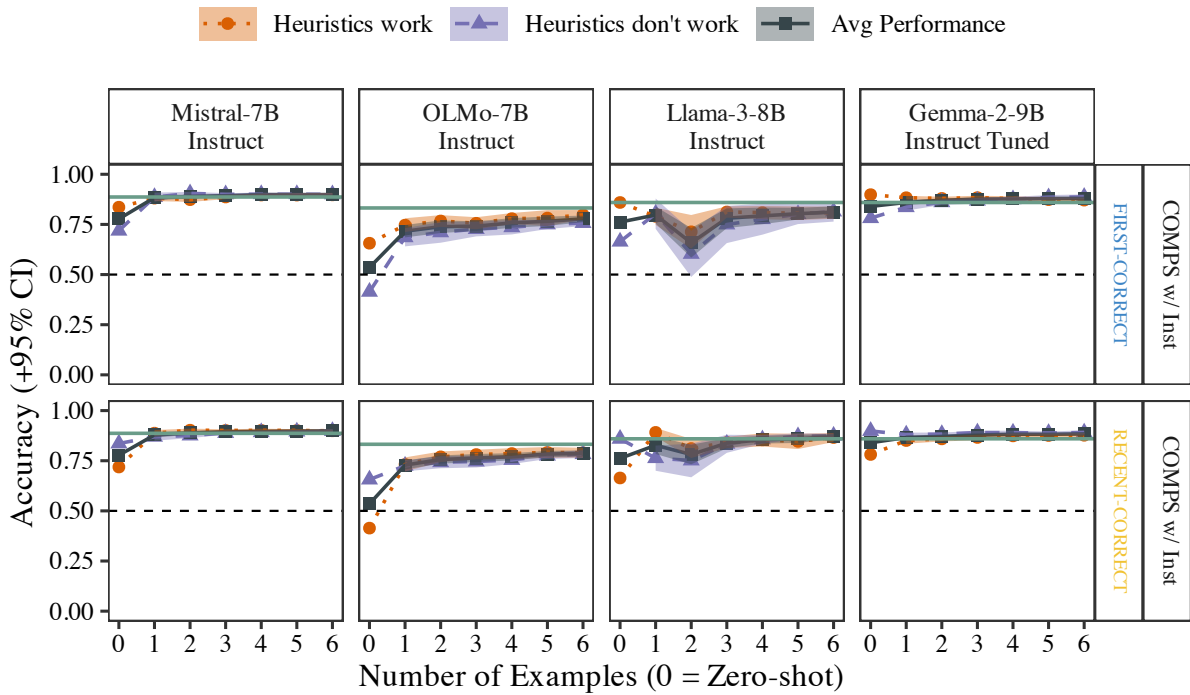


Figure 7: Fine-grained results for instruct tuned LMs on COMPS as a function of the number of in-context examples (with and without instructions).

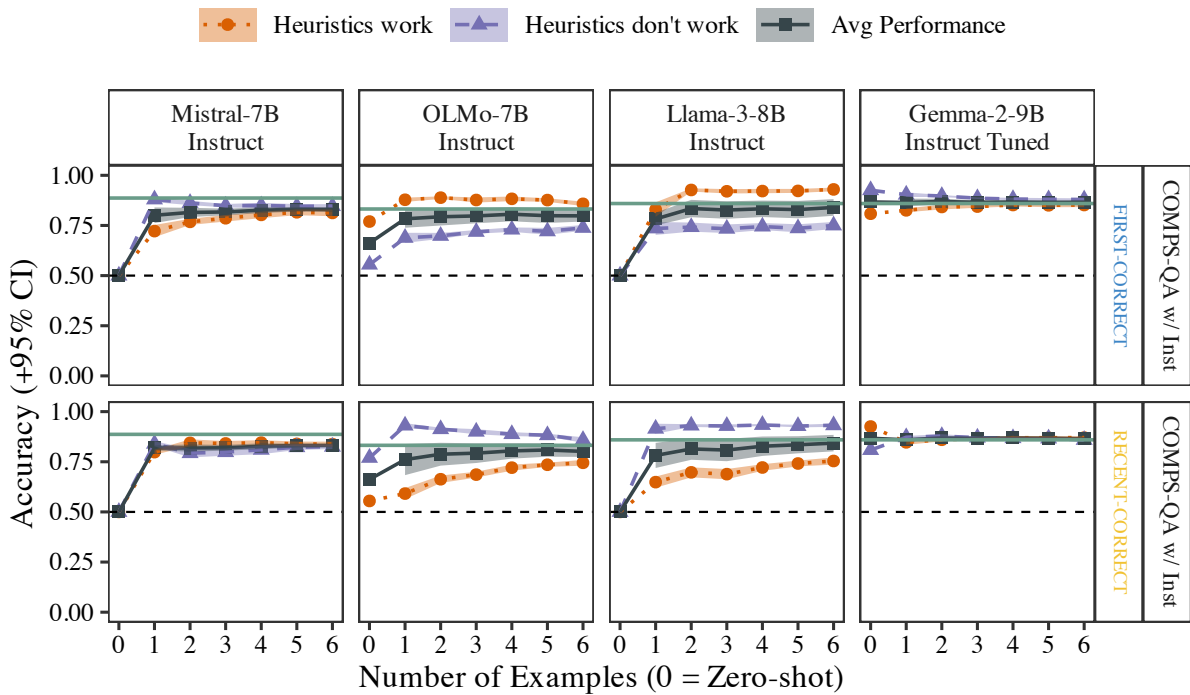


Figure 8: Fine-grained results for instruct tuned LMs on COMPS-QA as a function of the number of in-context examples (with and without instructions).

Model	Params	Pre-training Tokens	Vocab size
GPT-2 XL (Radford et al., 2019)	1.5B	8B	50,257
OPT-6.7b (Zhang et al., 2022)	6.7B	180B	50,272
Llama-2-7B (Touvron et al., 2023)	7B	1.8T	32,000
Mistral-7B (Jiang et al., 2023)	7B	?	32,000
OLMo-7B (Groeneveld et al., 2024)	7B	2.46T	50,304
Llama-3-8B (Dubey et al., 2024)	8B	15T	128,000
Gemma-2-9B (Team et al., 2024)	9B	8T	256,000
Llama-2-13B (Touvron et al., 2023)	13B	1.8T	32,000

Table 1: Overview of the non-instruction tuned LMs used in this work. ‘?’ indicates that the given value was not made available in the LM’s release. In addition to these, we also report results on instruction-tuned versions of **OLMo-7B**, **Mistral-7B**, **Llama-3-8B**, and **Gemma-2-9B**.

COMPS version	Instruction Template
COMPS	<p>Given a pair of statements that introduce novel entities as types of real world animals, write a true statement about the properties of the novel entities:</p> <p>{examples} (omitted in zero-shot) {test-stimulus}</p>
COMPS-QA	<p>Given a pair of statements that introduce novel entities as types of real world animals, answer the question that follows:</p> <p>{examples} (omitted in zero-shot) {test-stimulus}</p>

Table 2: Instructions for COMPS and COMPS-QA with instruction type: “minimal”

COMPS version	Instruction Template
COMPS	<p>Some aliens have come to earth, and it turns out they have their own language for talking about our animals here on Earth. Your job is to help the aliens learn about our Earthling animals by giving them some information about the animals.</p> <p>Let’s get started: {examples} (omitted in zero-shot) {test-stimulus}</p>
COMPS-QA	<p>Some aliens have come to earth, and it turns out they have their own language for talking about our animals here on Earth. Your job is to help the aliens learn about our Earthling animals by answering some questions about them.</p> <p>Let’s get started: {examples} (omitted in zero-shot) {test-stimulus}</p>

Table 3: Instructions for COMPS and COMPS-QA with instruction type: “aliens”

COMPS version	No. of shots	Instruction Template
COMPS	Zero-shot	<p>It is important to know and reason about the properties of entities in the world. The following is a pair of premise statements that introduce novel entities as new types of real world animals. Your task is to make a conclusion about the properties of one of the entities by reasoning over the premise statements.</p> <p>{test_stimulus}</p>
	Few-shot	<p>It is important to know and reason about the properties of entities in the world. The following example(s) show a pair of premise statements that introduce novel entities as new types of real world animals, followed by another statement that attributes a property to one of the entities introduced in the premise statements.</p> <p>Examples: {examples}</p> <p>Here is another pair of premise statements. Your task is to make a conclusion about the properties of one of the entities by reasoning over the premise statements.</p> <p>{test_stimulus}</p>
COMPS-QA	Zero-shot	<p>It is important to know and reason about the properties of entities in the world. The following is a pair of premise statements that introduce novel entities as new types of real world objects. The statements are followed by a question that asks which novel entity in the premise can a specific property can be attributed to. Answer the question by reasoning over the premise statements.</p> <p>{test_stimulus}</p>
	Few-shot	<p>It is important to know and reason about the properties of entities in the world. The following example(s) show a pair of premise statements that introduce novel entities as new types of real world animals. The statements are followed by a question that asks which novel entity in the premise can a specific property can be attributed to, and the answer to the question, obtained by reasoning over the premise statements.</p> <p>Examples: {examples}</p> <p>Here is another pair of premise statements. Answer the question that follows.</p> <p>{test_stimulus}</p>

Table 4: Instructions for COMPS and COMPS-QA with instruction type: “Detailed-1”

COMPS version	No. of shots	Instruction Template
COMPS	Zero-shot	<p>It is important to know and reason about the properties of entities in the world. The following is a pair of premise statements that introduce novel entities as new types of real world animals. Your task is to write a true statement about the properties of the novel entities.</p> <p>{test_stimulus}</p>
	Few-shot	<p>It is important to know and reason about the properties of entities in the world. The following example(s) show a pair of premise statements that introduce novel entities as new types of real world animals, followed by another statement that attributes a property to one of the entities introduced in the premise statements.</p> <p>Examples: {examples}</p> <p>Here is another pair of premise statements. Your task is to write a true statement about the properties of the novel entities.</p> <p>{test_stimulus}</p>
COMPS-QA	Zero-shot	<p>It is important to know and reason about the properties of entities in the world. The following is a pair of premise statements that introduce novel entities as new types of real world animals. The statements are followed by a question that asks which of the introduced entities a specific property can be attributed to. Answer the question by reasoning over the premise statements.</p> <p>{test_stimulus}</p>
	Few-shot	<p>It is important to know and reason about the properties of entities in the world. The following example(s) show a pair of premise statements that introduce novel entities as new types of real world animals. The statements are followed by a question that asks which of the introduced entities a specific property can be attributed to, and the answer to the question, obtained by reasoning over the premise statements.</p> <p>Examples: {examples}</p> <p>Here is another pair of premise statements. Answer the question that follows.</p> <p>{test_stimulus}</p>

Table 5: Instructions for COMPS and COMPS-QA with instruction type: “Detailed-2”