

An Incomplete Loop: Deductive, Inductive, and Abductive Learning in Large Language Models

Emmy Liu & Graham Neubig
 Language Technologies Institute
 Carnegie Mellon University
 emmy@cmu.edu, gneubig@cs.cmu.edu

Jacob Andreas
 CSAIL
 Massachusetts Institute of Technology
 jda@mit.edu

Abstract

Modern language models (LMs) can learn to perform new tasks in different ways: in *instruction following*, the target task is described explicitly in natural language; in *few-shot prompting*, the task is specified implicitly with a small number of examples; in *instruction inference*, LMs are presented with in-context examples and are then prompted to generate a natural language task description before making predictions. Each of these procedures may be thought of as invoking a different form of reasoning: instruction following involves deductive reasoning, few-shot prompting involves inductive reasoning, and instruction inference involves abductive reasoning. How do these different capabilities relate? Across four LMs (from the gpt and llama families) and two learning problems (involving arithmetic functions and machine translation) we find a strong dissociation between the different types of reasoning: LMs can sometimes learn effectively from few-shot prompts even when they are unable to explain their own prediction rules; conversely, they sometimes infer useful task descriptions while completely failing to learn from human-generated descriptions of the same task. Our results highlight the non-systematic nature of reasoning even in some of today’s largest LMs, and underscore the fact that very different learning mechanisms may be invoked by seemingly similar prompting procedures.¹

1 Introduction

Suppose a friend is teaching you to cook. You watch them place a pan on the stove and heat olive oil at low heat, adding minced garlic and chili flakes to the olive oil once it gets hot. Later, you decide to make the recipe yourself, but you are out of olive oil. You hypothesize that the olive oil served to cook the garlic without burning it, and so substitute butter for olive oil, as it should fulfill the same function. Here, you learned a generalizable cooking procedure by reasoning **abductively**—finding the rule that best explains your experience (Frankfurt, 1958; Peirce, 1965). This is only one way of learning: the friend could have instead provided a recipe, from which you could have reasoned **deductively** about how to apply it in your kitchen. If you had remembered other times when you put but-

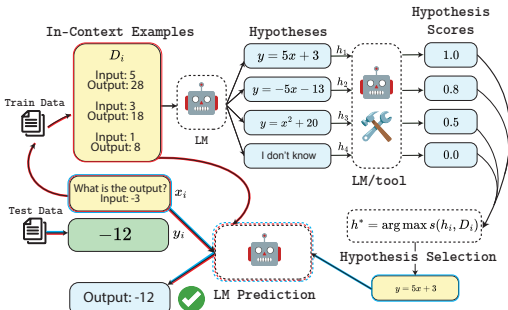


Figure 1: Diagram of abductive reasoning for an LM. Red arrows show data flow in inductive reasoning (few-shot prompting), while blue arrows show data flow in deductive reasoning (instruction following). Black arrows indicate data flow unique to abductive reasoning (instruction inference). Instruction inference generally improves on few-shot prompting and zero-shot chain of thought. However, success at inductive reasoning and success at instruction inference are not related.

¹Code and data, including generated hypotheses may be found at https://github.com/nightingal3/rule_induction

Reasoning type	Analogue in LMs	Citations
Deductive	Instruction following	Wei et al. (2021); Sanh et al. (2021)
Inductive	In-context learning	Brown et al. (2020)
Abductive	Chain-of-thought (one hypothesis)	Wei et al. (2023); Kojima et al. (2022)
Abductive	Externally aided instruction inference (many hypotheses)	Qiu et al. (2023); Wang et al. (2023); Zhu et al. (2023); Yang et al. (2024),
Abductive	LM-guided instruction inference (many hypotheses)	This paper

Table 1: Summary of reasoning types and analogues in language models. Citations to instruction following, in-context learning, and chain-of-thought are limited to the original paper due to the high number of papers on these topics, while we have tried to list all currently available papers on instruction inference (for abductive reasoning) with LMs.

ter on a hot pan before cooking, without explicitly reasoning about why, you could have cooked the same meal **inductively**.

Each form of reasoning has a close analogue in current procedures for steering language models (LMs). In order to induce LMs to perform new tasks, we may condition them on explicit commands (**instruction following**; Wei et al., 2021; Sanh et al., 2021), or a collection of examples from the task of interest (**few-shot prompting**; Brown et al., 2020); recently, several methods have been proposed that prompt LMs to *generate* textual commands from examples before conditioning on commands during prediction (**instruction inference**; Andreas et al., 2018; Honovich et al., 2022). But it is often unclear when to prefer one of these procedures over the other, and more generally how these different capabilities relate in current LMs. Does the ability to learn effectively from few-shot prompts imply the ability to perform instruction inference, or vice-versa? Are there tasks that can be learned from few-shot prompts, but not instructions? In this paper, we examine these questions through two tasks: learning numerical functions (§4.1) and translation models (§4.2, §4.3). We ask two questions:

RQ1: When does (abductive) instruction inference improve LM performance over ordinary (inductive) few-shot prompting?

Finding: Instruction inference improves over few-shot prompting in simple cases (linear function learning and simple artificial language learning), but suffers from incorrect hypotheses in a more complex case (low-resource machine translation).

RQ2: How does the ability to learn from instructions (deductively) relate to the ability to learn from in-context examples (inductively or abductively)?

Finding: The ability to learn through abduction (proposing hypotheses) is generally not related to learning through induction (few-shot learning).

Our results highlight the diverse capabilities (and diversity of different reasoning mechanisms) triggered by different prompts and examples in current LMs; future work may investigate how these reasoning types can be combined or made consistent to enhance problem solving in LMs.

2 Three Types of Reasoning in Language Models

We first give a formalized definition of instruction following, in-context learning, and instruction inference, relating these processes to deductive reasoning, inductive reasoning, and abductive reasoning respectively. Figure 1 gives a schematic of the reasoning loop, while Table 1 gives examples of work falling in each category.² Throughout this section, we assume

²We note that there is some disagreement in the philosophy literature about the exact distinction about inductive and abductive reasoning. Here we use commonly-cited definitions, which happen to distinguish in-context learning from the more explicit process of verbalizing and testing hypotheses.

that we have an input x , and we want to produce an output y with an autoregressive LM conditioned on some additional piece of data D that specifies the target task: $p_{LM}(y | x, D)$. We use in-context learning of linear functions as a running example.

2.1 Instruction Following

In instruction following, we want to map each input x to a y according to a general instruction or prediction rule D that is specified in the input. This may be viewed as a kind of deductive reasoning, which begins with one or more premises, and applies logically valid rules to reach a conclusion (Shapiro & Kouri Kissel, 2024). For instance, the instruction D might name a general function, such as *Apply this function to the input: $y = 5x + 3$* , followed by a specific query: *Input: -3*. This is illustrated in Figure 1.

2.2 Few-Shot Prompting

Few-shot prompting, by contrast, specifies the target task implicitly through examples. For each input x , the task specification D consists of a set of k training examples $D = \{(x_1, y_1), \dots, (x_k, y_k)\}$. We then sample from $p_{LM}(y | x, D)$. Few-shot learning requires LMs to perform inductive reasoning (Hawthorne, 2021). Unlike deductive reasoning, there is no explicit premise stated, but the model must complete the task in a similar way to the examples. In Figure 1, this is pairs of inputs and outputs, i.e. *Input: 5, Output: 28*. These examples help specify the task as a numeric prediction task, as well as the identity of the specific target function.

2.3 Instruction inference

Instruction inference connects instruction following and in-context learning: given examples $D = \{(x_i, y_i)\}_{i=1}^k$ and input x , we can instruct the model to generate a hypothesis h about the identity of the task, i.e. an instruction describing the task associated with D . We may sample one hypothesis and immediately condition on it (a form of chain-of-thought prompting), or sample several and select the most promising one.

Compared to chain-of-thought, fewer studies have explored multi-hypothesis instruction induction; those that do typically rely on an external validation model rather than using the LM itself to evaluate. Our approach (in §3) has the following high-level form: after sampling n hypotheses h_1, \dots, h_n from $p_{LM}(h | t, x)$, we evaluate each hypothesis by assigning a score $\text{score}(h_i, t)$ based on the hypothesis and in-context examples. We choose the best one $h^* = \arg \max \text{score}(h_i, D)$, and feed it back into the context as an instruction.³

Abductive reasoning is often called “inference to the best explanation” (Lipton, 2001; Douven, 2021). Suppose that the model generates the hypotheses shown in Figure 1. We may parse hypotheses and apply them back to the in-context examples, then ranking them by prediction error. Then the model simply has to follow the instruction as in Section 2.1.

3 Methods

We describe the implementation of instruction inference with multiple hypotheses in this section. As the other settings (few-shot, zero-shot chain-of-thought, etc) are already well understood, we do not explain them further, but include the exact prompts used to implement methods in Appendix B, F, and J.

Instruction inference [Abduction with many hypotheses] After generating n hypotheses $\mathcal{H} = \{h_1, \dots, h_n\}$, we explore methods for reranking them. For all experiments, $n = 5$. Generally, these reranking methods capture “fit” to training data. External validator reranking

³Several recent papers have recently proposed similar processes, but called this inductive reasoning. We believe that abductive reasoning may be a more apt term in any process that includes hypothesis evaluation and selection. See Section 6 for more discussion.

Domain	In-context examples (D_i)	Query ex-ample (x_i)	Expected answer (y_i)	Hypothesis example (h_i)
Functions	(-10, -213), (9, 167), (4, 67), ...	15	287	$f(x) = 20x - 13$
Colours	(lug dax, blue green), (lug zup, blue yellow), (lug bluf, blue blue), ...	lug walm dax bluf	blue blue blue green green	lug \rightarrow blue
Kalamang vocab (ek)	(That is their place., Tompat ma me muin), (I'm getting pandandus, I want to make a mat.), ...	Sakina is pouching guavas.	Sakina sarimara lawat.	guava \rightarrow sarim
Kalamang grammar (ek)	""	-	-	Order of subject and verb: SV

Table 2: Examples of in-context examples, queries, and hypotheses in each domain.

was used only in the functions domain (see Section 4.2). The other reranking methods used a language model. Given scores of each hypothesis, we choose:

$$h^* = \arg \max_{h_i \in \mathcal{H}} \text{score}(h_i, D_i).$$

Settings with instruction inference are referred to collectively as `instruction_inference`. We detail the score functions below:

Verbalized confidence We directly prompt the model to estimate the probability of the hypothesis given D_i . $\text{score}_{\text{verbal}}(h_i, D_i)$ is set to the model’s probability estimation.

P(data) We use a separate LM,⁴ to generate log probabilities for in-context examples given the hypothesis. $\text{score}_{\text{P(data)}}(h_i, D_i)$ is the sum of log-probabilities of tokens of D_i .

P(answer) This is similar to P(data) and uses the same template, but $\text{score}_{\text{P(answer)}}(h_i, D_i)$ only sums log-probabilities of tokens in answers y_i in the in-context examples.

External validator For structured hypotheses, it is possible to parse them and apply them back to in-context examples, with the score as the negative error. It may not be possible to parse all hypotheses, if they are in natural language or inconsistent formats.

4 Domains and Evaluation

We investigate LM behavior in three domains: linear function inference, an artificial language learning task, and vocabulary + typological feature learning in the Kalamang language. We refer to the underlying task we would like to improve (output prediction, translation) as the **base task**, and the task of inferring an explicit natural language hypothesis from few-shot examples as the **abductive task**. Examples of each task are in Table 2. We also evaluate hypotheses themselves in each domain. For the exact prompts used in each domain, refer to Appendix B, F and J. Models used across all domains consisted of gpt family models (gpt-3.5-turbo; Brown et al., 2020, gpt-4-turbo, OpenAI et al., 2024) and llama family models (llama-2-7b-chat, llama-2-70b-chat, Touvron et al., 2023).

4.1 Linear Functions

Following investigations of language models’ in-context learning abilities (Garg et al., 2023; Akyürek et al., 2023), we construct a dataset of 40 linear functions $f(x) = ax + b$, where a and b are uniformly sampled from the integers $[-20, 20]$, along with 5 test examples for each function, yielding a 200 test examples. For each test example, we randomly generate 5 in-context examples of the function $(x_i, y_i)_{i=1}^5$ and one test example. The test example and inputs for in-context examples are also uniformly sampled from the integers $[-20, 20]$. We refer to this as the **functions domain**. Prompts used are given in Appendix B.

⁴text-davinci-002 was used as the logprobs-generating LM for all models

Hypotheses In this domain, hypotheses are proposed using all in-context examples for each test example. The model was presented with functions (and instructed to write them) in the form $y = ax + b$. For external validator reranking, we parse the generated hypothesis, and score it by its negative mean-squared error (MSE) when applied to the in-context examples.⁵ If $h_i(x_{ik})$ represent executing the parsed hypothesis represented by h_i on example x_{ik} :

$$s_{\text{ground.truth}}(h_i, D_i) = \sum_{j=1}^k (h_i(x_{ij}) - f(x_{ij}))^2$$

Evaluation To evaluate the base task, we use 0-1 accuracy, as well as median squared errors. To evaluate hypotheses, we examine the accuracy of model-proposed a and b coefficients, as well as the Spearman correlation between proposed coefficients and the ground truth.

4.2 Simple artificial languages

Inspired by compositional instruction-following datasets, we generate a simple dataset where the inputs are nonce words such as *lug*, and the outputs are colour terms such as *blue* (Lake & Baroni, 2017; Lake et al., 2019). We call the expanded dataset the **colours domain**. The ruleset for this domain can be found in Appendix E. We generate 200 test examples and 800 training examples.⁶ Prompts used are in Appendix F.

Hypotheses Following the original miniscan, we create a fixed minimal set of 5 in-context examples that contains each nonce word at least once, and from which the meaning of each nonce word can be reasonably inferred. During instruction inference, we isolate one nonce word at a time from the sentence, and have the model try to induce the vocabulary mapping for that word from 5 retrieved in-context examples containing that word. The ground truth grammar was written by the author, and consisted of production rules for each nonce term. Prompts used in the abductive and base tasks can be found in Appendix F

Evaluation We also use 0-1 accuracy to evaluate the base task of nonce word translation, as well as corpus-level chrF (Popović, 2015). To evaluate the quality of the hypotheses themselves, we extract the production rules proposed by models and compare 0-1 accuracy against the ground truth. As a lenient evaluation on the “repeat” terms (see Appendix E for nonce terms and meanings), we marked a hypothesis for “repeat” terms as correct if it contained the term *repeat* or the numerals 2 and 3 respectively.

4.3 Kalamang translation

For low-resource translation, we use the Machine Translation with One Book (MTOB) dataset, an English–Kalamang dataset with a grammar book (Tanzer et al., 2024). Kalamang is an extremely low-resource language with fewer than 200 speakers, and virtually no text on the web. The base task is to perform sentence-level translation in both directions, while the abductive task is to infer correct grammar features and vocabulary mappings. The dataset consisted of 100 test sentences (50 in each direction) and 400 train sentences. A ground-truth bilingual dictionary was provided (Visser, 2020). A grammar book was included as well, but to check correctness of high-level grammar inferences, we compiled a high level grammar sketch instead from WALS and GramBank features (Dryer & Haspelmath, 2013; Skirgård et al., 2023). More details about the grammar sketch can be found in Appendix K. Otherwise, we use the same experimental settings as the baseline, summarized in Appendix I.

Hypotheses We split hypotheses into vocabulary and grammar. The ground truth instruction included retrieved examples from the wordlist and the grammar sketch. To induce the grammar sketch, we posed each grammar feature as a question, for instance: *What is the order of subject and verb in Kalamang?*, and sampled 5 sentence pairs at a time with the

⁵If a model produced an unparsable hypothesis (for instance, *I don't know*, or a generic answer like $y = ax + b$), that hypothesis was assigned a score of $-\infty$.

⁶We generate this dataset instead of using the original miniscan to mitigate memorization of the original nonce words, as well as to gain more test examples, as the original dataset has only 10.

Domain	Deductive reasoning works on average	Abductive reasoning works on average	Hypothesis proposal works on average	Abductive reasoning related to inductive reasoning
Functions	✓	✓	✗	✗
Colours	✓	✓	✗	✗
Kalamang	✗	✗	✗	✗

Table 3: Summary of results. A checkmark indicates that the property held for all or almost all language models, a half-checkmark indicates a partial success for all or almost all language models, while an X-mark represents lack of success for most language models.

requisite parts of speech until the model proposed an answer to the question. If the model responded that it was unclear, this would repeat up to a maximum of 10 iterations. We only performed this process once, for GPT-3.5-turbo and GPT-4-turbo respectively. Each model used its self-induced grammar sketch at test time.⁷ For vocabulary induction, we followed a similar process as in the colours domain.⁸

Evaluation We use corpus-level chrF. To evaluate the grammar sketch and vocabulary induction, we respectively compare to the ground truth wordlist and grammar sketch.

5 Results

In this section, we first evaluate models’ concrete performance across different domains (**RQ1**:). We highlight the significant improvement instruction inference offers in some cases in synthetic tasks, yet despite this, improvements are not uniform across tasks, and not attested in challenging domains like Kalamang. Our analysis also highlights an inconsistent correlation between the quality of generated hypotheses and few-shot learning success (**RQ2**:), meaning that the ability to generate or follow instructions doesn’t reliably predict task mastery or vice versa. These results (Table 3) suggest that while structured instructions can boost performance in simpler scenarios, their impact is less predictable in complex settings. Furthermore, the relationship between instruction induction, instruction following, and in-context learning is complex, and each capability may rely on separate unknown aspects of model architecture, training procedures, or data.

5.1 When does instruction inference improve over in-context learning?

How effective is using the true instruction? In the domains we study, the ground-truth instruction tends to yield accurate results. Figure 3 displays mean accuracy in the linear functions and colours domains over six trials, sampled at different temperatures (see Appendix A) for details. Notably, in linear functions (3A), GPT-4-turbo’s accuracy increases to 96% from a baseline of 30%, with GPT-3.5 also notably improving. In the colours domain, (3B), we see the true instruction also helps all models except for Llama-2-7b. However, this trend does not extend to the Kalamang task, where most models struggled to leverage the provided wordlist and grammar sketch effectively, indicated by chrF scores in Figure 4.

How effective are models’ induced instructions? Self-generated instructions also improve on the baseline in many cases, with variations by domain. In linear functions, models’ hypothesis induction markedly surpasses the few-shot baseline, with both the verbalized confidence and log-probability based reranking methods yielding comparable improvements (see Figure 3(A) again). For the colours domain (3(B)), instruction inference benefits performance, though not as strongly as in linear functions. chrF scores follow a similar trend, and are depicted in Appendix G. Interestingly, for gpt-4-turbo and Llama-2-7b, using the

⁷Llama-2 models were found to be unable to propose grammar features with our prompts, so we used the GPT-3.5-turbo induced grammar sketch for these models.

⁸For Kalamang, due to computational costs, we cache the first parseable hypothesis proposed for each word and reuse it on subsequent sentences containing that word.

models’ self-proposed hypotheses benefits performance more than using the ground-truth grammar for the colours language, despite the fact that self-proposed are not always correct. Unlike the two synthetic domains, inducing grammar and vocabulary items for Kalamang does not improve translation metrics in most cases.

5.2 How does the ability to induce instructions relate to in-context learning?

How accurate are model-generated hypotheses? In assessing accuracy of models’ self-generated hypotheses across different domains, our findings reveal significant variations in accuracy. In the linear functions domain, we plot hypotheses generated by two models in Figure 2, and list the Spearman ρ (Spearman, 1904) as well as p-values of model-predicted coefficients in Table 4. GPT-3.5-turbo and GPT-4-turbo’s proposed coefficients are positively correlated with the real coefficients, and this is statistically significant. However, this level of accuracy is not observed with Llama-2 models, indicating a disparity in model capabilities. In the colours domain, most models, except GPT-4-turbo, tend to generate inaccurate hypotheses. The exact accuracy of proposed hypotheses for colours is shown in Appendix H. Mappings of simple vocabulary items such as *lug* tend to more accurate, with GPT-4-turbo achieving an 87% mean accuracy for hypotheses about this word. On the other hand, the difficulty increases with the terms which involved repeats, with GPT-3.5-turbo only achieving a 15% mean accuracy on *bluf*, the “repeat twice” term.

However, when extending these analyses to Kalamang, all models’ performance in predicting grammar features is relatively poor, with GPT-3.5-turbo predicting 5/18 and 4/18 features correct respectively. Appendix L shows the correctness of each grammar feature. Vocabulary induction accuracy is also generally low, falling between 10-20% for most models in both the English to Kalamang as well as Kalamang to English directions. See Appendix M for details, as well as averaged segment-level chrF for the vocabulary hypotheses.

Is abductive reasoning related to in-context learning ability? In our final analysis, we examine whether the ability to induce correct instructions correlates with success in in-context learning. Specifically, we focus on the final hypotheses selected by the gpt model family, given that many hypotheses generated by llama models were incoherent or not formatted correctly. We use the point-biserial correlation (Pearson, 1895) to assess the relationship between the accuracy of a model’s final hypothesis and its success in in-context learning.

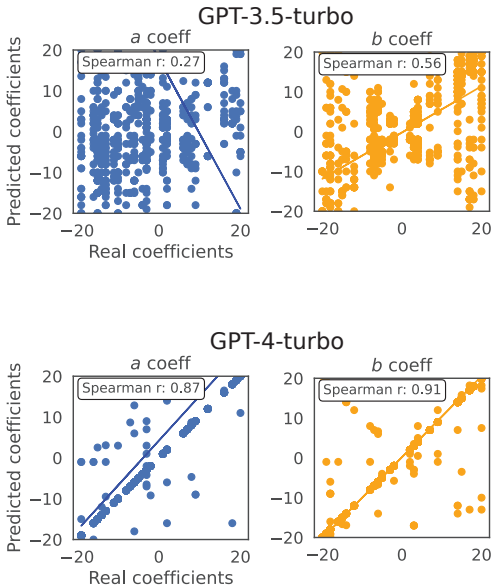


Figure 2: Real coefficients of linear functions and relationship to hypothesized coefficients for GPT-3.5-turbo and GPT-4-turbo. Remaining models can be found in Appendix D. The x-axis has been truncated for visualization purposes (as there are some large outlier hypotheses). GPT-4-turbo is able to induce a reasonable function in-context, but other models struggle.

Model	<i>b</i> corr.	<i>b</i> p-val	<i>a</i> corr	<i>a</i> p-val
GPT-3.5-turbo	0.2685	5.013×10^{-15}	0.5570	2.747×10^{-67}
GPT-4-turbo	0.8485	2.153×10^{-292}	0.9132	0.0
Llama-2-7b	-0.0069	0.9012	0.1455	0.0081
Llama-2-70b	0.1393	0.0595	-0.0175	0.8222

Table 4: Spearman correlation and p-values of true vs predicted coefficients for each model in the functions domain.

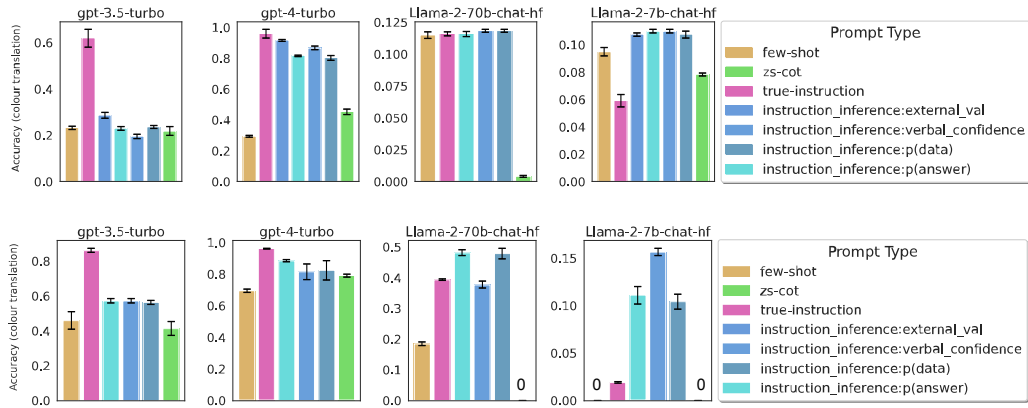


Figure 3: Accuracy of models in synthetic domains with and without hypothesis generation. Error bars indicate standard error. The top row shows results for the functions domain, while the bottom row shows results for the colours domain. Results are aggregated across 6 runs, and zero values are marked with ‘0’.

In linear functions, we examine GPT-3.5-turbo because of GPT-4-turbo’s consistent accuracy in selecting correct hypotheses.⁹ The analysis reveals a chance-level agreement (0.0013), suggesting GPT-3.5-turbo may be able to predict outputs without identifying the underlying function well, or vice versa. This reveals a dissociation between prediction accuracy and instruction induction.

In the colours domain, we examine both gpt models and conduct a similar analysis for each nonce word. We divide the test examples into examples containing each word, and examine the point-biserial correlation between accuracy of induced word meanings and correct translations in the few-shot context. This correlation is generally low, and few p-values are significant after correcting for multiple comparisons with false discovery rate (FDR) (Benjamini & Hochberg, 1995). See Appendix H for details.

In Kalamang, we repeat the process of computing point-biserial correlation between vocabulary induction correctness with segment-level chrF in the few-shot translations. Unlike the other domains, there is a small positive correlation between correct vocabulary hypotheses and chrF in gpt models. See Appendix M for details. We note that chrF is a more fine-grained measure than accuracy, and the initial scores were low enough that copying some correct vocabulary items may have had a slight impact on otherwise completely incorrect translations.

6 Related Work

Hypothesis Proposal with Language Models Recent work explores hypothesis proposals to improve language model performance in synthetic tasks, namely ACRE, the original MiniSCAN, ListOps, and versions of the ARC dataset (Qiu et al., 2023; Wang et al., 2023). These methods often rely on domain-specific interpreters or code generation by LMs, akin to our functions domain’s ground-truth reranker. We additionally explore probability-based reranking for rule selection across different domains and assesses the accuracy of model-induced rules prior to reranking.

Language models have also been used to automate hypothesis discovery as an end in itself, to discover distributional differences in text (Zhong et al., 2022; 2023). In this case, the hypothesis proposer LM is paired with a validator trained to filter out irrelevant hypotheses. We similarly find that hypotheses generated by language models themselves may not be very accurate inherently.

⁹This means agreement cannot be computed with the point-biserial correlation

Abductive Reasoning in Language Models Datasets for abductive commonsense reasoning (Wang et al., 2019; Bhagavatula et al., 2020; Zhang et al., 2020) and logical reasoning (Sinha et al., 2019; Yang et al., 2024) have been proposed. Bhagavatula et al. (2020) focuses on the ability for a model to select the more plausible hypothesis from pairs of hypotheses, while Yang et al. (2024) focuses on proposing natural language rules and evaluating them against ground-truth human rules. Zhao et al. (2023) use contrastive explanations to tune a model to recognize fluent ones. Comparatively, we do not focus on commonsense reasoning, and examine relationships between reasoning types on the same task.

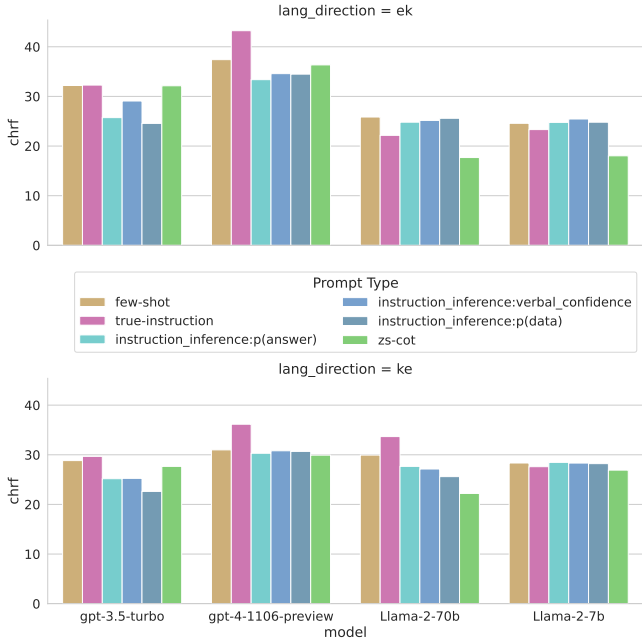


Figure 4: chrF scores for Kalamang under different methods, in English to Kalamang direction (top row) and Kalamang to English direction (bottom row)

Program Synthesis and Library Learning The general approach we outline for abductive reasoning can be considered a soft form of program synthesis with language models. Library learning, in which a set of reusable tools is learned from examples, is related (Ellis et al., 2020). A similar approach to ours using LLMs is proposed by Zhu et al. (2023), who use a similar two-step process to learn a library of rules for an arithmetic domain as well as a previous synthetic dataset (Sinha et al., 2019). We find similar gains on synthetic domains, and further examine challenges in applying abductive reasoning in complex domains.

LM-generated Instructions Instruction backtranslation is a related concept (Honovich et al., 2022; Li et al., 2024), in which an LM generates instructions for textual data. However, it differs in that the proposed instructions are not directly used at the same time that they are proposed, and it does not focus on generating rules.

7 Conclusion

We have examined the interplay between deductive, inductive, and abductive reasoning in LMs through the tasks of hypothesis proposal, in-context learning, and self-generated instruction following. Across three domains (linear function learning, artificial language translation, and Kalamang translation), we show that instruction inference is able to improve over few-shot prompting in simple synthetic domains, but that the relationship between these types of reasoning is complex, and they may not work together as expected presently when models are solving complex tasks. As abductive reasoning seems to be a relatively weaker capability in current language models as compared to instruction following, future work could develop more advanced mechanisms for natural-language hypothesis verification and correction. Enhancements in these areas could accelerate progress towards models capable of autonomous learning and self-improvement without the need for fine-tuning.

Acknowledgements

Thank you to Kiril Gashteovski, Ekin Akyürek, Shuyan Zhou, and Linlu Qiu for helpful discussions on this project. Additional thanks to David Mortensen for suggestions on grammar features and linguistic consultation. Thank you to NEC Labs Europe for supporting this project through the Student Research Fellowship program. We also acknowledge the support of the National Sciences and Engineering Research Council of Canada (NSERC) through a postgraduate fellowship given to EL (PGSD). We also appreciate support from Intel and the National Science Foundation through grants CCF-2217064 and IIS-2238240.

References

- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. What learning algorithm is in-context learning? investigations with linear models, 2023.
- Jacob Andreas, Dan Klein, and Sergey Levine. Learning with latent language. In *Proceedings of the Annual Meeting of the North American Chapter of the Association for Computational Linguistics*, 2018.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(1):289–300, 1995.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen tau Yih, and Yejin Choi. Abductive commonsense reasoning, 2020.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
- Igor Douven. Abduction. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition, 2021.
- Matthew S. Dryer and Martin Haspelmath (eds.). *WALS Online (v2020.3)*. Zenodo, 2013. doi: 10.5281/zenodo.7385533. URL <https://doi.org/10.5281/zenodo.7385533>.
- Kevin Ellis, Catherine Wong, Maxwell Nye, Mathias Sable-Meyer, Luc Cary, Lucas Morales, Luke Hewitt, Armando Solar-Lezama, and Joshua B. Tenenbaum. Dreamcoder: Growing generalizable, interpretable knowledge with wake-sleep bayesian program learning, 2020.
- H. Frankfurt. Peirce’s notion of abduction. *Journal of Philosophy*, 55:593–596, 1958.
- Shivam Garg, Dimitris Tsipras, Percy Liang, and Gregory Valiant. What can transformers learn in-context? a case study of simple function classes, 2023.
- James Hawthorne. Inductive Logic. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2021 edition, 2021.
- Or Honovich, Uri Shaham, Samuel R. Bowman, and Omer Levy. Instruction induction: From few examples to natural language task descriptions, 2022.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- Brenden M. Lake and Marco Baroni. Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. In *International Conference on Machine Learning*, 2017. URL <https://api.semanticscholar.org/CorpusID:46761158>.

Brenden M. Lake, Tal Linzen, and Marco Baroni. Human few-shot learning of compositional instructions. In *Annual Meeting of the Cognitive Science Society*, 2019. URL <https://api.semanticscholar.org/CorpusID:58006558>.

Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Omer Levy, Luke Zettlemoyer, Jason Weston, and Mike Lewis. Self-alignment with instruction backtranslation, 2024.

Peter Lipton. *Inference to the Best Explanation*. Routledge, 2001.

OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024.

- Karl Pearson. Notes on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London*, 58:240–242, 1895.
- Charles Sanders Peirce. *Collected Papers of Charles Sanders Peirce, Volume 5*, volume 5. Harvard University Press, 1965. URL <http://www.hup.harvard.edu/catalog.php?isbn=9780674138001>.
- Maja Popović. chrF: character n-gram F-score for automatic MT evaluation. In Ondřej Bojar, Rajan Chatterjee, Christian Federmann, Barry Haddow, Chris Hokamp, Matthias Huck, Varvara Logacheva, and Pavel Pecina (eds.), *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pp. 392–395, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/W15-3049. URL <https://aclanthology.org/W15-3049>.
- Linlu Qiu, Liwei Jiang, Ximing Lu, Melanie Sclar, Valentina Pyatkin, Chandra Bhagavatula, Bailin Wang, Yoon Kim, Yejin Choi, Nouha Dziri, and Xiang Ren. Phenomenal yet puzzling: Testing inductive reasoning capabilities of language models with hypothesis refinement, 2023.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2021.
- Stewart Shapiro and Teresa Kouri Kissel. Classical Logic. In Edward N. Zalta and Uri Nodelman (eds.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Spring 2024 edition, 2024.
- Koustuv Sinha, Shagun Sodhani, Jin Dong, Joelle Pineau, and William L. Hamilton. Clutrr: A diagnostic benchmark for inductive reasoning from text, 2019.
- Hedvig Skirgård, Hannah J. Haynie, Damián E. Blasi, Harald Hammarström, Jeremy Collins, Jay J. Latache, Jakob Lesage, Tobias Weber, Alena Witzlack-Makarevich, Sam Passmore, Angela Chira, Luke Maurits, Russell Dinnage, Michael Dunn, Ger Reesink, Ruth Singer, Claire Bower, Patience Epps, Jane Hill, Outi Vesakoski, Martine Robbeets, Noor Karolin Abbas, Daniel Auer, Nancy A. Bakker, Giulia Barbos, Robert D. Borges, Swintha Danielsen, Luise Dorenbusch, Ella Dorn, John Elliott, Giada Falcone, Jana Fischer, Yustinus Ghanggo Ate, Hannah Gibson, Hans-Philipp Göbel, Jemima A. Goodall, Victoria Gruner, Andrew Harvey, Rebekah Hayes, Leonard Heer, Roberto E. Herrera Miranda, Nataliia Hübler, Biu Huntington-Rainey, Jessica K. Ivani, Marilen Johns, Erika Just, Eri Kashima, Carolina Kipf, Janina V. Klingenberg, Nikita König, Aikaterina Koti, Richard G. A. Kowalik, Olga Krasnoukhova, Nora L.M. Lindvall, Mandy Lorenzen, Hannah Lutzenberger, Tônia R.A. Martins, Celia Mata German, Suzanne van der Meer, Jaime Montoya Samamé, Michael Müller, Saliha Muradoğlu, Kelsey Neely, Johanna Nickel, Miina Norvik, Cheryl Akinyi Oluoch, Jesse Peacock, India O.C. Pearey, Naomi Peck, Stephanie Petit, Sören Pieper, Mariana Poblete, Daniel Prestipino, Linda Raabe, Amna Raja, Janis Reimringer, Sydney C. Rey, Julia Rizaew, Eloisa Ruppert, Kim K. Salmon, Jill Sammet, Rhiannon Schembri, Lars Schlabach, Frederick W.P. Schmidt, Amalia Skilton, Wikaliler Daniel Smith, Hilário de Sousa, Kristin Sverredal, Daniel Valle, Javier Vera, Judith Voß, Tim Witte, Henry Wu, Stephanie Yam, Jingting Ye, Maisie Yong, Tessa Yuditha, Roberto Zariquiey, Robert Forkel, Nicholas Evans, Stephen C. Levinson, Martin Haspelmath, Simon J. Greenhill, Quentin D. Atkinson, and Russell D. Gray. Grambank reveals global patterns in the structural diversity of the world’s languages. *Science Advances*, 9, 2023. doi: 10.1126/sciadv.adg6175.
- Charles Spearman. The proof and measurement of association between two things. *American Journal of Psychology*, 15:72–101, 1904.
- Garrett Tanzer, Mirac Suzgun, Eline Visser, Dan Jurafsky, and Luke Melas-Kyriazi. A benchmark for learning to translate a new language from one grammar book, 2024.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas

- Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- Eline Visser. Kalamang dictionary. *Dictionaria*, (13):1–2737, 2020. URL <https://dictionaria.clld.org/contributions/kalamang>.
- Cunxiang Wang, Shuailong Liang, Yue Zhang, Xiaonan Li, and Tian Gao. Does it make sense? and why? a pilot study for sense making and explanation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4020–4026, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1393. URL <https://aclanthology.org/P19-1393>.
- Ruocheng Wang, Eric Zelikman, Gabriel Poesia, Yewen Pu, Nick Haber, and Noah D. Goodman. Hypothesis search: Inductive reasoning with language models, 2023.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
- Zonglin Yang, Li Dong, Xinya Du, Hao Cheng, Erik Cambria, Xiaodong Liu, Jianfeng Gao, and Furu Wei. Language models as inductive reasoners, 2024.
- Hongming Zhang, Xinran Zhao, and Yangqiu Song. Winowhy: A deep diagnosis of essential commonsense knowledge for answering winograd schema challenge, 2020.
- Wenting Zhao, Justin Chiu, Claire Cardie, and Alexander Rush. Abductive commonsense reasoning exploiting mutually exclusive explanations. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14883–14896, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.831. URL <https://aclanthology.org/2023.acl-long.831>.
- Ruiqi Zhong, Charlie Snell, Dan Klein, and Jacob Steinhardt. Describing differences between text distributions with natural language, 2022.
- Ruiqi Zhong, Peter Zhang, Steve Li, Jinwoo Ahn, Dan Klein, and Jacob Steinhardt. Goal driven discovery of distributional differences via language descriptions, 2023.
- Zhaocheng Zhu, Yuan Xue, Xinyun Chen, Denny Zhou, Jian Tang, Dale Schuurmans, and Hanjun Dai. Large language models can learn rules, 2023.

Appendix

A Model Inference Settings

Generation settings, as well as other details, used in the three domains are detailed here. Each model was tested 6 times for each setting (few-shot, zs-cot . . .) at different temperatures, and the aggregate results are shown in figures and tables throughout the paper. The exact model versions used for gpt models were gpt-3.5-turbo-0613 and gpt-4-1106-preview.

Functions When answering questions in the base task, gpt models were tested 3 times each at temperature $T = \{0, 1\}$ (note that $T = 0$ is nondeterministic in gpt models because of hardware). No max number of tokens was set for the generation. For llama models, $T = \{0.1, 1\}$ was used, also with no max number of tokens.

When generating hypotheses, we always used the same model as performed the base task. That is to say, gpt-3.5-turbo would generate hypotheses used by gpt-3.5-turbo, and so on. Hypotheses were always generated with a temperature above 0 to encourage generation of diverse hypotheses. For gpt models, hypotheses were generated with $T = 1$ and for llama models, with $T = 1.0625$. The specific value for llama models was because llama would usually generate the same hypothesis 5 times at $T = 1$, but higher values greatly increased the number of nonsensical and badly formatted hypotheses.

When self-evaluating hypotheses, the verbalized confidence score was generated at $T = 0$. When using the log-probabilities from text-davinci-002 to rerank hypotheses, the model was also set to $T = 0$.

Colours As in the functions domain, gpt models and llama models were respectively tested 3 times each at temperature $T = \{0, 1\}$ and $T = \{0.1, 1\}$ with no max number of tokens.

When generating hypotheses, $T = 1$ was used for all models.

When self-evaluating hypotheses, settings were the same as in the functions domain.

Kalamang For the base translation task, the same settings were used for generation as in Tanzer et al. (2024). Due to cost constraints, we ran only once in each translation direction on each setting with $T = 0.05$.

Vocabulary hypotheses for all models were generated with $T = 1$. Grammar feature hypotheses were generated independently of translations, and the first non-null hypothesis was chosen due to cost constraints. $T = 0.7$ was used for gpt grammar hypotheses, and $T = 1$ was used for llama grammar hypotheses.

When self-evaluating hypotheses, $T = 0$ was once again used for all models.

B Prompts for Linear Functions Domain

Table 5: Prompts for the functions domain. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The wording for the “prompt with self-induced hypothesis” and “zero-shot chain-of-thought prompt” were slightly changed from the few-shot examples prompt because models were sometimes confused by long-winded hypotheses, and responded with a long-winded answer in return, causing failures to parse their answers. In comparison, models mostly returned outputs in correct formats in other settings.

Prompt Type	Usage	Prompt Text
Base system prompt	For reasoning with in-context examples	You are a problem solving system. Your job is to use the input-output pairs to solve the problem as well as you can.
Hypothesis proposal system prompt	For proposing hypotheses based on in-context examples	You are a pattern recognition system. Your job is to come up with a function that describes the data as well as you can.
Instruction following system prompt	For applying a proposed hypothesis or ground-truth hypothesis to the input	You are a problem solving system. Your job is to apply the function to the data in order to produce an answer.
Few-shot examples prompt	For reasoning with in-context examples only	<p>Return the output preceded by 'Output:'</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Input: {query input}</p>
Prompt with ground-truth hypothesis	Used when prompting the model to directly apply the correct hypothesis to the input. In-context examples are also included.	<p>Function:</p> <p>{The real function}</p> <p>Examples:</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Return the output preceded by 'Output:'</p> <p>Input: {query input}</p>
Prompt for hypothesis induction	Used to have the model propose a single hypothesis for the function based on in-context examples	<p>Write the function that captures the relationship between inputs and outputs.</p> <p>You should write it in the form $y = ax^0 + bx^1$.</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Function (please write explicitly in the exact form "Output: $y = ax^0 + bx^1$):</p>

Table 5: Prompts for the functions domain. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The wording for the “prompt with self-induced hypothesis” and “zero-shot chain-of-thought prompt” were slightly changed from the few-shot examples prompt because models were sometimes confused by long-winded hypotheses, and responded with a long-winded answer in return, causing failures to parse their answers. In comparison, models mostly returned outputs in correct formats in other settings.

		Use this function to apply to the input example to get the correct output.
		{ model’s hypothesis }
Prompt with a self-induced hypothesis	Used similarly to the “prompt with ground-truth hypothesis”, except with a self-generated hypothesis. The wording is slightly changed.	<p>However, just write the output like what’s shown in these examples.</p> <p>Input: {input1} Output: {output1}</p> <p>Input: {input2} Output: {output2}</p> <p>...</p> <p>Return the output preceded by ‘Output:’</p> <p>Input: {query input}</p>
		Return the output preceded by ‘Final Output:’
		Input: {input1} Output: {output1}
		Input: {input2} Output: {output2}
		...
Prompt for zero-shot chain-of-thought	Used to encourage the model to generate a chain of thought.	<p>Let’s think step by step about what the function could be. Remember to write down ‘Final Output:’ before your final answer.</p> <p>Input: {query input}</p>
		How likely is this hypothesis about the function to be true given the data?
		Examples:
		Input: {input1} Output: {output1}
		Input: {input2} Output: {output2}
		...
Prompt for hypothesis probability estimate	Used to prompt a language model directly for reranking hypotheses	Function explanation: {model’s hypothesis}
		Please give a probability between 0 and 1 inclusive, and only answer with a number.
		Probability:
		These are examples of applying this function:
		{model’s hypothesis}
		Examples:
		Input: {input1} Output: {output1}
		Input: {input2} Output: {output2}
		...

C Prediction Fit of Other Models On Linear Functions

Figure 5 shows predictions made by each model when using a few-shot prompt, true instruction, or self-induced instruction.

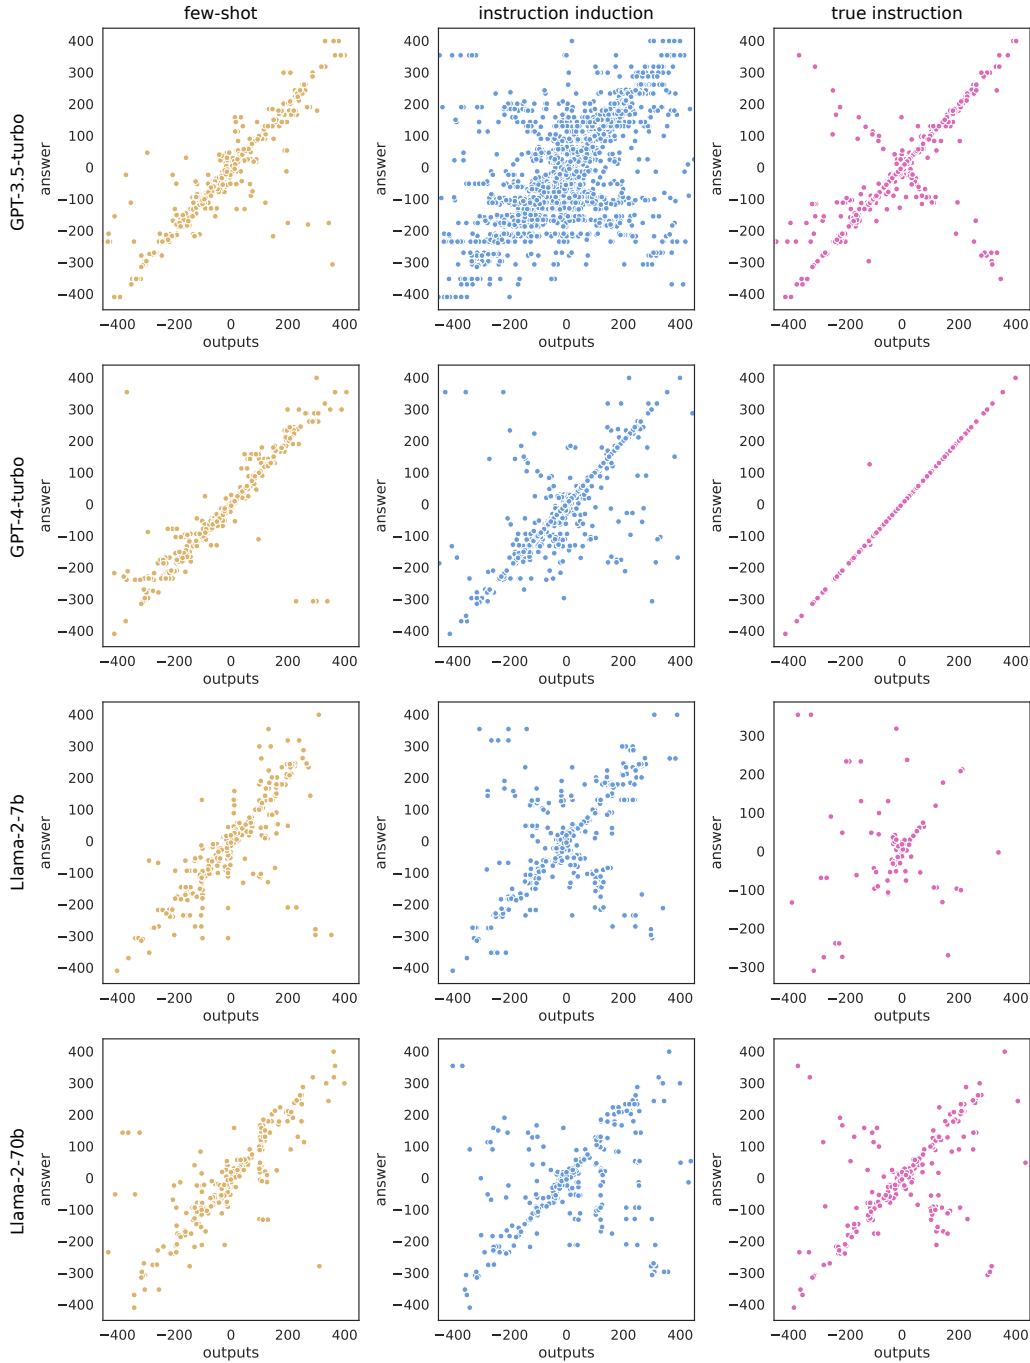


Figure 5: Model predictions plotted against true function output for all models. Range is restricted to the $[-400, 400]$ range for visualization purposes, although there are large outlier values for all models.

D Predicted Coefficients compared to Real Coefficients in Linear Functions

Figure 6 plots model hypotheses about the coefficients of x^0 and x^1 in a linear function against the actual coefficients

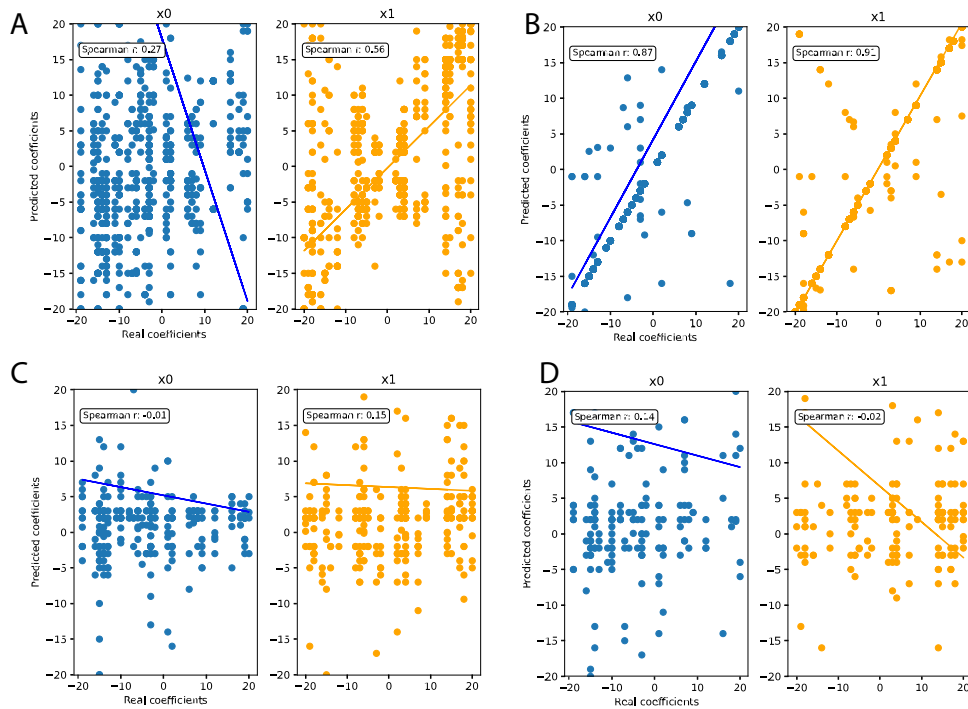


Figure 6: Model predictions plotted against true function output for all models. Range is restricted to the $[-400, 400]$ range for visualization purposes, although there are large outlier values for all models.

E Grammar and Details of Colours Domain

The rules of the colours domain are expressed through production rules below. Models were found to respond better to verbal instructions than formal ones in initial testing (e.g. a verbal statement "repeat twice" rather than $[[x]]b1uf \rightarrow [[x]][[x]]$)

Listing 1: Grammar of the colours language, as presented to LMs.

```

lug -> blue
dax -> green
wif -> red
zup -> yellow
bluf -> repeat the last action twice
walm -> repeat the last action three times

```

Training and test data was automatically generated by generating sentences of up to 5 nonce words on the source side, with shorter sentences being more likely (the respective probabilities for sentence lengths from 1 to 5 are $[0.4, 0.3, 0.15, 0.1, 0.05]$). Each colour could also be repeated with the repeat actions, and whether repeat nonce words were inserted was also random, though skewed toward no repetition (the probabilities were respectively $[0.8,$

0.1, 0.1] for no repeats, one repeat, or two repeats). A repeat term never followed another repeat term, and the same colour word never appeared twice consecutively to make it easier to learn the repeat terms.

Additionally, we provided a fixed set of few-shot examples that covered all the nonce terms:

Listing 2: Few-shot examples in the colours translation task

```

Input: lug dax
Output: blue green

Input: wif zup
Output: red yellow

Input: lug bluf
Output: blue blue

Input: wif walm
Output: red red red

Input: lug walm dax bluf
Output: blue blue blue green green

```

F Prompts for Colours Domain

Table 6: Prompts for the colours domain. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The “prompt with a self-induced hypothesis” was slightly modified from the base prompt in order to encourage models to follow the correct formatting, while the “prompt for zero-shot chain-of-thought” was slightly modified to encourage models to generate a concrete chain of thought, and also to follow the correct formatting.

Prompt Type	Usage	Prompt Text
Base system prompt	For reasoning with in-context examples	You are a problem solving system. Your job is to use the input-output pairs to solve the problem as well as you can.’’
Hypothesis proposal system prompt	For proposing hypotheses based on in-context examples	You are a rule induction system. Your job is to figure out the rules underlying a problem and report on them. Use the examples to guide your thinking.
Instruction following system prompt	For applying a proposed hypothesis or ground-truth hypothesis to the input	You are a parser. Carefully use the grammar to parse inputs to determine the correct output.
Few-shot examples prompt	For reasoning with in-context examples only	Return the output preceded by 'Output:' Input: {input1} Output: {output1} Input: {input2} Output: {output2} ... Input: {query input} Remember to start your answer with 'Output:'

Table 6: Prompts for the colours domain. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The “prompt with a self-induced hypothesis” was slightly modified from the base prompt in order to encourage models to follow the correct formatting, while the “prompt for zero-shot chain-of-thought” was slightly modified to encourage models to generate a concrete chain of thought, and also to follow the correct formatting.

Prompt with ground-truth hypothesis	Used when prompting the model to directly apply the correct hypothesis to the input. In-context examples are also included.	<p>Use this grammar to parse the input example to get the correct output.</p> <p>Grammar:</p> <p>{colours domain grammar}</p> <p>Examples:</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Return the output preceded by 'Output:'</p> <p>Input: {query input}</p>
Prompt for hypothesis induction	Used to have the model propose a single hypothesis for the translation of a word.	<p>The below examples contain the nonce word {word}.</p> <p>Using the examples, deduce what {word} means.</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Write your answer like this: word -> meaning.</p> <p>Meaning can be a word or a general rule dependent on the context.</p> <p>Rule:</p>
Prompt with a self-induced hypothesis	Used similarly to the “prompt with ground-truth hypothesis”, except with a self-generated hypothesis. The wording is slightly changed.	<p>Use this grammar to parse the input example to get the correct output.</p> <p>{ model’s hypothesis grammar }</p> <p>However, just write the output like what’s shown in these examples.</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Return the output preceded by 'Output:'</p> <p>Input: {query input}</p>

Table 6: Prompts for the colours domain. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The “prompt with a self-induced hypothesis” was slightly modified from the base prompt in order to encourage models to follow the correct formatting, while the “prompt for zero-shot chain-of-thought” was slightly modified to encourage models to generate a concrete chain of thought, and also to follow the correct formatting.

Prompt for zero-shot chain-of-thought	Used to encourage the model to generate a chain of thought.	<p>Return the output preceded by 'Final Output:'</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>...</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Let's think step by step about what the translation could be. Work through your answer step by step and show your work. Remember to write down 'Final Output:' before your final answer.</p> <p>Input: {query input}</p>
Prompt for hypothesis probability estimate	Used to prompt a language model directly for reranking hypotheses	<p>These are examples of the translation of the word {word}.</p> <p>Examples:</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>...</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p> <p>Given these examples, how likely is this hypothesis about the meaning of {word}? {model's word hypothesis}</p> <p>Please give a probability between 0 and 1 inclusive, and only answer with a number.</p> <p>Probability:</p>
Prompt for data logprobs given hyp estimate	Used to rerank hypotheses based on logprobs.	<p>These are examples of applying this function: {model's word hypothesis}</p> <p>Examples:</p> <p>Input: {input1}</p> <p>Output: {output1}</p> <p>...</p> <p>Input: {input2}</p> <p>Output: {output2}</p> <p>...</p>

G chrF For Colours Domain

Figure 7 depicts the mean chrF score over 6 trials for each model in each setting. Reranking methods are averaged for instruction-induction.

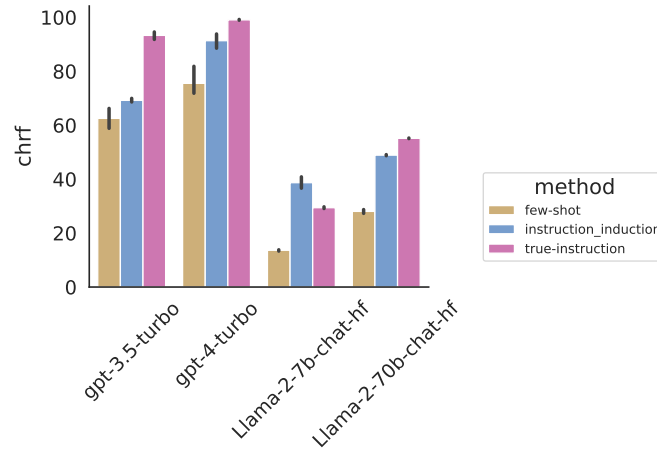


Figure 7: chrF scores for the colours domain under different methods. All reranking methods are averaged.

H Vocabulary Induction Accuracy for Colours Domain

We show in Table 7 the mean accuracy for each colour term for gpt models, computed over all selected (highest-ranked) word hypotheses. Point-biserial correlation between hypothesis correctness (in the instruction-induction case) and few-shot correctness is also shown. P-values are corrected with FDR in the trials for each model. We performed the correction per model rather than across all models because we were separately interested in the behaviour of each model. We parsed the production rule generated by the model in order to determine hypothesis correctness.

Table 7: Vocabulary induction correctness and relation to few-shot translation correctness for gpt models. Significant correlations after correction for multiple comparisons are bolded.

Model	Word	method	Mean Correctness	Correlation	p-value	p-value corrected
gpt-3.5-turbo	lug	verbal.conf	0.604859335038363	-0.0370824643671235	0.300350893877687	0.52561406
gpt-3.5-turbo	dax	verbal.conf	0.233072916666666	0.0669078409779957	0.0638449077681366	0.29564626
gpt-3.5-turbo	wif	verbal.conf	0.329268292682926	0.0360864838621108	0.646415727402187	0.82271093
gpt-3.5-turbo	zup	verbal.conf	0.61	0.154599744147798	0.0288284171902121	0.20179892
gpt-3.5-turbo	bluf	verbal.conf	0	-	-	-
gpt-3.5-turbo	walm	verbal.conf	0	-	-	-
gpt-4-turbo	lug	verbal.conf	0.875763747454175	0.0439535251603629	0.168735841765082	3.58e-1
gpt-4-turbo	dax	verbal.conf	0.921487603305785	-0.0418339829001646	0.193444842357041	3.58e-1
gpt-4-turbo	wif	verbal.conf	0.865853658536585	0.0447789368805147	0.569118855838926	7.32e-1
gpt-4-turbo	zup	verbal.conf	0.87	0.0579317588377811	0.415170703498878	5.75e-1
gpt-4-turbo	bluf	verbal.conf	0.153846153846153	0.130289326661761	0.255554806693343	3.83e-1
gpt-4-turbo	walm	verbal.conf	0.171428571428571	0.0489476613672	0.687394764355678	7.55e-1
gpt-3.5-turbo	lug	p_data	0.28772378516624	0.0591580864392629	0.0983076746979401	0.29564626
gpt-3.5-turbo	dax	p_data	0.10546875	0.00996669750420649	0.782730731634431	0.84294079
gpt-3.5-turbo	wif	p_data	0.195121951219512	0.24388397206115	0.00164950240178991	0.02309303
gpt-3.5-turbo	zup	p_data	0.4	-0.114775306453679	0.105587948530317	0.29564626
gpt-3.5-turbo	bluf	p_data	0.0769230769230769	-0.1673596703488	0.143038746170574	0.33375707
gpt-3.5-turbo	walm	p_data	0	-	-	-
gpt-4-turbo	lug	p_data	0.731160896130346	-0.0803208788270973	0.0118067271484766	5.31e-2
gpt-4-turbo	dax	p_data	0.733471074380165	-0.0411584340901287	0.200743028038944	3.58e-1
gpt-4-turbo	wif	p_data	0.768292682926829	-0.0255095864281423	0.745759257358691	7.55e-1
gpt-4-turbo	zup	p_data	0.76	0.0221788843765235	0.755246389745868	7.55e-1
gpt-4-turbo	bluf	p_data	0.102564102564102	-0.140859042454752	0.218662613923779	3.58e-1
gpt-4-turbo	walm	p_data	0.142857142857142	-0.263591223386182	0.0274677998646683	9.89e-2
gpt-3.5-turbo	lug	p_answer	0.436061381074168	-0.0128717673177528	0.719304414221268	0.83918848
gpt-3.5-turbo	dax	p_answer	0.190104166666666	0.00399825404003988	0.9119154429019	0.91191544
gpt-3.5-turbo	wif	p_answer	0.353658536585365	-0.0743944887571608	0.343771340895407	0.53475542
gpt-3.5-turbo	zup	p_answer	0.59	-0.0879883939992145	0.21536453378319	0.43072907
gpt-3.5-turbo	bluf	p_answer	0.0512820512820512	-0.0551410966570355	0.631586163340137	0.82271093
gpt-3.5-turbo	walm	p_answer	0	-	-	-
gpt-4-turbo	lug	p_answer	0.70468431771894	-0.0661341939950676	0.0382590601605542	1.15e-1
gpt-4-turbo	dax	p_answer	0.712809917355371	0.188125007589795	3.67e-9	6.61e-8
gpt-4-turbo	wif	p_answer	0.817073170731707	-0.249727938065423	0.00126030696185871	1.13e-2
gpt-4-turbo	zup	p_answer	0.68	-0.142141439177509	0.0446656422362514	1.15e-1
gpt-4-turbo	bluf	p_answer	0.23076923076923	-0.294147299400922	0.00894893196851883	5.31e-2
gpt-4-turbo	walm	p_answer	0.0571428571428571	-0.0578005412233204	0.634584373040489	7.55e-1

I MTOB Experimental Settings

We generally used the same experimental setup as in `mtob`. For the few-shot condition, we used two reference sentences for each word on the source side, selected via longest subsequence. For the `true-instruction` setting, we used reference sentences, in addition to the wordlist and true grammar sketch. For each word in the source sentence, the most similar word from the wordlist was also retrieved based on the longest common substring. In the `instruction-inference` settings, the model’s self-induced grammar sketch was substituted for the true grammar sketch, and the model was also prompted to first create hypotheses about the translation of each word in the source sentence.

At the time that we conducted experiments, the Kalamang to English training set was not available, so we created this training set by reversing the source and target in the English to Kalamang training set. This may have slightly harmed translations in this direction, as this caused the training set to be different from the one reported in the benchmark results.

J MTOB Prompts

Table 8: Prompts for Kalamang translation with MTOB. We follow the same prompt formatting as in Tanzer et al. (2024), with the exception of replacing retrieved grammar book passages with the grammar sketch in Appendix K. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like `{this}`. The English to Kalamang direction is depicted in the examples. ‘’’ indicates that the prompt is the same as in Appendix F.

Prompt Type	Usage	Prompt Text
Base system prompt	For reasoning with in-context examples	’’’
Hypothesis proposal system prompt (Grammar)	For proposing grammar feature hypotheses based on in-context examples	’’’
Hypothesis proposal system prompt (Vocab)	For proposing vocabulary mappings based on in-context examples	’’’
Instruction following system prompt	For applying a proposed hypothesis or ground-truth hypothesis to the input	’’’

Table 8: Prompts for Kalamang translation with MTOB. We follow the same prompt formatting as in Tanzer et al. (2024), with the exception of replacing retrieved grammar book passages with the grammar sketch in Appendix K. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The English to Kalamang direction is depicted in the examples. ‘’’ indicates that the prompt is the same as in Appendix F.

Few-shot examples prompt	For reasoning with in-context examples only	<p>Human: Kalamang is a language spoken on the Karas Islands in West Papua. Translate the following sentence from English to Kalamang: {query input}</p> <p>To help with the translation, here is a translated sentence with words similar to ‘{word}’ in a list of translated Kalamang-English reference sentences:</p> <p>English sentence: {eng sentence} Kalamang translation: {kgv sentence} <Repeated for all words> Now write the translation. English: {query input} Kalamang translation:</p>
Prompt with ground-truth hypothesis	Used when prompting the model to directly apply the correct hypothesis to the input. In-context examples are also included.	<p>Human: Kalamang is a language spoken on the Karas Islands in West Papua. Translate the following sentence from English to Kalamang: {query input}</p> <p>To help with the translation, here is one of the closest entries to ‘{word}’ in the Kalamang=English bilingual dictionary: English word or phrase: ‘{word}’ Kalamang translation: ‘{translation}’ <Repeated for all words></p> <p>To help with the translation, here is a translated sentence with words similar to ‘{word}’ in a list of translated Kalamang-English reference sentences:</p> <p>English sentence: {eng sentence} Kalamang translation: {kgv sentence} <Repeated for all words></p> <p>To help with the translation, here’s a grammar sketch of Kalamang: {grammar sketch}</p> <p>Now write the translation. English: {query input} Kalamang translation:</p>
Prompt for hypothesis induction	Used to have the model propose a single hypothesis for the translation of a word.	<p>The following sentences contain the word ‘{word}’ in English.</p> <p>Examples:</p> <p>English sentence: {eng sentence} Kalamang translation: {kgv sentence} ...</p> <p>What is the Kalamang translation of the word ‘{word}’? Write your answer like this: {English word} -> {Kalamang translation}.</p>

Table 8: Prompts for Kalamang translation with MTOB. We follow the same prompt formatting as in Tanzer et al. (2024), with the exception of replacing retrieved grammar book passages with the grammar sketch in Appendix K. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The English to Kalamang direction is depicted in the examples. ‘’’ indicates that the prompt is the same as in Appendix F.

Prompt with a self-induced hypothesis	Used similarly to the “prompt with ground-truth hypothesis”, except with a self-generated hypothesis. The wording is slightly changed.	<p>Human: Kalamang is a language spoken on the Karas Islands in West Papua. Translate the following sentence from English to Kalamang: {query input}</p> <p>To help with the translation, here is one of the closest entries to ‘{word}’ in the Kalamang=English bilingual dictionary: English word or phrase: ‘{word}’ Kalamang translation: ‘{hypothesis translation}’ <Repeated for all words></p> <p>To help with the translation, here is a translated sentence with words similar to ‘{word}’ in a list of translated Kalamang-English reference sentences:</p> <p>English sentence: {eng sentence} Kalamang translation: {kgv sentence} <Repeated for all words></p> <p>To help with the translation, here’s a grammar sketch of Kalamang: {hypothesis grammar sketch}</p> <p>Now write the translation. English: {query input} Kalamang translation:</p>
Prompt for zero-shot chain-of-thought	Used to encourage the model to generate a chain of thought.	<p>Human: Kalamang is a language spoken on the Karas Islands in West Papua. Translate the following sentence from English to Kalamang: {query input}</p> <p>To help with the translation, here is a translated sentence with words similar to ‘{word}’ in a list of translated Kalamang-English reference sentences:</p> <p>English sentence: {eng sentence} Kalamang translation: {kgv sentence} <Repeated for all words></p> <p>Let’s think step by step about what the translation could be. Now write the translation. English: {query input} Kalamang translation:</p>

Table 8: Prompts for Kalamang translation with MTOB. We follow the same prompt formatting as in Tanzer et al. (2024), with the exception of replacing retrieved grammar book passages with the grammar sketch in Appendix K. Newlines are depicted visually for ease of reading. Variables that are substituted depending on the question are marked like {this}. The English to Kalamang direction is depicted in the examples. ‘’’ indicates that the prompt is the same as in Appendix F.

		These are examples of the translation of the word {word}.
		Examples: {example translation pairs} ...
Prompt for hypothesis probability estimate	Used to prompt a language model directly for reranking hypotheses	Given these examples, how likely is this hypothesis about the meaning of {word}? {model’s word hypothesis}
		Please give a probability between 0 and 1 inclusive, and only answer with a number.
		Probability:
		These are examples of sentences that contain the word {word}: This is the translation of the word: {model’s word hypothesis}
Prompt for data logprobs given hyp estimate	Used to rerank hypotheses based on logprobs.	Examples: {example translation pairs}

K Kalamang Grammar Sketch

Listing 3: Grammar sketch for Kalamang, compiled from WALS and Grambank features.

```

=== Start of grammar sketch ===
Basic Word Order: SV (Subject-Verb), OV (Object-Verb)
Noun Phrase Construction: Postpositional, Genitive-Noun, Noun-Adjective, Noun-
    Demonstrative, Noun-Num, Possessed-Possessor
Articles: No definite/specific or indefinite articles
Morphological Marking: No productive singular, dual, or plural marking on nouns
    ; Possession marked by suffix on the possessed noun; Tense marked by
    auxiliary particle
Syntactic Alignment: Accusative
Negation: Standard negation marked clause-finally; Distinct imperative negation
Reduplication: Both verbs and nouns can be reduplicated
=== End of grammar sketch ===

```

L Predicted Kalamang Grammar Sketches and Accuracy

Listing 4: Grammar sketch for Kalamang, generated by GPT-3.5-turbo. Each feature was predicted individually based on randomly retrieved examples from the training data.

```

=== Start of grammar sketch ===
Basic Word Order: SV (Subject-Verb), OV (Object-Verb)
Noun Phrase Construction: Adposition-Noun Phrase, Genitive-Noun, Adjective-Noun
    , Noun-Demonstrative, Noun-Num, Possessor-Possessed
Articles: Unsure if there are definite/specific or indefinite articles
Morphological Marking: Unsure if singular, dual, or plural marking on nouns;
    Unsure if possession marked by suffix on the possessed noun; Unsure if
    tense marked by auxiliary particle

```

Syntactic Alignment: Unsure
 Negation: Standard negation marked clause-initially; Unsure if distinct imperative negation
 Reduplication: Unsure if verbs and nouns can be reduplicated
 === End of grammar sketch ===

Listing 5: Grammar sketch for Kalamang, generated by GPT-4-turbo. Each feature was predicted individually based on randomly retrieved examples from the training data.

=== Start of grammar sketch ===
 Basic Word Order: SV (Subject-Verb), VO (Verb-Object)
 Noun Phrase Construction: Unsure on order of adposition and noun phrase, Genitive-Noun, Adjective-Noun, Unsure on order of noun and demonstrative, Num-Noun, Possessor-Possessed
 Articles: No definite/specific or indefinite articles
 Morphological Marking: Unsure if productive singular, dual, or plural marking on nouns; Unsure if Possession marked by suffix on the possessed noun; Unsure if tense marked by auxiliary particle
 Syntactic Alignment: Unsure
 Negation: Standard negation marked clause-initially; Unsure if distinct imperative negation
 Reduplication: Unsure if verbs and nouns can be reduplicated
 === End of grammar sketch ===

Model	Grammar feature accuracy
GPT-3.5-turbo	27.78%
GPT-4-turbo	22.22%
Llama-2-7B	0%
Llama-2-70B	0%

Table 9: Overall grammar sketch accuracy of models. Unsure answers were marked as incorrect.

M Vocabulary Induction Accuracy for MTOB

The overall vocabulary induction accuracy for gpt models is shown in Table 10. Accuracy and chrF were calculated overall when exclude morphology is marked No, and morphologically rich words in the dictionary (marked with * or with - to represent suffixes/prefixes) were excluded when this value was marked Yes. For words with multiple translations in the dictionary, the translation was marked correct if it matched any of the possible translations. Words without an entry in the dictionary were skipped. If a null hypothesis such as *I don't know* was proposed, it was marked as incorrect.

Due to hypotheses for this task being more difficult to automatically evaluate, an author also annotated the correctness of results for a single setting (instruction-inference: verbalized confidence) in order to calculate the association between translation quality and hypothesis correctness more accurately. Results can be found in Table 11.

Table 10: Accuracy and segment-level chrF for vocabulary hypotheses in Kalamang translation. If morphology was excluded, all words with a * or - symbol in the translation (corresponding to prefixes/suffixes) were excluded, whereas otherwise we matched whether or not the model’s hypothesis began/ended with the correct characters.

direction	model	rerank method	exclude morphology	acc	chrF
ek	gpt-4-turbo	p_data.given.hyp.guess	No	0.1528	20.33
ek	gpt-4-turbo	p_data.given.hyp.guess	Yes	0.0813	11.86
ek	gpt-4-turbo	p_answer.given.hyp.logprobs	No	0.16	20.69
ek	gpt-4-turbo	p_answer.given.hyp.logprobs	Yes	0.08621	12.18
ek	gpt-4-turbo	p_data.given.hyp.logprobs	No	0.166	20.79
ek	gpt-4-turbo	p_data.given.hyp.logprobs	Yes	0.1015	13.21
ke	gpt-3.5-turbo	p_answer.given.hyp.logprobs	No	0.1033	15.33
ke	gpt-3.5-turbo	p_answer.given.hyp.logprobs	Yes	0.8244	13.11
ke	gpt-3.5-turbo	p_data.given.hyp.guess	No	0.11	16.42
ke	gpt-3.5-turbo	p_data.given.hyp.guess	Yes	0.09018	14.33
ke	gpt-3.5-turbo	p_data.given.hyp.logprobs	Yes	0.0738	12.34
ke	gpt-3.5-turbo	p_data.given.hyp.logprobs	No	0.1014	15.28
ke	Llama-2-70b	p_data.given.hyp.guess	Yes	0.1344	21.69
ke	Llama-2-70b	p_data.given.hyp.guess	No	0.155	23.62
ek	gpt-3.5-turbo	p_data.given.hyp.guess	Yes	0.0922	15.05
ek	gpt-3.5-turbo	p_data.given.hyp.guess	No	0.1217	19.29
ke	gpt-4-turbo	p_answer.given.hyp.logprobs	Yes	0.08127	13.54
ke	gpt-4-turbo	p_answer.given.hyp.logprobs	No	0.0979	15.25
ke	gpt-4-turbo	p_data.given.hyp.logprobs	Yes	0.07641	12.82
ke	gpt-4-turbo	p_data.given.hyp.logprobs	No	0.0924	14.46
ke	Llama-2-70b	p_answer.given.hyp.logprobs	Yes	0.1339	21.92
ke	Llama-2-70b	p_answer.given.hyp.logprobs	No	0.1447	22.92
ke	Llama-2-70b	p_data.given.hyp.logprobs	Yes	0.1334	21.36
ke	Llama-2-70b	p_data.given.hyp.logprobs	No	0.1442	22.38

Model	Correlation	p-val
GPT-3.5-turbo	0.1746	0.006331
GPT-4-turbo	0.1201	0.05088

Table 11: Point-biserial correlation between segment-level chrF of each translation with the correctness of vocabulary hypotheses induced for that question (in the instruction-inference:verbalized confidence setting). No correction for multiple comparisons was done because we examined hypotheses at the model level (see Appendix H for how this was done in the colours domain)

N Example Kalamang Translations

Table 12 shows example translations made by GPT-3.5-turbo in both directions.

Table 12: Example translations by the GPT-3.5-turbo model in different settings.

Source	Reference	Translation (few-shot)	Translation (instruction-induction: verbalized confidence)	Translation (true-instruction)
Who asked you? kiemara paruo?	Namana kat gerket?	Naman mua kiemara kowam?	Naman era?	Naman mua
The rope is wrapped around it.	Karek mat kokarap.	Kainasu esun kiwat semda bo dalangdalang.	Sar tok esun et dalangdalang na.	Karek mambon guanggarien ma he me.
Before, they said Japan and The Netherlands were at war.	Wiseme mu toni Japangbon Beladarbon nausair.	Wiseme mu Jafaris emun Nabil emun kaden kiesun ning.	Wise me mu salai Japan an Sar Beladargara din jambat mara.	Tikninda, termus siram panci tubak Beladar tamatko pang.
Mu bo Tamisen-gga som dande yecie.	They are going (on a return trip) to bury someone in Antalisa.	They go to Tamisengga and return.	They go to Antalisa, someone died, and they return.	They will return to Tamisengga until they bury the person.
Suo me kene arkon metko.	At Suo is a kene tree.	I washed my body in the river.	Back me gives birth canoe.	Wash your face near the Narkon tree.
An istrat kahendengoa marmar ba temun.	I'm walking on a long wide road.	I put the long stick on the ground.	I have a long drying rack but I am walking.	I walked far down the street to my uncle.