

# EVALUATING THE INSTRUCTION-FOLLOWING ABILITIES OF LANGUAGE MODELS USING KNOWLEDGE TASKS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In this work, we focus our attention on developing a benchmark for instruction-following where it is easy to verify both task performance as well as instruction-following capabilities. We adapt existing knowledge benchmarks and augment them with instructions that are a) conditional on correctly answering the knowledge task or b) use the space of candidate options in multiple-choice knowledge-answering tasks. This allows us to study model characteristics, such as their change in performance on the knowledge tasks in the presence of answer-modifying instructions and distractor instructions. In contrast to existing benchmarks for instruction following, we not only measure instruction-following capabilities but also use LLM-free methods to study task performance. We study a series of openly available large language models of varying parameter sizes (1B-405B) and closed source models namely GPT-4o-mini, GPT-4o. We find that even large-scale instruction-tuned LLMs fail to follow simple instructions in zero-shot settings. We release our dataset, the benchmark, code, and results for future work.

## 1 INTRODUCTION

The growth of increasingly powerful large language models has resulted in the development of end-user applications including assistants for coding and software engineering (Ozkaya, 2023; Zhang et al., 2023; Ross et al., 2023), workflow and business automations (Grohs et al., 2023; Wornow et al., 2024), self-help assistants (Zhou et al., 2020; Shuster et al., 2022) and more. The need for highly accurate and controllable systems that follow precise instructions have led to the development of methods to improve reliability and consistency in the output for LLMs. Such methods include few-shot prompting (Gao et al., 2020; Kojima et al., 2022b), reasoning with explanations (Wei et al., 2022; Huang & Chang, 2022), checking for consistency/self-consistency (Wang et al., 2022), use of intermediate evaluators or LLMs operating as judges (Zheng et al., 2023), etc.

While there has been a lot of focus on assessing the knowledge of LLMs (Brown et al., 2020; Heizerling & Inui, 2021), logical reasoning (Hendrycks et al., 2021; Wei et al., 2022; Ma et al., 2024), programmatic ability (Dakhel et al., 2023; Chen et al., 2021), problem solving ability (Lightman et al., 2024), etc, the study of their ability to follow precise instructions is relatively nascent; works such as FoFo (Xia et al., 2024), InFoBench (Qin et al., 2024), RuleBench (Sun et al., 2024), IFEval (Zhou et al., 2023b) attempt to address this gap. While FoFo Xia et al. (2024) assesses the ability of models to generate outputs conforming to existing real-world output formats such as the HL7-CDA format used in Healthcare applications, RuleBench (Sun et al., 2024) assesses a model’s capabilities on inferential rule-following using rules which can be encoded in instructions and first-order-logic. On the other hand, benchmarks such as InFoBench (Qin et al., 2024) and IFEval (Zhou et al., 2023b) assess the ability of LLMs to follow arbitrary task specific instructions though neither InFoBench nor IFEval provide easy ways of verifying (i) task success and (ii) instruction following capabilities simultaneously (see Section 2 for a detailed discussion). **Constraints such as formatting style, length are harder to verify along with task performance (not just instruction performance) while assessing instruction-following with verifiable is limited.**

In this work, we focus our attention on developing a benchmark for instruction-following where it is easy to verify both task performance as well as instruction following capabilities. We adapt existing

Table 1: Comparison with existing Instruction Following benchmarks

Benchmark	Task	Deterministic Outputs	Content Verification	QA-Conditioned Instructions	Evaluator
FoFo Eval (Xia et al., 2024)	Format following	✓	✗	✗	LLM
IFEval (Zhou et al., 2023b)	Instruction Following	✓	✗	✗	Direct
InFoBench (Qin et al., 2024)	Instruction Following	✗	✓	✗	LLM
RuleBench (Qin et al., 2024)	Inferential Rule Following	✓	✓	✓	Direct
This work	Instruction Following	✓	✓	✓	Direct

commonly used knowledge benchmarks including MMLUPro (Wang et al., 2024), MathQA Amini et al. (2019), Winogrande (Sakaguchi et al., 2021), BoolQ (Clark et al., 2019), PIQA Bisk et al. (2020) and augment them with two broad classes of instructions: (i) Instructions that are conditional on the answer to the question (ii) Instructions that are applied uniformly regardless of the answer or task. We include a detailed study of multiple LLMs and find that even the largest models have trouble following relatively simple instructions. [Our list of instructions demonstrate: 1. Simple changes of return text instead of labels results in a drop. 2. Simple tasks of counting, concatenation, conditional exclusion/inclusion/application as well as distracting instructions all result in significant drop in performance. To the best of our knowledge, there is no prior work that demonstrates this with verifiable results and bench-marking of current models.](#)

Our contributions are as follows: (i) We release the first benchmark that assesses the zero-shot instruction-following performance of models using knowledge and reasoning question-answering (QA) tasks. (ii) We employ multiple QA-conditioned instructions to examine instruction-following performance across different instruction classes, including those dependent on answer-type. (iii) We include instruction instances that serve as distractors for the original knowledge-tasks (iv) Unlike previous studies, we use LLM-free evaluation metrics to assess both knowledge and instruction-following abilities. (v) We offer automated error analysis measures, pre-classifying likely errors for each instruction instance. (vi) Our benchmark creation method is easy to extend to new instructions and datasets.

## 2 RELATED WORK

Evaluating the capabilities of large language models (LLMs) has been a significant area of research, with studies focusing on various aspects of LLM performance. Researchers have developed multiple benchmarks to assess factual knowledge Petroni et al. (2019); Roberts et al. (2020); Lin et al. (2022), logical reasoning abilities Wei et al. (2022); Zhou et al. (2023a); Saparov et al. (2023), general problem-solving capabilities Kojima et al. (2022a) and more.

Recently there have also been studies on instruction-following - for instance, FoFo Xia et al. (2024) evaluates models on format-following tasks and studies the ability of LLMs to generate outputs in existing real-world formats. In a similar vein, IFEval Zhou et al. (2023b) assesses LLMs’ ability to follow arbitrary task-specific instructions (e.g.) based on response length, casing, etc, focusing primarily on whether the instructions are followed rather than the correctness of the output for the task. InFoBench Qin et al. (2024) advances this research by introducing a metric known as the ‘Decomposed Requirements Following Ratio’ (DRFR) which is based on each aspect of an instruction that needs to be met. Along with 500 diverse instructions and 2,250 decomposed questions, InFoBench offers performance evaluation using OpenAI’s GPT4, across multiple constraint categories and highlights key areas where advanced LLMs can improve in complex instruction-following tasks. LLMBAR (Zeng et al., 2024) is another contribution to this area, as it provides a meta-evaluation benchmark specifically designed to test an LLM evaluator’s ability to discern instruction-following outputs. The benchmark consists of 419 manually curated pairs of outputs, where one output adheres to instructions and the other, while potentially more engaging or deceptive, does not. [Li et al. \(2024\) propose a method to evaluate instruction following ability via verbalizer manipulation. Specifically, they modify the classification task labels with different verbalizers which may or may not be semantically relevant to the task. They observe that all models fail to follow instructions when they instruct the model to flip the labels \(unnatural setting\). They evaluate the framework on mostly traditional NLP tasks like Sentiment Analysis, textual entailment etc.](#)

Our work builds upon these efforts by developing a benchmark that allows for easy verification of both task performance and instruction-following capabilities simultaneously. We augment existing knowledge benchmarks by creating instructions that are *conditional* on answering the QA-based knowledge task correctly. We also include instructions that are applied on the candidate space of answers provided in these knowledge tasks. Our approach of applying instructions on knowledge tasks provides an easy way of measuring performance. Further, it also allows us to study the interactions between knowledge and instruction following, and to investigate whether instructions serve as distractors for the original knowledge task when the instructions should result in no change to the original answer of the knowledge task.

### 3 INSTRUCTION-FOLLOWING EVALUATION DATASET

We now describe the process for creating our evaluation dataset.

**Design Principles:** We develop our instructions keeping the following design principles in mind: (i) We would like instructions to be unambiguous and be presented in a way that can be communicated clearly - if humans cannot follow the instructions and agree on the same output, LLMs should and likely would not be able to. (ii) We would like them to be easy to follow and not require complex reasoning abilities to follow so that models at all scales have a fair chance of success, (iii) The instructions need to have deterministic outputs that use the original answers of the knowledge-task or the candidate space of answers, or both, so that they can be evaluated easily with instruction specific scorers. (iv) We would like our benchmark to be based on a diverse mix of knowledge tasks, and be easily extensible to new ones.

#### 3.1 KNOWLEDGE AND REASONING TASKS

We select the following knowledge tasks that are commonly used in LLM evaluations as the basis for our instruction-following benchmark. These datasets involve either binary classification or multiple-choice-questions (MCQs) spanning different reasoning and problem-solving skills.

(i) **MMLUPro** (Wang et al., 2024): MMLUPro extends the MMLU dataset to make it more challenging by a) increasing the number of options from four to ten and b) increasing problem difficulty by focusing on more reasoning oriented problems. We consider all 14 subjects in the MMLUPro benchmark. We cap the maximum number of samples for each subject to be 150 samples.

(ii) **MathQA** (Amini et al., 2019): MathQA dataset consists of math word problems presented as Multiple-Choice-Questions (MCQs). Given a math question and four options, the model has to select the correct answer.

We also select a few common-sense and reasoning datasets:

(iii) **BoolQ** (Clark et al., 2019): BoolQ is a boolean question-answering dataset. Given a passage and a boolean question around the passage, the model has to select either *True* or *False*.

(iv) **PIQA** (Bisk et al., 2020): Physical Interaction: Question Answering (PIQA) involves answering questions that involve commonsense reasoning around physical objects. Given a question and two options, the model has to select the most plausible option.

(v) **Winogrande** (Sakaguchi et al., 2021): Winogrande involves a fill-in-the-blank task with binary options, the model has to select the correct option for a given sentence. The task involves reasoning for pronoun resolution.

We select a subset of 1500 samples randomly from each of the above datasets.

#### 3.2 INSTRUCTION CATEGORIES

Unlike datasets that require open-ended generation for answering, our selected tasks have a structured answer-space. This allows us to craft instructions using these answer-spaces in a way that can be verified easily. We define the following instruction categories.

Table 2: Categories of instructions and the number of instances of each in the Full and Lite subsets.

Instruction Group	Name	Definition	# Instances	
			Full	Lite
String Manipulation	<code>alternate_case_correct_answer</code>	Print the text corresponding to the correction candidate answer of knowledge task in alternate case	7867	950
	<code>capitalize_correct_answer</code>	Print the text corresponding to the correct candidate answer of the knowledge task in upper case.	7867	950
	<code>reverse_correct_answer_alternate_case</code>	Reverse the text corresponding to the correct candidate answer of the knowledge task and print it in alternate case.	9573	1383
	<code>reverse_correct_answer</code>	Print the text corresponding to the correct answer in reverse	7868	951
Format Correct Answer	<code>numformat_numeric_answer</code>	Apply a specified decimal formatting the correct answer if it is a numeric quantity, otherwise print the correct answer as is.	11336	1600
	<code>print_correct_answer_in_words</code>	If the correct answer is a numeric quantity, display the numeric quantity in words, otherwise print the correct answer as is.	9874	1320
	<code>print_correct_answer_append_string</code>	Append a pre-specified string to the text associated with the correct candidate answer.	7867	950
Operations on List (Conditional on Correct Answer)	<code>increment_incorrect_numeric_answers_by_one</code>	If the candidate answer values are numeric quantities increment them by one and show them as a list. Other value types are not modified.	7117	825
	<code>sort_only_incorrect_answers</code>	Sort the candidate answers that are incorrect in ascending order	7867	950
	<code>use_incorrect_options_to_create_string</code>	Sort the incorrect candidates in ascending order and take the last character of the text associated with each incorrect option to create a string	7868	951
Operations on List	<code>sort_options_to_create_string</code>	Sort all candidate answers in ascending order and use the last character of the text associated with each incorrect candidate to create a string.	7867	950
Numeric Manipulation	<code>increment_correct_numeric_answer_by_one</code>	If the correct answer is a numeric quantity, increment it by one, otherwise print the correct answer as is.	9757	1352

(i) **String Manipulation:** This operation involves manipulating the characters within the correct answer. We apply simple transformations like changing the case of the answer text or reversing the answer text, etc.

(ii) **Format Correct Answer:** This operation involves displaying the correct answer in the specified format. This involves printing any numeric answers in words or appending a string to the correct answer, etc.

(iii) **Numeric Manipulation:** This instruction involves incrementing a numeric quantity by one and has no effect on non-numeric answer text.

(iv) **Operations on list (Conditional):** This operation involves conditionally manipulating the candidate answer space – for instance, incrementing incorrect answers by one, sorting the incorrect answers, etc.

(v) **Operations on list:** These are simple instructions that do not depend on the correct answer of the original knowledge tasks. Examples, include - sorting all candidate options, concatenating characters from each candidate option, etc.

For each instruction category, we create multiple instructions. Table 2 presents the 13 instructions we have included in our work. The task prompts (instructions) for each of the 13 instruction types with an example are available in the Appendix (Section A.2).

**Instruction Creation:** To create each instruction, the authors iteratively refined them until all the authors had complete agreement in the output when they followed them manually. Examples of aspects of iterative improvement include - explicitly making clear what is not to be included in the output, how the output is to be presented, etc. We then asked 2 computer science researchers to follow and generate the output for 75 instructions across all our instruction types and datasets. We found that both the researchers were able to follow our instructions successfully and generated the same response for 93.33% of the instances. The first annotator generated the correct response for 98.67% of the instances, while the second annotator for 94.67% of the instances. Upon analyzing their responses, we found the only instruction-following error was rounding off the decimal number when truncating to two decimal places. We also found very few human errors in the annotator’s response, specifically for instructions like `reverse_correct_answer_alternate_case` on datasets with long output text such as PIQA.

**Answering baseline-instructions:** We additionally develop two baseline instructions – (1) printing the correct answer option<sup>1</sup> from the candidate space (`print_correct_answer_label`), and (2) printing only the text associated with the correct answer option (`print_correct_answer_text`).

<sup>1</sup>We use ‘label’ and ‘option’ interchangeably to denote the candidates in a multiple-choice QA task.

**Instructions with no-effect:** Certain instructions may be inapplicable for some knowledge tasks. For example, in the MathQA dataset, some instances have *none of these* as the correct answer and are not numeric. Here, instructions such as *numformat\_numeric\_answer* or *increment\_correct\_numeric\_answer\_by\_one* will not affect the existing answer of the knowledge-task. We refer to these instructions as “*distractor*” instances and expect that in these instances, models should perform as well as they do on the answering baseline-instructions. We include details and statistics of such instructions in the Appendix (Section A.4 Tables 7 and 9).

### 3.3 METRICS

**Exact Match:** We report the model performance as exact match under two settings - *strict* and *loose*. In the **strict** setting, we perform basic string parsing (removing beginning and ending whitespaces, quotations, etc.) and compare the model prediction to the expected output for the applied instruction.

However, we observe that models often make errors when following the primary instruction. These could be minor copying errors, such as missing a period or comma, or even fixing typos within the provided options. On the other hand, they could also be instruction following mistakes, where for instance the option label is added to the response even when the prompt explicitly states otherwise. Given that we do not expect models to make such mistakes given clear instructions, we use the strict metric in the majority of our evaluations.

However, we also define a relaxed version of the exact match called **loose** exact match, allowing for a Levenshtein distance Levenshtein (1966) of two edit operations between the prediction and ground truth. Additionally, we also perform whitespace-free matching as part of our loose criterion. Similar to Zhou et al. (2023b), we consider our loose match as a complement to the strict one.

### 3.4 BENCHMARK DATASET

We create two versions of our benchmark dataset - ‘Full’ and ‘Lite’ (for lower inference costs).

**Full Benchmark:** We select a subset of 1500 samples randomly from each datasets and apply each applicable instruction on the same. For MMLUPro, we consider a subset of 150 samples per subject and apply each applicable instruction.

**Lite Benchmark:** We select a subset of 150 samples randomly from the full version created above for each dataset and apply each applicable instruction on the same. For MMLUPro, we consider a subset of 25 samples per subject and apply each applicable instruction. Statistics for the above two versions are available in presented in Table 2 in the appendix. Detailed statistics for each dataset and the instruction types are provided in the appendix section A.4. Additionally, each benchmark includes a set of instances when instructions have no effect (called the no-effect subset).

#### 3.4.1 BENCHMARK RANKING

An effective instruction-following model should not only be capable of following a variety of instructions across different knowledge-tasks but should also be unaffected by instructions when they are inapplicable i.e, they should be robust to ‘distractors’. Therefore, we define an overall benchmark score for a model as its arithmetic mean of the following:

**Exact-Match Score** ( $\mu_{EM}$ ): We compute the micro-average of the exact-match scores using all instances of every instruction type in the benchmark.

**Instruction Category Score (IC Score):** We compute the micro-average exact-match scores for every instance per instruction category and then compute the arithmetic mean.

**Knowledge Task Subset Score: (KTS Score):** We compute the micro-average exact-match scores for every instance per knowledge-task, and then compute the arithmetic mean.

**Exact Match Score on ‘Instructions with no-effect’** ( $\mu'_{EM}$ ): We compute the micro-average of all instruction instances in the benchmark that have no effect on the original knowledge-task answers (i.e.) ‘distractors’.

Table 3: List of Models evaluated on our benchmark.

Small (< 7B parameters)	Medium (7 – 30B parameters)	Large (> 30B parameters)	Frontier
Llama-3.2-1B-Instruct (1B)	Mistral-7B-Instruct-v0.3 (7B)	Qwen2.5-32B-Instruct (32B)	Llama-3.1-405B-Instruct (405B)
Qwen2.5-1.5B-Instruct (1.5B)	Qwen2.5-7B-Instruct (7B)	Llama-3.1-70B-Instruct (70B)	GPT-4o-mini-2024-07-18
Llama-3.2-3B-Instruct (3B)	Phi-3-small-8k-instruct (7B)	Qwen2.5-72B-Instruct (72B)	GPT-4o-2024-08-06
Qwen2.5-3B-Instruct (3.0B)	Llama-3.1-8B-Instruct (8B)		
Phi-3.5-mini-instruct - (3.8B)	Gemma-2-9b-it (9B)		
	Phi-3-medium-4k-instruct (14B)		
	Qwen2.5-14B-Instruct (14B)		
	Gemma-2-27b-it (27B)		

## 4 EVALUATION

We present an evaluation on our benchmark using a variety of models and study the following research questions: (i) Do models display a difference in performance on the two simple baseline instruction tasks? (ii) Do models display a variation in performance across our different instruction categories? (iii) Are models robust to, or get distracted by instructions that do not apply to the task? (iv) Does the size of a model impact its instruction-following capability?

### 4.1 MODELS AND INFERENCE

We evaluate our benchmark on a range of open instruction-tuned models and parameter sizes. For ease of presentation, we categorize them based on their parameter count as shown in Table 3. Our inference code uses vLLM Kwon et al. (2023) for running the evaluations. We use greedy decoding for generations and `bf16` as floating point precision. We generate a maximum of 1024 tokens per instance. We use A100 80GB GPUs for running inference. We use an instance hosted by a cloud provider for Llama-3.1-405B-Instruct, while we use OpenAI APIs for GPT4-o and GPT4-o-mini models.

In all our experiments, we perform zero-shot Chain-of-Thought (CoT) Wei et al. (2024) reasoning. Models see the same prompt based on prompt guides for the original knowledge tasks in `lm-evaluation-harness` framework Gao et al. (2024) and OpenAI evals.<sup>2</sup> We instruct the model to generate reasoning first and then the answer (See examples in Appendix Section A.2). We write custom post-processing scripts to extract the model’s answer as described in the next section.

### 4.2 OUTPUT POST-PROCESSING

All our task prompt templates, as shown in Appendix A.2, explicitly instruct the model to provide their final response after a ‘`Response:`’ keyword. As part of our *strict* evaluation metric (Section 3.3), we search for and extract the response after this keyword while computing the metrics.

However, we observe that models may not always follow this, and can instead generate a wide range of other keywords (e.g.) *(the final answer is, the output is, etc)*, or no keyword at all. Given the diverse possible responses, we make a good-faith attempt to capture these patterns as part of our *loose* evaluation to classify a wider range of model responses. In our subsequent results, we use the loose evaluation for error analysis, and denote the specific type of strategy elsewhere.

### 4.3 RESULTS

We begin this section by first presenting our results on the answering baseline-instructions and then proceed to our results on instruction-following for the different categories. We then look at the impact of distractors and knowledge-task characteristics on model performance.

#### 4.3.1 PRINTING THE CORRECT ANSWER

We begin our experiments with the simplest task – given a multiple-choice question with option labels and their texts, we instruct the models to print the text associated with the correct answer instead of the answer label. From a knowledge perspective, this task is no harder than selecting the

<sup>2</sup><https://github.com/openai/simple-evals>

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377




Figure 1: Left: Average exact match performance across all tasks for the *print\_correct\_answer* and *print\_correct\_answer\_label* instructions. Right: Knowledge and instruction following (IF) errors across all tasks for the *print\_correct\_answer* instruction. A lower error is better. Both results shown using Full Benchmark data. Lite Benchmark results can be found in Appendix Figure 5.

answer label. However, as shown on the left in Figure 1, we observe a significant drop ( $\sim 20\%$  on average) in knowledge-task performance when instructing the model to respond with the text associated with the answer instead of its label. The pattern is consistent for frontier models like *GPT-4o* on the Lite Benchmark (Figure 5).

We hypothesize that this drop in performance could, in part, be due to the training process resulting in models being over-fit to certain input/output task formats, resulting in worse instruction following for other formats. Some common issues we observed include models outright ignoring the instruction and continuing to generate labels, or generating only Chain-of-Thought reasoning without a final answer, missing the output keyword specified in the prompt, etc, reflected by the knowledge and instruction following errors in Figure 1 on the right.

We observe that the errors decrease as the inherent model capability and size increases. Note that incorrect answers could correspond to both knowledge and instruction following errors. The figure also shows that we capture most errors. Figures 26-28 show the error analysis for different model families, where we observe larger models making fewer errors (20% – 80% reduction) for the Llama and Qwen models. The Phi model family however does not show this trend, calling for a closer look at their instruction training methodology. Figures 29-32 takes a deep-dive at the error distribution for each instruction category across model scales. We observe that models make the most errors on string manipulation tasks, and model scale does little to mitigate this. For the other categories, errors reduce as the model size increases. Inspired by this, and to further illustrate the challenges that LLMs face on simple instruction-following, we study their performance when the final output requires first inferring the correct answer, and then applying operations specified by the instructions on the correct answer.

#### 4.3.2 ANSWER-CONDITIONED INSTRUCTION PERFORMANCE

We present results from different model scales across our five instruction categories in Figures 2a-2d. We compare this to their corresponding performance on the baseline task of *print\_correct\_answer* (PCA).

**Small-Scale and Medium-Scale models:** We observe a 10% – 40% drop in performance compared to the baseline across all instruction categories (Figures . 2a and 2b). In particular, we notice that all models struggle on our set of string and numeric manipulation instructions, suggesting a bias towards certain input/output format instructions (see Figure 24a and 24b for instruction-specific results for each model).

We further notice that models such as Qwen-2.5-14B and Llama-3.1-8B exhibit good loose evaluation scores for numeric manipulations and formatting the correct answer, but suffer a large drop

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

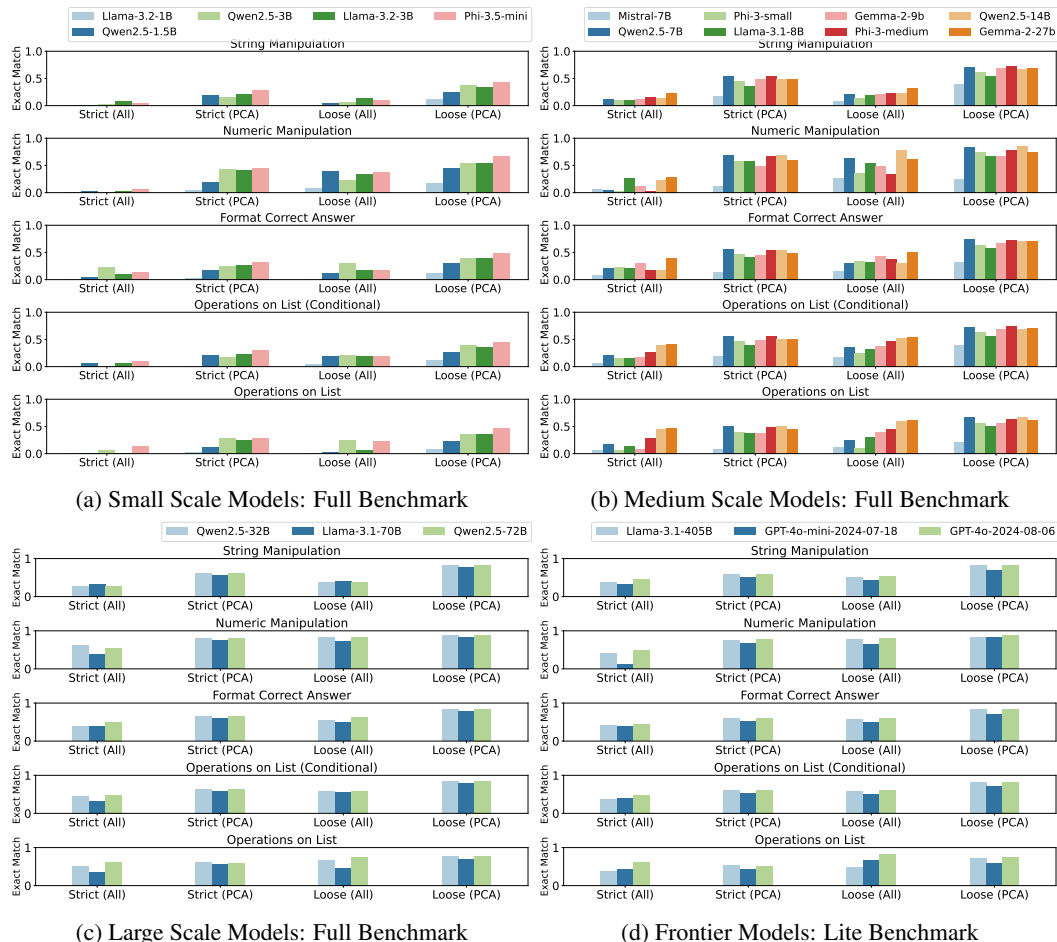


Figure 2: Performance variation (strict and loose) of exact match scores across the different answer-conditioned instruction categories (All) from Table 2 compared to corresponding performance on *print\_correct\_answer* (PCA).

in the corresponding strict evaluation. This difference suggests that these models are able to grasp the expectations from the instructions, but fail to follow them precisely. Examples of this include incrementing an answer by 0.1 when asked to increment by one, or only returning the special string when instructed to append it to the correct answer text, or even adding/missing characters from the provided options when instructed to return them as is. Finally, all models also find the operations on list categories to be challenging – where interestingly the performance of models across both sets - conditional on correct answers vs. not, is similar.

**Large-scale and Frontier models:** The improved capabilities of these larger models are evident from the absolute improvement in performance as shown in Figures 2c and 2d. We also observe a smaller drop in performance between their loose evaluation and strict evaluation scores, reflecting more precise instruction following. However, the trend of performance deterioration (5% – 40%) across instructions, compared to their respective baseline knowledge-answering tasks, still persists in these models, demonstrating opportunities to improve their instruction following.

### 4.3.3 EFFECT OF PARAMETER SIZE WITHIN A MODEL FAMILY

We report the performance on Full Benchmark for models from the Llama family and Qwen family of models in Figures 33 and 34 in the appendix. We observe a consistent pattern of improvements in instruction following-ability with increase in model capacity for the Llama family. However, this is not the case for Qwen family of models. Specifically, for some instructions



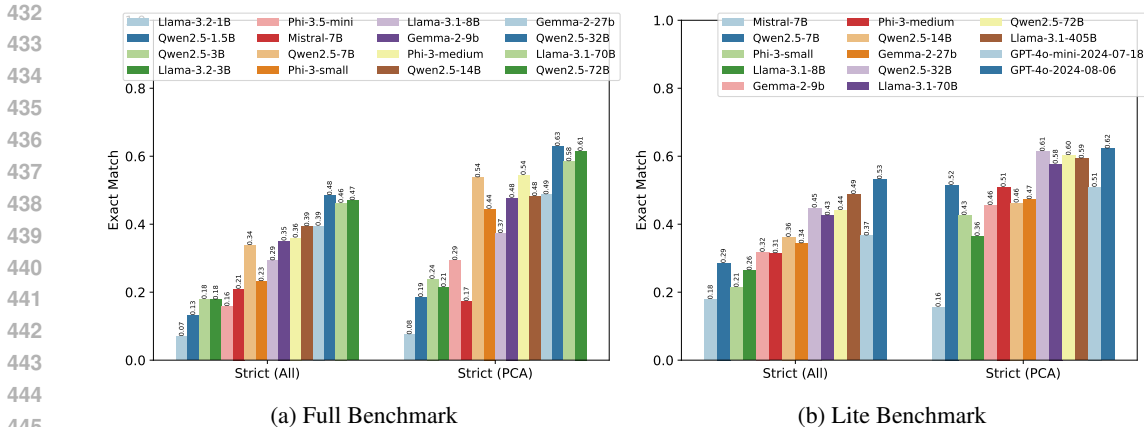


Figure 3: Impact of distractor instructions on exact match performance across tasks and instructions, compared to its corresponding *print\_correct\_answer* performance. A drop indicates the model getting distracted by an inapplicable instruction.

like *print\_correct\_answer*, *print\_correct\_answer\_label*, *sort\_only\_incorrect\_answers* the Qwen 1.5B model outperforms 3B model. Qwen 3B model is better than Qwen 7B and 14B variants for the *print\_correct\_answer\_append\_string* instruction. We consistently see 32B and 72B variants outperforming other models by a significant margin.

#### 4.3.4 INSTRUCTIONS AS DISTRACTORS

Our dataset also includes instructions that apply only when certain properties of a knowledge-task answer are fulfilled. For instance, instructions for incrementing the correct answer by one if numeric, formatting numeric values, and printing any numeric answers in words, do not apply on tasks with textual answers. They serve as distractors, and we expect model performance to be unaffected since these instructions are not applicable and do not alter the original knowledge-task answer.

However, from Figure 3, we observe that there is a 5-20% drop in small, medium, large, and frontier scale models. In figures 39a, 39b, 39c we report details of how different model families (Llama, Qwen and Phi) are affected by distractors, at different scales. We find that the Llama family and Phi of models are extremely distracted by instructions that require reversing and casing text (even though the instruction is inapplicable on numerical data), and report a drop of nearly 75-78% while Qwen family of models (at all scales) is relatively robust to such distractors. On the other hand, distractor instructions that are based on numeric operations lead to a minor drop in performance in Llama and Qwen models but still affect Phi family of models significantly. While model failures in the presence of distractors have been studied before (Shi et al., 2023; Feng et al., 2024), to the best of our knowledge this is the first work to study them in an instruction-following setting.

#### 4.3.5 KNOWLEDGE-TASK CHARACTERISTICS AND INSTRUCTION-FOLLOWING

As seen in Figure 4, the performance drop for models for an instruction category can also be dependent on the nature of the knowledge-task. For instance, models appear to have a larger relative drop on MathQA as compared to MMLUPro for the numeric manipulation instruction category. Models also struggle more on string manipulation operations on PIQA - probably because of the long sentences that are part of answer candidates. Other knowledge-task and model scales have been presented in Appendix Section A.6.

### 4.4 BENCHMARK

We report the *strict* scores of the medium, large and frontier models on the Lite Benchmark in Table 4. Unsurprisingly, GPT4o model performs the best on our benchmark data while large and medium-scale models like Llama-3.1 405B, Qwen2.5 72B, and, Qwen2.5 32B models appear to be better than other openly available models including Llama-3.1-70B-instruct and the Gemma family of models. We also include the results on the full benchmark in Appendix Table 5. We note that the ranking of models is largely consistent and that small models are much weaker than larger models.

486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

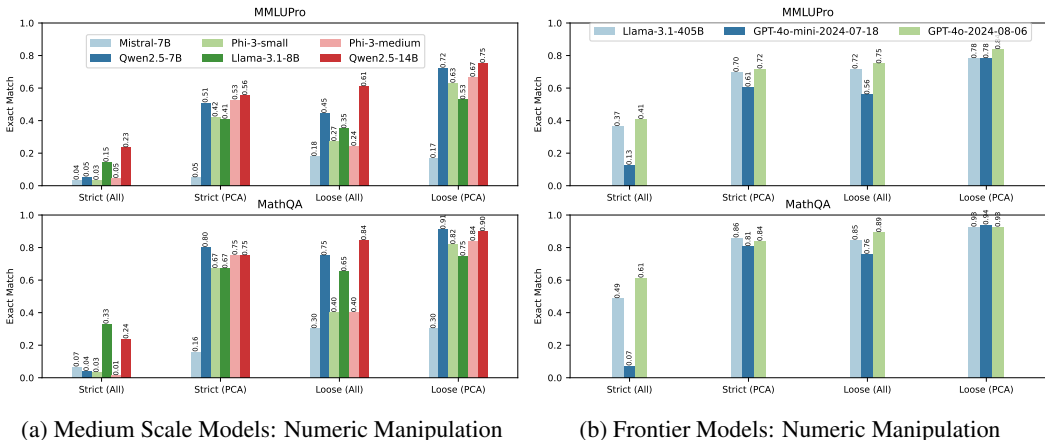


Figure 4: Performance variation (strict and loose) of exact match scores for the Numeric Manipulation instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

Table 4: Performance of the Medium, Large and Frontier Models on our Lite Benchmark - models ranked in order of performance using the average score (higher is better).

Models	$\mu_{EM}$	IC Score	KTS Score	$\mu_{EM}$	Average Score
GPT-4o-2024-08-06	0.4790	0.4990	0.5543	0.5318	0.5161
Llama-3.1-405B	0.4236	0.4537	0.4920	0.4883	0.4644
Qwen2.5-72B	0.4021	0.4690	0.4548	0.4410	0.4417
Qwen2.5-32B	0.3710	0.4402	0.4311	0.4481	0.4226
Llama-3.1-70B	0.3394	0.3832	0.3946	0.4253	0.3856
GPT-4o-mini-2024-07-18	0.3601	0.3327	0.4299	0.3659	0.3722
Gemma-2-27b	0.3254	0.3673	0.3902	0.3430	0.3565
Qwen2.5-14B	0.2508	0.2996	0.2980	0.3620	0.3026
Phi-3-medium	0.2056	0.2250	0.2512	0.2932	0.2437
Gemma-2-9b	0.1716	0.1952	0.2133	0.3092	0.2223
Qwen2.5-7B	0.1700	0.1860	0.2029	0.2849	0.2109
Llama-3.1-8B	0.1568	0.1996	0.1840	0.2637	0.2010
Phi-3-small	0.1418	0.1535	0.1780	0.1970	0.1676
Mistral-7B	0.0566	0.0786	0.0755	0.1789	0.0974

## 5 DISCUSSION & CONCLUSION

In this work, we demonstrated how modern LLMs fail to follow simple instructions. We took a novel approach to studying instruction-following by grounding instructions on existing knowledge tasks. Our approach has the advantage of being easily extendable for new instruction types and domains, while also enabling LLM-free evaluations with some degree of automated error analysis. We demonstrated that not only do models fail to follow simple instructions (e.g.) printing the answer text instead of the label, but their performance drops further when compound but simple, instructions are included. Even when instructions that should have no effect on the knowledge-tasks are used, models at all scales report a drop in performance, though the extent of deterioration varies. As models are increasingly being viewed as agents and assistants, it is crucial that models have better guarantees of following user instructions. As our benchmark demonstrates, there is a lot of scope for improvement and we hope the community finds it helpful in improving the current state-of-the-art.

Lastly, before concluding, we would like to re-emphasize the choice of the strict measure to study performance - if instructions specify how the task is to be completed then models should not add extraneous text, respond by rephrasing the question as part of the final response, make copying errors, etc. The nature of errors made by models as reflected in the difference between loose and strict scores, the automated error analysis sets and the large amount of unclassified errors highlights that instruction-tuning of LLMs requires special focus on *instruction-following*.

## REFERENCES

- 540  
541  
542 Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh  
543 Hajishirzi. MathQA: Towards interpretable math word problem solving with operation-based  
544 formalisms. In Jill Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019*  
545 *Conference of the North American Chapter of the Association for Computational Linguistics:*  
546 *Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2357–2367, Minneapolis,  
547 Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1245.  
548 URL <https://aclanthology.org/N19-1245>.
- 549 Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning  
550 about physical commonsense in natural language. In *Thirty-Fourth AAAI Conference on Artificial*  
551 *Intelligence*, 2020.
- 552 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhari-  
553 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agar-  
554 wal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh,  
555 Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz  
556 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec  
557 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In  
558 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neu-  
559 ral Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc.,  
560 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/  
561 file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfcb4967418bfb8ac142f64a-Paper.pdf).
- 562 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared  
563 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large  
564 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.
- 565 Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina  
566 Toutanova. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In Jill  
567 Burstein, Christy Doran, and Thamar Solorio (eds.), *Proceedings of the 2019 Conference of*  
568 *the North American Chapter of the Association for Computational Linguistics: Human Lan-*  
569 *guage Technologies, Volume 1 (Long and Short Papers)*, pp. 2924–2936, Minneapolis, Min-  
570 nesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1300. URL  
571 <https://aclanthology.org/N19-1300>.
- 572 Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C Desmarais,  
573 and Zhen Ming Jack Jiang. Github copilot ai pair programmer: Asset or liability? *Journal of*  
574 *Systems and Software*, 203:111734, 2023.
- 575 Wanyong Feng, Jaewook Lee, Hunter McNichols, Alexander Scarlatos, Digory Smith, Simon  
576 Woodhead, Nancy Ornelas, and Andrew Lan. Exploring automated distractor generation for  
577 math multiple-choice questions via large language models. In Kevin Duh, Helena Gomez, and  
578 Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL 2024*,  
579 pp. 3067–3082, Mexico City, Mexico, June 2024. Association for Computational Linguistics.  
580 doi: 10.18653/v1/2024.findings-naacl.193. URL [https://aclanthology.org/2024.  
581 findings-naacl.193](https://aclanthology.org/2024.findings-naacl.193).
- 582 Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Fos-  
583 ter, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muen-  
584 nighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lin-  
585 tang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework  
586 for few-shot language model evaluation, 07 2024. URL [https://zenodo.org/records/  
587 12608602](https://zenodo.org/records/12608602).
- 588 Tianyu Gao, Adam Fisch, and Danqi Chen. Making pre-trained language models better few-shot  
589 learners. *arXiv preprint arXiv:2012.15723*, 2020.
- 590 Michael Grohs, Luka Abb, Nourhan Elsayed, and Jana-Rebecca Rehse. Large language models can  
591 accomplish business process management tasks. In *International Conference on Business Process*  
592 *Management*, pp. 453–465. Springer, 2023.

- 594 Benjamin Heinzerling and Kentaro Inui. Language models as knowledge bases: On entity rep-  
595 resentations, storage capacity, and paraphrased queries. In Paola Merlo, Jorg Tiedemann, and  
596 Reut Tsarfaty (eds.), *Proceedings of the 16th Conference of the European Chapter of the Asso-*  
597 *ciation for Computational Linguistics: Main Volume*, pp. 1772–1791, Online, April 2021. As-  
598 sociation for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.153. URL <https://aclanthology.org/2021.eacl-main.153>.  
599
- 600 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn  
601 Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset.  
602 In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks*  
603 *Track (Round 2)*, 2021. URL <https://openreview.net/forum?id=7Bywt2mQsCe>.  
604
- 605 Jie Huang and Kevin Chen-Chuan Chang. Towards reasoning in large language models: A survey.  
606 *arXiv preprint arXiv:2212.10403*, 2022.
- 607 Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
608 language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*,  
609 volume 35, pp. 22199–22213, 2022a.
- 610 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
611 language models are zero-shot reasoners. *Advances in neural information processing systems*,  
612 35:22199–22213, 2022b.
- 613 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.  
614 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model  
615 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating*  
616 *Systems Principles*, 2023.
- 617 V Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Proceedings*  
618 *of the Soviet physics doklady*, 1966.
- 619 Shiyang Li, Jun Yan, Hai Wang, Zheng Tang, Xiang Ren, Vijay Srinivasan, and Hongxia Jin.  
620 Instruction-following evaluation through verbalizer manipulation. In Kevin Duh, Helena Gomez,  
621 and Steven Bethard (eds.), *Findings of the Association for Computational Linguistics: NAACL*  
622 *2024*, pp. 3678–3692, Mexico City, Mexico, June 2024. Association for Computational Linguis-  
623 tics. doi: 10.18653/v1/2024.findings-naacl.233. URL [https://aclanthology.org/](https://aclanthology.org/2024.findings-naacl.233)  
624 [2024.findings-naacl.233](https://aclanthology.org/2024.findings-naacl.233).  
625
- 626 Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan  
627 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let’s verify step by step. In *The Twelfth*  
628 *International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=v8L0pN6EOi>.  
629
- 630 Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring how models mimic human  
631 falsehoods. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings*  
632 *of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*  
633 *Papers)*, pp. 3214–3252, Dublin, Ireland, May 2022. Association for Computational Linguis-  
634 tics. doi: 10.18653/v1/2022.acl-long.229. URL <https://aclanthology.org/2022.acl-long.229>.  
635
- 636 Qianli Ma, Haotian Zhou, Tingkai Liu, Jianbo Yuan, Pengfei Liu, Yang You, and Hongxia Yang.  
637 Let’s reward step by step: Step-level reward model as the navigators for reasoning, 2024. URL  
638 <https://openreview.net/forum?id=RSQL6xvUYW>.  
639
- 640 Ipek Ozkaya. Application of large language models to software engineering tasks: Opportunities,  
641 risks, and implications. *IEEE Software*, 40(3):4–8, 2023.
- 642 Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and  
643 Alexander Miller. Language models as knowledge bases? In Kentaro Inui, Jing Jiang, Vincent  
644 Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Nat-*  
645 *ural Language Processing and the 9th International Joint Conference on Natural Language Pro-*  
646 *cessing (EMNLP-IJCNLP)*, pp. 2463–2473, Hong Kong, China, November 2019. Association for  
647 Computational Linguistics. doi: 10.18653/v1/D19-1250. URL <https://aclanthology.org/D19-1250>.

- 648 Yiwei Qin, Kaiqiang Song, Yebowen Hu, Wenlin Yao, Sangwoo Cho, Xiaoyang Wang, Xuansheng  
649 Wu, Fei Liu, Pengfei Liu, and Dong Yu. Infobench: Evaluating instruction following ability in  
650 large language models. 2024.
- 651 Adam Roberts, Colin Raffel, and Noam Shazeer. How much knowledge can you pack into the  
652 parameters of a language model? In Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu  
653 (eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Process-*  
654 *ing (EMNLP)*, pp. 5418–5426, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.437. URL [https://aclanthology.org/2020.](https://aclanthology.org/2020.emnlp-main.437)  
655 [emnlp-main.437](https://aclanthology.org/2020.emnlp-main.437). URL [https://aclanthology.org/2020.](https://aclanthology.org/2020.emnlp-main.437)  
656 [emnlp-main.437](https://aclanthology.org/2020.emnlp-main.437).
- 657 Steven I Ross, Fernando Martinez, Stephanie Houde, Michael Muller, and Justin D Weisz. The  
658 programmer’s assistant: Conversational interaction with a large language model for software de-  
659 velopment. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*,  
660 pp. 491–514, 2023.
- 661 Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. Winogrande: an ad-  
662 versarial winograd schema challenge at scale. *Commun. ACM*, 64(9):99–106, aug 2021. ISSN  
663 0001-0782. doi: 10.1145/3474381. URL <https://doi.org/10.1145/3474381>.
- 664 Abulhair Saparov, Richard Yuanzhe Pang, Vishakh Padmakumar, Nitish Joshi, Mehran Kazemi,  
665 Najoung Kim, and He He. Testing the general deductive reasoning capacity of large language  
666 models using OOD examples. In *Thirty-seventh Conference on Neural Information Processing*  
667 *Systems*, 2023. URL <https://openreview.net/forum?id=MCVfX7HgPO>.
- 668 Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli,  
669 and Denny Zhou. Large language models can be easily distracted by irrelevant context, 2023.  
670 URL <https://arxiv.org/abs/2302.00093>.
- 671 Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung,  
672 Moya Chen, Kushal Arora, Joshua Lane, et al. Blenderbot 3: a deployed conversational agent that  
673 continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*, 2022.
- 674 Wangtao Sun, Chenxiang Zhang, Xueyou Zhang, Ziyang Huang, Haotian Xu, Pei Chen, Shizhu He,  
675 Jun Zhao, and Kang Liu. Beyond instruction following: Evaluating inferential rule following of  
676 large language models, 2024. URL <https://arxiv.org/abs/2407.08440>.
- 677 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-  
678 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.  
679 *arXiv preprint arXiv:2203.11171*, 2022.
- 680 Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming  
681 Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi  
682 Fan, Xiang Yue, and Wenhu Chen. Mmlu-pro: A more robust and challenging multi-task language  
683 understanding benchmark, 2024. URL <https://arxiv.org/abs/2406.01574>.
- 684 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi,  
685 Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language  
686 models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Ad-*  
687 *vances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/](https://openreview.net/forum?id=_VjQlMeSB_J)  
688 [forum?id=\\_VjQlMeSB\\_J](https://openreview.net/forum?id=_VjQlMeSB_J).
- 689 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi,  
690 Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language  
691 models. In *Proceedings of the 36th International Conference on Neural Information Processing*  
692 *Systems*, NIPS ’22, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088.
- 693 Michael Wornow, Avanika Narayan, Krista Opsahl-Ong, Quinn McIntyre, Nigam H Shah,  
694 and Christopher Re. Automating the enterprise with foundation models. *arXiv preprint*  
695 *arXiv:2405.03710*, 2024.
- 696 Congying Xia, Chen Xing, Jiangshu Du, Xinyi Yang, Yihao Feng, Ran Xu, Wenpeng Yin, and  
697 Caiming Xiong. Fofu: A benchmark to evaluate llms’ format-following capability. *arXiv preprint*  
698 *arXiv:2402.18667*, 2024.
- 699  
700  
701

- 702 Zhiyuan Zeng, Jiatong Yu, Tianyu Gao, Yu Meng, Tanya Goyal, and Danqi Chen. Evaluating large  
703 language models at evaluating instruction following. In *International Conference on Learning*  
704 *Representations (ICLR)*, 2024.
- 705
- 706 Quanjun Zhang, Tongke Zhang, Juan Zhai, Chunrong Fang, Bowen Yu, Weisong Sun, and Zhenyu  
707 Chen. A critical review of large language model on software engineering: An example from  
708 chatgpt and automated program repair. *arXiv preprint arXiv:2310.08879*, 2023.
- 709 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang,  
710 Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and  
711 chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623, 2023.
- 712
- 713 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schur-  
714 mers, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. Least-to-most prompting  
715 enables complex reasoning in large language models. In *The Eleventh International Confer-*  
716 *ence on Learning Representations*, 2023a. URL <https://openreview.net/forum?id=WZH7099tgfM>.
- 717
- 718 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny  
719 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint*  
720 *arXiv:2311.07911*, 2023b.
- 721
- 722 Li Zhou, Jianfeng Gao, Di Li, and Heung-Yeung Shum. The design and implementation of xiaoice,  
723 an empathetic social chatbot. *Computational Linguistics*, 46(1):53–93, 2020.
- 724

## 725 A APPENDIX

726

727 We describe how we automatically classify errors in section A.1. We list all instructions with ex-  
728 ample input, ground truth, and expected instruction output in A.2. We the report results on the  
729 Full Benchmark in A.3. The detailed statistics of Full ad Lite Benchmark are presented in A.4.  
730 Section A.5 presents the comparison between model’s performance on print correct answer and  
731 print correct answer labels tasks on the Lite Benchmark. Section A.6 presents performance of dif-  
732 ferent models for each instruction category in comparison with its corresponding performance on  
733 *print\_correct\_answer* (PCA).

734

### 735 A.1 AUTOMATED CLASSIFICATION OF ERRORS

736

737 For each instruction in the ‘String Manipulation’, ‘Format Correct Answer’, and ‘Numeric Manipu-  
738 lation’ instruction categories, we create a set of error classes that are based on the incorrect answers  
739 to the original knowledge task and the subsequent application of the instruction on that incorrect  
740 answer. We create ‘instruction-following’ error sets (IFError) and ‘knowledge-error sets’ (Knowl-  
741 edgeError) as follows:

742 **IFError:** We inspected some representative model outputs and created a set of errors that LLMs  
743 could plausibly make for each instruction. This set includes errors such as: (i) answering with an  
744 answer label (option) when requiring/using the text, (ii) answering with any candidate answer text  
745 without applying the instruction, (iii) a combination of these, (iv) instruction specific errors based  
746 on these – for example, for the instruction *alternate\_case\_correct\_answer* such error cases could  
747 include alternate casing starting with a lowercase character (the instruction text explicitly requires  
748 that it start with upper case).

749 **KnowledgeError:** For every instruction instance, this set includes: (i) Answering with the incorrect  
750 candidate answer – regardless of whether it is its label or text, (ii) Applying the instruction (when  
751 applicable) on the incorrect candidate answer.

752 Note that these error sets are not mutually exclusive – a model response to the alternate casing in-  
753 struction that starts with lowercase and is applied on the incorrect answer candidate is an ‘IFError’  
754 as well as a ‘KnowledgeError’. Further, these error sets are meant to be *high-precision* to aid anal-  
755 ysis and cannot reasonably be expected to anticipate and cover everything LLMs might respond with.  
Hence, there will be model errors that we cannot classify and we intend to continue to expand the

756 benchmark’s error classification capabilities. We would also like to highlight that if models make a  
757 lot of errors that are not classified, it is likely that those are more severe instruction-following errors.  
758

## 759 760 761 762 763 A.2 PROMPTS

764  
765  
766 We now list the prompts for each instruction and dataset combination defined in our benchmark.  
767  
768  
769  
770

771 Dataset: Winogrande

772 Instruction: print\_correct\_answer\_append\_string

773  
774 Given a sentence with a blank and 2 options namely 'A' and 'B', print the text associated  
775 with the option label that answers the question correctly. Do not print the option label. With-  
776 out including any additional characters, also append the following text string at the end of  
777 the output: 'ENDOFRESPONSE'. Do not add any special characters including quotations,  
778 spaces, new lines etc in the response. Think step by step and in the end, finish your response  
779 with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output ex-  
780 pected.

781 Sentence: It is an article of faith that the paper is more important than the exam , even though  
782 the \_ weighs less heavily on the grade.

783 **Options:**

784 A. paper

785 B. exam

786 **Ground Truth:** A

787 **Instruction Output:** Response:paperENDOFRESPONSE

788  
789  
790  
791  
792  
793  
794  
795 Dataset: Winogrande

796 Instruction: alternate\_case\_correct\_answer

797  
798 Given a sentence with a blank and 2 options namely 'A' and 'B', answer the question by  
799 printing the text associated with the correct option label in alternate case. The first character  
800 should be in uppercase and the following characters should alternate between lowercase and  
801 uppercase. Do not print the option label. Think step by step and in the end, finish your  
802 response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final  
803 output expected.

804 Sentence: The wooden doors at my friends work are worse than the wooden desks at my  
805 work, because the \_ material is stronger.

806 **Options:**

807 A. doors

808 B. desks

809 **Ground Truth:** B

**Instruction Output:** Response:DeSkS

810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821  
822  
823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835  
836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863

Dataset: MathQA  
Instruction: numformat\_numeric\_answer

Given a mathematical question and 5 options namely 'a', 'b', 'c', 'd', and 'e', as candidate answers, print the text associated with the option label that answers the question correctly. If the answer is numeric print it in two decimal places as long as it contains no other string or units of measurement. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** a man walking at the rate of 5 km / hr crosses a bridge in 15 minutes . the length of the bridge ( in meters ) is :

**Options:**

- a. 600
- b. 750
- c. 1000
- d. 1250
- e. none of these

**Ground Truth:** d

**Instruction Output:** Response:1250.00

Dataset: MathQA  
Instruction: sort\_options\_to\_create\_string

Given a mathematical question and 5 options namely 'a', 'b', 'c', 'd', and 'e', as candidate answers, sort the list of options using their values, in alphabetical order. Use only the text associated with the option labels and not the option labels while sorting. Then, create a string by concatenating the last character of the text associated with each option value. If the last character is a special character (such as period, comma, quotation, etc) use the previous character. Print only the final string and not the sorted list. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** marts income is 50 percent more than tims income and tims income is 40 percent less than juans income . what percentage of juans income is marts income

**Options:**

- a. 124 %
- b. 120 %
- c. 96 %
- d. 90 %
- e. 64 %

**Ground Truth:** d

**Instruction Output:** Response:40604

Dataset: PIQA  
Instruction: reverse\_correct\_answer

Given a question and two answer candidates 'A' and 'B', answer the question by printing the text associated with the correct option label, in reverse. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** Butcher Shop

**Options:**

- A. will decimate fish from the ocean into digestible pieces
- B. will decimate a full cow into digestible pieces

**Ground Truth:** B

**Instruction Output:** Response:seceip elbitsegid otni woc lluf a etamiced lliw



864  
865  
866  
867  
868  
869  
870  
871  
872  
873  
874  
875  
876  
877  
878  
879  
880  
881  
882  
883  
884  
885  
886  
887  
888  
889  
890  
891  
892  
893  
894  
895  
896  
897  
898  
899  
900  
901  
902  
903  
904  
905  
906  
907  
908  
909  
910  
911  
912  
913  
914  
915  
916  
917

Dataset: PIQA  
Instruction: print\_correct\_answer

Given a question and two answer candidates 'A' and 'B', answer the question by selecting the value associated with the option label corresponding to the correct answer. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** how to avoid paint spill when adding paint to your brush

**Options:**

A. Put a rubber band on your paint can to get rid of that excess glue on your paint brush, this will prevent spilling paint on the paint stir stick where the lid is.

B. Put a rubber band on your paint can to get rid of that excess glue on your paint brush, this will prevent spilling paint on the edge where the lid is.

**Ground Truth:** B

**Instruction Output:** Response:Put a rubber band on your paint can to get rid of that excess glue on your paint brush, this will prevent spilling paint on the edge where the lid is.

Dataset: Winogrande  
Instruction: use\_incorrect\_options\_to\_create\_string

Given a sentence with a blank and 2 options namely 'A' and 'B', create a string by concatenating the last character of every option value, excluding the option value corresponding to the correct answer. Use only the text associated with the option labels and not the option labels while sorting. If the last character is a special character (such as period, comma, quotation, etc) use the previous character. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

Sentence: Fiction books were interesting and easy to read for Logan but not Brett because \_ enjoyed real tales.

**Options:**

A. Logan

B. Brett

**Ground Truth:** B

**Instruction Output:** Response:n

Dataset: PIQA  
Instruction: reverse\_correct\_answer\_alternate\_case

Given a question and two answer candidates 'A' and 'B', reverse the text associated with the answer label that correctly answers the question. Print this reversed text in alternate case starting with upper case. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** What else should I add to a peanut butter sandwich?

**Options:**

A. Take some pasta sauce and put it onto the other piece of bread with a knife.

B. Take some jelly and put it onto the other piece of bread with a knife.

**Ground Truth:** B

**Instruction Output:** Response:..EfInK A Htlw dAeRb fO EcEiP ReHtO EhT OtNo tI TuP DnA YlLeJ EmOs eKaT

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

Dataset: Winogrande  
Instruction: use\_options\_to\_create\_string

Given a sentence with a blank and 2 options namely 'A' and 'B', create a string by concatenating the last character of every option value (not option label). If the last character is a special character (such as period, comma, quotation, etc) use the previous character. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

Sentence: Megan focused less on proper posture than Lindsey because \_ wanted to become a model.

**Options:**

- A. Megan
- B. Lindsey

**Ground Truth:** B

**Instruction Output:** Response:ny

Dataset: MathQA  
Instruction: print\_correct\_answer\_label

Given a mathematical question and 5 options namely 'a', 'b', 'c', 'd', and 'e', as candidate answers, answer the question by selecting the option label corresponding to the correct answer. Do not include the text associated with the option label in the answer. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** a reduction of 20 % in the price of salt enables a lady to obtain 2 kgs more for rs . 100 , find the original price per kg ?

**Options:**

- a. 12.6
- b. 12.1
- c. 12.5
- d. 12.4
- e. 12.7

**Ground Truth:** c

**Instruction Output:** Response:c

Dataset: PIQA  
Instruction: increment\_correct\_numeric\_answer\_by\_one

Given a question and two answer candidates 'A' and 'B', print the text associated with the option label that answers the question correctly. Note that if the correct answer is a numeric quantity, including dollar values and percentages but contains no other string or units of measurement, print the value after increasing its value by 1. Dollar values should be prefixed with '\$'. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** how to winterize windows

**Options:**

- A. put weather stripping around them to stop air from escaping and air from coming in
- B. put weather stripping around them to stop air from escaping and air from coming into the dishwasher

**Ground Truth:** A

**Instruction Output:** Response:put weather stripping around them to stop air from escaping and air from coming in

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

Dataset: MathQA  
Instruction: sort\_only\_incorrect\_answers

Given a mathematical question and 5 options namely 'a', 'b', 'c', 'd', and 'e', as candidate answers, excluding the option that answers the question correctly, print a sorted list (ascending order) of the incorrect options. Do not print the option labels. Use the text associated with the option labels and not the option labels while sorting and printing. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** the sector of a circle has radius of 21 cm and central angle 108 o . find its perimeter ?

**Options:**

- a. 81.6 cm
- b. 85.9 cm
- c. 90 cm
- d. 92 cm
- e. 95 cm

**Ground Truth:** a

**Instruction Output:** Response:['85.9 cm', '90 cm', '92 cm', '95 cm']

Dataset: PIQA  
Instruction: print\_correct\_answer\_in\_words

Given a question and two answer candidates 'A' and 'B', print the text associated with the option label that answers the question correctly. However, if the correct answer is a numeric value with no additional text (including percentages, currency, units of measurement etc), print the numeric answer in words. For example, if the answer is '32' print 'thirty-two' without quotes. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** How do I make the pattern for the baby leather shoes?

**Options:**

- A. Create a template on a piece of paper by placing your babies shoe on the paper and drawing around it.
- B. Create a template on a piece of paper by placing your babies foot on the paper and drawing around it.

**Ground Truth:** A

**Instruction Output:** Response:Create a template on a piece of paper by placing your babies shoe on the paper and drawing around it.

1026  
1027  
1028  
1029  
1030  
1031  
1032  
1033  
1034  
1035  
1036  
1037  
1038  
1039  
1040  
1041  
1042  
1043  
1044  
1045  
1046  
1047  
1048  
1049  
1050  
1051  
1052  
1053  
1054  
1055  
1056  
1057  
1058  
1059  
1060  
1061  
1062  
1063  
1064  
1065  
1066  
1067  
1068  
1069  
1070  
1071  
1072  
1073  
1074  
1075  
1076  
1077  
1078  
1079

Dataset: BoolQ  
Instruction: increment\_incorrect\_numeric\_answers\_by\_one

Given a passage and a boolean question, and the possible answer candidates 'A' or 'B', print the list of incorrect answers (not the answer label). Increase each value by 1 while printing if it is a numeric quantity including dollar values, percentages but contains no other string or units of measurement. Do not print the option labels. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

Passage: A Star Is Born is an upcoming American musical romantic drama film produced and directed by Bradley Cooper, in his directorial debut. Cooper also wrote the screenplay with Will Fetters and Eric Roth. A remake of the 1937 film of the same name, it stars Cooper, Lady Gaga, Andrew Dice Clay, Dave Chappelle, and Sam Elliott, and follows a hard-drinking country musician (Cooper) who discovers and falls in love with a young singer (Gaga). It marks the third remake of the original 1937 film (which featured Janet Gaynor and Fredric March), which was adapted into a 1954 musical (starring Judy Garland and James Mason) and then remade as a 1976 rock musical with Barbra Streisand and Kris Kristofferson.

**Question:** is bradley cooper a star is born a remake

**Options:**

A. True

B. False

**Ground Truth:** A

**Instruction Output:** Response:['False']

Dataset: PIQA  
Instruction: capitalize\_correct\_answer

Given a question and two answer candidates 'A' and 'B', answer the question by printing the text associated with the correct option label in uppercase. Do not print the option label. Think step by step and in the end, finish your response with 'Response:\$RESPONSE' where \$RESPONSE (without quotes) is the final output expected.

**Question:** wool

**Options:**

A. can be used to line cookie tins

B. can be used to line pants

**Ground Truth:** B

**Instruction Output:** Response:CAN BE USED TO LINE PANTS

### A.3 FULL BENCHMARK LEADERBOARD

We report *strict* scores of the small, medium, and, large models on the Full Benchmark in Table 5. We observe that Qwen2.5 72B and 32B variants outperforms all other models. Llama-3.1 70B is ranked third with a significant gap between the second best model Qwen2.5 32B. There is a significant drop in performance as the parameter size decreases.

### A.4 ADDITIONAL BENCHMARK STATISTICS

The following sections reports detailed statistics for the Full and Lite Benchmark. We report statistics for both instruction following and Instructions with no-effect subsets. We observe that for some dataset (knowledge tasks) and instruction combinations, the corresponding entries are zero indicating that there is no single instance where the instruction gets applied (Instructions with no-effect) or there is no single instance where the instruction doesn't get applied (instruction follow subset).

1080  
 1081  
 1082  
 1083  
 1084  
 1085  
 1086  
 1087  
 1088  
 1089  
 1090  
 1091  
 1092  
 1093  
 1094  
 1095  
 1096  
 1097  
 1098  
 1099  
 1100  
 1101  
 1102  
 1103  
 1104  
 1105  
 1106  
 1107  
 1108  
 1109  
 1110  
 1111  
 1112  
 1113  
 1114  
 1115  
 1116  
 1117  
 1118  
 1119  
 1120  
 1121  
 1122  
 1123  
 1124  
 1125  
 1126  
 1127  
 1128  
 1129  
 1130  
 1131  
 1132  
 1133

Table 5: Performance of the Small, Medium, and Large Models on our Full Benchmark - models ranked in order of performance using the average score (higher is better).

Models	$\mu_{EM}$	IC Score	KTS Score	$\mu_{EM}$	Average Score
Qwen2.5-72B	0.4157	0.4827	0.4343	0.4697	0.4506
Qwen2.5-32B	0.3822	0.4501	0.4108	0.4846	0.4319
Llama-3.1-70B	0.3389	0.3532	0.3582	0.4623	0.3781
Gemma-2-27b	0.3392	0.3589	0.3727	0.3947	0.3664
Qwen2.5-14B	0.2545	0.2803	0.2767	0.3932	0.3012
Phi-3-medium	0.2013	0.1835	0.2167	0.3406	0.2355
Gemma-2-9b	0.1633	0.1583	0.1850	0.3409	0.2119
Qwen2.5-7B	0.1658	0.1515	0.1753	0.3366	0.2073
Llama-3.1-8B	0.1460	0.1741	0.1573	0.2927	0.1925
Phi-3-small	0.1352	0.1171	0.1568	0.2129	0.1555
Phi-3.5-mini	0.0883	0.0967	0.0932	0.1397	0.1045
Llama-3.2-3B	0.0704	0.0571	0.0780	0.1632	0.0922
Mistral-7B	0.0458	0.0541	0.0543	0.2093	0.0909
Qwen2.5-3B	0.0663	0.0704	0.0725	0.1090	0.0796
Qwen2.5-1.5B	0.0382	0.0347	0.0436	0.1223	0.0597
Llama-3.2-1B	0.0039	0.0037	0.0041	0.0184	0.0075

1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166

Table 6: Full Benchmark: Instruct Follow Stats

	MMLUPro													PIQA	MathQA	Winogrande			
	BooQ	Physics	Health	Economics	Law	Philosophy	Business	Other	Chemistry	Psychology	History	Computer Science	Biology				Math	Engineering	
numformat_numeric_answer	0	150	53	150	150	44	150	150	150	150	47	150	150	150	150	553	1500	1500	0
increment_incorrect_numeric_answers_by_one	1500	150	150	150	150	0	150	150	150	150	0	150	150	150	150	1500	1500	1500	1267
sort_only_incorrect_answers	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
use_options_to_create_string	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
print_correct_answer_label	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
reverse_correct_answer_alternate_case	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
alternate_case_correct_answer_by_one	0	150	22	42	1	4	150	150	150	150	11	86	13	150	53	0	1500	540	0
print_correct_answer	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
print_correct_answer_append_string	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
print_correct_answer_in_words	0	142	14	10	0	3	81	31	108	10	0	84	8	150	35	0	1500	1500	0
sort_options_to_create_string	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
reverse_correct_answer	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
use_incorrect_options_to_create_string	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267
capitalize_correct_answer	1500	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1500	1500	1500	1267

Table 7: Full Benchmark: Instructions with no-effect

	MMLUPro													MathQA	PIQA	BooQ	Winogrande	
	Health	Economics	Math	Psychology	Law	Computer Science	Physics	Other	Business	Chemistry	Engineering	Biology	History					Philosophy
numformat_numeric_answer	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1337	1285	1500	1267
print_correct_answer_in_words	150	150	150	150	150	150	150	150	150	150	150	150	150	150	1331	1500	1500	1267
increment_correct_numeric_answer_by_one	150	150	150	150	150	150	150	150	150	150	150	150	150	150	928	1500	1500	1267
reverse_correct_answer_alternate_case	30	48	150	13	3	122	150	150	150	150	33	2	15	1500	0	0	0	0
reverse_correct_answer	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0
use_incorrect_options_to_create_string	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0

1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181  
1182  
1183  
1184  
1185  
1186  
1187  
1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199

Table 8: Lite Benchmark: Instruct Follow Stats

MMLUPro																	
BoolQ	chemistry	other	physics	math	biology	philosophy	psychology	economics	history	health	law	engineering	business	computer science	PIQA	MathQA	Winogrande
150	66	82	65	53	53	50	54	69	49	60	38	74	73	54	253	352	150
150	66	82	65	53	53	50	54	69	49	60	38	74	73	54	253	352	150
0	25	25	25	25	13	4	11	25	1	22	1	25	25	25	0	150	0
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
0	25	25	25	25	8	3	10	10	0	14	0	25	25	25	0	150	0
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
0	25	25	25	25	25	25	25	25	25	25	25	25	25	25	0	150	0
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150
150	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150

Table 9: Lite Benchmark: Instructions with no-effect

MMLUPro																		
BoolQ	biology	health	law	engineering	chemistry	math	business	physics	history	psychology	other	computer science	economics	philosophy	MathQA	Winogrande	BoolQ	PIQA
57	47	41	56	41	56	46	48	52	29	44	50	63	57	43	293	150	150	198
57	47	41	56	41	56	46	48	52	29	44	50	63	57	43	293	150	150	198
25	25	3	25	3	25	25	25	25	2	13	25	25	25	15	150	0	0	0
25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150	150
25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150	150
25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150	150
25	25	25	25	25	25	25	25	25	25	25	25	25	25	25	150	150	150	150
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

### A.5 PRINTING THE CORRECT ANSWER

We present the comparison between model’s performance on print correct answer and print correct answer labels tasks on the Lite Benchmark in Table 5. We observe that all models show a drop in performance when instructed to print correct answer instead of the label.

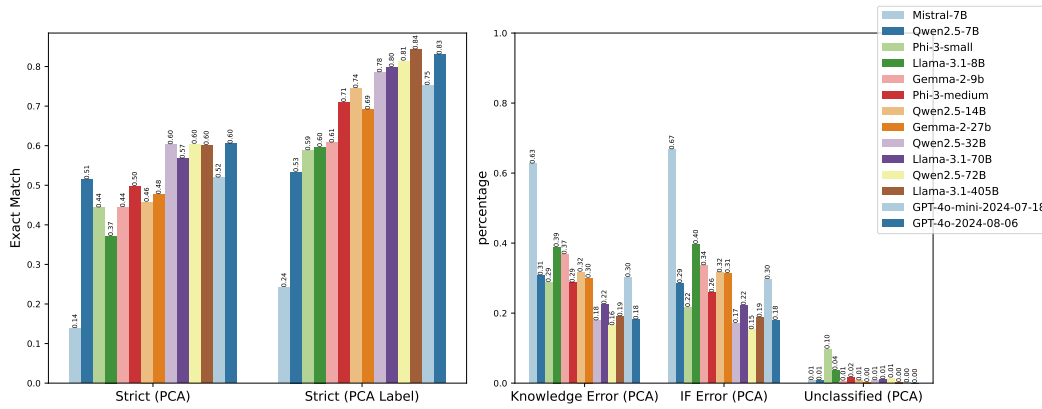


Figure 5: Lite Benchmark: Performance of LLMs on Printing the correct answer task and error comparison. PCA refers to *print\_correct\_answer* instruction and PCA label refers to *print\_correct\_answer\_label*.

### A.6 KNOWLEDGE-TASK CHARACTERISTICS AND INSTRUCTION-FOLLOWING

We now present performance of different models for each instruction category in comparison with its corresponding performance on *print\_correct\_answer* (PCA). The patterns remains consistent.




Figure 6: Small Scale Models: Performance variation (strict and loose) of exact match scores for the Numeric Manipulation instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).






Figure 7: Large-Scale Models: Performance variation (strict and loose) of exact match scores for the Numeric Manipulation instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

## A.7 INSTRUCTION SPECIFIC RESULTS

In this section, we report results at the individual instruction level across knowledge tasks.

## A.8 ERROR CLASSIFICATION

### A.8.1 INFLUENCE OF PARAMETER SIZE

We report the performance on Full Benchmark for models from the Llama family and Qwen family of models (Figures 33 and 34). We observe a consistent pattern of improvements in instruction following-ability with increase in model capacity for the Llama family. However, this is not the case for Qwen family of models. Specifically, for some instructions like *print\_correct\_answer*, *print\_correct\_answer\_label*, *sort\_only\_incorrect\_answers* the Qwen 1.5B model outperforms 3B model. Qwen 3B model is better than Qwen 7B and 14B variants for the *print\_correct\_answer\_append\_string* instruction. We consistently see 32B and 72B variants outperforming other models by a significant margin.

## A.9 INFLUENCE OF DISTRACTORS

1308  
1309  
1310  
1311  
1312  
1313  
1314  
1315  
1316  
1317  
1318  
1319  
1320  
1321  
1322  
1323  
1324  
1325  
1326  
1327  
1328  
1329  
1330  
1331  
1332  
1333  
1334  
1335  
1336  
1337  
1338  
1339  
1340  
1341  
1342  
1343  
1344  
1345  
1346  
1347  
1348  
1349  
1350  
1351  
1352  
1353  
1354  
1355  
1356  
1357  
1358  
1359  
1360  
1361

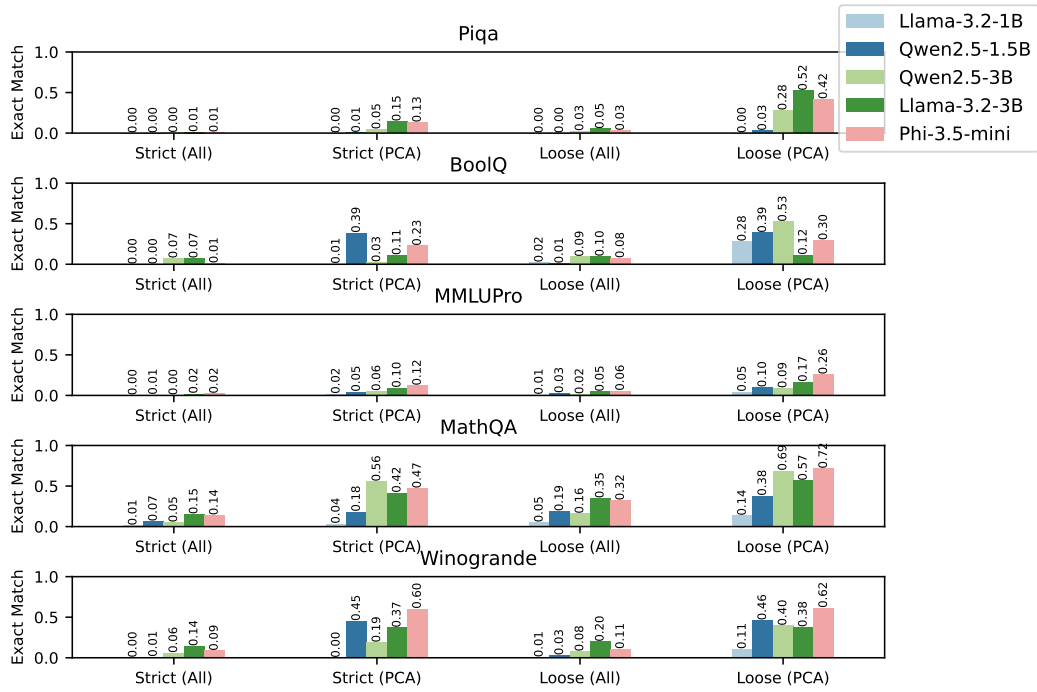


Figure 8: Small-Scale Models: Performance variation (strict and loose) of exact match scores for the String Manipulation instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

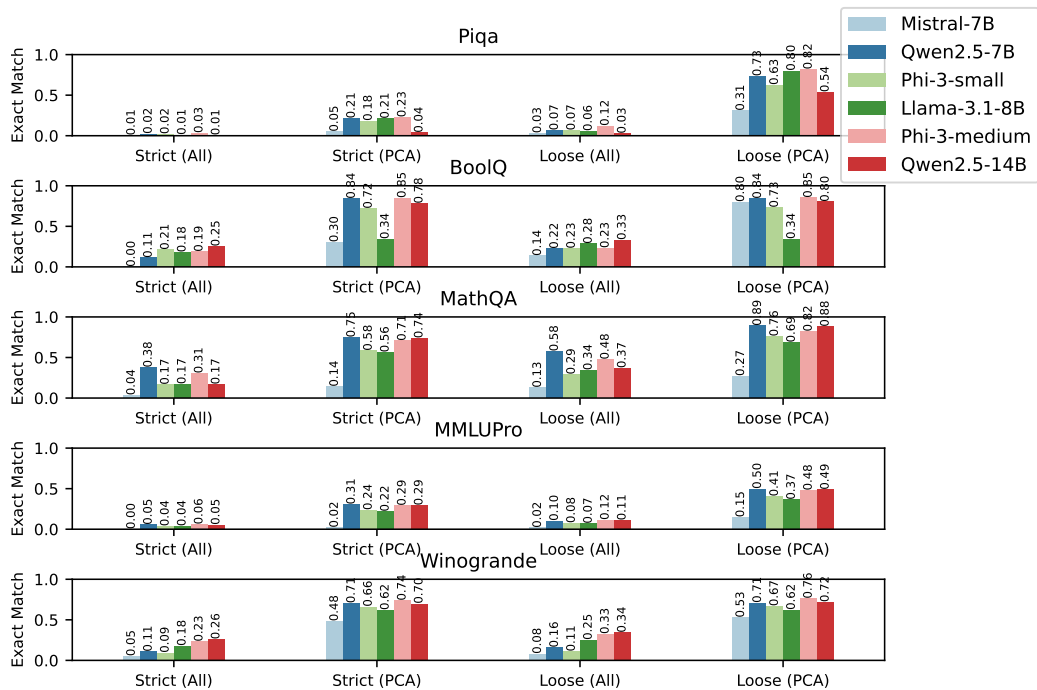



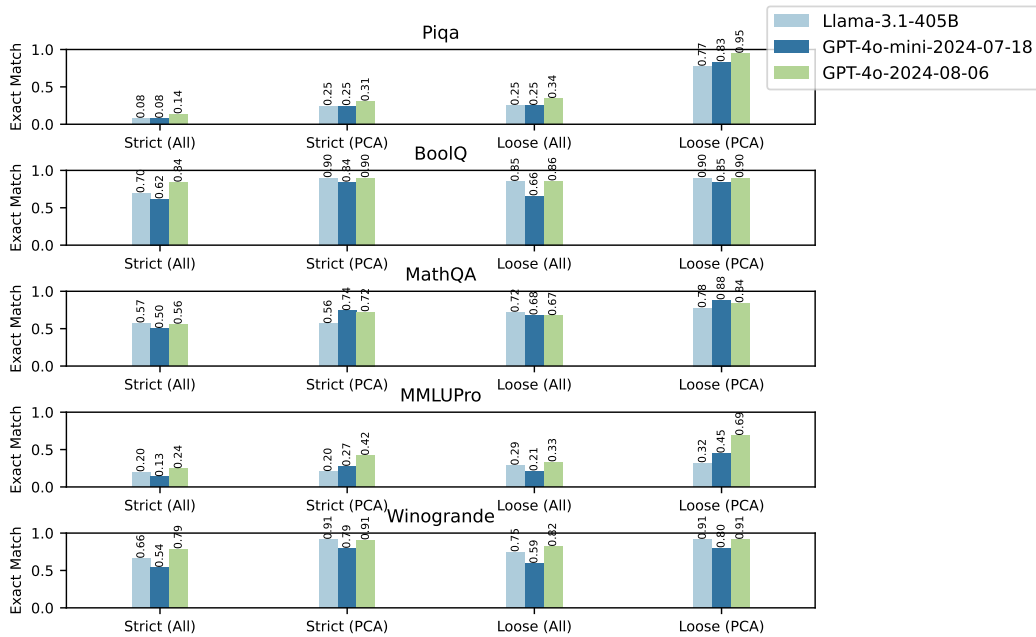
Figure 9: Medium-Scale Models: Performance variation (strict and loose) of exact match scores for the String Manipulation instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1362  
1363  
1364  
1365  
1366  
1367  
1368  
1369  
1370  
1371  
1372  
1373  
1374  
1375  
1376  
1377  
1378  
1379  
1380  
1381  
1382  
1383  
1384  
1385



1386 Figure 10: Large-Scale Models: Performance variation (strict and loose) of exact match scores  
1387 for the String Manipulation instruction category compared to its corresponding performance on  
1388 *print\_correct\_answer* (PCA).  
1389

1390  
1391  
1392  
1393  
1394  
1395  
1396  
1397  
1398  
1399  
1400  
1401  
1402  
1403  
1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415



1412 Figure 11: Frontier Models: Performance variation (strict and loose) of exact match scores  
1413 for the String Manipulation instruction category compared to its corresponding performance on  
1414 *print\_correct\_answer* (PCA).  
1415

1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439

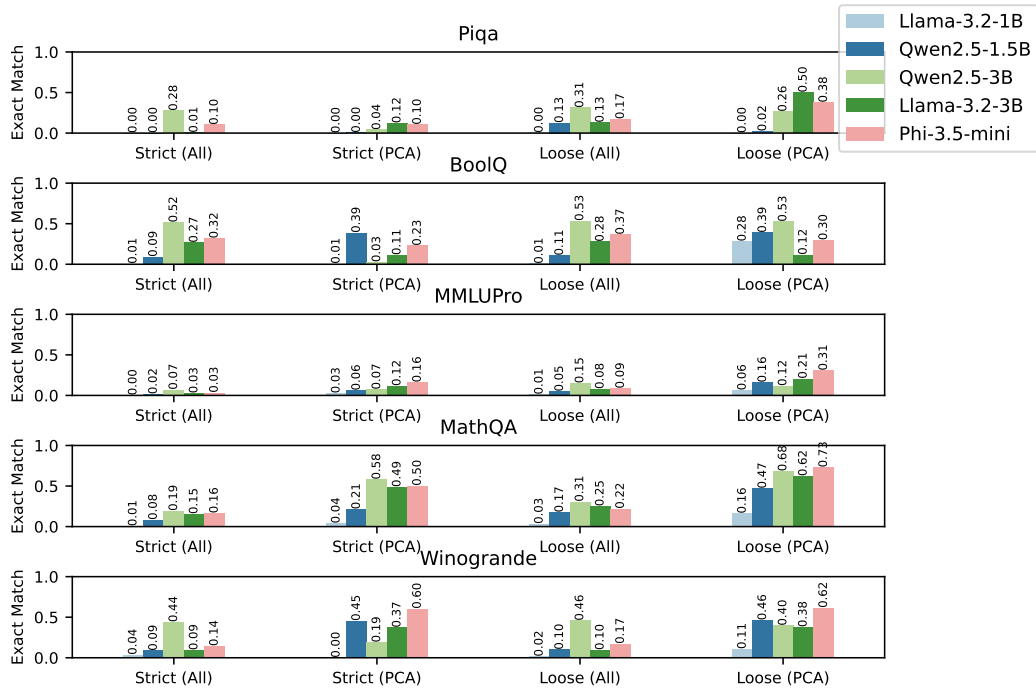


Figure 12: Small-Scale Models: Performance variation (strict and loose) of exact match scores for the Format Correct Answer instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457  
1458  
1459  
1460  
1461  
1462  
1463  
1464  
1465  
1466  
1467  
1468  
1469

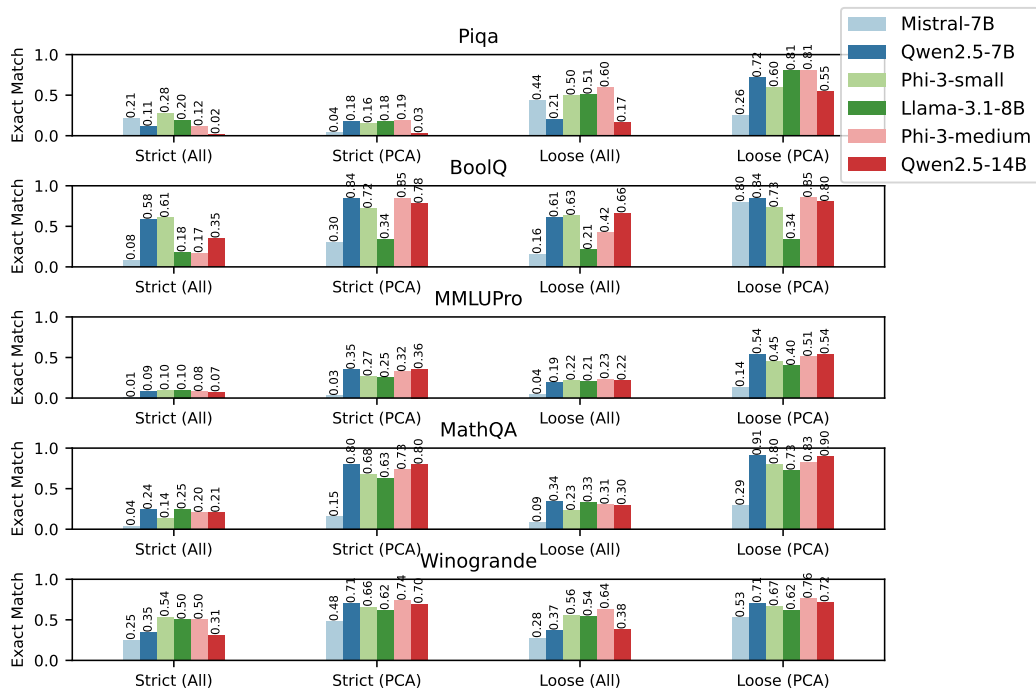


Figure 13: Medium-Scale Models: Performance variation (strict and loose) of exact match scores for the Format Correct Answer instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1470  
1471  
1472  
1473  
1474  
1475  
1476  
1477  
1478  
1479  
1480  
1481  
1482  
1483  
1484  
1485  
1486  
1487  
1488  
1489  
1490  
1491  
1492  
1493




Figure 14: Large-Scale Models: Performance variation (strict and loose) of exact match scores for the Format Correct Answer instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1494  
1495  
1496  
1497  
1498  
1499  
1500  
1501  
1502  
1503  
1504  
1505  
1506  
1507  
1508  
1509  
1510  
1511  
1512  
1513  
1514  
1515  
1516  
1517  
1518  
1519  
1520  
1521  
1522  
1523




Figure 15: Frontier-Scale Models: Performance variation (strict and loose) of exact match scores for the Format Correct Answer instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1524  
1525  
1526  
1527  
1528  
1529  
1530  
1531  
1532  
1533  
1534  
1535  
1536  
1537  
1538  
1539  
1540  
1541  
1542  
1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551  
1552  
1553  
1554  
1555  
1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565  
1566  
1567  
1568  
1569  
1570  
1571  
1572  
1573  
1574  
1575  
1576  
1577




Figure 16: Small-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).




Figure 17: Medium-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

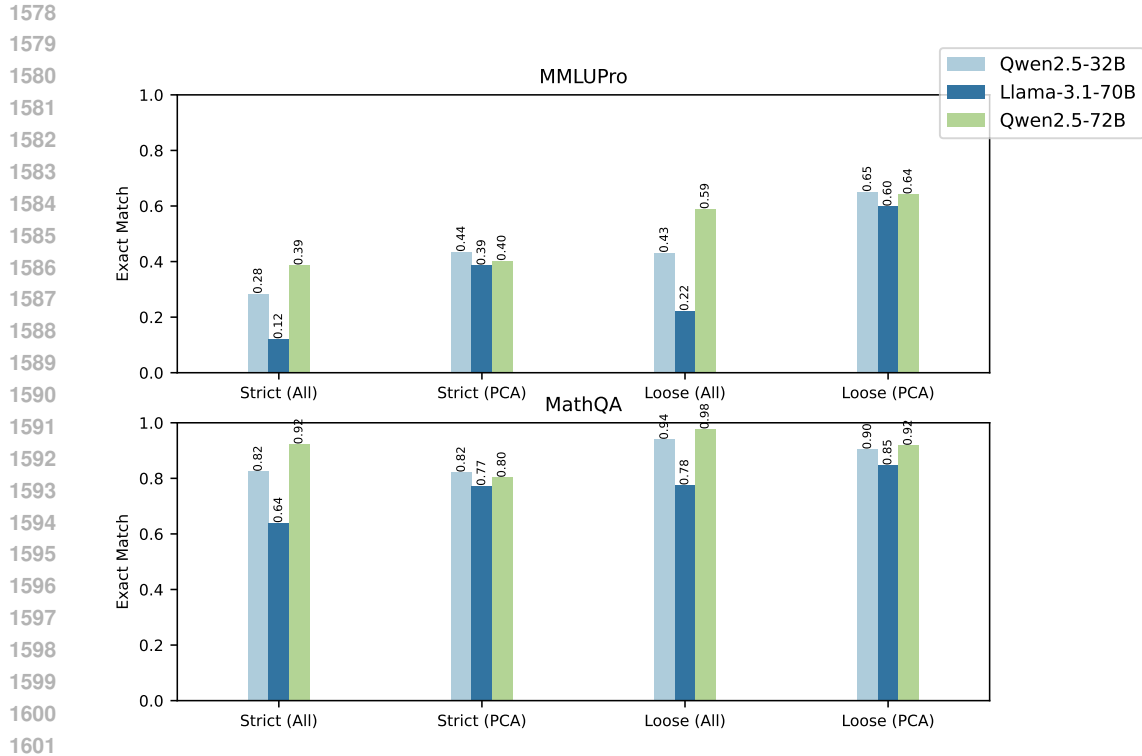


Figure 18: Large-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).




Figure 19: Frontier Models: Performance variation (strict and loose) of exact match scores for the Operations on List instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654

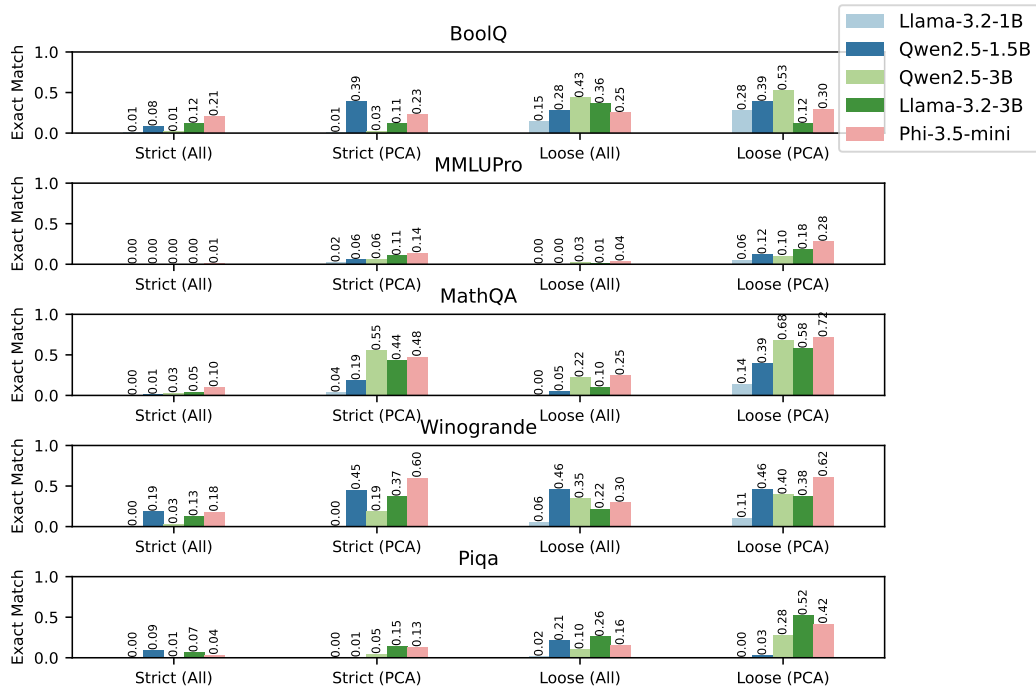


Figure 20: Small-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List (Conditional) instruction category compared to its corresponding performance on *print.correct\_answer* (PCA).

1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673  
1674  
1675  
1676  
1677  
1678  
1679  
1680  
1681  
1682  
1683  
1684  
1685

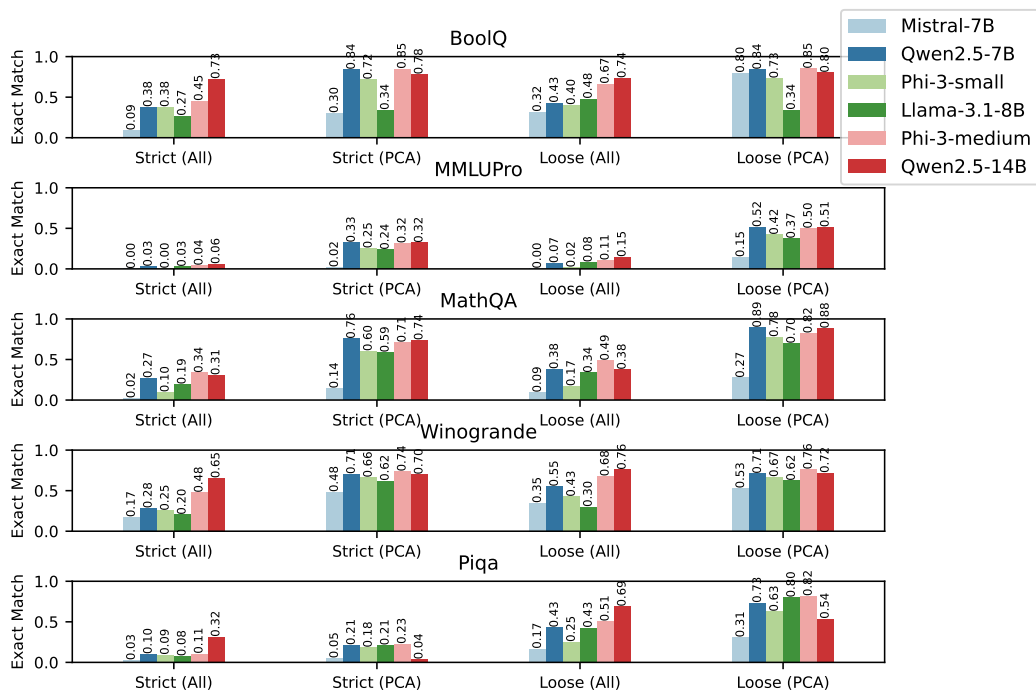


Figure 21: Medium-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List (Conditional) instruction category compared to its corresponding performance on *print.correct\_answer* (PCA).



1686  
1687  
1688  
1689  
1690  
1691  
1692  
1693  
1694  
1695  
1696  
1697  
1698  
1699  
1700  
1701  
1702  
1703  
1704  
1705  
1706  
1707  
1708  
1709

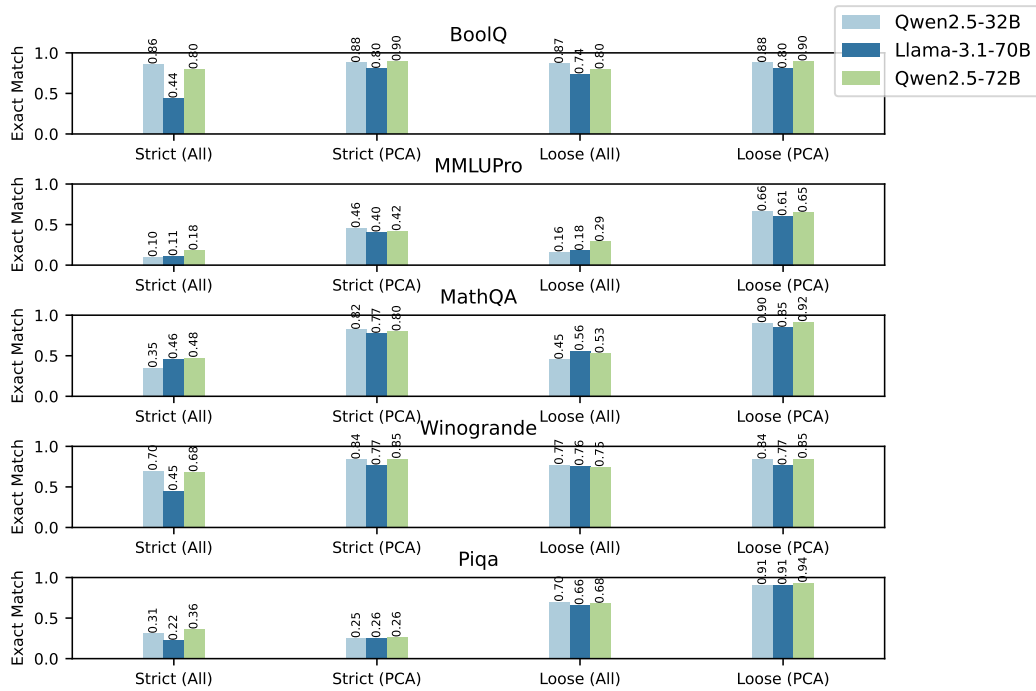


Figure 22: Large-Scale Models: Performance variation (strict and loose) of exact match scores for the Operations on List (Conditional) instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1710  
1711  
1712  
1713  
1714  
1715  
1716  
1717  
1718  
1719  
1720  
1721  
1722  
1723  
1724  
1725  
1726  
1727  
1728  
1729  
1730  
1731  
1732  
1733  
1734  
1735  
1736  
1737  
1738  
1739

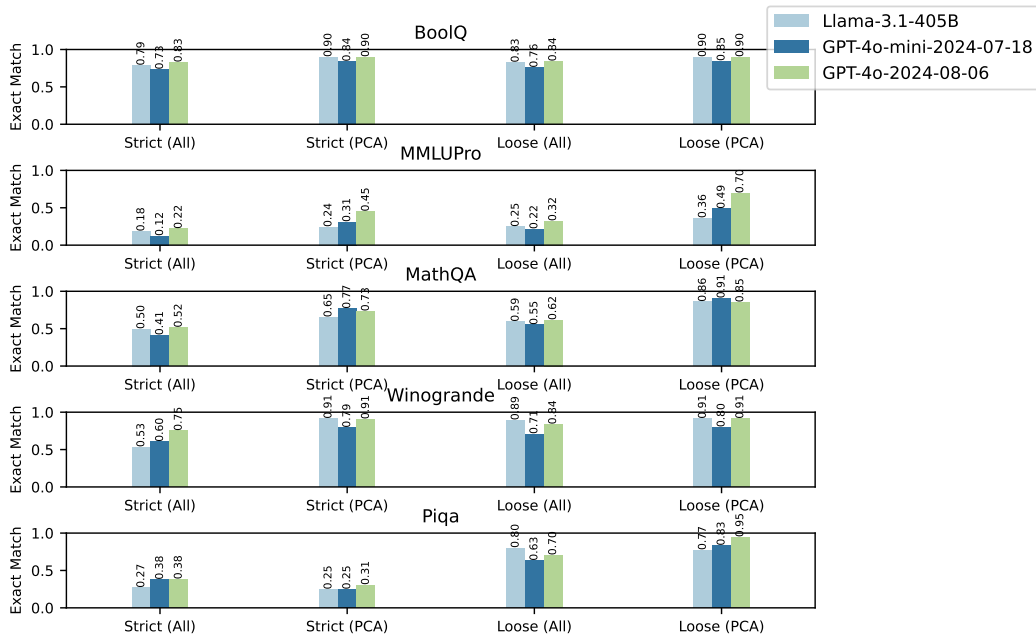




Figure 23: Frontier Models: Performance variation (strict and loose) of exact match scores for the Operations on List (Conditional) instruction category compared to its corresponding performance on *print\_correct\_answer* (PCA).

1740  
1741  
1742  
1743  
1744  
1745  
1746  
1747  
1748  
1749  
1750  
1751  
1752  
1753  
1754  
1755  
1756  
1757  
1758  
1759  
1760  
1761  
1762  
1763  
1764  
1765  
1766  
1767  
1768  
1769  
1770  
1771  
1772  
1773  
1774  
1775  
1776  
1777  
1778  
1779  
1780  
1781  
1782  
1783  
1784  
1785  
1786  
1787  
1788  
1789  
1790  
1791  
1792  
1793




(a) Exact Match (Strict)




(b) Exact Match (loose)

Figure 24: Performance variation of exact match scores on individual instructions across models on Full Benchmark

1794  
1795  
1796  
1797  
1798  
1799  
1800  
1801  
1802  
1803  
1804  
1805  
1806  
1807  
1808  
1809  
1810  
1811  
1812  
1813  
1814  
1815  
1816  
1817  
1818  
1819  
1820  
1821  
1822  
1823  
1824  
1825  
1826  
1827  
1828  
1829  
1830  
1831  
1832  
1833  
1834  
1835  
1836  
1837  
1838  
1839  
1840  
1841  
1842  
1843  
1844  
1845  
1846  
1847



(a) Exact Match (Strict)



(b) Exact Match (Loose)

Figure 25: Performance variation of exact match scores on individual instructions across models on Lite Benchmark

1848  
1849  
1850  
1851  
1852  
1853  
1854  
1855  
1856  
1857  
1858  
1859  
1860  
1861  
1862  
1863  
1864  
1865  
1866  
1867  
1868  
1869  
1870  
1871  
1872  
1873  
1874  
1875  
1876  
1877  
1878  
1879  
1880  
1881  
1882  
1883  
1884  
1885  
1886  
1887  
1888  
1889  
1890  
1891  
1892  
1893  
1894  
1895  
1896  
1897  
1898  
1899  
1900  
1901




Figure 26: Llama model family: Knowledge Errors and IFErrors




Figure 27: Qwen model family: Knowledge Errors and IFErrors

1902  
1903  
1904  
1905  
1906  
1907  
1908  
1909  
1910  
1911  
1912  
1913  
1914  
1915  
1916  
1917  
1918  
1919  
1920  
1921  
1922  
1923  
1924  
1925  
1926  
1927  
1928  
1929  
1930  
1931  
1932  
1933  
1934  
1935  
1936  
1937  
1938  
1939  
1940  
1941  
1942  
1943  
1944  
1945  
1946  
1947  
1948  
1949  
1950  
1951  
1952  
1953  
1954  
1955

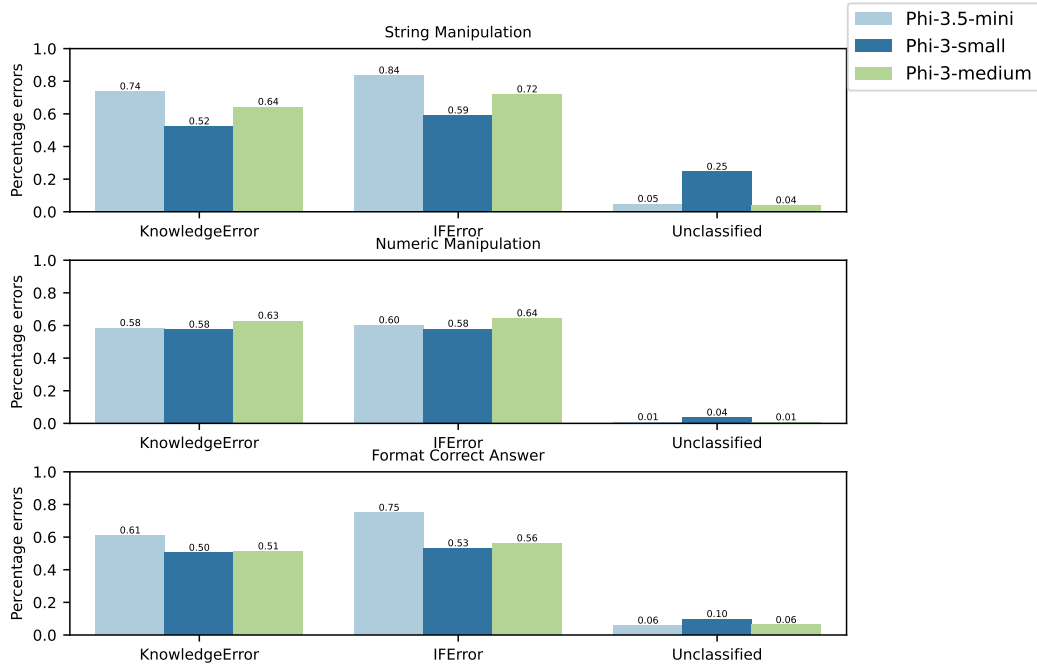


Figure 28: Phi model family: Knowledge Errors and IFErrors




Figure 29: Small Scale Models: Knowledge Errors and IFErrors

1956  
1957  
1958  
1959  
1960  
1961  
1962  
1963  
1964  
1965  
1966  
1967  
1968  
1969  
1970  
1971  
1972  
1973  
1974  
1975  
1976  
1977  
1978  
1979  
1980  
1981  
1982  
1983  
1984  
1985  
1986  
1987  
1988  
1989  
1990  
1991  
1992  
1993  
1994  
1995  
1996  
1997  
1998  
1999  
2000  
2001  
2002  
2003  
2004  
2005  
2006  
2007  
2008  
2009

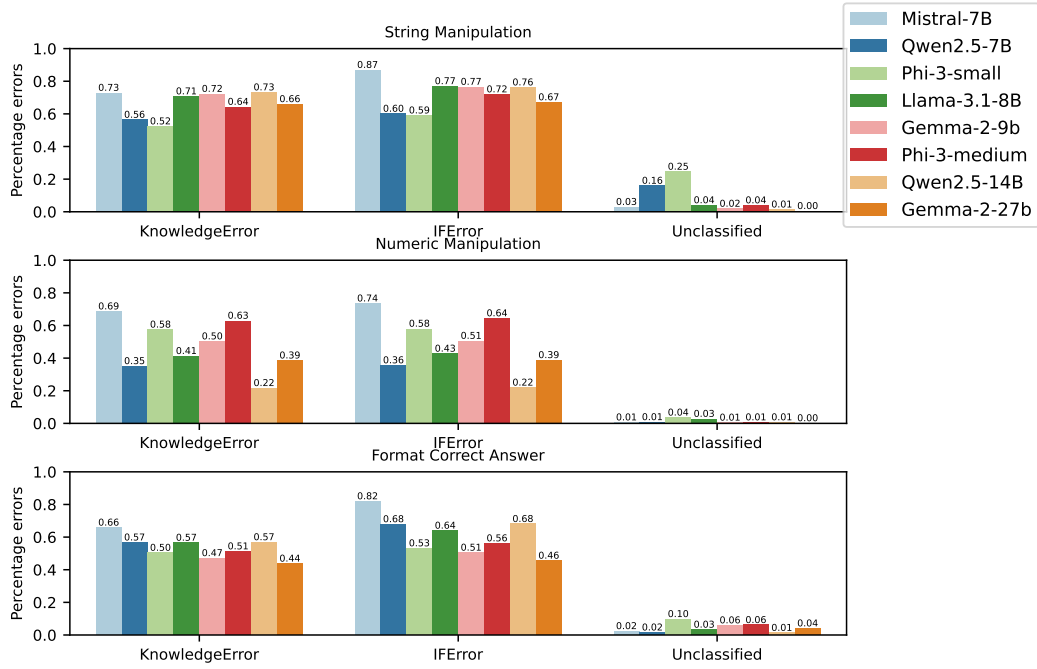


Figure 30: Medium Scale Models: Knowledge Errors and IFErrors




Figure 31: Large Scale Models: Knowledge Errors and IFErrors

2010  
2011  
2012  
2013  
2014  
2015  
2016  
2017  
2018  
2019  
2020  
2021  
2022  
2023  
2024  
2025  
2026  
2027  
2028  
2029  
2030  
2031  
2032  
2033  
2034  
2035  
2036  
2037  
2038  
2039  
2040  
2041  
2042  
2043  
2044  
2045  
2046  
2047  
2048  
2049  
2050  
2051  
2052  
2053  
2054  
2055  
2056  
2057  
2058  
2059  
2060  
2061  
2062  
2063

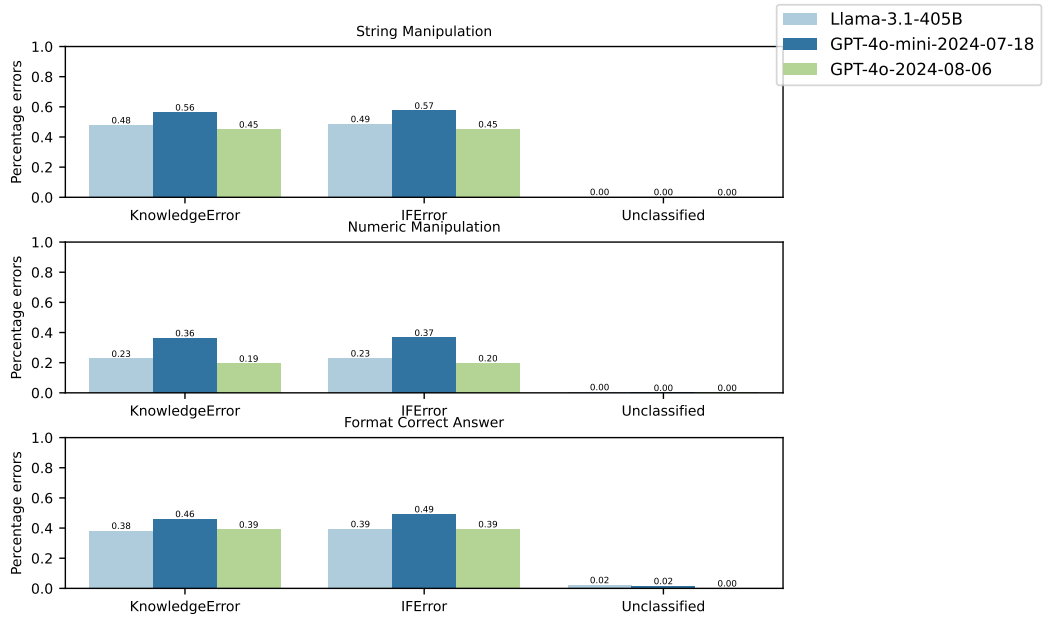


Figure 32: Frontier Models: Knowledge Errors and IFErrors

2064  
2065  
2066  
2067  
2068  
2069  
2070  
2071  
2072  
2073  
2074  
2075  
2076  
2077  
2078  
2079  
2080  
2081  
2082  
2083  
2084  
2085  
2086  
2087  
2088

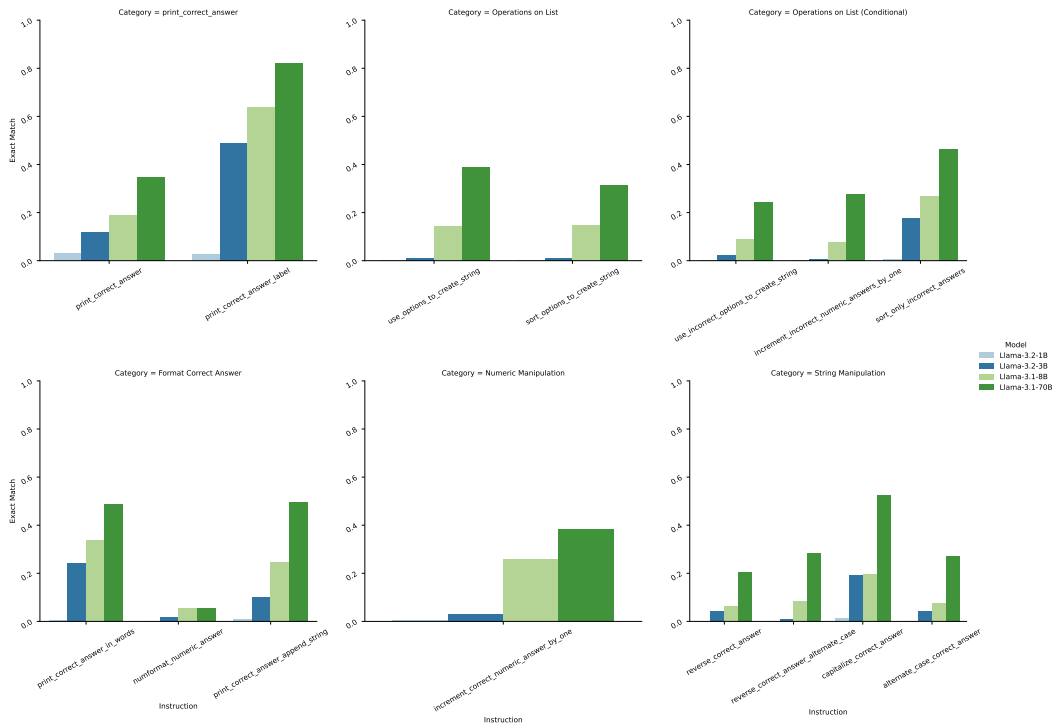


Figure 33: Performance variation (strict) of exact match scores for different instruction categories for Llama family of models

2091  
2092  
2093  
2094  
2095  
2096  
2097  
2098  
2099  
2100  
2101  
2102  
2103  
2104  
2105  
2106  
2107  
2108  
2109  
2110  
2111  
2112  
2113  
2114

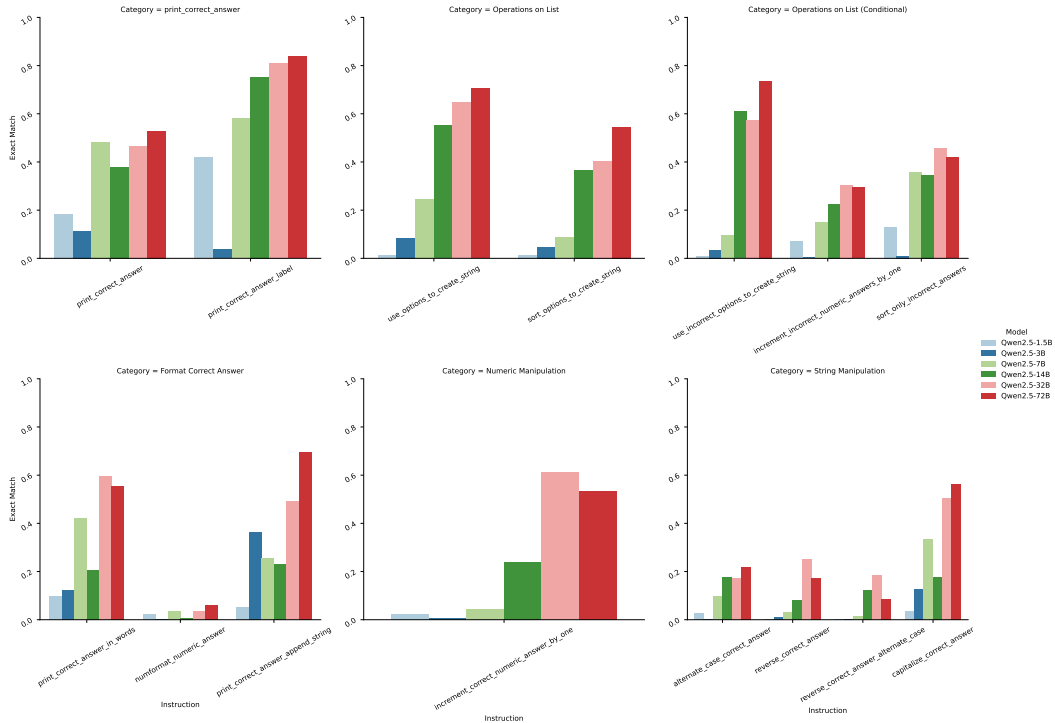


Figure 34: Performance variation (strict) of exact match scores for different instruction categories for Qwen family of models



2118  
 2119  
 2120  
 2121  
 2122  
 2123  
 2124  
 2125  
 2126  
 2127  
 2128  
 2129  
 2130  
 2131  
 2132  
 2133  
 2134  
 2135  
 2136  
 2137  
 2138  
 2139  
 2140  
 2141  
 2142  
 2143  
 2144  
 2145  
 2146  
 2147  
 2148  
 2149  
 2150  
 2151  
 2152  
 2153  
 2154  
 2155  
 2156  
 2157  
 2158  
 2159  
 2160  
 2161  
 2162  
 2163  
 2164  
 2165  
 2166  
 2167  
 2168  
 2169  
 2170  
 2171

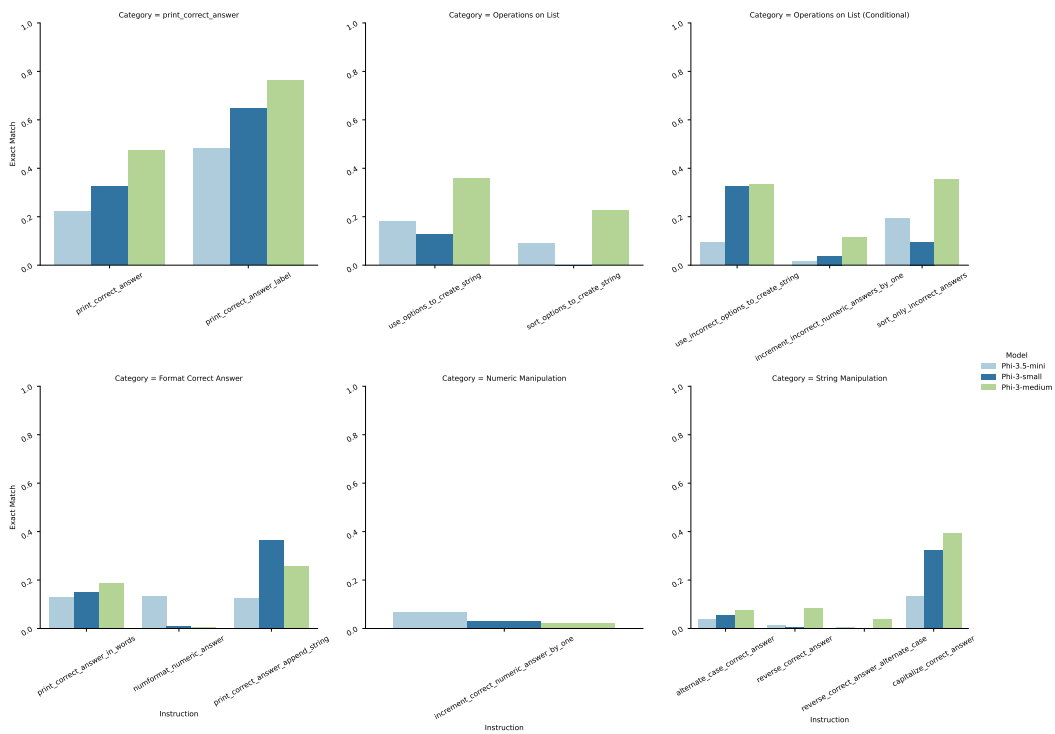


Figure 35: Performance variation (strict) of exact match scores for different instruction categories for Phi family of models

2172  
2173  
2174  
2175  
2176  
2177  
2178  
2179  
2180  
2181  
2182  
2183  
2184  
2185  
2186  
2187

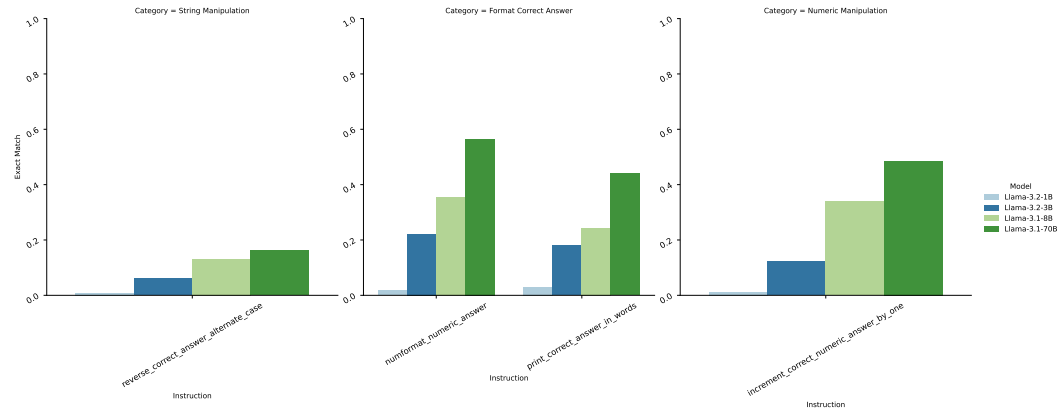


Figure 36: Performance variation (strict) of exact match scores for different instruction categories for Llama family of models on No Instruction following subset

2190  
2191  
2192  
2193  
2194  
2195  
2196  
2197  
2198  
2199  
2200  
2201  
2202




Figure 37: Performance variation (strict) of exact match scores for different instruction categories for Qwen family of models on No Instruction following subset

2203  
2204  
2205  
2206  
2207  
2208  
2209  
2210  
2211  
2212  
2213  
2214  
2215  
2216  
2217  
2218  
2219  
2220  
2221  
2222






Figure 38: Performance variation (strict) of exact match scores for different instruction categories for Phi family of models on No Instruction following subset

2223  
2224  
2225


2226  
2227  
2228  
2229  
2230  
2231  
2232  
2233  
2234  
2235  
2236  
2237  
2238  
2239  
2240  
2241  
2242  
2243  
2244  
2245  
2246  
2247  
2248  
2249  
2250  
2251  
2252  
2253  
2254  
2255  
2256  
2257  
2258  
2259  
2260  
2261  
2262  
2263  
2264  
2265  
2266  
2267  
2268  
2269  
2270  
2271  
2272  
2273  
2274  
2275  
2276  
2277  
2278  
2279



(a) Distractors on Llama Family of Models



(b) Distractors on Qwen Family of Models



(c) Distractors on Phi Family of Models