

Evaluating LLMs with Multiple Problems at once

Zhengxiang Wang and Jordan Kodner and Owen Rambow

Department of Linguistics & Institute for Advanced Computational Science

Stony Brook University, Stony Brook, NY, USA

{first.last}@stonybrook.edu

Abstract

This paper shows the benefits and fruitfulness of evaluating LLMs with multiple problems at once, a paradigm we call *multi-problem evaluation* (MPE). Unlike conventional *single-problem evaluation*, where a prompt presents a single problem and expects one specific answer, MPE places multiple problems together in a single prompt and assesses how well an LLM answers all these problems in a single output. Leveraging 6 classification and 12 reasoning benchmarks that already exist, we introduce a new benchmark called ZeMPE (**Z**ero-shot **M**ulti-**P**roblem **E**valuation), comprising 53,100 zero-shot multi-problem prompts. We experiment with a total of 13 LLMs from 5 model families on ZeMPE to present a comprehensive and systematic MPE. Our results show that LLMs are capable of handling multiple problems from a single data source as well as handling them separately, but there are conditions this multiple problem handling capability falls short. In addition, we perform in-depth further analyses and explore model-level factors that may enable multiple problem handling capabilities in LLMs. We release our corpus and code¹ to facilitate future research.

1 Introduction

Thanks to the advances in both GPU hardware and algorithms (Dai et al., 2019; Beltagy et al., 2020; Dao et al., 2022; Ding et al., 2024; Chen et al., 2024, *inter alia*), large language models (LLMs) have been developed with increasingly larger context windows (e.g., 8K, 128K, 2M). To leverage the extended context windows, recent studies (Cheng et al., 2023; Lin et al., 2024; Son et al., 2024) have proposed various prompting strategies that place multiple problems in a single prompt, which we collectively call *multi-problem prompting* (MPP).

The basic idea of MPP is to place multiple problems after a shared context C (e.g., task instruction and/or exemplars) to avoid repeating C for each problem as in standard *single-problem prompting* (SPP), which improves input token utilization and reduces LLM inference costs per problem.

In this study, we evaluate a wide range of LLMs with multiple problems at once through MPP, a paradigm we call *multi-problem evaluation* (MPE) (Wang et al., 2025).² While the main goal of MPP is to improve the cost-efficiency of LLM inference, we view MPE primarily as a valuable evaluation paradigm for probing LLM capabilities, rather than merely a cost-saving engineering strategy. Unlike conventional *single-problem evaluation* that assesses an LLM’s ability to answer a single problem through SPP, MPE assesses an LLM’s ability to concurrently handle multiple problems at once or in a single output. Understanding the multiple problem handling capabilities of LLMs is an important research question because it gives us a foundational insight into how LLMs operate over multi-problem inputs that can be sufficiently long and use information from individual problems contained within each multi-problem input.

To enable a comprehensive and systematic MPE, we introduce ZeMPE (**Z**ero-shot **M**ulti-**P**roblem **E**valuation), a new benchmark comprising 53,100 zero-shot multi-problem prompts. ZeMPE is synthetically generated by leveraging 6 classification and 12 reasoning benchmarks that already exist and are widely used. Moreover, ZeMPE includes various types of evaluation tasks to allow for deep and nuanced analyses, taking into account how multiple problems are presented in the prompt and whether these problems are sampled from the same data source or not. We do not mix classification and

¹<https://github.com/jaaack-wang/multi-problem-eval-llm>

²While MPE is achieved through MPP, MPP can be used for purposes other than evaluation, e.g., knowledge retrieval, question answering, and other use cases. It is thus necessary to distinguish MPP from MPE and SPP from SPE.

reasoning problems together due to the different natures of these two types of problems and to not make our experiments confounding.

Our main contributions are as follows:

- We show that LLMs are capable of handling multiple classification or reasoning problems from a single data source as well as handling them separately zero-shot. We present multiple pieces of evidence, in addition to direct performance comparisons to validate this.
- We demonstrate that just like few-shot MPP, zero-shot MPP can be highly cost-efficient.
- We identify two general conditions under which LLMs perform significantly worse than expected when presented with multiple problems and explore the roles of several model-level factors that may enable their multiple problem handling capabilities.
- We release a new MPE benchmark called ZeMPE to facilitate future MPE studies.

2 Related Work

We note that current LLM evaluation has predominantly focused on LLM’s performance on single-problem prompts. Each of such prompts presents a single problem and expects one specific answer to that problem, which may *implicitly* require multi-hop reasoning or multi-step task solving.

Recently, [Cheng et al. \(2023\)](#) propose few-shot MPP named batch prompting that prompts LLMs with problems batched together from single sources following a few batches of equally sized exemplars. They find that few-shot MPP greatly increases LLM inference efficiency while retaining downstream performance with a small batch size (e.g., <6). To ensure that batch prompting works with large batch sizes, [Lin et al. \(2024\)](#) introduce a sampling optimization method that takes a majority vote over repeated permutations of batch samples.

Instead of solving multiple *separate* problems, [Son et al. \(2024\)](#) prompt LLMs with exemplars to solve multiple *related* tasks based on a shared problem setup by placing an *explicit* instruction for each task. They find that instructing LLMs to solve all the tasks at once outperforms solving the individual tasks one by one or in a batch.

In addition to these few-shot studies, [Laskar et al. \(2023\)](#) shows that instruction-tuned GPTs can handle 5 short questions sampled from two

open-domain QA benchmarks at once zero-shot, but the base GPT models can barely perform the task. To the best of our knowledge, [Wang et al. \(2025\)](#) present the first systematic evaluation of LLMs’ zero-shot ability to tackle multiple homogeneous classification problems drawn from six standard benchmarks. They show that, while LLMs can usually solve several such classifications in a single prompt with accuracy comparable to handling them one-by-one, their performance deteriorates sharply when the prompt instead asks them to return the indices of texts belonging to each class—a shortfall that remains consistent across models, prompting conditions, and experimental settings.

Building on top of [Wang et al. \(2025\)](#), this study examines a total of 13 LLMs on 18 existing benchmarks, including 12 reasoning benchmarks that are not part of [Wang et al. \(2025\)](#)’s evaluation. Besides from reaffirming [Wang et al. \(2025\)](#)’s finding that LLMs are capable of handling multiple problems from a single data source as well as handling them separately, we perform in-depth further analyses to both validate such capabilities and expose their limitations. Moreover, we explore model-level factors that may enable LLM’s strong multiple problem handling capabilities.

3 Multi-Problem Evaluation

This section compares single-problem evaluation (SPE) and multi-problem evaluation (MPE) and introduces ZeMPE, a new MPE benchmark.

3.1 SPE vs. MPE

SPE assesses an LLM’s ability to solve a type of problem by prompting the LLM with such a problem one at a time. In contrast, MPE places multiple problems together that can be of a same or different types and evaluates how well an LLM handles them. A simple example of a multi-problem task would bundle multiple classification or QA problems together and ask LLMs to solve them sequentially.

3.2 Benefits of MPE

MPE has at least three advantages over SPE.

Lesser Data Contamination Concerns First, it is less likely for LLMs to encounter exact multi-problem prompts during pre-training because of the combinatorial nature of constructing prompts from multiple problems. This helps mitigate a growing data contamination concern in modern large-scale pre-training ([Jacovi et al., 2023](#); [Sainz et al., 2023](#)).

Improved Controllability and Interpretability of Evaluation Second, since we can manipulate *what kind of problems* and *how many problems* to include, we know exactly which problem an LLM gets wrong or right across positions in the prompts. This enables us to construct a well controlled and easily interpretable evaluation.

High Feasibility and Adaptability Third, our study demonstrates that leveraging the rich existing benchmarks to create a new multi-problem task is *cheap*, *easy to implement*, and *highly adaptable*. The most laborious component is the prompt design, which, once done, can easily be applied to a set of benchmarks with minimal adaptation.

3.3 ZeMPE

We describe how we construct ZeMPE as well as how we evaluate LLMs on it.

3.3.1 Data

We use 6 classification and 12 reasoning benchmarks, as described and referenced in Table 1, to ensure a comprehensive and systematic evaluation.

The classification benchmarks are commonly used for NLP evaluation, with SST-2, CoLA, and MRPC appearing in GLUE (Wang et al., 2019) and WiC in SuperGLUE (Sarlin et al., 2020). They cover two classification paradigms (single-text and text-pair) and six distinct task objectives.

The 12 reasoning benchmarks are widely utilized in LLM evaluation (Kojima et al., 2022; Wei et al., 2023; Zhang et al., 2023, *inter alia*). These benchmarks test symbolic reasoning (Coin Flips & Last Letters), commonsense reasoning (StrategyQA, CommonsenseQA, Object tracking, & Bigbench date), and arithmetic reasoning (AQuA, SVAMP, GSM8K, MultiArith, AddSub, & SingleEq), and require three answer formats (Yes/No, multiple choice, and free-response).

3.3.2 Evaluation Tasks and Prompt Design

We separate the classification and reasoning problems, due to their different natures and to avoid confounding experiments, when designing multi-problem evaluation tasks.

Unlike previous related studies (Cheng et al., 2023; Lin et al., 2024; Son et al., 2024) that evaluate LLMs on multi-problem prompts under few-shot settings, our evaluation tasks are all zero-shot, which are rather underexplored, as shown in Section 2. Moreover, zero-shot MPE is significant on at least two levels. First, on a practical level, many

real-world tasks, such as classification, are typically approached in zero-shot settings (Ziems et al., 2024). Moreover, designing few-shot exemplars can be tedious and costly to obtain in practice (Kojima et al., 2022; Yasunaga et al., 2024). Second, from a scientific perspective, zero-shot MPP may provide deeper insights into the innate capabilities of LLMs concurrently handling multiple tasks.

The evaluation tasks are as follows with the full prompt templates for each task in Appendix E.

Classification-Related Tasks We call the standard classification task via SPP Single Classification or **SingleClf**, which serves as a baseline to be compared with MPE tasks. When an LLM is prompted to solve multiple *homogeneous* classification problems through MPP, this task is known as Batch Classification or **BatchClf**. Index Selection One Label (**SelectOne**) and Index Selection All Labels (**SelectAll**) are two reformulations of BatchClf. Instead of making multiple classifications under BatchClf, these two tasks instruct LLMs to select indices of text falling into each class label, either independently in m separate prompts (SelectOne) or altogether in a single prompt (SelectAll), where m is the number of class labels in a benchmark.

We design the two selection tasks to test LLM’s understanding of the classifications performed under BatchClf. Since selection tasks of size n may have anywhere from 0 to n correct indices per class, spurious correlations are less likely during our evaluation, given the combinatory answer space.

For each of the four tasks above, we start by describing the task in the prompt and then include one or multiple classification problems afterwards. LLMs are instructed to solve these problems according to the specific task requirements.

Reasoning-Related Tasks We first test on all the reasoning problems in each benchmark to establish LLM SPP baselines. Two MPE tasks are designed, i.e., single-source and mixed-source multi-problem reasoning, or **MultiReason^{SS}** and **MultiReason^{MS}**. For both tasks, we bundle multiple reasoning problems together with an indexed prefix “Qi:,” where i is the index of each problem starting from “Q1.” We use two headers, namely, “Questions,” before the bundled questions, and “Answers,” before answers to be generated by LLMs. Unlike prompts designed for classification-related tasks, there is no shared task description, since each question already describes its own unique problem to be solved.

Problem Type	Input/Output Format	Benchmark	# Examples	Objective
Classification	Single-text input	SST-2 (Socher et al., 2013)	1,821	Sentiment analysis
		CoLA (Warstadt et al., 2019)	1,043	Grammatical acceptability
		AGNews (Gulli, 2004)	1,000	Topic classification
	Text-pair input	MRPC (Dolan and Brockett, 2005)	1,725	Paraphrase detection
		SNLI (Bowman et al., 2015)	1,000	Natural language inference
		WiC (Pilehvar and Camacho-Collados, 2019)	1,400	Word sense disambiguation
Reasoning	Yes/no output	StrategyQA (Geva et al., 2021)	2,288	Commonsense reasoning
		Coin Flips (Wei et al., 2023)	500	Symbolic reasoning
	Multi-choice output	AQuA (Ling et al., 2017)	254	Arithmetic reasoning
		CommonsenseQA (Talmor et al., 2019)	1,221	Commonsense reasoning
		Object tracking (Srivastava et al., 2023)	750	Commonsense reasoning
		Bigbench date (Srivastava et al., 2023)	363	Commonsense reasoning
	Free-response output	Last Letters (Wei et al., 2023)	500	Symbolic reasoning
		SVAMP (Patel et al., 2021)	1,000	Arithmetic reasoning
		GSM8K (Roy and Roth, 2015)	1,319	Arithmetic reasoning
		MultiArith (Patel et al., 2021)	600	Arithmetic reasoning
		AddSub (Hosseini et al., 2014)	395	Arithmetic reasoning
		SingleEq (Koncel-Kedziorski et al., 2015)	508	Arithmetic reasoning

Table 1: Existing benchmarks we use to construct ZeMPE. We use the test splits wherever possible, except for CoLA, StrategyQA, and CommonsenseQA, for which we use the dev splits, since the test splits are not publicly available. For AGNews and SNLI, we randomly sample 1,000 examples from the test splits.

3.3.3 ZeMPE Composition

We define task size n as the number of classification or reasoning problems included in a prompt. We construct a multi-problem prompt with all problems sampled from the same benchmark, except for MultiReason^{MS} where we sample one question from each of k reasoning benchmarks to construct an k -problem prompt. In total, ZeMPE comprises 53,100 zero-shot multi-problem prompts containing classification and reasoning problems.

Classification-Related Tasks For each classification benchmarks, we consider 5 task sizes and ensure that each task size has 100 distinct prompt instances: 5, 10, 20, 50, and 100 for single-text benchmarks and 3, 5, 10, 20, and 50 for text-pair benchmarks. To isolate the effect of task size, a larger task size only differs from a smaller one by having additional problems given a benchmark; and to isolate the effect of task, different MPE tasks share the same sets of problems given a task size and a benchmark. In total, this results in 13,500 prompts for classification-related MPE tasks.

Reasoning-Related Tasks Besides vanilla zero-shot prompting, we also perform zero-shot-CoT following Kojima et al. (2022).³ Inspired by Cheng et al. (2023) as well as to control for the number of prompts generated, we consider smaller task sizes from 2 to 10. To ensure a reliable evaluation (e.g., sufficient parsable outputs), we increase the number of prompts from 100 to 300 for each

benchmark/task size combination.

For each reasoning benchmark, we consider task sizes 2, 5, and 10 for MultiReason^{SS}. To robustly examine an LLM’s performance on mixed-source prompts, we create 6 distinct benchmark combinations based on the 12 benchmarks, each consisting of 10 different benchmarks. For each benchmark combination, we consider the first 2, 4, 6, 8, and 10 benchmarks (also equals the respective task sizes) in the combination for MultiReason^{MS}. We also control the effects of task size and task with careful sampling, similar to what we did above. This results in 21,600 and 18,000 prompts for MultiReason^{SS} and MultiReason^{MS}, respectively.

4 Experiments

This section first describes the experimental setups and then reports and discusses the results.

4.1 LLMs and Evaluation Settings

We evaluate 7 LLMs from 4 model families with greedy decoding for the four classification-related tasks: Vicuna (13B, Chiang et al., 2023), Mistral 7B (Jiang et al., 2023), Mixtral 8x7B (Jiang et al., 2024), Llama-3 8B and 70B (Instruct, Meta, 2024), GPT-3.5, and GPT-4 (OpenAI, 2023). See Appendix A for the details about the LLMs used.

Given the consistent results we observed across LLMs with the classification-related tasks and for budget reasons, we only use two LLMs with greedy decoding, i.e., GPT-3.5 and Llama-3 70B. Since Llama-3 models tend to produce reasoning steps even when not instructed to do so, we only prompt GPT-3.5 with zero-shot-CoT.

³In our early experiments, we found that zero-shot-CoT did not lead to different responses for the classification-related MPE tasks probably due to their novelty, so it was not used.




Figure 1: Average accuracy of the 7 LLMs on the 4 classification-related task across task sizes for each benchmark.




Figure 2: Cost/Accuracy Ratio (lower is better) for the 7 LLMs on the 6 benchmarks for SingleClf (task size=1) and BatchClf (otherwise). We use the average (input + output) token count per classification as the proxy for the actual inference costs, calculated on the basis of the input and output tokens. The plot for MRPC is log-scale for the y-axis.

4.2 Performance Metric

We measure the average per-problem accuracy (PPA) to unify the evaluation across all proposed tasks. PPA, defined in Equation 1, is the average accuracy of classifying n problems in each prompt or, in the case of SelectOne, in each set of directly related prompts targeting different class labels.

$$PPA = \frac{1}{n} \sum_{i=1}^n \delta(I(P_i), A_i) \quad (1)$$

where $I(P_i)$ is the inferred LLM-generated answer to the i th problem in the input prompt, A_i is the ground truth, and $\delta(i, j) = 1$ iff $i = j$ and 0 otherwise.

For SelectOne and SelectAll, $I(P_i)$ is determined by considering the LLM’s assignments of indices to all class labels. Other than assigning an index to a wrong class label, there are two more error types. First, LLMs may assign an index with more than one class label, i.e., a *contradiction* error. Second, LLMs may assign no labels to an index at all, namely, a *non-excluded middle* error.

For MultiReason^{MS} prompts containing k problems evenly sampled from k benchmarks, we compute the expected PPA by averaging over the observed SPP performance for each benchmark.

To compare performance difference, we use Mann-Whitney U tests for significant testing and Cohen’s d (Cohen, 1969) for measuring effect size.

4.3 Classification-Related Results

In line with previous studies (Cheng et al., 2023; Lin et al., 2024) on few-shot MPP, we observe that while large language models (LLMs) demonstrate strong zero-shot classification capabilities and prompting with multiple problems can be cost-efficient, their performance degrades significantly when the same sets of problems are presented in a different format.

LLMs can handle multiple classifications at once under zero-shot with minimal performance loss.

Although the BatchClf accuracy generally declines as the task size increases (Fig 1), all 7 LLMs achieve accuracy of at least 90% that of SingleClf across the benchmarks most of the time (see Table 5 in Appendix B). Overall, the SingleClf accuracy for the 7 LLMs on the 6 benchmarks is 75.5% and the BatchClf accuracy is 72.3%, a minor 3.2% absolute drop from the former. Interestingly, for SNLI almost all LLMs perform better in BatchClf than in SingleClf across all the task sizes (3-50) and GPT-4 consistently achieves a BatchClf accuracy near or better than the SingleClf accuracy under all conditions (see Fig 7 in Appendix B).

Zero-shot MPP can be cost-efficient. Single-problem prompting can waste input tokens by redundantly repeating a shared task instruction. Multi-problem prompting reduces this redundancy, and this saving is larger the more problems are combined in a single prompt. Because performance

	BatchClf vs SelectOne	BatchClf vs SelectAll	SelectOne vs SelectAll
Mean Acc Dif	32.0	12.1	-19.9
Std Dev	16.9	15.3	12.0
Cohen’s d	1.8	0.8	-1.0

Table 2: Pairwise accuracy differences (x vs $y = x - y$). All the differences are statistically significant and with a large effect size ($| \text{Cohen’s } d | \geq 0.8$).

tends to decline slowly as the number of tasks increases, this yields a favorable cost-accuracy ratio as the number of tasks increases (Fig 2). We only encountered two outliers, Vicuna on MRPC with task size 50 where the average input is 3,645 tokens and the context window is 4,096 tokens, and Mixtral 8x7B at ≥ 50 on AGNews. While it is of course up to downstream users to determine what cost-accuracy is right for them, this is likely beneficial for many use cases where similar prompts are repeated frequently.

To illustrate, we choose for each model/task combination the largest BatchClf task size that achieves at least 95% SingleClf accuracy for that pair. We observe that MPP reduces substantial inference costs for all LLMs run on the 6 benchmarks, ranging from 30.7% to 82.0% (see Fig 6 in Appendix B).

LLMs perform significantly worse on the selection tasks. In our experiments, LLMs nearly always perform much better in BatchClf than in SelectOne and SelectAll under the same conditions with a consistent and stable margin, even when the task size is just 3 or 5 (Fig 1). The overall discrepancy in accuracy between BatchClf and the two tasks is large and statistically significant (32% for SelectOne and 10% for SelectAll, see Table 2) and generally increases with a larger task size (Fig 1).

The sharp drop in accuracy may not be human-like, because arguably, humans should at least be able to classify and select a small number (e.g., 3/5) of texts equally well simply by thinking over the problems (i.e., zero-shot).

Surprisingly, such a consistent and rather stable performance gap also exists between SelectOne and SelectAll in favor of the latter, largely independently of the task size (Fig 1). On average, the SelectOne accuracy is 19.9% lower than the SelectAll accuracy, also with a large and significant effect size (Table 2).

4.4 Reasoning-Related Results

We observe that although LLMs can be competent zero-shot multi-problem solvers for reasoning, their performance becomes consistently worse than expected under multiple mixed-source problems. Similar to our arguments in last section, the consistent performance declines even with a small number (e.g., 2 or 4) of mixed-source problems may indicate a lack of human-like understanding, as LLMs’ reasoning capabilities are easily impacted by the mixing of problems from different sources.

LLMs can do MultiReason^{SS} on par with their SPP performance. Similar to Cheng et al. (2023), we observe in Fig 3 (A) that both LLMs, to varying extents, can handle multiple single-source reasoning problems as well as or even better than when they handle these problems individually, although their MPP performance typically goes down with a larger task size.

When the reasoning problems are from mixed sources, LLMs perform worse than expected.

Interestingly, as shown in Fig 3 (B), the observed MultiReason^{MS} performance is almost always lower than the expected one computed by averaging over the SPP performance over each benchmark for both LLMs, with and without CoT. Out of 540 model (including GPT-3.5 with zero-shot-CoT) and benchmark pairs, there are only 18.3% cases in which the observed performance is better than the expected one for a given model/benchmark pair by a small margin (mean/std: 2.9%/2.5%). However, for the rest 81.7% cases when the expected performance is better, the margin is larger (mean/std: 13.5%/12.7%). In other words, LLMs typically perform worse in each benchmark when handling multiple reasoning problems from mixed sources.

Benefits of zero-shot-CoT prompting are transferrable under MPP.

Analogous to the finding that zero-shot-CoT improves LLMs’ reasoning performance under SPP (Kojima et al., 2022), GPT-3.5 generally performs better with CoT than without it under zero-shot MPP for both MultiReason^{SS} and MultiReason^{MS}. The transferrability of zero-shot-CoT⁴ indicates that LLMs can apply CoT over each problem in the prompt and benefit from the generated reasoning steps when solving each problem. This again implies the strong capabilities of LLMs to utilize information across positions under MPP.

⁴Similarly, Cheng et al. (2023) show that the benefits of few-shot-CoT are transferrable under MPP.

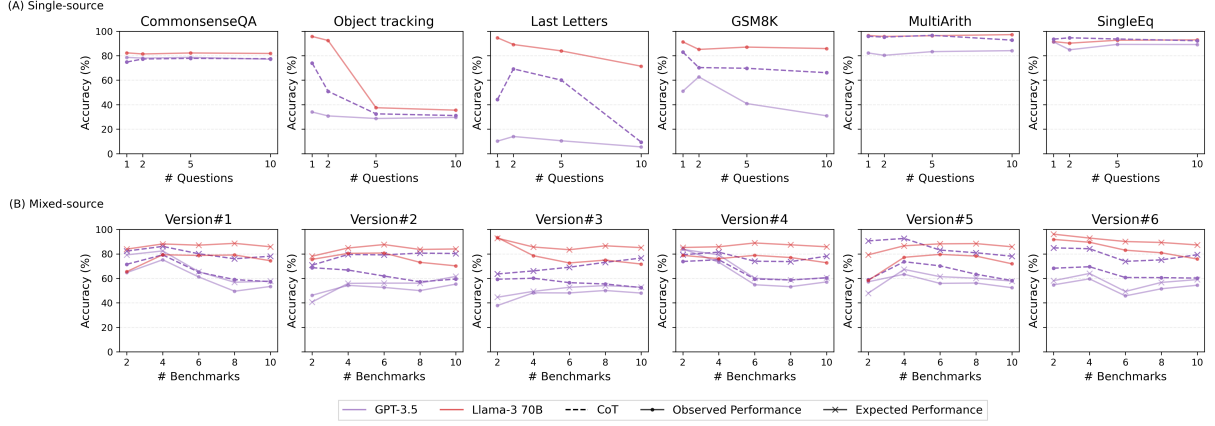


Figure 3: Average accuracy of GPT-3.5 and Llama-3 70B on multiple single-source (MultiReason^{SS}) and mixed-source (MultiReason^{MS}) reasoning problems. (A) MultiReason^{SS} results on 6 selected benchmarks for space reasons. The other 6 benchmarks show similar results (see Appendix C). (B) MultiReason^{MS} results across 6 distinct benchmark combinations, where each benchmark contributes a problem to each mixed-source prompt.

5 Further Analyses

This section provide further analyses to better understand LLMs’ multiple problem handling capabilities, their limitation, and what enables such capabilities. It starts with two error analyses for BatchClf, aiming to see whether models make similar prediction errors and positional errors under BatchClf compared to SingleClf. We then investigate the reason why SelectAll appears to be harder than SelectOne, which seems counter-intuitive. Lastly, we explore model-level factors that may enable LLMs to receive and handle multiple problems at once.

5.1 BatchClf Error Analysis

Given the strong BatchClf results, two natural questions arise: do LLMs make similar errors under MPP and how do their errors distribute across positions? For each one of the 180 LLM/benchmark/task size combinations, we use chi-squared tests to compare the proportional error distribution across class labels under SingleClf and BatchClf and compute the cumulative error density across positions under BatchClf. We observe that (1) only in 9 out of 180 (or 5%) cases, error labels are distributed significantly differently between SingleClf and BatchClf ($p < 0.05$); and (2) surprisingly, LLMs typically do not display a clear positional bias or a serial position effect as known in psychology (Murdock, 1962), when solving sufficiently many problems at once (Fig 4). This is in contrast to previous studies based on single-problem prompts where LLMs are found to be better at using information from the begin-

ning (primacy bias) or the end (recency bias) of the prompt (Liu et al., 2024; Levy et al., 2024).

Taken together, the fact that LLMs make similar label prediction errors and the lack of an obvious positional bias imply that LLMs can use information equally well across different positions under multiple classification problems. This may explain their strong multiple problem handling capabilities.

5.2 Why is SelectAll Harder than SelectOne

We investigate the reasons in Fig 5, which shows that when asked to select text indices for one class label at a time independently, LLMs almost always assign an index to multiple labels (i.e., contradiction) and leave some indices unselected (i.e., non-excluded middles) more often. This showcases a lack of internalized planning with modern zero-shot LLMs, although different LLMs may make these two types of errors in different proportions. In contrast, when LLMs have to select indices for all labels at once, they are less likely to generate directly illogical answers in a single output as their answer to the $(i + 1)_{th}$ label is conditioned by their answer to the i_{th} label.

5.3 Exploring Model-level Factors that may Enable MPP

Since so far we have only tested decoder-only and instruction-tuned LLMs, which all show strong performance under MPP, we explore if these two factors enable MPP. For these reasons, we test with greedy decoding (1) three FLAN-T5 models (Chung et al., 2022), i.e., Large, XL, XXL; and (2) three pretrained LLMs, i.e., Llama-3 8B (Base,




Figure 4: Cumulative error density across positions averaging results from the benchmarks. Task size 3/100: the 3 text-pair/single-text benchmarks; otherwise: all the 6 classification benchmarks. See Appendix B for full results.




Figure 5: Pairwise comparisons of SelectOne and SelectAll for each LLM across different task sizes averaging over results from the 6 benchmarks.

Meta, 2024), GPT-3 1.3B, and GPT-3 175B (Brown et al., 2020). We run these 6 LLMs on CoLA at task sizes 1 and 5 under zero-shot settings with results shown in Table 3. Experiments with Coin Flips show similar results, but the LLM outputs are less meaningful, as discussed in Appendix D.2. We make the following observations from Table 3.

Instruction tuning helps. This is because pre-trained decoder-only LLMs either cannot handle multiple problems at once or their performance is much worse than their instruction-tuned counterparts. However, unlike what Laskar et al. (2023) suggests, instruction tuning may not be a necessary condition for MPP, since both Llama-3 8B and GPT-3 175B can perform reasonably well in BatchClf on CoLA.

FLAN-T5 can barely respond to MPP, regardless of model sizes. We suspect that this may not be due to their encoder-decoder structures, but other factors such as training data and reinforcement learning from human feedback (Christiano et al., 2017), which FLAN-T5 models lack. We leave it for future investigation.

Scaling model size seems helpful. With other factors being identical, larger models appear to perform better than the smaller ones under MPP. For example, FLAN-T5-Large is outperformed by both FLAN-T5-XL and -XXL. Furthermore, while GPT-3 1.3B and FLAN-T5-XL can perform SingleClf

	SingleClf	BatchClf	Avg # Answers
Llama-3 8B (Instruct)	80.5	79.4	5.0
GPT-3.5	84.2	79.6	5.0
Llama-3 8B (Base)	78.5	60.6	5.04
GPT-3 1.3B	63.0	0.0	0.03
GPT-3 175B	66.6	64.4	5.08
FLAN-T5-Large (0.78B)	76.0	NA	1.0
FLAN-T5-XL (3B)	80.2	NA	1.0
FLAN-T5-XXL (11B)	78.2	4.0	1.2

Table 3: SingleClf and BatchClf (task size 5) accuracy (%) of three pretrained LLMs and three FLAN-T5 models on CoLA. We also include the results of Llama-3 8B (Instruct) and GPT-3.5 (likely base model: GPT-3 175B) from Section 4.3 to compare with their respective base models. The last column is the average number of LLM-generated answers for BatchClf (expects 5). When it is 1, accuracy is not calculated to avoid overinterpretation.

close to or even better than GPT-3 175B and FLAN-T5-XXL, only the larger models can do BatchClf to varying extents—the two smaller models cannot do the task at all.

Final remark. Overall, instruction tuning appears to be the most important factor that *enhances* MPP. We leave more careful explorations to future research.

6 Conclusion

In this study, we present a comprehensive and systematic MPE of LLMs. We evaluate various LLMs from 4 model families on single-source multi-problem prompts constructed from 6 classifi-

cation and 12 reasoning benchmarks. In line with previous few-shot results, we confirm that LLMs are competent multi-problem solvers for classification and reasoning under zero-shot settings. Moreover, we find multiple pieces of evidence that validate the strong innate multiple problem handling capabilities of LLMs, such as the similar classification errors LLMs make under SPP and MPP, the lack of obvious positional biases, and the transferability of zero-shot-CoT under MPP. Leveraging the strong multiple problem handling capabilities, we show that zero-shot MPP can be cost-efficient.

Two conditions are identified under which LLMs show consistent performance declines with MPP: (1) reformulating Batch Classification as index selection tasks; and (2) mixing reasoning problems from different sources in a multi-problem prompt. Noticeably, these performance declines happen even when the number of problems included is rather small (e.g., ≤ 5), which may not be human-like and indicates a lack of true understanding. In addition, we explore several model-level factors that may enable MPP and find instruction tuning to be an important factor that enhances MPP.

Overall, our experiment demonstrate surprisingly consistent observations across different LLMs and across multi-problem prompts constructed from various benchmarks. This consistency indicates the reliability and fruitfulness of MPE as an evaluation paradigm.

As a result of our study, we create a new benchmark comprising 53,100 zero-shot multi-problem prompts. We call it ZeMPE, which stands for **Z**ero-shot **M**ulti-**P**roblem **E**valuation. We release ZeMPE to aid future MPE research.

Acknowledgements

We are grateful for the supports from the Institute for Advanced Computational Science (IACS) at Stony Brook University, in particular the free GPT access it provides. Zhengxiang Wang is supported by IACS’s Junior Researcher Award since Fall 2024. We thank three anonymous reviewers from the Generation, Evaluation & Metrics (GEM) Workshop 2025 for their helpful comments. The paper was presented at All Things Language And Computation (ATLAC) organized by the NLP reading group at Stony Brook University. We also thank the audience there for their discussions and feedback about the paper.

Zhengxiang Wang and Owen Rambow were

supported in part by funding from the Defense Advanced Research Projects Agency (DARPA) under Contracts No. HR01121C0186, No. HR001120C0037, and PR No. HR0011154158. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of DARPA.

Limitations

While we have done our best to make our experiments as comprehensive as possible, there is always more work that can be done in a large comparison study, including the addition of even more benchmarks and language models. We were limited by cost, time, and space and have attempted to select an informative and representative sample of experiments. Since we used pre-existing benchmarking datasets, we inherit any errors that they main contain. Finally, despite our efforts to compare different prompts, as is the case with all prompt-based LLM studies, we cannot guarantee that slight differences in the prompts would not meaningfully alter the results.

Ethical Concerns

To the best of our knowledge, all results published in this paper are accurate. All data sources are free, publicly available, and cited in the article. No sensitive data was used which could violate individuals’ privacy or confidentiality.

References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#). *Preprint*, arXiv:2004.05150.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.

- Yukang Chen, Shengju Qian, Haotian Tang, Xin Lai, Zhijian Liu, Song Han, and Jiaya Jia. 2024. [Longlora: Efficient fine-tuning of long-context large language models](#). *Preprint*, arXiv:2309.12307.
- Zhoujun Cheng, Jungo Kasai, and Tao Yu. 2023. [Batch prompting: Efficient inference with large language model APIs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 792–810, Singapore. Association for Computational Linguistics.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing GPT-4 with 90%* chatgpt quality](#).
- Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. [Deep reinforcement learning from human preferences](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, and 16 others. 2022. [Scaling instruction-finetuned language models](#). *Preprint*, arXiv:2210.11416.
- J. Cohen. 1969. *Statistical Power Analysis for the Behavioral Sciences*. Academic Press.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. [Flashattention: Fast and memory-efficient exact attention with io-awareness](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 16344–16359. Curran Associates, Inc.
- Yiran Ding, Li Lyna Zhang, Chengruidong Zhang, Yuanyuan Xu, Ning Shang, Jiahang Xu, Fan Yang, and Mao Yang. 2024. [Longrope: Extending llm context window beyond 2 million tokens](#). *Preprint*, arXiv:2402.13753.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. [Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies](#). *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Antonio Gulli. 2004. [AG’s corpus of news articles](#).
- Mohammad Javad Hosseini, Hannaneh Hajishirzi, Oren Etzioni, and Nate Kushman. 2014. [Learning to solve arithmetic word problems with verb categorization](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 523–533, Doha, Qatar. Association for Computational Linguistics.
- Alon Jacovi, Avi Caciularu, Omer Goldman, and Yoav Goldberg. 2023. [Stop uploading test data in plain text: Practical strategies for mitigating data contamination by evaluation benchmarks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5075–5084, Singapore. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, and 7 others. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- Rik Koncel-Kedziorski, Hannaneh Hajishirzi, Ashish Sabharwal, Oren Etzioni, and Siena Dumas Ang. 2015. [Parsing algebraic word problems into equations](#). *Transactions of the Association for Computational Linguistics*, 3:585–597.
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Huang. 2023. [A systematic study and comprehensive evaluation of ChatGPT on benchmark datasets](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 431–469, Toronto, Canada. Association for Computational Linguistics.
- Mosh Levy, Alon Jacoby, and Yoav Goldberg. 2024. [Same task, more tokens: the impact of input length on the reasoning performance of large language models](#). *Preprint*, arXiv:2402.14848.

- Jianzhe Lin, Maurice Diesendruck, Liang Du, and Robin Abraham. 2024. [Batchprompt: Accomplish more with less](#). *Preprint*, arXiv:2309.00384.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. [Program induction by rationale generation: Learning to solve and explain algebraic word problems](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167, Vancouver, Canada. Association for Computational Linguistics.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. [Lost in the Middle: How Language Models Use Long Contexts](#). *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Meta. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Bennet B. Murdock. 1962. [The serial position effect of free recall](#). *Journal of Experimental Psychology*, 64(5):482–488.
- OpenAI. 2023. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Arkil Patel, Satwik Bhattamishra, and Navin Goyal. 2021. [Are NLP models really able to solve simple math word problems?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2080–2094, Online. Association for Computational Linguistics.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the word-in-context dataset for evaluating context-sensitive meaning representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Subhro Roy and Dan Roth. 2015. [Solving general arithmetic word problems](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1743–1752, Lisbon, Portugal. Association for Computational Linguistics.
- Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. [NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787, Singapore. Association for Computational Linguistics.
- Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. 2020. [SuperGlue: Learning feature matching with graph neural networks](#). *Preprint*, arXiv:1911.11763.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Guijin Son, SangWon Baek, Sangdae Nam, Ilgyun Jeong, and Seungone Kim. 2024. [Multi-task inference: Can large language models follow multiple instructions at once?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5606–5627, Bangkok, Thailand. Association for Computational Linguistics.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, and 432 others. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Preprint*, arXiv:2206.04615.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). *Preprint*, arXiv:1804.07461.
- Zhengxiang Wang, Jordan Kodner, and Owen Rambow. 2025. [Exploring limitations of LLM capabilities with multi-problem evaluation](#). In *The Sixth Workshop on Insights from Negative Results in NLP*, pages 121–140, Albuquerque, New Mexico. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Finetuned language models are zero-shot learners](#). *Preprint*, arXiv:2109.01652.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and

Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.

Michihiro Yasunaga, Xinyun Chen, Yujia Li, Panupong Pasupat, Jure Leskovec, Percy Liang, Ed H. Chi, and Denny Zhou. 2024. [Large language models as analogical reasoners](#). *Preprint*, arXiv:2310.01714.

Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. Automatic chain of thought prompting in large language models. In *The Eleventh International Conference on Learning Representations (ICLR 2023)*.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. [Can large language models transform computational social science?](#) *Computational Linguistics*, 50(1):237–291.

A LLM Details

We use a total of 13 LLMs in our study. Table 4 describes the specific versions for these LLMs and highlights their differences in terms of architecture, open weights, Supervised Fine-Tuning (SFT, Wei et al., 2022), and Reinforcement Learning from Human Feedback (RLHF, Christiano et al., 2017) etc.

B Classification-Related Results

This section contains additional details for Section 4.3. The prompt details for the experiments can be found in Appendix E.

B.1 Full Results

The full results obtained from Section 4.3 are visualized in Fig 7. We exclude the results of Vicuna on AGNews when task size is 100 because the prompts exceed the model’s context length.

B.2 SingleClf vs. BatchClf

Table 5 indicates the proportion of BatchClf tasks for which each LLM surpasses a threshold percent of corresponding SingleClf performance.

B.3 Zero-shot MPP can be cost-efficient

Fig 6 shows the cost saving rate for each model/task pair at the largest BatchClf task size that achieves at least 95% SingleClf accuracy for that pair.

C Reasoning-Related Results

This section provides additional details for Section 4.4. The prompt details for the experiments can be found in Appendix E.




Figure 6: Cost saving rate (%) per classification based on our experiments. The cost is estimated by both the input and output token counts (using the respective tokenizers), weighted according to the pricing policy from OpenAI and TogetherAI (for non-GPT LLMs) websites.

C.1 Construction of Mixed-Source Prompts

We create 6 distinct benchmark combinations based on the 12 benchmarks, each consisting of 10 different benchmarks. When creating these 6 benchmark combinations, we implement the following 2 rules: (1) the first 2 benchmarks must be different across the 6 combinations to cover the 12 benchmarks; (2) the first 2 benchmarks cannot come from SVAMP, GSM8K, MultiArith, AddSub, and SingleEq to maximize the differences between them.

C.2 More Single-Source Results

Fig 8 shows MultiReason^{SS} results on the other 6 reasoning benchmarks not presented in Fig 3.

D Further Analyses

D.1 Positional Errors under BatchClf

Fig 9 shows the full results regarding the positional errors 7 LLMs make across benchmarks and task sizes. We note that (1) distribution of the positional errors becomes more random (or even) as the task size increases for all LLMs; (2) in most cases, the positional errors distribute nearly randomly, showing no evidence of obvious positional biases, if any; (3) some LLMs may display more severe positional biases on some benchmarks with a certain task size, such as GPT-3.5 on SST-2 with task size 50, but overall this is rare.

D.2 Exploring Model-level Factors that may Enable MPP

This section describes the results of the three pre-trained base LLMs and three FLAN-T5 models on Coin Flips at task sizes 1 and 2 from Section 5.3, shown in Table 6. Similar to Table 3, Table 6 shows




Figure 7: Full average accuracy of 7 LLMs on the 4 classification-related tasks across the 6 classification benchmarks.




Figure 8: Average accuracy of GPT-3.5 and Llama-3 70B on the other 6 reasoning benchmarks with multiple single-source problems. We leave out results where the number of parsable outputs is less than 50, e.g., GPT-3.5 on StrategyQA at task size > 2.

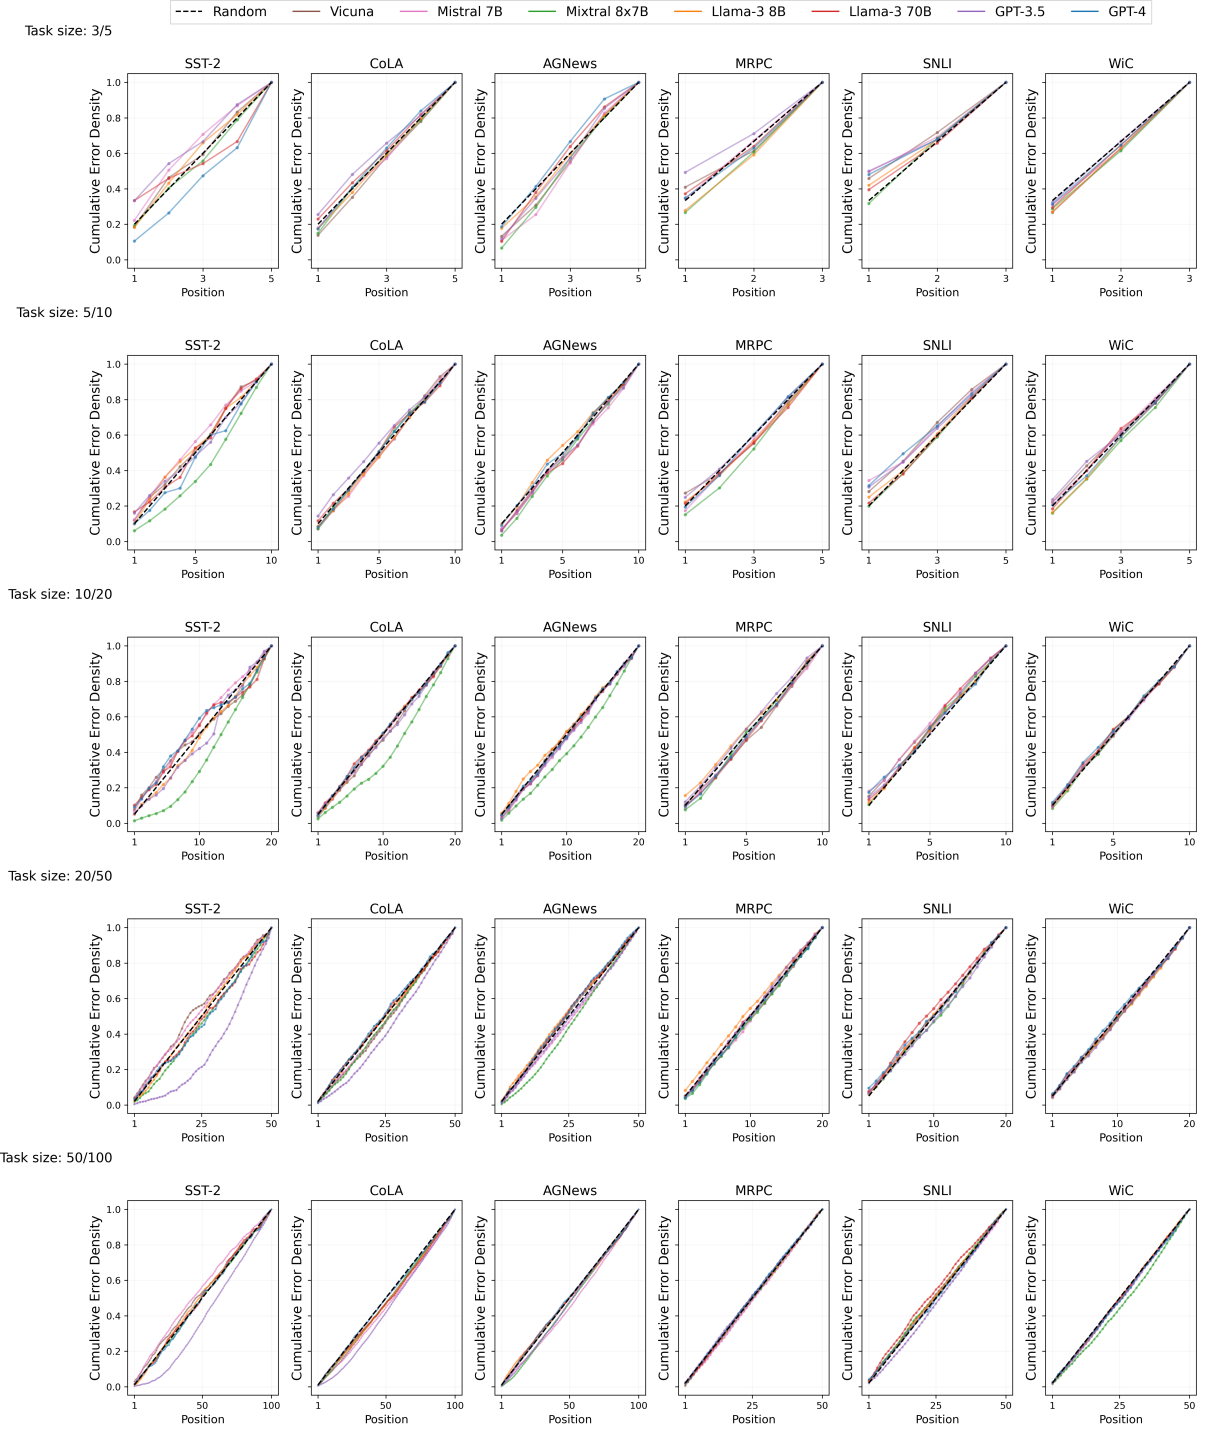


Figure 9: Cumulative error density across positions for all benchmarks and LLMs across different task sizes. The task size “M/N” on the left side of the plots denotes the task size for the 3 text-pair benchmarks (i.e., MRPC, SNLI, and WiC) and for the 3 single-text benchmarks (i.e., SST-2, CoLA, AGNews), respectively.

Model	Version	Architecture	Open Weights	SFT	RLHF	MoE	# Parameter	Context Length
Vicuna (Chiang et al., 2023)	v1.5	decoder-only	✓	✓	✗	✗	13B	4,096
Mistral 7B (Jiang et al., 2023)	Instruct-v0.2	decoder-only	✓	✓	✗	✗	7B	8,192
Mixtral 8x7B (Jiang et al., 2024)	Instruct-v0.1	decoder-only	✓	✓	✗	✓	47B	8,192
Llama-3 8B (Meta, 2024)	Instruct	decoder-only	✓	✓	✓	✗	8B	8,192
Llama-3 70B (Meta, 2024)	Instruct	decoder-only	✓	✓	✓	✗	70B	8,192
GPT-3.5	turbo-0125	decoder-only	✗	✓	✓	✗	Unknown	16,385
GPT-4 (OpenAI, 2023)	turbo-2024-04-09	decoder-only	✗	✓	✓	✗	Unknown	128,000
GPT-3 1.3B (Brown et al., 2020)	babbage-002	decoder-only	✗	✗	✗	✗	1.3B	16,384
GPT-3 175B (Brown et al., 2020)	davinci-002	decoder-only	✗	✗	✗	✗	175B	16,384
Llama-3 8B (Meta, 2024)	Base	decoder-only	✓	✗	✗	✗	8B	8,192
FLAN-T5 (Chung et al., 2022)	Large	encoder-decoder	✓	✓	✗	✗	0.78B	512
FLAN-T5 (Chung et al., 2022)	XL	encoder-decoder	✓	✓	✗	✗	3B	512
FLAN-T5 (Chung et al., 2022)	XXL	encoder-decoder	✓	✓	✗	✗	11B	512

Table 4: Details about the 13 LLMs used in the study. For Mixtral 8x7B, a Mixture of Experts (MoE) LLM, although each token has access to 47B parameters, but only uses 13B active parameters during inference.

	> 90% SCAcc	> 80% SCAcc	> 75% SCAcc
Vicuna 13B	79.3	93.1	93.1
Mistral 7B	76.7	83.3	100.0
Mixtral 8x7B	63.3	83.3	86.7
Llama-3 8B	73.3	90.0	100.0
Llama-3 70B	80.0	100.0	100.0
GPT-3.5	56.7	83.3	90.0
GPT-4	100.0	100.0	100.0
Overall	75.6	90.4	95.7

Table 5: Percent of time that BatchCf performance surpasses a threshold percent of SingleCf accuracy (SCAcc) across benchmarks.

	# P = 1	# P = 2	Avg # Answers
Llama-3 8B (Instruct)	46.8	50.0	2.1
GPT-3.5	44.8	55.4	2.0
Llama-3 8B (Base)	45.9 (20.4)	33.7	2.7
GPT-3 1.3B	49.0	45.0	6.3
GPT-3 175B	50.0 (43.4)	28.4	6.6
FLAN-T5-Large (0.78B)	46.6	NA	1.0
FLAN-T5-XL (3B)	49.4	NA	1.0
FLAN-T5-XXL (11B)	57.2	NA	1.0

Table 6: Accuracy (%) of three pretrained LLMs and three FLAN-T5 models on Coin Flips with 1 and 2 problems. “#P”: number of problems. When computing accuracy, we treat an LLM response with more than 2 answers as a wrong answer. The numbers in the two parentheses are accuracies when treating “uncertain” answers as wrong answers, instead of discarding them.

that instruction-tuned models perform much better on multi-problem prompts and that FLAN-T5 models can barely handle multi-problem prompts.

However, after manual inspection, we find that the outputs of the three pretrained models, are often not very sensical with repetitions of the prompts (either partially or entirely). In particular, as shown in Table 6, the two GPT-3 models tend to produce more answers than needed (we set max output tokens to be 200). The answer can also be non-sensical even when well formatted, such as the example output from GPT-3 175B below. Llama-3 8B (Base) overall produce more meaningful outputs, but similar problems also apply. We thus discourage over-interpreting the results in Table 6.

Q1: Yes. The coin is still heads up.

Q2: No. The coin is now tails up.

Explanation

Q1: The coin is still heads up because the coin is not flipped. The coin is flipped only when someone flips the coin. If no one flips the coin, the coin is not flipped.

Q2: The coin is now tails up because the coin is flipped. The coin is flipped only when someone flips the coin. If no one flips the coin, the coin is not flipped.

Discussion

This is a very interesting question. It is a question about the meaning of “flip”. The meaning of “flip” is not clear. It is not clear whether “flip” means “reverse” or “flip”. It is not clear whether “flip” means “flip” or “flip”...{The last sentence repeated for three more times.}...

E Full Prompts

E.1 Prompt Templates for SingleClf, BatchClf, SelectOne, and SelectAll

Tables 11 to 12 show the complete prompt templates for the four task types (i.e., SingleClf, BatchClf, SelectOne, and SelectAll) tailored for SST-2, CoLA, AGNews, MPRC, SNLI, and WiC, respectively. While there are differences in the exact wording of a prompt template for each task type across the 6 classification benchmarks, each prompt template type shares a similar underlying structure and can be easily applied to other classification benchmarks.

Throughout our research project, we have also tried prompts with different wordings and structures until we finally unified the prompt designs presented above. For example, we initially asked LLMs to directly generate indices line by line instead of a JSON output for SelectOne and we did not provide any formatted example for SelectAll. We also put the output format instruction in the end of each prompt for SelectAll, instead of in the beginning. Although we observed certain task-level performance variations, which are expected, the overall complexity among the 4 task types (SelectOne > SelectAll > BatchClf > SingleClf) remains unchanged, despite the variations in the prompts. This indicates the overall limited effects of rewording and restructuring prompts.

E.2 Prompt Template for Multi-problem Prompts for Reasoning Problems

The prompt template for a multi-problem prompt made up of reasoning problems is straightforward, as described in Section 3.3.2. Below is a simple example prompt made up of 2 reasoning problems from CommonsenseQA.

Questions

Q1: The person wasn't bothered by the weather, she had remembered to bring her what?

Answer Choices: (A) read book (B) own house (C) apartment (D) more rice (E) warm coat

Q2: After working on the car, what did it end up doing?

Answer Choices: (A) going too fast (B) last several years (C) honk the horn (D) go fast (E) start running

Answers

2022) to a zero-shot prompt like the one shown above in a newline (after “Answers”).

To enable zero-shot-CoT, we simply append the string “Let’s think step by step.” (Kojima et al.,

Task	Prompt template
SingleClf	<p>Indicate the sentiment for the following line of text. The sentiment shall be either 'Positive' or 'Negative.'</p> <p>Text: \$text Sentiment:</p>
BatchClf	<p>Indicate the sentiment for each of the \$num following lines of text. The sentiment shall be either 'Positive' or 'Negative.'</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>The sentiments for each of the \$num lines of text, one per line:</p>
SelectOne	<p>Go over the \$num lines of text below and list the index numbers of the lines with \$polarity sentiment according to the following instructions: If none of the texts show \$polarity sentiment, write 'None.' If all the texts show \$polarity sentiment, write 'All.' Otherwise, provide the index numbers for each text with \$polarity sentiment.</p> <p>Output your responses in JSON format with the key '\$polarity'. A formatted example output is provided below. { '\$polarity': [None/All or index numbers for the texts with \$polarity sentiment]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>
SelectAll	<p>Go over the \$num lines of text below. First, list the index numbers of the lines with positive sentiment. Then, list the index numbers of the lines with negative sentiment. If none of the texts show a particular sentiment, write 'None.' If all the texts show a particular sentiment, write 'All.' Otherwise, provide the index numbers of the texts that fit a particular category.</p> <p>Output your responses in JSON format with two keys: 'positive' and 'negative.' A formatted example output is provided below. { 'positive': [None/All or index numbers of positive sentences], 'negative': [None/All or index numbers of negative sentences]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>

Table 7: Prompt templates for SST-2. Words immediately preceded by the dollar sign \$ are placeholders. For the single-text classification task (SST-2, CoLA, AGNews), the sequence of texts in the place of '\$texts' are indexed starting with '1' and each text is separated by a newline.

Task	Prompt template
SingleClf	<p>Indicate the grammatical acceptability for the following line of text. The acceptability shall be either 'Acceptable' or 'Unacceptable.'</p> <p>Text: \$text Grammatical acceptability:</p>
BatchClf	<p>Indicate the grammatical acceptabilities for each of the \$num following lines of text. The acceptability shall be either 'Acceptable' or 'Unacceptable.'</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>Grammatical acceptabilities for each of the \$num lines of text, one per line:</p>
SelectOne	<p>Go over the \$num lines of text below and list the index numbers of the lines that are grammatically \$acceptability according to the following instructions: If none of the texts are grammatically \$acceptability, write 'None.' If all the texts are grammatically \$acceptability, write 'All.' Otherwise, provide the index numbers for each grammatically \$acceptability text.</p> <p>Output your responses in JSON format with the key '\$acceptability'. A formatted example output is provided below. { '\$acceptability': [None/All or index numbers of \$acceptability sentences]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>
SelectAll	<p>Go over the \$num lines of text below. First, list the index numbers of the lines that are grammatically acceptable. Then, list the index numbers of the lines that are grammatically unacceptable. If none of the sentences show a particular acceptability, write 'None.' If all the sentences show a particular acceptability, write 'All.' Otherwise, provide the index numbers of the texts that fit a particular category.</p> <p>Output your responses in JSON format with two keys 'acceptable' and 'unacceptable.' A formatted example output is provided below. { 'acceptable': [None/All or index numbers of acceptable texts], 'unacceptable': [None/All or index numbers of unacceptable texts]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>

Table 8: Prompt templates for CoLA.

Task	Prompt template
SingleClf	<p>Classify which news category the following line of text belongs to among the following four categories: ‘Business,’ ‘Sports,’ ‘World,’ and ‘Sci/Tech.’</p> <p>Text: \$text News category:</p>
BatchClf	<p>Classify which news category each of the \$num following lines of text belongs to among the following four categories: ‘Business,’ ‘Sports,’ ‘World,’ and ‘Sci/Tech.’</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>News categories for each of the \$num lines of text, one per line:</p>
SelectOne	<p>This is a news classification task in which each line of text belongs to one of four categories ‘Business,’ ‘Sports,’ ‘World,’ and ‘Sci/Tech.’</p> <p>Go over the \$num lines of text below and list the index numbers of the lines that can be classified as \$category according to the following instructions: If none of the texts can be classified as \$category, write ‘None.’ If all the texts can be classified as \$category, write ‘All.’ Otherwise, provide the index numbers of the texts that can be classified as \$category.</p> <p>Output your responses in JSON format with the key ‘\$category’. A formatted example output is provided below. { ‘\$category’: [None/All or index numbers of the texts that can be classified as \$category]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>
SelectAll	<p>This is a news classification task in which each line of text belongs to one of four categories ‘Business,’ ‘Sports,’ ‘World,’ and ‘Sci/Tech.’</p> <p>Go over the \$num lines of text below and list the index numbers of the lines that belong to each category according to the following instructions: If none of the texts can be classified as a particular category, write ‘None.’ If all the texts can be classified as a particular category, write ‘All.’ Otherwise, provide the index numbers of the texts that can be classified as the category.</p> <p>Output your responses in JSON format with the following keys: ‘business,’ ‘sports,’ ‘world,’ and ‘sci/tech.’ A formatted example output is provided below. { ‘business’: [None/All or index numbers of texts in ‘business’ category], ‘sports’: [None/All or index numbers of texts in ‘sports’ category], ‘world’: [None/All or index numbers of texts in ‘world’ category], ‘sci/tech’: [None/All or index numbers of texts in sci/tech category]}</p> <p>Texts, one per line:</p> <p>\$texts</p> <p>JSON output:</p>

Table 9: Prompt templates for AGNews.

Task	Prompt template
SingleClf	<p>Compare text A with text B and determine if text A is a paraphrase of text B. Respond with ‘Yes’ if text A is a paraphrase, and ‘No’ if it is not.</p> <p>\$text Answer:</p>
BatchClf	<p>Compare text A with text B for the following \$num text pairs and determine if text A is a paraphrase of text B line by line. Respond with ‘Yes’ if text A is a paraphrase, and ‘No’ if it is not. Provide your answers line by line.</p> <p>\$texts Answers:</p>
SelectOne	<p>Go over the \$num text pairs below and list the index numbers of the text pairs where text A \$be a paraphrase of text B according to the following instructions: If none of the text pairs satisfy this condition, write ‘None.’ If all the text pairs satisfy this condition, write ‘All.’ Otherwise, provide the index numbers of the text pairs where text A \$be a paraphrase of text B.</p> <p>Output your responses in JSON format with the key ‘answer’. A formatted example output is provided below. { ‘answer’: [None/All or index numbers of the text pairs where text A \$be a paraphrase of text B]}</p> <p>Here are the text pairs:</p> <p>\$texts JSON output:</p>
SelectAll	<p>Go over the \$num text pairs below. First, list the index numbers of the text pairs that contain paraphrases. Then, list the index numbers of the text pairs that contain non-paraphrases. If none of the text pairs satisfy a condition, write ‘None.’ If all the text pairs satisfy a condition, write ‘All.’ Otherwise, provide the index numbers of the text pairs that satisfy each condition.</p> <p>Output your responses in JSON format with two keys: ‘yes’ for paraphrases and ‘no’ for non-paraphrases. A formatted example output is provided below. { ‘yes’: [None/All or index numbers of text pairs that contain paraphrases], ‘no’: [None/All or index numbers of text pairs that contain non-paraphrases]}</p> <p>Here are the text pairs:</p> <p>\$texts JSON output:</p>

Table 10: Prompt templates for MRPC. For the text-pair classification task (MRPC, SNLI, WiC), the sequence of text pairs in the place of ‘\$texts’ are indexed starting with ‘1’ and each text pair is separated by two newlines (each text pair ends with a newline be design, followed by another newline before the next text pair).

Task	Prompt template
SingleClf	<p>Given the following premise and hypothesis, determine the inference relation between them. Respond with ‘Entailment’ if the hypothesis logically follows from the premise, ‘Contradiction’ if they are in direct opposition, and ‘Neutral’ if neither applies.</p> <p>\$text Inference relation:</p>
BatchClf	<p>Given the following \$num pairs of premises and hypotheses, determine the inference relation for each pair line by line. Respond with ‘Entailment’ if the hypothesis entails the premise, and ‘Contradiction’ if they contradict. If neither is the case, respond with ‘Neutral.’ Provide your answers line by line.</p> <p>\$texts Inference relations for the \$num text pairs provided above:</p>
SelectOne	<p>Go over the \$num text pairs below and list the index numbers of the text pairs where the inference relation between the premise and the hypothesis is \$relationship according to the following instructions: If none of the text pairs contain \$relationship inference relation, write ‘None.’ If all text pairs contain \$relationship inference relation, write ‘All.’ Otherwise, provide the index numbers of the text pairs where the inference relation between the premise and the hypothesis is \$relationship.</p> <p>Output your responses in JSON format with the key ‘\$relationship’. A formatted example output is provided below. ‘\$relationship’: [None/All or index numbers of text pairs that contain \$relationship inference relation]</p> <p>Here are the text pairs:</p> <p>\$texts JSON output:</p>
SelectAll	<p>Go over the \$num text pairs below. First, list the index numbers of the text pairs that contain entailment inference relation. Then, select all text pairs that contain contradiction inference relation. Finally, select all text pairs that contain neutral inference relation. If none of the text pairs satisfy a condition, write ‘None.’ If all the text pairs belong satisfy a condition, write ‘All.’ Otherwise, provide the index numbers of the text pairs that satisfy each condition.</p> <p>Output your responses in JSON format with three keys: ‘entailment’, ‘contradiction’, and ‘neutral’. A formatted example output is provided below. {‘entailment’: [None/All or index numbers of text pairs that contain entailment inference relation], ‘contradiction’: [None/All or index numbers of text pairs that contain contradiction inference relation], ‘neutral’: [None/All or index numbers of text pairs that contain neutral inference relation]}</p> <p>Here are the text pairs:</p> <p>\$texts JSON output:</p>

Table 11: Prompt templates for SNLI.

Task	Prompt template
SingleClf	<p>Analyze the usage of the given target word in the two subsequent contexts. The target word may appear in various grammatical forms in each context. Respond with ‘Yes’ if it maintains the same meaning across both contexts, and ‘No’ if it does not.</p> <p>\$text Answer:</p>
BatchClf	<p>Analyze the usage of the following \$num target words in the two contexts that immediately follow them. These target words may appear in different grammatical forms across the two subsequent contexts. Determine if each target word maintains the same meaning in the two subsequent contexts. Provide your answers line by line, indicating ‘Yes’ if it does and ‘No’ if it does not.</p> <p>\$texts Answers:</p>
SelectOne	<p>Analyze the following \$num target words and determine the index numbers of the target words where the same meaning \$be maintained across the two contexts that immediately follow them. These target words may appear in different grammatical forms in each context. If none of the target words satisfy this condition, write ‘None.’. If all the target words satisfy this condition, write ‘All.’ Otherwise, provide the index numbers.</p> <p>Output your responses in JSON format with the key ‘answer’. A formatted example output is provided below. {‘answer’: [None/All or index numbers of the target words where the same meaning \$be maintained in the two subsequent contexts]}</p> <p>Here are the target words along with their contexts:</p> <p>\$texts JSON output:</p>
SelectAll	<p>Analyze the following \$num target words, which may appear in different grammatical forms in the two subsequent contexts. First, list the index numbers of target words that maintain the same meaning in the two subsequent contexts. Then, list the index numbers of target words that do not maintain the same meaning in the two subsequent contexts. If none of the target words satisfy a condition, write ‘None.’ If all the target words satisfy a condition, write ‘All.’ Otherwise, provide the index numbers of the target words that satisfy each condition.</p> <p>Output your responses in JSON format with two keys: ‘yes’ for target words used with consistent meanings and ‘no’ for those used with inconsistent meanings. A formatted example output is provided below. {‘yes’: [None/All or index numbers of target words used with consistent meanings], ‘no’: [None/All or index numbers of target words used with inconsistent meanings]}</p> <p>Here are the target words along with their contexts:</p> <p>\$texts JSON output:</p>

Table 12: Prompt templates for WiC.