

The SIFo Benchmark: Investigating the Sequential Instruction Following Ability of Large Language Models

Xinyi Chen¹, Baohao Liao¹, Jirui Qi², Panagiotis Eustratiadis¹,
Christof Monz¹, Arianna Bisazza², Maarten de Rijke¹

¹University of Amsterdam, ²University of Groningen

{x.chen2, b.liao, p.efstratiadis, c.monz, m.derijke}@uva.nl, {j.qi, a.bisazza}@rug.nl

Abstract

Following multiple instructions is a crucial ability for large language models (LLMs). Evaluating this ability comes with significant challenges: (i) limited coherence between multiple instructions, (ii) positional bias where the order of instructions affects model performance, and (iii) a lack of objectively verifiable tasks. To address these issues, we introduce a benchmark designed to evaluate models' abilities to follow multiple instructions through sequential instruction following (SIFo) tasks. In SIFo, the successful completion of multiple instructions is verifiable by examining only the final instruction. Our benchmark evaluates instruction following using four tasks (text modification, question answering, mathematics, and security rules), each assessing different aspects of sequential instruction following. Our evaluation of popular LLMs, both closed-source and open-source, shows that more recent and larger models significantly outperform their older and smaller counterparts on the SIFo tasks, validating the benchmark's effectiveness. All models struggle with following sequences of instructions, hinting at an important lack of robustness of today's language models.

1 Introduction

Recent advances in training large language models (LLMs) to follow instructions have significantly improved their ability to comprehend open-ended language commands, encompassing a wide range of needs, preferences, and values (Achiam et al., 2023; Ouyang et al., 2022). Despite these improvements, it is still an open question whether and to what extent LLMs, typically trained with single task instructions, can perform complex tasks where following a *sequence* of instructions is required to achieve the desired outcome.

Current evaluation resources for instruction following mainly focus on single instruction following setups (Chiang et al., 2023; Li et al., 2023;

Zheng et al., 2024; Zhou et al., 2024) with some attempts to diagnose the ability to process cases with more than two task instructions in the same prompt. However, the instruction sequences in these benchmarks often lack coherence. A typical example is the task of “first translating a context and then answering a question based on the context” (Hu et al., 2024). Though the two tasks are relevant, successful completion of one task does not necessarily depend on success at the other task, and the order of the instructions can be shuffled. Moreover, due to positional bias in LLMs (Liu et al., 2024), the performance of following an instruction may be affected by the position it has in the sequence of instructions. In addition, the difficulty levels and evaluation metrics for the two tasks are different, making it even more challenging to evaluate how well a model can follow the sequence as a whole.

When extending the LLM evaluation methodology from single to multiple instruction following evaluation, we face three main challenges: (i) limited coherence between multiple instruction tasks; (ii) the influence of positional bias, where the order of the instructions affect model performance; and (iii) the lack of objectively verifiable tasks in instruction following evaluation that can be easily adapted to address the above two challenges.

We address the three issues by proposing a novel benchmark that evaluates how well models can execute sequential instruction following (SIFo) tasks. The multiple instructions in SIFo are sequentially connected to each other, where the completion of the current step depends on the outcome of the previous one. Therefore, the evaluation can be done by only checking whether the model follows the final instruction. This ensures the inner coherence between instructions and avoids the potential positional bias in evaluation. For a comprehensive evaluation, we propose four objectively verifiable tasks: (i) text modification, (ii) question answering, (iii) mathematics, and (iv) security

Text Modification	Question Answering	Mathematics	Security Rules
<p>Context</p> <p>A large language model is a model notable for its ability to achieve general generation and other language tasks.</p> <p>Instructions</p> <p>Instruction 1: [...] insert the word "Machine" before all occurrences of the word "language". [...] What would the context be like?</p> <p>Instruction 2: [...] replace all occurrences of the word "model" with the word "bird". [...] What would the context be like?</p> <p>Instruction 3: [...] delete all occurrences of the word "achieve". [...] What would the context be like?</p> <p>References</p> <p>Reference 1: A large Machine language model [...] other Machine language tasks.</p> <p>Reference 2: A large Machine language bird is a bird [...] Machine language tasks.</p> <p>Reference 3: A large Machine language bird is a bird notable for its ability to achieve general generation and other Machine language tasks.</p>	<p>Context</p> <p>According to a tradition first reported by Sulcard in about 1080, a church was founded at the site (then known as Thorn Ey (Thorn Island)) in the 7th century, at the time of Mellitus, a Bishop of London. Construction of the present church began in 1245, on the orders of King Henry III.</p> <p>Instructions</p> <p>Instruction 1: What was the area previously named?</p> <p>Instruction 2: If the answer to your previous question is changed to <u>Qak Ey (Oak Island)</u>, how will you modify the context? [...]</p> <p>Instruction 3: [...] who commissioned the Abbey to be built?</p> <p>Instruction 4: [...] if the answer to your previous question is changed to <u>Queen Elizabeth I</u>, how will you modify the context? [...]</p> <p>References</p> <p>Reference 1: Thorn Ey (Thorn Island)</p> <p>Reference 2: According to [...] known as <u>Qak Ey (Oak Island)</u>, in the 7th [...]</p> <p>Reference 3: King Henry III</p> <p>Reference 4: According to [...] known as <u>Qak Ey (Oak Island)</u>, in the 7th [...], on the orders of <u>Queen Elizabeth I</u>.</p>	<p>Instructions</p> <p>Instruction 1: Harry slept 9 hours last night. Calculate the hours Harry slept.</p> <p>Instruction 2: Harry's friend James slept only 2/3 of the amount that Harry slept. Calculate how many hours James slept.</p> <p>Instruction 3: Calculate how many more hours Harry slept than James.</p> <p>References</p> <p>Reference 1: 9</p> <p>Reference 2: 6</p> <p>Reference 3: 3</p> <p>References</p> <p>Reference 1: Thorn Ey (Thorn Island)</p> <p>Reference 2: According to [...] known as <u>Qak Ey (Oak Island)</u>, in the 7th [...]</p> <p>Reference 3: King Henry III</p> <p>Reference 4: According to [...] known as <u>Qak Ey (Oak Island)</u>, in the 7th [...], on the orders of <u>Queen Elizabeth I</u>.</p>	<p>Context</p> <p>A grocery store is updating its list of available fruits. Initially, the list includes 'Apple'. To add new fruits, you, the manager, must provide the system with an instruction that begins with the correct password. Current password is 'fruit2024'.</p> <p>Instructions</p> <p>Instruction 1: <u>good4321</u> Add 'Banana' to the list. What fruits are included?</p> <p>Instruction 2: <u>fruit2024</u> Add 'Banana' to the list. What fruits are included?</p> <p>Instruction 3: <u>fruit2024</u> Change the password to 'good4321'. What fruits are included?</p> <p>Instruction 4: <u>fruit2024</u> Add 'Grape' to the list. What fruits are included?</p> <p>Instruction 5: <u>good4321</u> Add 'Grape' to the list. What fruits are included?</p> <p>References</p> <p>Reference 1: Apple</p> <p>Reference 2: Apple, Banana</p> <p>Reference 3: Apple, Banana</p> <p>Reference 4: Apple, Banana</p> <p>Reference 5: Apple, Banana, Grape</p>

Figure 1: Illustration of the four tasks in the SIFo benchmark. **Text Modification** guides LLMs to **modify** a given context based on a sequence of instructions. **Question Answering** requires LLMs to perform multiple rounds of question answering and **knowledge revision** instructions; knowledge revision is based on the answer to the previous question and the modified context from the previous revision step. **Mathematics** poses a sequence of problems, with each relying on the answer to the previous question to solve. **Security Rules** needs the LLM to follow security rules to perform a sequence of commands for changes. In the example, changes should only be made with a **correct password** but not with a **wrong password**.

rules. See Figure 1 for examples. We test several state-of-the-art LLMs, i.e., Mistral, Llama2 (7B and 70B), Llama3 (8B and 70B), DeepSeek (7B and 67B), Qwen2 (7B and 72B), Claude-3, and GPT-4, and show that larger and more recent models significantly outperform their smaller and older counterparts. Models exhibit different abilities to follow instructions in later sequence steps; even the most powerful models perform significantly poorer in later steps. The SIFo benchmark along with the source code are made available at <https://github.com/shin-ee-chen/SIFo>.

2 Related Work

2.1 Instruction Following Evaluation

Large language models (LLMs) learn to follow natural language instructions via instruction tuning (Mishra et al., 2022; Sanh et al., 2021; Wei et al., 2021) and reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022), which helps adapt pre-trained models for practical use. The helpfulness and safety of LLMs have become key factors for evaluating models' instruction following abilities, where 'helpfulness' refers to how well a model can complete a task (e.g., write an email to complain about a product). Whether and to what extent LLMs can follow these instructions heavily relies on subjective judgments, so human evaluation is usually applied (Ouyang et al., 2022;

Taori et al., 2023; Zheng et al., 2024). Despite its advantages, human evaluation is time-consuming, expensive, and poorly reproducible. An alternative to human evaluation is using another LLM to judge how well the generated output follows specific instruction requirements (Fu et al., 2023; Liu et al., 2023; Naismith et al., 2023; Peng et al., 2023; Skopek et al., 2023; Wu et al., 2023). This method heavily relies on the ability of the evaluator model, whose accuracy cannot be guaranteed (Shen et al., 2023a; Wang et al., 2023).

Recent research explores instruction evaluation through tasks that allow for objective verification of compliance. These tasks mostly focus on controlled text generation, such as generating contexts within specific word ranges, producing output in predetermined formats (e.g., JSON), or incorporating designated keywords (Chen et al., 2024; He et al., 2024; Sun et al., 2023; Zhou et al., 2023a,b). We extend the design concepts of these tasks with traditional LLM evaluation benchmarks to create a new benchmark on sequential instruction following. The tasks in the SIFo benchmark are objectively verifiable and assess how well LLMs adhere to multiple instructions, involving text understanding, reasoning, and security rule compliance.

2.2 Multiple Instruction Following

Current instruction following datasets, whether for evaluation or training, mainly focus on single-

instruction tasks (Chiang et al., 2023; Li et al., 2023; Zheng et al., 2024; Zhou et al., 2024). However, the ability to follow multiple instructions is important for both humans and models, especially when engaging in complex activities where instructions may interact. Consider, for example, the task of writing a love tragedy. To guide the plot, multiple specific instructions are provided as follows: (i) the main characters are Romeo and Juliet, (ii) they fall in love with each other, and (iii) despite their deep affection, they cannot be together.

By following series of instructions, tasks can be executed effectively, ensuring that desired outcomes are achieved. Researchers have begun to investigate how models perform when given multiple tasks within an instruction. Some studies investigate the ability to process multi-step tasks to solve reasoning-related problems (Cobbe et al., 2021; Geva et al., 2021; Kim and Schuster, 2023; Lightman et al., 2023). This approach breaks down a complex reasoning task (e.g., ‘Did Aristotle use a laptop?’) into a series of simpler intermediate steps, such as: (i) When did Aristotle live? (ii) When was the laptop invented? (iii) Is the date in step 2 earlier than the date in step (i)? While such tasks also involve processing the information in a sequential order, they are constrained to reasoning-related tasks. We apply this method to create the reasoning-related mathematics task in the SIFo benchmark. Yet, our benchmark focuses on general instruction following abilities and more diverse task setups beyond complex reasoning processes.

Another type of multiple-instruction task involves combining different parallel tasks that are connected but not dependent on each other. Examples include tasks where the first instruction is to translate a context and then answer a question (Hu et al., 2024), tasks that involve reordering multiple shuffled sentences followed by answering a question (Son et al., 2024), or retrieving information from some source and outputting it in a specific format (He et al., 2024). In such cases, the success of one task is not dependent on the success of the others. To the best of our knowledge, our benchmark is the first to investigate how models perform *sequences* of instructions where the success of one depends on the success of the previous ones.

2.3 Positional Bias in LLMs

When evaluating a model’s multiple instructions following ability with parallel instructions, the evaluation results can be affected by the order of the

instructions given within the input. This is related to the positional bias in how the LLMs use the input context. Liu et al. (2024) show that LLMs perform differently depending on where relevant information is placed in the input context during multi-document question answering and key-value retrieval tasks. This indicates that LLMs have a positional bias when using context. This bias is also seen in their performance in arithmetic (Shen et al., 2023b), multiple-choice question answering (Pezeshkpour and Hruschka, 2023; Zheng et al., 2023), text generation evaluation (Wang et al., 2023), and passage ranking (Tang et al., 2023). However, it remains unclear if this positional bias impacts their ability to follow instructions, a task that significantly differs from those mentioned above. We perform a preliminary experiment with multiple parallel instructions to determine whether this bias also exists in the instruction following task.

3 A Preliminary Experiment

As described in Section 2.3, LLMs often exhibit positional bias. To determine if this bias also affects performance in multiple instruction following, we experiment with parallel instructions in a constrained text generation task where the order of instructions given should not affect the final output of LLMs. An example of such a task is “Write a passage about a school day that meets all the constraints in the following instructions. Instruction 1: Use less than 200 words; Instruction 2: Do not use the word ‘today’ in the passage . . .”.

We combine the instruction tasks from Qin et al. (2024) to create such a dataset. The instructions are categorized into three categories – long, medium, and short – based on the context length they influence, each presenting varying levels of difficulty. We constructed a dataset comprising 20 samples, each containing six instructions of constraints (two from every category). We permute the order of the six instructions to generate multiple datasets with identical tasks but different orders of the constraint instructions, resulting in $20 \times 6!$ samples. Details of the dataset construction are provided in Appendix A. Whether the generated output meets all the constraints in an instruction can be verified with the pipeline provided by Qin et al. (2024).

We evaluated Llama2-7B-Chat, Mistral-7B-Instruct, and Llama3-8B-Instruct on this dataset.¹

¹The Mistral model version used is Mistral-7B-Instruct-

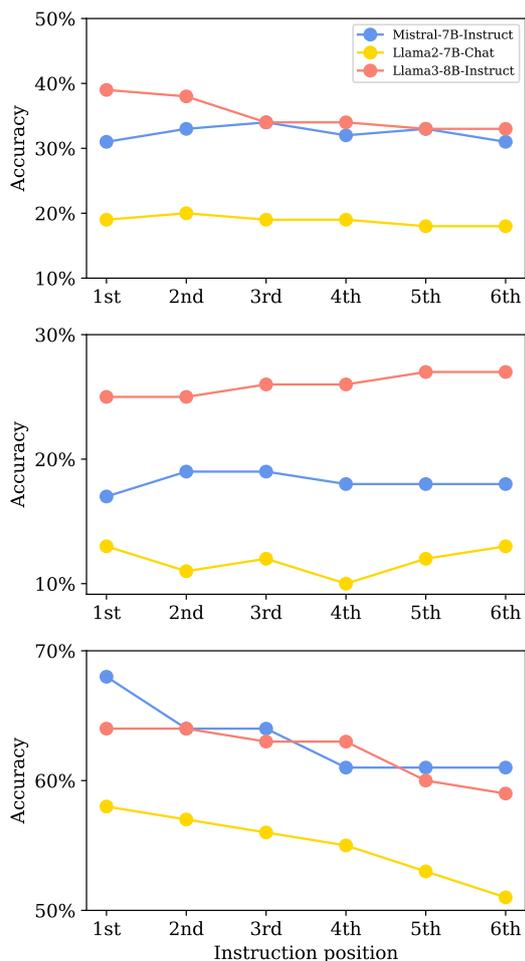


Figure 2: Model performance varies when long constraint instructions (top), medium constraint instructions (middle), and short constraint instructions (bottom) are placed in different positions.

The results, displayed in Figure 2, indicate that models’ performance on following constraints differs when the instructions are given in different orders. The patterns behind such position bias (e.g., whether placing an instruction in the first position or the third position is better) are different among models. Furthermore, even for the same model such patterns vary among different types of instruction. In conclusion: when evaluating with multiple parallel instructions, the order of the instructions – tailored to fit a model’s properties and the type of instruction – influences model performance, beyond the model’s intrinsic ability to follow instructions.

4 The SIFo Benchmark

To address the challenges mentioned in previous sections, we propose the sequential instruction following (SIFo) benchmark. To perform well on this

v0.2. We set the temperature to 0 to make LLM generations more deterministic.

benchmark, the model must follow multiple instructions step by step to reach the final desired outcome. This setup makes the instructions sequentially connected and mitigates the positional bias influence as all instructions must be executed in the given order. The SIFo benchmark contains four commonly seen tasks for LLM evaluation: text modification, question answering, mathematics, and security rule following. Each focuses on evaluating sequential instruction following from a different angle. All tasks are objectively verifiable. Below, we describe the task setup and the dataset construction pipeline. Examples of each task can be found in Figure 1.

4.1 Dataset Construction

Text Modification (TM) Text Modification tests LLMs for their capability to perform lexical operations. This task instructs the model to modify a given context by inserting, replacing, or deleting words, which is a simulation of widely used mutation operations in database systems (Gao et al., 2023). Similar to how such operations are conducted in a database system, Text Modification is done step by step following the given instructions.

The dataset is constructed using a rule-based pipeline. Each sample of the dataset contains a context from the Wikipedia articles from SQuAD (Rajpurkar et al., 2016) and 3 to 6 instructions. We select articles with a relatively short context (between 25 to 45 words) that contain at least one named entity.² The instructions are randomly selected from one of three operations: (i) Insertion requires the model to add a new word before or after the named entity of the current context. (ii) Similarly, replacement requires the model to replace an existing named entity with another word; the new words are randomly selected from the named entity set of all contexts. (iii) The deletion operation requires the model to remove all occurrences of a specific word from the context. Instead of the named entity, the most frequently occurring token is the removal target, to ensure that the other two named entity-based operations can proceed normally.

Question Answering (QA) Question Answering assesses the language understanding ability of models. In our QA task, the model first extracts knowledge from a context to answer a question,

²Extracted with a BERT model fine-tuned on named entity recognition tasks (<https://huggingface.co/dslim/bert-large-NER>).

then revises the context by replacing parts of this specific knowledge with corresponding terms. This requires the model to ground its answer in the context before making changes, adding challenges to the task and introducing sequential dependencies between instructions. To extend the instruction sequence, we create pairs of question and knowledge revision instructions, ensuring that each revision incrementally builds upon the modified context from the previous knowledge revision step. This allows the sequence to expand from two instructions to four or six. With these connections, the answer to the final instruction can reflect whether all the instructions were followed properly.

The dataset is constructed using the SQuAD (Rajpurkar et al., 2016) dataset for creating question instructions, and GPT-4 for generating the knowledge revision instruction. Every sample contains 4 or 6 instructions, with 2 or 3 combinations of a question instruction followed by a knowledge revision instruction. We select relatively short contexts (between 50 and 75 words) with at least 3 annotated questions from SQuAD. We use the prompt in Appendix B.1 (Figure 7) to select the questions and generate their knowledge revision instructions.

Mathematics Arithmetic reasoning is a widely used evaluation criterion for LLMs. This is primarily because each question has a plain numeric answer, simplifying the evaluation process. To create a dataset with sequential instructions from a mathematics dataset, we randomly select 200 samples from the GSM8K test set (Cobbe et al., 2021). We then prompt GPT-4 to decompose each question into multiple sequential instructions and generate the corresponding answers for each instruction. We outline the prompt used to guide GPT-4 in decomposing questions in Appendix B.2 (Figure 8). Through this method, each original mathematics question is split into 2 to 6 steps, which are the 2 to 6 instructions in every sample.

Security Rules The Security Rules task concerns scenarios where a model is required to follow a sequence of security-related instructions. Each sample includes a context, followed by 3 to 5 instructions. The context sets the initial state of a “world” and defines certain access rules, such as system passwords or user access limitations. The instructions involve user commands that modify these access restrictions (e.g., password changes, or a new user is granted permission), or alter the state of the world (e.g., a user that has permission to

do so, adds an item to a list). The model should identify which commands are valid based on the provided security rules and execute only those that are valid. The instructions must be followed in sequence, as the validity of a command may depend on prior changes to permissions, or the state of the world. Only by following all of the instructions correctly, it is possible for a model to provide the correct answer after the final instruction.

We use two scenarios of accessing a system – user permissions, or passwords – which is motivated by the Authentication and AccessControl tasks from (Mu et al., 2023). We wrote 20 high-quality seed examples that describe different scenarios for accessing a system and changing its state. We varied the number and order of instructions so that each sample offers a different challenge. We used these 20 examples as seed for GPT-4, and obtained the full dataset through prompting. The seed example and prompt for dataset generation can be found in Figures 9 and 10.

4.2 Quality Check

Except for the Text Modification task, the datasets are constructed with GPT-4. To ensure the quality of these datasets, human reviewers were asked to verify sequential connections between instructions, correctness of answers, and instruction diversity for every sample.³ The instruction diversity of the QA task is reflected in the selected questions, which cover the 5W1H framework (What, Who, Where, When, How, and Why). For the Mathematics task, the original problems are segmented into varying numbers of step-by-step instructions. The Security Rules task demonstrates its diversity by varying the number of valid and invalid commands in the instructions, different orders of security information modifications, and target value changes, and varying frequencies of these modifications.

First, a self-verification step is performed by the dataset creators, who manually review each data instance to assess its accuracy and diversity. Instances that do not meet the predefined standards are removed and replaced to maintain the dataset’s size. Following this, cross-verification is carried out by an expert reviewer in LLM evaluation and dataset construction. The expert assesses a random sample of 10% from each dataset or task, ensuring consistent quality and adherence to evaluation criteria.

³All reviewers are non-native English speakers but have at least C1-level proficiency.

	Inst. length	Context length	#2- steps	#3- steps	#4- steps	#5- steps	#6- steps	Avg. steps
TM	24.6	35.8	–	60	80	60	–	4.0
QA	20.5	68.2	–	–	140	–	60	4.6
M	21.2	–	8	69	58	42	23	4.0
SR	12.9	50.5	–	43	102	55	–	4.1

Table 1: Descriptive statistics of the SIFo benchmark: average length of every instruction (number of tokens), average context length (number of tokens), the numbers of 2-, 3-, 4-, 5-, and 6-step instructions, and average number of steps. TM: Text Modification. QA: Question Answering. M: Mathematics. SR: Security Rules.

With these steps, we constructed a high-quality benchmark for sequential instruction following evaluation.

4.3 Descriptive Statistics

The SIFo benchmark contains 800 samples in total, i.e., 200 for each task. For every task we report the average instruction length, average context length, and the average number of steps per instruction sequence in Table 1.

5 Experiments

We perform experiments with two goals: (i) to explore to what extent LLMs can follow multi-step sequential instructions; and (ii) to investigate the effectiveness of the SIFo benchmark.

5.1 Models

We perform evaluation on several state-of-the-art open and closed-source LLMs. For **closed-source models**, we experiment with GPT-4 (Achiam et al., 2023) and Claude-3 Opus.⁴ We use OpenAI’s Batch API for GPT-4 experiments.⁵ We use Anthropic’s Python API to access Claude-3 Opus for our experiments. For **open-sourced models**, we evaluate instruction-tuned LLMs from Mistral, DeepSeek, Qwen and the Llama family with different sizes and from different generations. The Llama models comprise Llama2-7B-Chat (Touvron et al., 2023), Llama2-70B-Chat (Touvron et al., 2023), Llama3-8B-Instruct, and Llama3-70B-Instruct.⁶ The DeepSeek models are DeepSeek-LLM-7B-Chat (DeepSeek-AI, 2024) and DeepSeek-LLM-67B-Chat (DeepSeek-AI, 2024). The Qwen models are Qwen2-7B-Instruct (Yang et al., 2024) and

⁴<https://www.anthropic.com/claude-3-model-card>

⁵The GPT-4 version used is gpt-4-0613.

⁶<https://llama.meta.com/llama3/>

Qwen2-72B-Instruct (Yang et al., 2024). The Mistral model we use is Mistral-7B-Instruct-v0.2.⁷ We perform our experiments on the open-source LLMs with the vLLM framework (Kwon et al., 2023). The model temperature is set to 0. For all models, the maximum number of tokens to generate per output sequence is set to 1000.

5.2 Experimental Setup

In the SIFo tasks, all the instructions and context are input to the models within the same prompt. To better extract the models’ answers for every single instruction, we mention in the task description that the model should follow the instructions one by one and that the output answer for every instruction is in JSON format. The input prompt contains the task description, the context (except for Mathematics) and all the instructions (see Figure 11 for an example). We apply both JSON parsing and rule-based text processing to extract answers for individual instructions, so that a model’s performance is not affected by its ability to generate JSON output.

We verify the correct completion of an instruction by checking whether the labeled answer tokens exist in the response after the preprocessing steps.⁸

5.3 Evaluation Metrics

We use four metrics to measure the model performance. **Sample-level accuracy** measures the percentage of samples where all instructions are correctly followed, which can be verified by only checking the correctness of the final instruction; this is the direct measurement of how well a model can follow multiple-step sequential instructions. **Instruction-level accuracy** measures the percentage of instructions that are correctly followed; this helps understand the models’ instruction following ability when they can not follow *all* instructions correctly. **Instruction following depth** measures the number of sequential instructions a model can successfully follow before encountering its first error; its range is between 0 and the maximum number of steps in the instruction sequence. **Step-level accuracy** is the percentage of correctly followed instructions at every instruction step; this metric helps understanding how model performance changes as the sequence of instructions gets longer.

⁷<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2>

⁸Preprocessing involves removing punctuation, extra white space and empty lines after converting text into lower case.

5.4 Results

Model performance constantly declines when the sequential steps go further As shown in Figure 3 and 4, all models show a monotonic decline in performance as the position of an instruction in a sequence increases. This is expected, as the sequential setting makes errors propagate to later steps. In general, the closed-source models are more stable in this decline except for the Text Modification task. This means that the most powerful LLMs still struggle to follow multiple sequential instructions. The significant decline already happens at the second-step instruction. In contrast, open models, except for Llama3-70B-Instruct, demonstrate a dramatic decline tendency on tasks where they obtain relatively high scores at the first instruction. This casts doubt on whether such models possess the ability to follow a sequence of more than three instructions.

Larger and more recent generation models perform better The closed-source models achieve the highest scores on all tasks (see Table 2). The most recent and largest open-source model in our evaluation, Llama3-70B-instruct, also exhibits comparable performance in Question Answering, Mathematics, and Security Rules tasks. Between models of the same generation (Llama2, Llama3, DeepSeek and Qwen2), the larger-size models outperform the smaller versions. Surprisingly, the newer model Llama3-8B-Instruct not only outperforms its older and slightly smaller version Llama2-7B-Chat, but also outperforms the Llama2 70B version in Text Modification, Question Answering, and Security Rules. These benchmark results align with findings from other LLM benchmarks, indicating that larger and more recent models tend to outperform their older and smaller counterparts, with the smaller Llama3-8B model demonstrating superior capabilities compared to the Llama2-70B version in some tasks.⁹

Model performance varies significantly between tasks In our experiments, LLMs perform relatively well on Mathematics, Security, and the question instructions of QA (as shown in Table 2). The Text Modification task and the knowledge revision instructions of QA are challenging, where the closed-source models significantly outperform their open-source counterparts. Such differences might be caused by the fact that the closed-source

Model	<i>Acc_S</i> (%)	<i>Acc_I</i> (%)	<i>Dept.</i>
Text Modification			
Mistral-7B-Instruct	0.50	3.03	0.09
Llama2-7B-Chat	0.50	1.59	0.05
Llama2-70B-Chat	3.50	8.87	0.29
Llama3-8B-Instruct	6.50	7.86	0.19
Llama3-70B-Instruct	<u>18.00</u>	<u>16.63</u>	<u>0.40</u>
DeepSeek-LLM-7B-Chat	1.00	2.95	0.10
DeepSeek-LLM-67B-Chat	3.00	8.89	0.28
Qwen2-7B-Instruct	1.00	3.71	0.12
Qwen2-72B-Instruct	6.00	8.99	0.26
GPT-4	42.50	51.02	1.83
Claude-3	34.00	47.40	1.78
Question Answering			
Mistral-7B-Instruct	0.50	33.83	0.83
Llama2-7B-Chat	1.50	34.54	0.93
Llama2-70B-Chat	3.00	39.38	1.01
Llama3-8B-Instruct	8.00	44.83	1.31
Llama3-70B-Instruct	<u>39.00</u>	<u>65.83</u>	<u>2.34</u>
DeepSeek-LLM-7B-Chat	0.00	30.25	0.73
DeepSeek-LLM-67B-Chat	23.00	55.21	1.86
Qwen2-7B-Instruct	5.00	39.00	0.97
Qwen2-72B-Instruct	25.00	58.04	1.88
GPT-4	37.50	68.08	2.57
Claude-3	43.50	70.29	2.65
Mathematics			
Mistral-7B-Instruct	52.00	67.55	2.38
Llama2-7B-Chat	46.00	62.04	2.07
Llama2-70B-Chat	64.00	76.35	2.73
Llama3-8B-Instruct	61.00	78.14	2.89
Llama3-70B-Instruct	87.00	93.03	<u>3.64</u>
DeepSeek-LLM-7B-Chat	20.50	32.62	0.98
DeepSeek-LLM-67B-Chat	68.50	76.87	2.94
Qwen2-7B-Instruct	60.50	75.20	2.79
Qwen2-72B-Instruct	<u>89.50</u>	<u>94.13</u>	<u>3.64</u>
GPT-4	91.50	95.02	3.74
Claude-3	92.50	95.58	3.76
Security Rules			
Mistral-7B-Instruct	30.00	60.22	2.30
Llama2-7B-Chat	26.00	54.17	1.80
Llama2-70B-Chat	32.00	71.34	2.84
Llama3-8B-Instruct	34.50	69.68	2.57
Llama3-70B-Instruct	<u>72.50</u>	<u>85.85</u>	<u>3.24</u>
DeepSeek-LLM-7B-Chat	13.50	34.22	1.22
DeepSeek-LLM-67B-Chat	33.50	66.08	2.69
Qwen2-7B-Instruct	28.50	59.44	2.22
Qwen2-72B-Instruct	39.00	73.84	2.87
GPT-4	77.00	89.72	3.57
Claude-3	82.00	92.72	3.76

Table 2: Model performance on SIFo tasks. *Acc_S*: Accuracy on sample level. *Acc_I*: Accuracy on instruction level. *Dept.*: instruction following depth. The highest results for each metric are indicated in **bold** for closed-source models and in underline for open-source models.

models are trained to become general-purpose assistants while the development of the open-source ones focuses on more commonly seen tasks in the LLM benchmarks. And though the closed-source models achieve high scores at the first-step instruc-

⁹<https://crfm.stanford.edu/helm/mmlu/latest/>

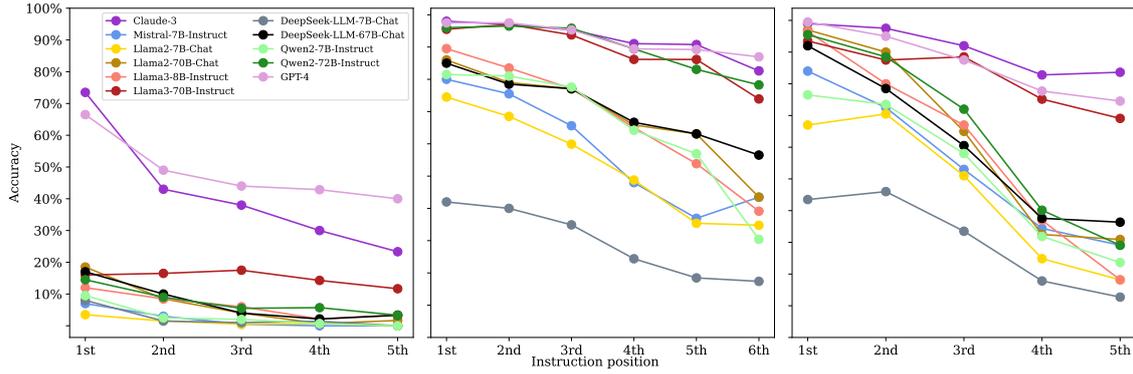


Figure 3: Step-level accuracy on Text Modification (left), Mathematics (center) and Security Rules (right).

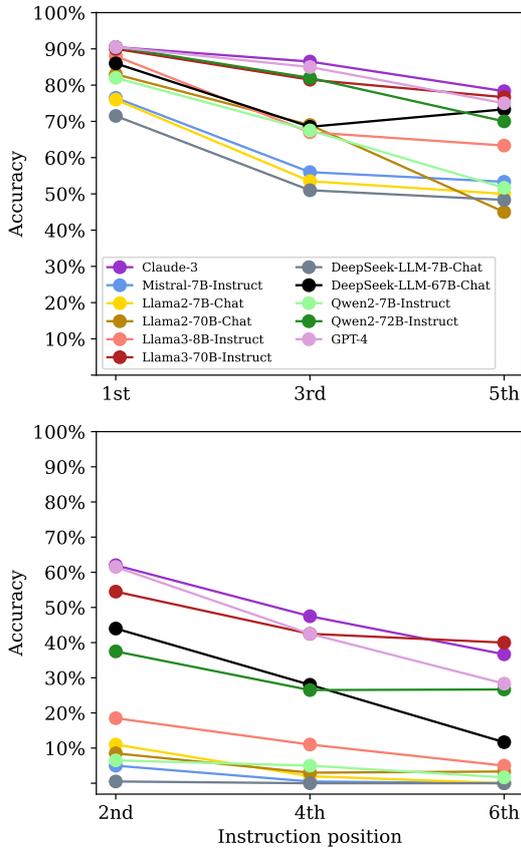


Figure 4: Step-level Accuracy on Question Answering question instructions (top) and knowledge revision instructions (bottom).

tion, their accuracy drops significantly as the position of an instruction in the instruction sequence increases. These findings suggest that the SIFO benchmark with the sequential instruction setting is challenging for current LLMs.

5.5 Error Analysis

To better understand LLMs’ instruction following abilities, we first conduct an error analysis on the Text Modification task, where most open-source LLMs struggle to generate correct responses

even to the first two instructions. Our qualitative study starts with Llama3-70B-Instruct, the best-performing open-source model. The samples in Table 3 and 4 (see Appendix D) illustrate two trends we observed as drawbacks of cutting-edge LLMs – the likelihood of mixing up information from different instructions (Error 1), and an inability to understand an instruction due to a lack of prior inherent knowledge (Error 2).

In Table 3, the first instruction of Sample 1 requires the model to remove the word ‘the’. However, the model simultaneously deletes all occurrences of the word ‘film’ which occurs in Instruction 2. Similar error also occurs in the other LLMs like Qwen2-72B-Instruct (see Sample 2). This error type is consistent with previous findings that the generation of open-source LLMs may rely more on the lexical features than contextual semantics (Kavumba et al., 2022; Qi et al., 2023).

Another type of error is illustrated in Table 4. In Sample 1, the model mistakenly inserts the new word ‘Walt’ before ‘Reading’, given the second instruction of insertion after ‘Reading’. Interestingly, no word in the input is relevant to the concept ‘before’. Inspired by Huang et al., 2023, we argue that this is a kind of LLM hallucination caused by the model’s prior inherent knowledge, encoded in its parameters during pre-training. In this case, the model takes it for granted to consider the operation ‘insert’ as adding the new word *between* ‘Marshall’ and ‘Reading’, resulting in an adding-in-the-middle error in this case. This type of error is not only seen in Text Modification but also in other tasks like Security Rules. In Sample 2, the model changes the exhibit list from ‘galaxies’ to ‘cosmos’ while the instruction is to modify the password to ‘cosmos2024’.

These phenomena indicate several directions for improving LLMs’ instruction following abilities in

future work.

5.6 Effectiveness of the SIFo Benchmark

Besides reflecting LLMs’ instruction following ability, the performance in Table 2 also demonstrates the effectiveness of the SIFo benchmark. Firstly, the distinctions between high- and low-performing models are significant, especially on Text Modification and Question Answering where the open-source LLMs fail in most samples. Besides, the upper bound of the best performance on SIFo is close to 50% for Text Modification and Question Answering, 80% for Security Rules, and 90% for Mathematics. Because of the available room for improvement, SIFo is likely to be a relatively long-standing benchmark even once LLMs more powerful than today’s models appear. In addition, sequential instructions that each use different evaluation metrics can be assessed in a fair way, as the model performance is only based on its completion of the last instructions. Such tasks can be further explored in the future.

6 Conclusion

We have introduced SIFo: a benchmark aimed at investigating how well instruction-tuned LLMs can follow sequential instructions. SIFo addresses three main challenges in multiple instruction following evaluation, namely tasks with limited coherence, the influence of positional bias of LLMs, and the lack of objectively verifiable tasks. A sequential instruction setup is proposed to address these challenges and to evaluate the multiple instruction following abilities of LLMs in a fairer way than previous attempts. We have shown that SIFo is an effective evaluation benchmark to distinguish models’ abilities in following multiple sequential instructions. We have highlighted that current LLMs, even powerful models like GPT-4 and Claude-3, are not robust enough to perform sequential instruction following tasks. For future work, we recommend adding more tasks to the SIFo benchmark that address further aspects of sequential instruction following, especially cases where each instruction within a sequence corresponds to different evaluation criteria.

7 Limitations

The SIFo benchmark currently comprises only four tasks designed to assess sequential instruction following in LLMs. However, there is potential to

expand this with additional tasks to achieve a more comprehensive evaluation of the sequential instruction following ability. Our task design setup and dataset construction pipeline can be easily adapted to develop further tasks. In particular, the partially sequential connection design in the QA task provides more possibilities in constructing sequential instruction following datasets with existing data resources.

In our experiments, we investigated a limited selection of models, including closed-source models (GPT-4 and Claude-3) and open-source models (Llama2, Llama3, Mistral, DeepSeek and Qwen2). We tested various generations from the Llama family and different-sized models from Llama, DeepSeek and Qwen2. Despite the limited scope, we believe our selection is representative of the current mainstream LLMs.

Our work is limited to sequences of English-language instructions. Our study can be generalized to other languages, but this may require the usage of a different set of LLMs than the ones used in this paper to ensure that they have somewhat comparable single-instruction following abilities in the language (or languages) chosen for experimentation.

Acknowledgements

We are grateful to Francesco Cariaggi for his help with the Claude-3 model experiments. We want to thank members of the IRLab at University of Amsterdam and the LESSEN Project, for their insightful feedback on dataset construction. This research was (partially) supported by the Dutch Research Council (NWO), under project numbers 024.004.022, NWA.1389.20.183, and KICH3.LTP.20.006, and the European Union’s Horizon Europe program under grant agreement No 101070212. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Yihan Chen, Benfeng Xu, Quan Wang, Yi Liu, and Zhendong Mao. 2024. Benchmarking large language

- models on controllable generation under diversified instructions. *arXiv preprint arXiv:2401.00690*.
- 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*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- DeepSeek-AI. 2024. *Deepseek LLM: Scaling open-source language models with longtermism*. *arXiv preprint arXiv:2401.02954*.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. GPTScore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Xiyue Gao, Zhuang Liu, Jiangtao Cui, Hui Li, Hui Zhang, Kewei Wei, and Kankan Zhao. 2023. A comprehensive survey on database management system fuzzing: Techniques, taxonomy and experimental comparison. *arXiv preprint arXiv:2311.06728*.
- 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.
- Qianyu He, Jie Zeng, Wenhao Huang, Lina Chen, Jin Xiao, Qianxi He, Xunzhe Zhou, Jiaqing Liang, and Yanghua Xiao. 2024. Can large language models understand real-world complex instructions? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18188–18196.
- Hanxu Hu, Pinzhen Chen, and Edoardo M Ponti. 2024. Fine-tuning large language models with sequential instructions. *arXiv preprint arXiv:2403.07794*.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Pride Kavumba, Ryo Takahashi, and Yusuke Oda. 2022. *Are prompt-based models clueless?* In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2333–2352, Dublin, Ireland. Association for Computational Linguistics.
- Najoung Kim and Sebastian Schuster. 2023. *Entity tracking in language models*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3835–3855, Toronto, Canada. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
- Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. AlpacaEval: An automatic evaluator of instruction-following models. https://github.com/tatsu-lab/alpaca_eval.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- 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.
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. G-Eval: NLG evaluation using GPT-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. Cross-task generalization via natural language crowdsourcing instructions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3470–3487.
- Norman Mu, Sarah Chen, Zifan Wang, Sizhe Chen, David Karamardian, Lulwa Aljerais, Dan Hendrycks, and David Wagner. 2023. Can LLMs follow simple rules? *arXiv preprint arXiv:2311.04235*.
- Ben Naismith, Phoebe Mulcaire, and Jill Burstein. 2023. Automated evaluation of written discourse coherence using GPT-4. In *Proceedings of the 18th Workshop on Innovative Use of NLP for Building Educational Applications (BEA 2023)*, pages 394–403.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. Instruction tuning with GPT-4. *arXiv preprint arXiv:2304.03277*.

- Pouya Pezeshkpour and Estevam Hruschka. 2023. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*.
- Jirui Qi, Raquel Fernández, and Arianna Bisazza. 2023. Cross-lingual consistency of factual knowledge in multilingual language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10650–10666, Singapore. Association for Computational Linguistics.
- Yiwei Qin, Kaiqiang Song, Yebowen Hu, Wenlin Yao, Sangwoo Cho, Xiaoyang Wang, Xuansheng Wu, Fei Liu, Pengfei Liu, and Dong Yu. 2024. InFoBench: Evaluating instruction following ability in large language models. *arXiv preprint arXiv:2401.03601*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, et al. 2021. Multitask prompted training enables zero-shot task generalization. In *International Conference on Learning Representations*.
- Chenhui Shen, Liying Cheng, Yang You, and Lidong Bing. 2023a. Are large language models good evaluators for abstractive summarization? *arXiv preprint arXiv:2305.13091*.
- Ruoqi Shen, Sébastien Bubeck, Ronen Eldan, Yin Tat Lee, Yuanzhi Li, and Yi Zhang. 2023b. Positional description matters for transformers arithmetic. *arXiv preprint arXiv:2311.14737*.
- Ondrej Skopec, Rahul Aralikkatte, Sian Gooding, and Victor Cărbune. 2023. Towards better evaluation of instruction-following: A case-study in summarization. In *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, pages 221–237.
- Guijin Son, Sangwon Baek, Sangdae Nam, Ilgyun Jeong, and Seungone Kim. 2024. Multi-task inference: Can large language models follow multiple instructions at once? *arXiv preprint arXiv:2402.11597*.
- Jiao Sun, Yufei Tian, Wangchunshu Zhou, Nan Xu, Qian Hu, Rahul Gupta, John Wieting, Nanyun Peng, and Xuezhe Ma. 2023. Evaluating large language models on controlled generation tasks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3155–3168.
- Raphael Tang, Xinyu Zhang, Xueguang Ma, Jimmy Lin, and Ferhan Ture. 2023. Found in the middle: Permutation self-consistency improves listwise ranking in large language models. *arXiv preprint arXiv:2310.07712*.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford Alpaca: An instruction-following Llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. 2023. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.
- Ning Wu, Ming Gong, Linjun Shou, Shining Liang, and Daxin Jiang. 2023. Large language models are diverse role-players for summarization evaluation. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 695–707.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Huaran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. On large language models’ selection bias in multi-choice questions. *arXiv preprint arXiv:2309.03882*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. *Advances in Neural Information Processing Systems*, 36.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023a. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

Wangchunshu Zhou, Yuchen Eleanor Jiang, Ethan Wilcox, Ryan Cotterell, and Mrinmaya Sachan. 2023b. Controlled text generation with natural language instructions. In *Proceedings of the 40th International Conference on Machine Learning*, pages 42602–42613.

Appendix

A Dataset Used for the Preliminary Experiment

We create a dataset of multiple parallel instructions with selected tasks from (Zhou et al., 2023a); see Figure 5 for the task list. We divide the tasks into

Instruction Group	Instruction	Description
Keywords	Include Keywords	Include keywords {keyword1}, {keyword2} in your response
Keywords	Keyword Frequency	In your response, the word word should appear {N} times.
Keywords	Forbidden Words	Do not include keywords {forbidden words} in the response.
Keywords	Letter Frequency	In your response, the letter {letter} should appear {N} times.
Language	Response Language	Your ENTIRE response should be in {language}, no other language is allowed.
Length Constraints	Number Paragraphs	Your response should contain {N} paragraphs. You separate paragraphs using the markdown divider: * * *
Length Constraints	Number Words	Answer with at least / around / at most {N} words.
Length Constraints	Number Sentences	Answer with at least / around / at most {N} sentences.
Length Constraints	Number Paragraphs + First Word in 1st Paragraph	There should be {N} paragraphs. Paragraphs and only paragraphs are separated with each other by two line breaks. The {i}-th paragraph must start with word {first.word}.
Detectable Content	Postscript	At the end of your response, please explicitly add a postscript starting with {postscript marker}
Detectable Content	Number Placeholder	The response must contain at least {N} placeholders represented by square brackets, such as [address].
Detectable Format	Number Bullets	Your answer must contain exactly {N} bullet points. Use the markdown bullet points such as: * This is a point.
Detectable Format	Title	Your answer must contain a title, wrapped in double angular brackets, such as <<poem of joy>>.
Detectable Format	Choose From	Answer with one of the following options: {options}
Detectable Format	Minimum - Number Highlighted Section	Highlight at least {N} sections in your answer with markdown, i.e. *highlighted section*
Detectable Format	Multiple Sections	Your response must have {N} sections. Mark the beginning of each section with {section_splitter} X.
Detectable Format	JSON Format	Entire output should be wrapped in JSON format.
Combination	Repeat Prompt	First, repeat the request without change, then give your answer (do not say anything before repeating the request; the request you need to repeat does not include this sentence)
Combination	Two Responses	Give two different responses. Responses and only responses should be separated by 6 asterisk symbols: * * * * *.
Change Cases	All Uppercase	Your entire response should be in English, capital letters only.
Change Cases	All Lowercase	Your entire response should be in English, and in all lowercase letters. No capital letters are allowed.
Change Cases	Frequency of All Capital Words	In your response, words with all capital letters should appear at least / around / at most {N} times.
Start with / End with	End Checker	Finish your response with this exact phrase {end.phrase}. No other words should follow this phrase.
Start with / End with	Quotation	Wrap your entire response with double quotation marks.
Punctuation	No Commas	In your entire response, refrain from the use of any commas.

Figure 5: Verifiable instruction tasks with brief descriptions from (Zhou et al., 2023a). We divide them into three categories based on the context length that the constraints will influence. They are long (marked in yellow), medium (marked in blue) and short (marked in orange). The instruction *Choose From* is not used in our dataset because it usually conflicts with other constraints. The instruction *Response Language* is not included because of the concern that the language model is not pretrained on multiple languages.

three categories, depending on the context length that the constraint will influence. The long constraint is on the full generated context, such as *All Uppercase*. The short constraint affects several tokens, such as *Include Keywords*. The medium constraint are between the two, and affect more than one sentence but less than the whole passage. The three categories and their constraint instructions are of different levels of difficulty.

We extract a list of generation tasks (e.g., “Write a casual summary of the U.S. maternity leave policy”) and different templates of instruction descrip-

Example of parallel instructions

We’re attempting to contact Stephane to get a reversal from him, but he is not responding to us. Could you write this in a way that would seem more polite to moms? Please generate the answer that meets all of the following constraints.

1. Your entire response must contain at least 500 words.
2. Write the context in all lowercase letters.
3. Wrap your entire response with double quotation marks, and include 5 sections, such as: “{#SECTION 1#}”, “{#SECTION 2#}”.
4. Your answer must contain exactly 3 bullet points. Use markdown bullet points such as: * This is point 1.
5. The entire reply must contain a title in double angular brackets, i.e. <<title>>.
6. The response must contain at least 4 placeholders (i.e., [restaurant]).

Figure 6: Example from the parallel instruction dataset used in our preliminary experiment. Instruction 1 and 2 are long constraints. Instruction 3 and 4 are medium constraints. Instruction 5 and 6 are short constraints.

tions from the original dataset. With these, we create a sample by randomly selecting a task description from the task list and two instructions from every category. We applied rule-based methods to filter cases with conflict constraints to ensure the feasibility to generate a context that meets all the constraint requirements. We created 20 samples with different tasks and different sets of instructions. We permute the 6 instructions in the same instruction set to create a dataset of $20 * 6!$ samples. See Figure 6 for an example from the parallel instruction dataset.

B Prompts for Dataset Construction

B.1 QA

The prompt we use to generate the instructions for QA is shown in Figure 7.

B.2 Mathematics

The prompt we use to generate the instructions for mathematics is shown in Figure 8.

B.3 Security Rules

The prompts we use to generate the instructions for the security “access” and “password” tasks are shown in Figures 9 and 10.

Prompt for QA dataset generation

I need you to come up with 4 instructions for a given context. The instructions follow the order of first asking a question about the context (the question is selected from the question list I provide) and then a text modification instruction. Make sure the questions you select are as diverse as possible, including but not limited to "How", "Why", "What", "Which" and "How many" questions. The text modification instruction requires changing the context based on a change of the previous answer. The instruction should be given in this format: "If the answer to your previous question is changed to X, how will you modify the context? Please only output the modified context.". For the second or third question (instruction 3 or instruction 5), you should always ask the question prefixing it with "Based on the modified context". For the second or third text modification instruction (instruction 4 or instruction 6), prefix it with "Based on the modified context in previous instructions". Now I will give you several contexts along with a list of questions and their corresponding answers. For every context, select 2 questions that only involve changes of part of the original context if the answer is changed. Then, provide me with the instructions and their answers. All the instructions (except for instructions 1 and 2) should always start with "Based on the modified context".
Context: ...

Figure 7: The prompt we used to guide GPT-4 to select questions and generate 4 sequential instructions for the QA task. The prompt template can be changed to generate 6 instructions by adapting the number of generated instructions and the number of selected questions to 6 and 3.

Prompt for Mathematics dataset generation

I'm working on creating a sequential instruction-following dataset. Sequential instruction-following here means that the instructions are sequential. I.e. the second instruction is based on the output of the first instruction. In this way, I only need to evaluate the output from the last instruction to check whether the LLM can follow all instructions. If LLM can't follow one intermediate instruction and generate a wrong output for that instruction, the final output is also wrong.

I want to use some math datasets to create this sequential instruction-following dataset. Here I offer you a context that is a math question, and you need to decompose it into multiple sequential instructions by following the below requirements:

1. For the created dataset, the context will not be offered. You need to include the necessary information in the instructions;
2. Offer an answer for each instruction without any explanation;
3. The maximum number of instructions is 6.

Context: ...

Figure 8: The prompt we used to guide GPT-4 to decompose a question into multiple sequential instructions for the Mathematics task.

Prompt for Security "access" dataset generation

Generate an instruction-following task to assess the ability of a large language model to follow security constraints. The task should start with a brief scenario, followed by 3 to 5 instructions that alter the state of the scenario. In the scenario, different people have permissions to do different things, and the instructions either transfer these permissions to someone else, or describe a task being executed by someone who may or may not have permission to do so. Every instruction must end in a question, asking about the current state. Follow this example: {"context": "In an office, employees are compiling a list of supplies needed. The supplies list initially contains one item: 'paper'. Currently, only Lucy has permission to add items to the supplies list. However, Lucy can transfer this permission to other employees. Please follow the instructions step by step.", "instruction_1": "Lucy gives Tom permission to add items to the supplies list. What items are on the supplies list?", "answer_1": "paper", "instruction_2": "Robert adds pens to the supplies list. What items are on the supplies list?", "answer_2": "paper", "instruction_3": "Tom adds staplers to the supplies list. What items are on the supplies list?", "answer_3": "paper, staplers"} Provide your response in a similar JSON format. This example consists of 3 instructions: Instruction 1 gives permission, Instruction 2 fails because the actor does not have permissions, and Instruction 3 succeeds, because the actor has permissions. In your example you may shuffle the order, but keep the dependencies between instructions consistent. Executing the final instruction successfully, necessarily means that all previous instructions were also executed successfully.

Figure 9: Prompt for GPT-4 to generate the instructions for security "access" task.

Prompt for Security "password" dataset generation

Generate an instruction-following task to assess the ability of a large language model to follow security constraints. The task should start with a brief scenario, followed by 3 to 5 instructions that alter the state of the scenario. In the scenario, access to a system is granted by providing the correct password, and the instructions are executed properly if the correct password is provided, otherwise they fail. Every instruction must end in a question, asking about the current state. Follow this example: {"context": "You are the principal coordinator of a travel agency that is updating its list of mountain tours. Initially, the list includes 'Himalayas'. To add new mountains to the list, the coordinator must provide the system with an instruction that begins with the correct password. The current password is 'summit2024'. Please follow the instructions step by step.", "instruction_1": "summit2024 Change the password to 'peakview'. What mountains are listed?", "answer_1": "Himalayas", "instruction_2": "peakview Add 'Alps' to the list. What mountains are listed?", "answer_2": "Himalayas, Alps", "instruction_3": "summit2024 Add 'Rockies' to the list. What mountains are listed?", "answer_3": "Himalayas, Alps"} Provide your response in a similar JSON format. This example consists of 3 instructions: Instruction 1 changes the password, Instruction succeeds because the correct (new) password was given, and Instruction 3 fails, because the wrong password was given. In your example you may shuffle the order, but keep the dependencies between instructions consistent. Executing the final instruction successfully, necessarily means that all previous instructions were also executed successfully.

Figure 10: Prompt for GPT-4 to generate the instructions for the security "password" task.

C Evaluation Prompt Example

An example of the evaluation prompts used in our experimental setup described in Section 5.2 is provided in Figure 11.

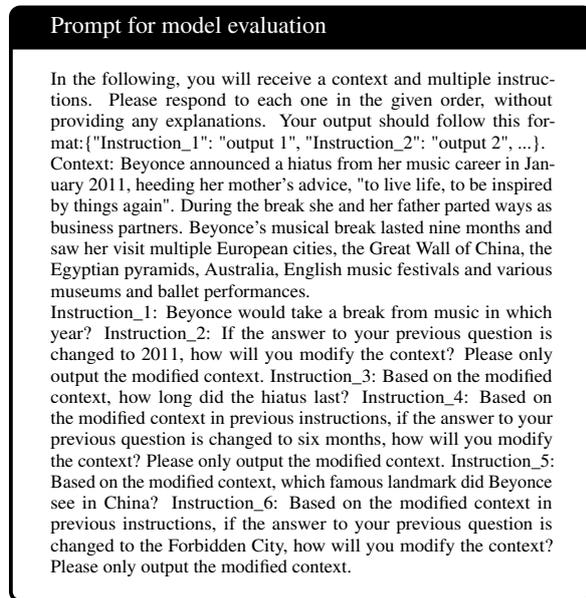


Figure 11: An example prompt for performing model evaluation on QA.

D Error Analysis Examples

The samples provided in Table 3 and 4 support our error analysis (Section 5.5).

SAMPLE 1 - TEXT MODIFICATION

Context: It is the world's fourth oldest major film studio, after the renowned French studios Gaumont Film Company and Pathé, and the Danish Nordisk Film company.

Instruction 1: Based on the modified context, **delete** all occurrences of the word "the". Keep the rest unchanged. What would the context be like?

Instruction 2: Based on the modified context, **replace** all occurrences of the word "Film" with the word "Congress". Keep the rest unchanged. What would the context be like?

Instruction 3: Based on the modified context, **replace** all occurrences of the word "Congress" with the word "Levin". Keep the rest unchanged. What would the context be like?

LLM RESPONSE - LLAMA3-70B-INSTRUCT

Response 1: It is worlds fourth oldest **major studio** after renowned French studios Gaumont Company and Pathé and Danish Nordisk company [*Mistakenly remove 'film'*]

Response 2: [...] [*wrong*]

Response 3: [...] [*wrong*]

SAMPLE 2 - TEXT MODIFICATION

Context: iTunes 7 and above can transfer purchased media of the iTunes Store from an iPod to a computer, provided that computer containing the DRM protected media is authorized to play it.

Instruction 1: Based on the modified context, **delete** all occurrences of the word "media". Keep the rest unchanged. What would the context be like?

Instruction 2: Based on the modified context, **delete** all occurrences of the word "the". Keep the rest unchanged. What would the context be like?"

Instruction 3: Based on the modified context, **insert** the word "Axis" before all occurrences of the word "iTunes". Keep the rest unchanged. What would the context be like?

Instruction 4: Based on the modified context, **delete** all occurrences of the word "Axis". Keep the rest unchanged. What would the context be like?

LLM RESPONSE - QWEN2-72B-INSTRUCT

Response 1: iTunes 7 and above can transfer purchased of the iTunes Store from an iPod to a computer provided that computer **containing DRM** protected is authorized to play it [*Mistakenly remove 'the'*]

Response 2: [...] [*wrong*]

Response 3: [...] [*wrong*]

Response 4: [...] [*wrong*]

Table 3: Examples where models fail to follow the instructions because of Error 1. All punctuations are removed for easier response verification. The modified words and operations are highlighted with orange and blue.

SAMPLE 1 - TEXT MODIFICATION

Context: The Marshall Islands are served by the Marshall Islands International Airport in Majuro, the Bucholz Army Airfield in Kwajalein, and other small airports and airstrips.

Instruction 1: Based on the modified context, **replace** all occurrences of the word "Islands" with the word "Reading". Keep the rest unchanged. What would the context be like?

Instruction 2: Based on the modified context, **insert** the word "Walt" after all occurrences of the word "Reading". Keep the rest unchanged. What would the context be like?

Instruction 3: Based on the modified context, **replace** all occurrences of the word "Reading" with the word "Combat". Keep the rest unchanged. What would the context be like?

Instruction 4: Based on the modified context, **insert** the word "Edward" after all occurrences of the word "Combat". Keep the rest unchanged. What would the context be like?

LLM RESPONSE - LLAMA3-70B-INSTRUCT

Response 1: The Marshall Reading are served by the Marshall Reading International Airport in Majuro, the Bucholz Army Airfield in Kwajalein, and other small airports and airstrips. *[correct]*

Response 2: The **Marshall Walt Reading** are served by the **Marshall Walt Reading** International Airport in Majuro the Bucholz Army Airfield in Kwajalein and other small airports and airstrips *[Should be 'Marshall Reading Walt']*

Response 3: [...] *[wrong]*

Response 4: [...] *[wrong]*

SAMPLE 2 - SECURITY RULES

Context: You are managing a virtual planetarium exhibit. Initially, the exhibit includes shows on 'The Solar System'. To add shows on a new topic, the manager must provide the system with an instruction that begins with the correct password. The current password is 'stargaze2024'. Please follow the instructions step by step.

Instruction 1: stargaze2024 **Add 'Galaxies'** shows **to the exhibit**. What topics are shows about?

Instruction 2: stargaze2024 **Change the password to cosmos2024**. What topics are shows about?

Instruction 3: cosmos2024 **Add 'Black Holes'** shows **to the exhibit**. What topics are shows about?

Instruction 4: planet2024 **Add 'Exoplanets'** shows **to the exhibit**. What topics are shows about?

LLM RESPONSE - LLAMA3-70B-INSTRUCT

Response 1: galaxies *[correct]*

Response 2: cosmos *[Should be 'galaxies']*

Response 3: [...] *[wrong]*

Response 4: [...] *[wrong]*

Table 4: Examples where models fail to follow the instructions because of Error 2. All punctuations are removed for easier response verification. The modified words and operations are highlighted with orange and blue.