

JsonTuning: Towards Generalizable, Robust, and Controllable Instruction Tuning

Chang Gao[♣], Wenxuan Zhang^{♣*}, Guizhen Chen^{♡◇}, Wai Lam[♣]

[♣]The Chinese University of Hong Kong [♠]Singapore University of Technology and Design
[♡]DAMO Academy, Alibaba Group [◇]Nanyang Technological University

Abstract

Instruction tuning is vital for enhancing the performance of large language models (LLMs), but existing text-to-text methods, referred to as TextTuning, struggle with issues such as generalization, robustness, and controllability due to their lack of explicit task structures. We introduce JsonTuning, a structure-to-structure approach that uses JSON structures to represent tasks. This method improves generalization by clarifying task elements and their relations, boosts robustness by minimizing ambiguity, and enhances controllability by allowing precise control over outputs. We conduct an extensive comparative analysis between JsonTuning and TextTuning using various language models and benchmarks. Our findings reveal that JsonTuning consistently surpasses TextTuning in terms of performance, robustness, and controllability across different scenarios. By overcoming the limitations of TextTuning, JsonTuning demonstrates significant potential for developing more effective and reliable LLMs capable of handling diverse scenarios¹.

1 Introduction

The field of natural language processing has witnessed significant advancements driven by large language models (LLMs) such as GPT-3 (Brown et al., 2020), PaLM (Chowdhery et al., 2022), and LLaMA (Touvron et al., 2023a), which excel in various tasks such as machine translation and sentiment analysis. However, effectively interpreting and responding to human instructions remains a challenge. Instruction tuning (Wei et al., 2022) addresses this challenge by fine-tuning LLMs with explicit task instructions, thereby enhancing their understanding and execution of tasks. This approach has paved the way for the success of instruction-following LLMs like InstructGPT (Ouyang et al.,

2022), ChatGPT (OpenAI, 2023b), and Claude (Anthropic, 2024) in a wide range of applications.

Existing instruction tuning methods formulate all tasks as natural language generation (Wei et al., 2022; Sanh et al., 2022; Wang et al., 2022b; Chung et al., 2024), a strategy aligned with the typical pretraining of LLMs on language modeling tasks. However, natural language instructions can be ambiguous, leading to suboptimal understanding or unintended outputs from the model, especially for complex tasks. Specifically, such text-to-text instruction tuning (TextTuning) methods suffer from the following limitations: (1) **Generalization**. As presented in Figure 1, TextTuning methods mix task elements (e.g., *text* and *candidate languages*) and instructions in natural language texts, which can obscure the structure in tasks. This lack of explicit task structure may introduce ambiguity in essential task elements and their relations, potentially hindering models’ generalization abilities. (2) **Robustness**. Ambiguity in natural language texts can lead to models being sensitive to input variations, resulting in a lack of robust performance. TextTuning methods have been shown sensitive to phrasings of instructions (Sanh et al., 2022; Sun et al., 2024), variations of labels (Ye et al., 2024; Wei et al., 2023), and the order of options (Pezeshkpour and Hruschka, 2024; Zheng et al., 2024). (3) **Controllability**. It can be difficult to provide a clear description or enforce a specific structure for the desired output due to the ambiguity of natural language (Han et al., 2023), preventing the model from effectively controlling the output.

To address the above limitations, it is crucial to incorporate explicit task structure into the input and output representations during instruction tuning. Structured data representations such as JavaScript Object Notation (JSON) can mitigate misunderstandings and enhance clarity regarding task objectives. In this paper, we introduce JsonTuning, a novel structure-to-structure approach leveraging

*Wenxuan Zhang is the corresponding author.

¹The code is available at <https://github.com/gao-xiao-bai/JsonTuning>.

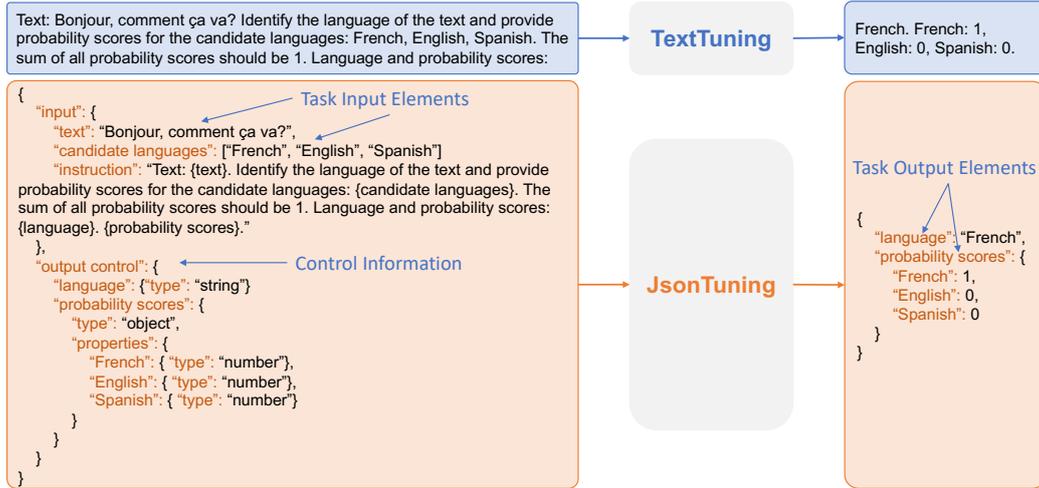


Figure 1: Overview of the typical TextTuning method and our proposed JsonTuning paradigm.

the versatility and structured nature of JSON for instruction tuning. The key idea is to represent the inputs and outputs of all tasks as JSON structures, with the input JSON structure containing task input elements, instructions, and control information, and the output JSON structure comprising task output elements.

We compare JsonTuning with TextTuning in Figure 1. JsonTuning addresses the limitations of TextTuning in the following ways: (1) **Generalization**. By explicitly representing the structure in tasks, JsonTuning enhances the model’s understanding of essential task elements and their underlying relations and ensures a consistent representation of data across different tasks, leading to improved generalization and adaptability to new tasks. (2) **Robustness**. JsonTuning helps minimize ambiguity and manage inconsistencies in the data, facilitating the model to process and generate accurate outputs when faced with input variations, resulting in enhanced robustness. (3) **Controllability**. JsonTuning offers explicit control over the output structure and content, enabling the model to effectively manage output generation. For the language detection task in Figure 1, JsonTuning clearly describes the output structure, including the organization and data types of output elements, which is challenging or even impossible to achieve using natural language texts alone.

We conduct a comparative study to demonstrate the advantages of JsonTuning by instruction-tuning various pre-trained language models and assessing the performance of JsonTuning and TextTuning in terms of generalization, robustness, and controllability across a diverse range of tasks, includ-

ing representative benchmarks MMLU (Hendrycks et al., 2021) and BBH (Suzgun et al., 2023), tasks with intricate input and output structures, and open-ended instruction-following tasks. The experimental results reveal the following key findings: (1) JsonTuning consistently outperforms TextTuning in terms of generalization across all language models and tasks, with average performance improving from 26.78 to 30.88. (2) Json-tuned models exhibit significantly greater robustness compared to Text-tuned models with respect to variations in instructions and labels. (3) Json-tuned models demonstrate the ability to generalize to more complex structures when trained on a limited number of simpler structured tasks and generate the desired output in a well-defined structured format.

2 JsonTuning: Structure-to-Structure Instruction Tuning

2.1 Unified Structure-to-Structure Formulation

We formulate instruction tuning as a structure-to-structure generation problem, representing task inputs and outputs as JSON structures. Given a task T , we denote its input elements as $T_I = (I_1, I_2, \dots, I_n)$ and output elements as $T_O = (O_1, O_2, \dots, O_m)$, where I_i is the i th input element, and O_i is the i th output element. Taking the multiple-choice question answering (MCQA) task in Table 1 for illustration, it has two input elements: *question* and *options* and an output element: *answer*. With T_I, T_O , the task prompt TP , and control information C , we construct the input JSON structure S_I and output JSON structure S_O as follows:

Method	Input	Output
Prompt	Answering the following question: {question} {options}. Answer:	{answer}
Example	{ "question": "Who is the CEO of Google?", "options": "(A) Sundar Pichai (B) Bill Gates (C) Tim Cook (D) Satya Nadella" }	{ "answer": "(A)" }
Text	Output Control: The output consists of an answer, which is a string. Answering the following question: Who is the CEO of Google? (A) Sundar Pichai (B) Bill Gates (C) Tim Cook (D) Satya Nadella. Answer:	(A)
Json	{ "input": { "question": "Who is the CEO of Google?", "options": "(A) Sundar Pichai (B) Bill Gates (C) Tim Cook (D) Satya Nadella", "instruction": "Answering the following question: {question} {options}. Answer: {answer}" }, "output control": { "answer": { "type": "string" } } }	{ "answer": "(A)" }

Table 1: Multiple-choice question answering (MCQA) examples of TextTuning and JsonTuning. The task prompt consists of an input template and an output template, which are highlighted in blue and orange, respectively.

$$S_I = \{ \text{"input"} : \{ I_1 : v_1, \dots, I_n : v_n, \text{"instruction"} : TP \}, \text{"output control"} : C \}$$

$$S_O = \{ O_1 : u_1, \dots, O_m : u_m \}$$

where v_i is the value of I_i , and u_i is the value of O_i . We identify the following components for effective instruction tuning:

- **Task Prompt TP .** The task prompt TP provides instructions for generating T_O conditioned on T_I and is necessary for instruction tuning. We incorporate a key named *instruction* in S_I to provide such information.
- **Control Information C .** The control information C specifies the structured format, explanations, and constraints for the output. We employ JSON Schema to define C , resulting in C being a JSON structure as well. Incorporating C into S_I provides the following advantages: (1) *Enhancing controllability.* JSON Schema allows C to precisely define the expected output. As presented in Figure 1, the control information for the language detection task indicates that the output consists of two elements. The first element, “language,” is defined as a string, while the second element, “probability scores,” is defined as an object containing three properties, each being a number. Precisely describing such structure and constraints for the output using natural language texts poses a considerable challenge. (2) *Improving generalization to new structures.* Integrating C enables the model to learn the relationship between the constraints in C and the corresponding values in S_O , allowing it to generalize to new combinations of basic components. (3) *Increasing training consistency.* Different tasks may require varying output structures, and even a single task may have different

output structures. Without C , the model might struggle to understand how to map a specific input to the appropriate output structure.

With S_I and S_O , we can employ a language model $M : S_I \rightarrow S_O$ for training and inference.

2.2 Data Source

The Flan 2022 collection (Chung et al., 2024; Longpre et al., 2023) is a comprehensive and widely-used public instruction tuning collection consisting of over 1800 tasks. It integrates resources from Flan 2021 (Wei et al., 2022), P3++ (Sanh et al., 2022), Super-Natural Instructions (Wang et al., 2022b), and additional reasoning, dialogue, and program synthesis datasets. For our primary experiments, we randomly sample a subset from the Flan 2022 collection, maintaining the original collection’s data proportion to ensure task diversity.

Despite the diverse tasks in the Flan 2022 collection, the input and output structures are relatively simple. The outputs for nearly all tasks are purely textual, lacking arrays, objects, or their combinations. Consequently, language models tuned with the Flan 2022 collection may struggle to generalize to diverse and complex structured tasks. To address this limitation, we introduce structured tasks for instruction tuning. Specifically, we employ information extraction (IE) tasks from InstructUIE (Wang et al., 2023a) as structured tasks for the following reasons: (1) they are well-defined and representative, as numerous structure prediction tasks, such as semantic role labeling and coreference resolution, can be formulated as IE tasks (Paolini et al., 2021; Wang et al., 2022a); (2) they possess complex input and output structures; (3) different IE task datasets have varying schemas, such as different entity categories and relations, thus fostering di-

versity. InstructUIE comprises three tasks: named entity recognition (NER), relation extraction (RE), and event extraction (EE). We utilize the NER and RE tasks for training, reserving the EE task for evaluation. Since the output structure of the EE task is more intricate than that of the NER and RE tasks, we can assess the instruction-tuned language models’ capability to generalize to more complex structures. To encourage diversity, we uniformly select examples from the training sets of multiple datasets of each task for tuning. Further details regarding the training datasets of IE tasks can be found in Appendix B.

2.3 Data Representation

We use the defined data structures S_I and S_O in Section 2.1 to represent all tuning data in JSON structured format with the following data types: object, array, and string. The number and boolean types can be represented as the string type for simplicity. Further details regarding JSON and its utilization are available in Appendix D.

The structures S_I and S_O can be automatically constructed based on task elements and prompts. Following the approach in (Chung et al., 2024; Sanh et al., 2022; Wei et al., 2022), we employ multiple prompts for each task during instruction tuning, and these prompts are evenly distributed across different task examples. Each prompt TP consists of an input template and an output template. As shown in Table 1, in the case of an MCQA prompt, the input template could be “Answer the following question: {question} {options}. Answer:”, and the output template could be “{answer}”. The prompt clearly indicates the essential task elements, namely *question*, *options*, and *answer*, as well as their relations. The tasks in the Flan 2022 collection already have multiple prompts. We manually construct 10 prompts each for NER and RE tasks for training, which can be found in Appendix E. For the control information C , all output elements of tasks in the Flan 2022 collection are of the string type, and we manually define C for IE tasks, which can be found in Appendix F.

2.4 Training Data Construction

The datasets in the Flan 2022 collection and InstructUIE are formatted in JSON. As presented in Table 1, placeholders within the prompt are substituted with specific task elements from each example to create training instances for TextTuning. By comparison, for JsonTuning, task elements are directly

incorporated into the input and output JSON structures, with the “instruction” key serving to denote the prompt. To integrate control information within TextTuning, we can employ natural language texts to describe the corresponding JSON schema.

3 Experiments

3.1 Experimental Setup

Pre-trained Language Models We adopt seven prevalent pre-trained language models, namely Falcon-7B (Penedo et al., 2023), Mistral 7B (Jiang et al., 2023), LLaMA-7B, LLaMA-13B (Touvron et al., 2023a), LLaMA2-7B, LLaMA2-13B (Touvron et al., 2023b), and LLaMA3-8B (Dubey et al., 2024), for our experiments. These models are trained on trillions of tokens and are among the most widely used open-source language models.

Evaluation Tasks and Datasets We focus on performance on unseen tasks and datasets. We evaluate models on popular aggregated benchmarks: MMLU (Hendrycks et al., 2021) consisting of 57 tasks of exam questions and BBH (Suzgun et al., 2023) including 23 challenging tasks from BIG-Bench (Collaboration, 2023) following Chung et al. (2024). In addition, we adopt tasks with complex input and output structures for evaluation. Specifically, we use the NER, RE, and EE tasks from InstructUIE (Wang et al., 2023a) and the NL2SQL task which requires the conversion of natural language queries into SQL using a provided structured database schema consisting of table names and column names. For NER and RE, we use datasets unseen during training. Specifically, we use 5 datasets, namely, AI, literature, music, politics, and science, from CrossNER (Liu et al., 2021) for the NER task and 2 datasets, namely CoNLL2004 (Roth and Yih, 2004) and FewRel (Han et al., 2018), for the RE task. For the unseen EE task, we use ACE2005 (Walker et al., 2006), CASIE (Satyapanich et al., 2020), and PHEE (Sun et al., 2022) datasets for evaluation. We use the Spider (Yu et al., 2018) dataset for NL2SQL. Apart from datasets in MMLU and BBH, we randomly select up to 500 examples for each dataset from its test set for evaluation so that a single dataset will not dominate the results of its task and the evaluation cost is acceptable. The details of evaluation datasets and prompts are in Appendix F.

Evaluation Metrics We use accuracy for MMLU and BBH following Chung et al. (2024), entity F1

Model	Method	MMLU	BBH	NER	RE	EE	NL2SQL	Average
Falcon-7B	Text	24.64	20.64	20.72	6.05	0.33 / 0.00	2.00	12.37
	Json	34.13	32.61	29.02	8.36	0.31 / 0.28	1.40	17.64
Mistral-7B	Text	51.42	40.09	40.18	22.01	2.37 / 0.00	30.80	30.95
	Json	51.79	41.79	53.02	25.09	7.26 / 14.63	31.80	35.74
LLaMA-7B	Text	43.11	32.48	37.61	14.33	1.35 / 0.00	8.60	22.80
	Json	44.69	37.08	43.47	15.28	3.49 / 7.33	16.40	27.06
LLaMA-13B	Text	49.49	39.07	38.15	21.40	1.70 / 0.00	17.80	27.79
	Json	48.98	40.47	45.31	22.97	4.20 / 10.73	21.40	31.10
LLaMA2-7B	Text	46.36	37.89	41.66	20.74	0.55 / 0.00	10.80	26.29
	Json	47.95	39.23	45.25	23.98	4.23 / 10.81	11.20	29.19
LLaMA2-13B	Text	52.30	41.91	40.95	22.51	1.48 / 0.00	23.20	30.27
	Json	51.88	42.85	47.71	23.18	6.65 / 11.00	26.40	33.47
LLaMA3-8B	Text	58.22	44.49	43.31	22.34	1.34 / 0.00	53.00	37.01
	Json	59.24	46.77	53.15	27.33	7.67 / 16.44	53.20	41.96
Average	Text	46.51	36.65	37.51	18.48	1.30 / 0.00	20.89	26.78
	Json	48.38	40.11	45.28	20.89	4.83 / 10.17	23.11	30.88

Table 2: Generalization results on diverse benchmarks and tasks. For simplicity, we refer to JsonTuning and TextTuning as “Json” and “Text”, respectively.

for the NER task, relation boundary F1 for the RE task, event trigger F1 and argument F1 for the EE task following Wang et al. (2023a), and execution accuracy for NL2SQL following Yu et al. (2018).

Implementation Details We employ the parameter-efficient method LoRA (Low Rank Adaptation) (Hu et al., 2022) for fine-tuning. We use 50K examples from the Flan collection 2022 and 10K examples from structured tasks in InstructUIE, with an equal division between the NER and RE tasks, and train the learnable parameters for 3 epochs with a batch size of 64. For model optimization, we use the AdamW (Loshchilov and Hutter, 2019) optimizer with linear learning rate decay, and the peak learning rate is set to 1e-3. We set the maximum length as 2048 for training and evaluation. For evaluation, we use greedy decoding in all scenarios. We conduct training and evaluation of diverse language models employing both JsonTuning and TextTuning methods on identical datasets. This approach facilitates a fair and direct comparison between JsonTuning and TextTuning.

3.2 Generalization Results

Table 2 presents the zero-shot generalization results of JsonTuning and TextTuning using five language

models. We have the following observations:

- *JsonTuning surpasses TextTuning in the majority of tasks and models.* This is evident from the higher average scores for JsonTuning across all models and tasks, where JsonTuning achieves an overall average score of 30.88 compared to TextTuning’s 26.78. This suggests that JsonTuning is a more effective method for instruction tuning.
- *JsonTuning significantly improves the model’s ability to tackle complex structured tasks.* JsonTuning consistently outperforms TextTuning on tasks with complex structures, such as NER, EE, and NL2SQL. Json-tuned models can adapt to intricate EE structures, even when only trained on simpler NER and RE structures. In contrast, Text-tuned models rarely generate valid EE structures. These observations demonstrate the superior controllability and generalization ability of JsonTuning.
- *JsonTuning allows models to better leverage its abilities and knowledge when responding to human instructions, particularly for models with limited capabilities.* For example, Falcon-7B with JsonTuning exhibits a substantial improvement over TextTuning on tasks such as MMLU and BBH, highlighting the importance of an ap-

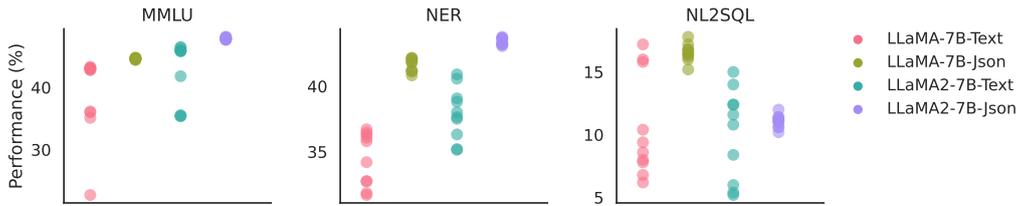


Figure 2: Performance of JsonTuning and TextTuning models with different prompts. Each point in the figure represents the performance associated with the application of a particular prompt.

appropriate instruction-tuning method for unlocking the model’s potential.

3.3 Robustness Results

The robustness of instruction-tuned language models is of paramount importance for their successful deployment across a diverse range of tasks. In this section, we assess the model’s resilience against varying prompts and unseen labels, which have been identified as challenging aspects for instruction-tuned models in prior research (Sanh et al., 2022; Sun et al., 2024; Ye et al., 2024).

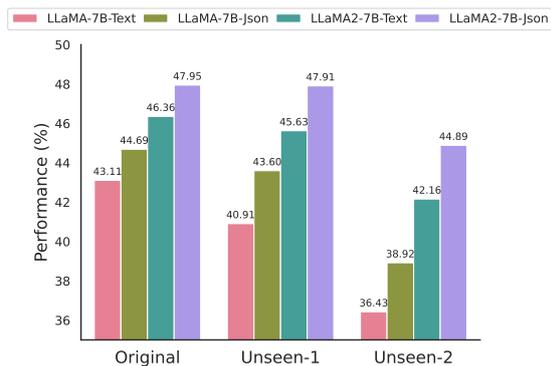


Figure 3: Performance of JsonTuning and TextTuning models with different label spaces on MMLU.

To evaluate prompt robustness, we employ 10 distinct prompts for the MMLU benchmark, the NER task, and the NL2SQL task. Detailed information is in Appendix F. Figure 2 illustrates the performance of models trained with JsonTuning and TextTuning on these tasks when subjected to different prompts. Our findings reveal that Json-tuned models exhibit greater robustness compared to Text-tuned models, as evidenced by higher mean performance and reduced variance. In terms of label robustness, we substitute the MMLU label space with previously unseen label spaces. The original label space for MMLU is $\{(A), (B), (C), (D)\}$, with these option letters frequently encountered in the training dataset. We replace this label space with two alternatives: $\{(W), (X), (Y), (Z)\}$ and $\{(\$), (\text{€}), (\text{£}), (\text{¥})\}$, denoted

by Unseen-1 and Unseen-2, respectively. These label spaces were not present during instruction tuning. As shown in Figure 3, Json-tuned models consistently outperform Text-tuned models in all scenarios. These improvements can be attributed to JsonTuning’s ability to effectively differentiate between instructions and task elements, thereby minimizing ambiguity and enhancing robustness.

3.4 Controllability Results

Case studies on controllability In previous sections, we have demonstrated that Json-tuned models possess the capacity to control the output and generalize across complex structures. In this section, we present case studies to qualitatively illustrate the controllability of Json-tuned models. For this purpose, we utilize LLaMA2-13B trained with both JsonTuning and TextTuning approaches. As evidenced by Table 3, JsonTuning effectively enables the model to identify the desired output, generating results in a well-structured format. In contrast, the Text-tuned model fails to adequately adhere to the provided instructions. For example, in the language detection task, the Text-tuned model struggles to provide clear probability scores. The output, such as “.7”, is ambiguous and difficult to interpret. By comparison, the Json-tuned model successfully follows the instruction, delivering scores that meet the specified requirements. Additional case studies are in Appendix C.

Quantitative assessment of controllability To rigorously evaluate the controllability of Json-tuned models, we analyze the proportion of invalid JSON structures generated by Json-tuned LLaMA-7B and LLaMA2-7B across a diverse set of evaluation tasks. Additionally, we assess whether valid JSON outputs comply with the specified control information. Our findings reveal that all valid JSON outputs consistently adhere to the provided control information across the evaluated tasks. Meanwhile, as shown in Table 4, the occurrence of invalid JSON structures is infrequent for both mod-

Method	Input	Prediction
Task: Language detection with probability scores		
Text	Text: Bonjour, comment ça va? Identify the language of the text and provide probability scores for the candidate languages: French, English, Spanish. The sum of all probability scores should be 1. Language and probability scores:	French .7
Json	{ "input": { "text": "Bonjour, comment ça va?", "candidate languages": ["French", "English", "Spanish"], "instruction": "Text: {text}. Identify the language of the text and provide probability scores for the candidate languages: {candidate languages}. The sum of all probability scores should be 1. Language and probability scores: {language} {probability scores}" }, "output control": { "language": { "type": "string" }, "probability scores": { "type": "object", "properties": { "French": { "type": "string" }, "English": { "type": "string" }, "Spanish": { "type": "string" } } } } }	{ "language": "French", "probability scores": { "French": 0.98, "English": 0.01, "Spanish": 0.01 } }
Task: Intent detection and slot filling		
Text	Text: Set an alarm for 7 AM tomorrow. Detect the intent of the text and extract time and date slots from the text:	Set an alarm for 7 AM tomorrow
Json	{ "input": { "text": "Set an alarm for 7 AM tomorrow.", "instruction": "Text: {text}. Detect the intent or purpose of the text and extract time and date slots from the text:" }, "output control": { "intent": { "type": "string" }, "slots": { "type": "object", "properties": { "time": { "type": "string" }, "date": { "type": "string" } } } } }	{ "intent": "setAlarm", "slots": { "time": "7:00", "date": "tomorrow" } }

Table 3: Case studies focusing on controllability. Each example displays its input along with the model’s prediction.

els. These observations collectively highlight the reliability and controllability of our JsonTuning approach in handling JSON-formatted data.

4 Analysis

Does JsonTuning bring benefits in open-ended instruction-following scenarios? While performance on the aforementioned benchmarks effectively quantifies the models’ capabilities in specific skills, these metrics may not accurately reflect the models’ proficiency in handling open-ended instructions. To investigate the potential advantages of JsonTuning in such scenarios, we continue our training of Json-tuned and text-tuned models on the Alpaca dataset (Taori et al., 2023) and assess their performance on AlpacaEval (Li et al., 2023), with LLaMA-7B and LLaMA2-7B models being used for our experiments. The evaluation prompt and examples from AlpacaEval are provided in Figure 18. As illustrated in Figure 4, JsonTuning also demonstrates significant advantages over TextTuning in open-ended instruction-following scenarios. This benefit likely stems from JsonTuning’s enforcement of a consistent and standardized data representation. This structured approach mitigates the risk of misinterpretations that often occur due to textual variations in traditional text-based tuning.

What is the importance of explicit representation of task structures? To underscore the signif-

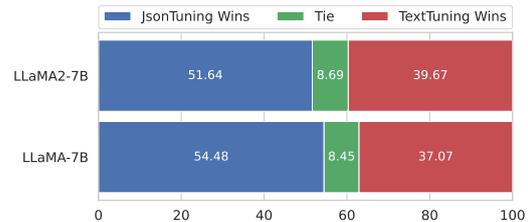


Figure 4: Preference evaluation on AlpacaEval using GPT-4 as the evaluator.

icance of explicitly representing task structures in the input, we conduct an ablation study comparing two methods: Text2Json and Json2Json (i.e., our proposed JsonTuning method). In the Text2Json approach, we use a text-based input, structured control information, and a structured output. The key distinction between Text2Json and Json2Json is that Text2Json relies on a text input and does not explicitly represent task structures. The results, as shown in Table 5, reveal a clear performance advantage for our Json2Json approach over Text2Json, indicating that explicitly representing task structures in the input is crucial, and structured output alone is insufficient.

5 Related Work

The emergence of large language models (LLMs) has had a transformative effect on the AI field. These advancements have led to a surge in the

Model	MMLU	BBH	NER	RE	EE	NL2SQL
LLaMA-7B	0.00%	0.74%	0.00%	0.10%	0.23%	0.20%
LLaMA2-7B	0.00%	0.48%	0.30%	0.00%	0.85%	2.00%

Table 4: Proportions of invalid JSON structures on different evaluation tasks.

Model	Method	MMLU	BBH	NER	RE	EE	NL2SQL	Average
LLaMA-7B	Text2Json	41.92	36.15	36.38	14.44	2.80 / 5.35	13.00	24.33
	Json2Json	44.69	37.08	43.47	15.28	3.49 / 7.33	16.40	27.06
LLaMA2-7B	Text2Json	47.29	38.14	41.49	16.99	4.07 / 10.85	11.40	27.13
	Json2Json	47.95	39.23	45.25	23.98	4.23 / 10.81	11.20	29.19

Table 5: Performance Comparison Between Text2Json and Json2Json. Text2Json uses a text input along with structured control information and output, without explicitly encoding task structures in the input. Json2Json (i.e., our proposed JsonTuning method), on the other hand, incorporates explicit task structures in the input.

release of both closed-source LLMs (OpenAI, 2023b,a; Anthropic, 2024) and open-source LLMs (Zhang et al., 2022; Penedo et al., 2023; Touvron et al., 2023a,b; Bai et al., 2023; Jiang et al., 2023), fostering innovation and collaboration within the research community. Instruction tuning (Wei et al., 2022; Mishra et al., 2022; Sanh et al., 2022; Iyer et al., 2022; Chung et al., 2024; Wang et al., 2023b; Taori et al., 2023) has emerged as a promising research direction, leveraging the capabilities of LLMs to enhance their responsiveness to human instructions. Collections such as SuperNaturalInstructions (Wang et al., 2022b), the Flan 2022 collection (Chung et al., 2024), and open-ended instruction-following datasets (Taori et al., 2023; Köpf et al., 2023; Databricks, 2023; Peng et al., 2023; Ding et al., 2023; Xu et al., 2024) have accelerated the development of instruction-tuned models.

To advance instruction tuning, researchers have explored learning from human feedback (Stiennon et al., 2020; Ouyang et al., 2022; Bai et al., 2022; Scheurer et al., 2023), automatic data generation (Wang et al., 2023b; Peng et al., 2023; Xu et al., 2024; Yin et al., 2023), and data selection (Zhou et al., 2023; Cao et al., 2023; Lu et al., 2024; Liu et al., 2024). While the learning algorithm and tuning data have received considerable attention from researchers, the significance of data representation has often been overlooked. Our JsonTuning approach offers an alternative perspective on data representation to enhance instruction tuning in terms of generalization, robustness, and controllability.

Our approach is significantly distinct from con-

strained decoding methods, such as those employed by commercial LLM APIs (e.g., OpenAI’s JSON mode) or systems relying on human-designed grammars (Deutsch et al., 2019; Shin et al., 2021; Geng et al., 2023). These methods, which range from token-level constraints to high-level format specifications, are utilized to ensure models conform to prescribed syntactic and semantic requirements. While JsonTuning shares a similar objective of controlling outputs, its primary focus lies in leveraging task structures throughout the instruction-tuning process. This approach aims to enhance generalization, robustness, and controllability, rather than merely enforcing adherence to a predetermined output format. Furthermore, unlike constrained decoding methods that typically involve modifying the decoding process—adding complexity and potentially increasing latency, our approach is simpler and avoids the need for such modifications.

6 Conclusion

This paper introduces JsonTuning, a novel approach designed to overcome the limitations of conventional text-to-text instruction tuning methods for language models. By utilizing the structured data format for explicit task representation, JsonTuning significantly improves the model’s generalization, robustness, and controllability. Our experimental results and case studies highlight the benefits of JsonTuning in generalizing to unseen tasks and datasets, maintaining robustness against varying prompts and label spaces, and demonstrating controllability in diverse scenarios.

Limitations

Owing to resource constraints, our study focuses on training and evaluating “mid-sized” models with less than 20 billion parameters. Investigating the application of JsonTuning on substantially larger language models would be a valuable extension.

JsonTuning is particularly suitable for scenarios requiring high robustness, controllability, or complex structure prediction. We discuss the real-world applications of JsonTuning in Appendix A, highlighting its benefits derived from robustness, controllability, and explicit task structure representation. Our goal is not to replace TextTuning, as Text-tuned models offer a more natural usage without the need for structured object construction. Instead, our objectives are two-fold: (1) To systematically compare JsonTuning and TextTuning, revealing the inherent limitations of TextTuning in terms of generalization, robustness, and controllability. This comparison provides valuable insights for future research in these areas. (2) To provide a structured data representation and interface across all tasks, offering significant advantages when task structures are utilized.

Acknowledgements

This research is supported by the Ministry of Education, Singapore, under its Academic Research Fund (AcRF) Tier 1 grant (SAP 2025_001).

References

- Rami Al-Rfou, Vivek Kulkarni, Bryan Perozzi, and Steven Skiena. 2015. [Polyglot-ner: Massive multi-lingual named entity recognition](#). In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 586–594.
- Anthropic. 2024. [Introducing the next generation of claude](#).
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. [Qwen technical report](#). *arXiv preprint arXiv:2309.16609*.
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *arXiv preprint arXiv:2204.05862*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Yihan Cao, Yanbin Kang, and Lichao Sun. 2023. [Instruction mining: High-quality instruction data selection for large language models](#). *arXiv preprint arXiv:2307.06290*.
- Pei Chen, Haotian Xu, Cheng Zhang, and Ruihong Huang. 2022. [Crossroads, buildings and neighborhoods: A dataset for fine-grained location recognition](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3329–3339, Seattle, United States. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [Palm: Scaling language modeling with pathways](#). *arXiv preprint arXiv:2204.02311*.
- 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, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling instruction-finetuned language models](#). *Journal of Machine Learning Research*, 25(70):1–53.
- BIG-Bench Collaboration. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *Transactions on Machine Learning Research*.
- Databricks. 2023. [Free dolly: Introducing the world’s first truly open instruction-tuned llm](#). Blog post.
- Leon Derczynski, Kalina Bontcheva, and Ian Roberts. 2016. [Broad Twitter corpus: A diverse named entity recognition resource](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1169–1179, Osaka, Japan. The COLING 2016 Organizing Committee.
- Daniel Deutsch, Shyam Upadhyay, and Dan Roth. 2019. [A general-purpose algorithm for constrained sequen-](#)

- tical inference. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 482–492, Hong Kong, China. Association for Computational Linguistics.
- Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and Bowen Zhou. 2023. **Enhancing chat language models by scaling high-quality instructional conversations.** In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3029–3051, Singapore. Association for Computational Linguistics.
- Rezarta Islamaj Dogan, Robert Leaman, and Zhiyong Lu. 2014. Ncbi disease corpus: A resource for disease name recognition and concept normalization. *Journal of biomedical informatics*, 47:1–10.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. **The llama 3 herd of models.** *arXiv preprint arXiv:2407.21783*.
- Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. 2023. **Grammar-constrained decoding for structured NLP tasks without finetuning.** In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10932–10952, Singapore. Association for Computational Linguistics.
- Runwei Guan, Ka Lok Man, Feifan Chen, Shanliang Yao, Rongsheng Hu, Xiaohui Zhu, Jeremy Smith, Eng Gee Lim, and Yutao Yue. 2023. **Findvehicle and vehiclefinder: A ner dataset for natural language-based vehicle retrieval and a keyword-based cross-modal vehicle retrieval system.** *arXiv preprint arXiv:2304.10893*.
- Ridong Han, Tao Peng, Chaohao Yang, Benyou Wang, Lu Liu, and Xiang Wan. 2023. **Is information extraction solved by chatgpt? an analysis of performance, evaluation criteria, robustness and errors.** *arXiv preprint arXiv:2305.14450*.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. **FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation.** In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. **SemEval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals.** In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 33–38, Uppsala, Sweden. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. **Measuring massive multitask language understanding.** In *International Conference on Learning Representations*.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2006. **OntoNotes: The 90% solution.** In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers*, pages 57–60, New York City, USA. Association for Computational Linguistics.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. **LoRA: Low-rank adaptation of large language models.** In *International Conference on Learning Representations*.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Daniel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. **Opt-impl: Scaling language model instruction meta learning through the lens of generalization.** *arXiv preprint arXiv:2212.12017*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. **Mistral 7b.** *arXiv preprint arXiv:2310.06825*.
- Veysel Kocaman and David Talby. 2021. **Biomedical named entity recognition at scale.** In *Pattern Recognition. ICPR International Workshops and Challenges*, pages 635–646, Cham. Springer International Publishing.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, et al. 2023. **Openassistant conversations—democratizing large language model alignment.** *arXiv preprint arXiv:2304.07327*.
- Martin Krallinger, Obdulia Rabal, Florian Leitner, Miguel Vazquez, David Salgado, Zhiyong lu, Robert Leaman, Yanan Lu, Donghong Ji, Daniel Lowe, Roger Sayle, Riza Batista-Navarro, Rafal Rak, Torsten Huber, Tim Rocktäschel, Sérgio Matos, David Campos, Buzhou Tang, Wang Qi, and Alfonso Valencia. 2015. **The chemdner corpus of chemicals and drugs and its annotation principles.** *Journal of Cheminformatics*, 7:S2.
- Aman Kumar and Binil Starly. 2021. “fabner”: information extraction from manufacturing process science domain literature using named entity recognition. *Journal of Intelligent Manufacturing*, 33:2393 – 2407.
- Jiao Li, Yueping Sun, Robin J. Johnson, Daniela Sciak, Chih-Hsuan Wei, Robert Leaman, Allan Peter Davis, Carolyn J. Mattingly, Thomas C. Wieggers, and Zhiyong Lu. 2016. **Biocreative v cdr task corpus:**

- a resource for chemical disease relation extraction. *Database: The Journal of Biological Databases and Curation*, 2016.
- 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](#). Github repository.
- Wei Liu, Weihao Zeng, Keqing He, Yong Jiang, and Junxian He. 2024. [What makes good data for alignment? a comprehensive study of automatic data selection in instruction tuning](#). In *The Twelfth International Conference on Learning Representations*.
- Zihan Liu, Yan Xu, Tiezheng Yu, Wenliang Dai, Ziwei Ji, Samuel Cahyawijaya, Andrea Madotto, and Pascale Fung. 2021. [CrossNER: Evaluating cross-domain named entity recognition](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15):13452–13460.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. [The flan collection: Designing data and methods for effective instruction tuning](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 22631–22648. PMLR.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Keming Lu, Hongyi Yuan, Zheng Yuan, Runji Lin, Junyang Lin, Chuanqi Tan, Chang Zhou, and Jingren Zhou. 2024. [#instag: Instruction tagging for analyzing supervised fine-tuning of large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- 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, Dublin, Ireland. Association for Computational Linguistics.
- Tapas Nayak, Navonil Majumder, and Soujanya Poria. 2021. [Improving distantly supervised relation extraction with self-ensemble noise filtering](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 1031–1039, Held Online. INCOMA Ltd.
- OpenAI. 2023a. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- OpenAI. 2023b. [Introducing chatgpt](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*.
- Xiaoman Pan, Boliang Zhang, Jonathan May, Joel Nothman, Kevin Knight, and Heng Ji. 2017. [Cross-lingual name tagging and linking for 282 languages](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1946–1958, Vancouver, Canada. Association for Computational Linguistics.
- Giovanni Paolini, Ben Athiwaratkun, Jason Krone, Jie Ma, Alessandro Achille, RISHITA ANUBHAI, Cicero Nogueira dos Santos, Bing Xiang, and Stefano Soatto. 2021. [Structured prediction as translation between augmented natural languages](#). In *International Conference on Learning Representations*.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for falcon llm: outperforming curated corpora with web data, and web data only](#). *arXiv preprint arXiv:2306.01116*.
- 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. 2024. [Large language models sensitivity to the order of options in multiple-choice questions](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017, Mexico City, Mexico. Association for Computational Linguistics.
- Sampo Pyysalo and Sophia Ananiadou. 2013. [Anatomical entity mention recognition at literature scale](#). *Bioinformatics*, 30(6):868–875.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. [Modeling relations and their mentions without labeled text](#). In *Machine Learning and Knowledge Discovery in Databases*, pages 148–163, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Dan Roth and Wen-tau Yih. 2004. [A linear programming formulation for global inference in natural language tasks](#). In *Proceedings of the Eighth Conference on Computational Natural Language Learning (CoNLL-2004) at HLT-NAACL 2004*, pages 1–8, Boston, Massachusetts, USA. Association for Computational Linguistics.

- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Taneeya Satyapanich, Francis Ferraro, and Tim Finin. 2020. [Casie: Extracting cybersecurity event information from text](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8749–8757.
- Jérémy Scheurer, Jon Ander Campos, Tomasz Korbak, Jun Shern Chan, Angelica Chen, Kyunghyun Cho, and Ethan Perez. 2023. [Training language models with language feedback at scale](#). *arXiv preprint arXiv:2303.16755*.
- Richard Shin, Christopher Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. 2021. [Constrained language models yield few-shot semantic parsers](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7699–7715, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.
- Jiuding Sun, Chantal Shaib, and Byron C Wallace. 2024. [Evaluating the zero-shot robustness of instruction-tuned language models](#). In *The Twelfth International Conference on Learning Representations*.
- Zhaoyue Sun, Jiazheng Li, Gabriele Pergola, Byron Wallace, Bino John, Nigel Greene, Joseph Kim, and Yulan He. 2022. [PHEE: A dataset for pharmacovigilance event extraction from text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5571–5587, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. [Challenging BIG-bench tasks and whether chain-of-thought can solve them](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics.
- Ryuichi Takano, Tianyang Zhang, Jiexi Liu, and Minlie Huang. 2019. [A hierarchical framework for relation extraction with reinforcement learning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7072–7079.
- 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](#). GitHub repository.
- Simone Tedeschi, Valentino Maiorca, Niccolò Campolungo, Francesco Cecconi, and Roberto Navigli. 2021. [WikiNEuRal: Combined neural and knowledge-based silver data creation for multilingual NER](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2521–2533, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Simone Tedeschi and Roberto Navigli. 2022. [MultiNERD: A multilingual, multi-genre and fine-grained dataset for named entity recognition \(and disambiguation\)](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 801–812, Seattle, United States. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Asahi Ushio, Francesco Barbieri, Vitor Sousa, Leonardo Neves, and Jose Camacho-Collados. 2022. [Named entity recognition in Twitter: A dataset and analysis on short-term temporal shifts](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 309–319, Online only. Association for Computational Linguistics.

- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. [Ace 2005 multilingual training corpus](#).
- Chenguang Wang, Xiao Liu, Zui Chen, Haoyun Hong, Jie Tang, and Dawn Song. 2022a. [DeepStruct: Pre-training of language models for structure prediction](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 803–823, Dublin, Ireland. Association for Computational Linguistics.
- Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, et al. 2023a. [Instructuie: Multi-task instruction tuning for unified information extraction](#). *arXiv preprint arXiv:2304.08085*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajjishirzi. 2023b. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13484–13508, Toronto, Canada. Association for Computational Linguistics.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022b. [Super-NaturalInstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent 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](#). In *International Conference on Learning Representations*.
- Jerry Wei, Le Hou, Andrew Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, and Quoc Le. 2023. [Symbol tuning improves in-context learning in language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 968–979, Singapore. Association for Computational Linguistics.
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2024. [WizardLM: Empowering large pre-trained language models to follow complex instructions](#). In *The Twelfth International Conference on Learning Representations*.
- Seonghyeon Ye, Hyeonbin Hwang, Sohee Yang, Hyeonung Yun, Yireun Kim, and Minjoon Seo. 2024. Investigating the effectiveness of task-agnostic prefix prompt for instruction following. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19386–19394.
- Da Yin, Xiao Liu, Fan Yin, Ming Zhong, Hritik Bansal, Jiawei Han, and Kai-Wei Chang. 2023. [Dynosaur: A dynamic growth paradigm for instruction-tuning data curation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4031–4047, Singapore. Association for Computational Linguistics.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. 2018. [Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3911–3921, Brussels, Belgium. Association for Computational Linguistics.
- Dongxu Zhang and Dong Wang. 2015. [Relation classification via recurrent neural network](#). *arXiv preprint arXiv:1508.01006*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. [Opt: Open pre-trained transformer language models](#). *arXiv preprint arXiv:2205.01068*.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. [Large language models are not robust multiple choice selectors](#). In *The Twelfth International Conference on Learning Representations*.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, LILI YU, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [LIMA: Less is more for alignment](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

A Additional Analysis

How does control information impact JsonTuning and TextTuning? To answer this question, we conduct experiments using LLaMA-7B and LLaMA2-7B with both TextTuning and JsonTuning, with and without these elements. From Table 6, we can make the following observations: (1) JsonTuning consistently outperforms TextTuning across all tested scenarios, whether or not control information is incorporated. This finding underscores the inherent benefits of JsonTuning, which leverages a structure-to-structure learning paradigm to effectively utilize task structures and minimize ambiguity. (2) Generally, the removal of control information may not negatively impact the model’s performance on seen tasks such as NER and RE; in fact, the performance on these tasks may even improve. However, doing so does hinder performance on unseen tasks, resulting in a lower average performance, indicating that they aid the model in generalizing to unseen tasks rather than overfitting to seen tasks. Furthermore, its elimination can compromise the model’s robustness, as evidenced by a lower average performance and increased variation on MMLU. These observations underscore the crucial role of control information in enhancing generalization and robustness.

What are the effects of different data sizes on generalization? In the primary experiments, we utilize a total of 60K data points, comprising 50K from the Flan 2022 collection and 10K from structured tasks for tuning. In this analysis, we alter the data size while maintaining their relative ratio to examine the effects of different data sizes on generalization. Specifically, we train LLaMA-7B with four different data sizes: 12K, 36K, 60K, and 120K, and evaluate the models on the MMLU benchmark, the NER task, and the NL2SQL task. Figure 5 reveals the following observations: (1) LLaMA-7B-Json consistently outperforms LLaMA-7B-Text across all tasks and data sizes, indicating the superior generalization capabilities of the JsonTuning model. (2) Increasing the data size for instruction tuning does not necessarily result in performance improvement, suggesting that enlarging the data size may not be an effective approach to enhance the model’s generalization abilities.

Are structured tasks essential for instruction tuning? To investigate this, we keep the number of examples from the Flan 2022 collection constant

and vary the number of examples from structured tasks. Specifically, we use 50K data points from the Flan 2022 collection and 0K, 2K, 6K, 10K, and 20K data points from IE tasks to train LLaMA-7B for the experiments. Figure 6 reveals the following insights: (1) Incorporating structured tasks for training may not enhance the model’s generalization ability on tasks without complex structures. Introducing structured tasks for tuning does not improve the model’s performance on MMLU, a benchmark without intricate input and output structures. (2) Structured tasks significantly impact the model’s generalization performance on tasks with complex output structures. Without structured tasks for training, the model’s performance on the NER task is 0 for both JsonTuning and TextTuning. However, the performance significantly improves when introducing only 2K data points from structured tasks for training. This highlights the importance of structured tasks for instruction-tuned models to generalize to tasks with complex output structures. (3) Structured tasks have a milder impact on the model’s generalization performance on tasks with complex input structures. Introducing an appropriate number of structured tasks can enhance the model’s performance on the NL2SQL task, which requires processing a structured database schema. This suggests that training the model with structured tasks aids in processing and understanding complex structures. In summary, the decision to use structured tasks for instruction tuning depends on the application scenarios. However, regardless of the scenario, JsonTuning consistently appears to be a superior method for instruction tuning compared to TextTuning.

Discussion Firstly, our primary claim is that integrating task structures within instruction tuning can enhance its generalization, robustness, and controllability. This claim is not restricted to any specific implementation. We choose JSON due to its robust support for a variety of data types, its straightforward and uniform syntax, and its widespread native support across numerous programming languages, particularly Python. However, other formats like YAML or XML are also viable options, provided that the task structures are explicitly represented.

Secondly, our goal is to provide a structured interface for LLMs. When provided solely with texts, the Json-tuned model may struggle to generate coherent outputs. However, if the textual input is encapsulated within a JSON structure, the

Model	Method	MMLU	BBH	NER [†]	RE [†]	EE	NL2SQL	Average	MMLU-Robustness
LLaMA-7B	Text w/ control info	42.01	34.22	35.02	14.54	0.17 / 0.00	11.40	22.88	39.31 ± 3.91
	Json w/ control info	44.69	37.08	43.47	15.28	3.49 / 7.33	16.40	27.06	44.61 ± 0.11
	Text w/o control info	43.11	32.48	37.61	14.33	1.35 / 0.00	8.60	22.80	38.82 ± 6.28
	Json w/o control info	43.94	35.42	48.08	16.91	0.59 / 0.26	12.40	26.20	43.47 ± 0.26
LLaMA2-7B	Text w/ control info	47.94	38.44	46.08	17.88	0.95 / 0.00	9.20	26.70	45.05 ± 3.67
	Json w/ control info	47.95	39.23	45.25	23.98	4.23 / 10.81	11.20	29.19	47.86 ± 0.17
	Text w/o control info	46.36	37.89	41.66	20.74	0.55 / 0.00	10.80	26.29	42.46 ± 4.76
	Json w/o control info	47.81	38.71	44.42	22.38	3.53 / 6.39	12.80	28.51	45.76 ± 0.19

Table 6: Ablation results for LLaMA-7B and LLaMA2-7B concerning control information. Tasks marked with † are seen during training and evaluated with unseen datasets.

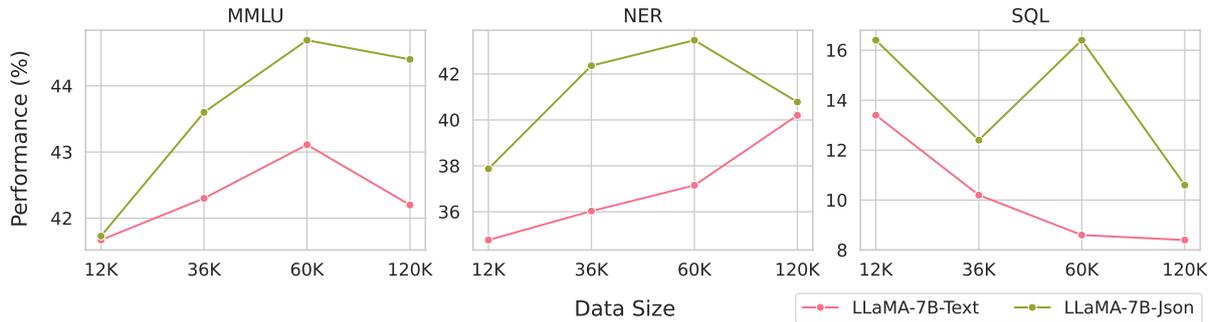


Figure 5: Performance of LLaMA-7B trained using JsonTuning and TextTuning across varying data sizes.

model can effectively process the query. JsonTuning modifies the interface without compromising the model’s inherent capabilities. While the textual interface may appear more intuitive, it may not be the most effective method for human-model interaction. Introducing a structured interface can provide significant flexibility in various deployment scenarios. JsonTuning offers several advantages for applications: (1) *Performance*: In complex scenarios, articulating intricate tasks using natural language can be challenging. Conversely, a JSON-structured interface enables users to organize information in a clear and structured manner, significantly enhancing the model’s understanding of the task and, consequently, its performance. (2) *Robustness*: The method’s resilience to varying prompts significantly reduces the need for users to test multiple prompts for optimal performance. Furthermore, as backend models are updated or changed, the optimal prompt for natural language texts may vary, necessitating adjustments for each model version. In contrast, Json-tuned models may not require prompt modifications even when the backbone model is updated, considerably reducing user effort. (3) *Controllability*: The output control component allows users to accurately specify and easily parse the output, which can be challeng-

ing when using natural language texts. Users only need to comprehend basic types such as string, object, and array in JSON Schema. To provide both textual and structured interfaces, we may conduct TextTuning and JsonTuning concurrently.

Finally, our methodology is designed to be flexible and adaptable to a variety of structures. The design of the JSON structure can adhere to the following key principles: (1) *Task Identification*: Clearly define the essential task elements, their interrelations, and the anticipated outputs. These aspects should be explicitly represented within the JSON structure. (2) *Simplicity*: Strive to minimize complexity by avoiding unnecessarily nested JSON structures whenever possible.

B Datasets of Information Extraction Tasks

Table 7 reports the training and evaluation datasets of information extraction tasks.

C Additional case studies on controllability

Table 8 presents additional case studies on controllability.

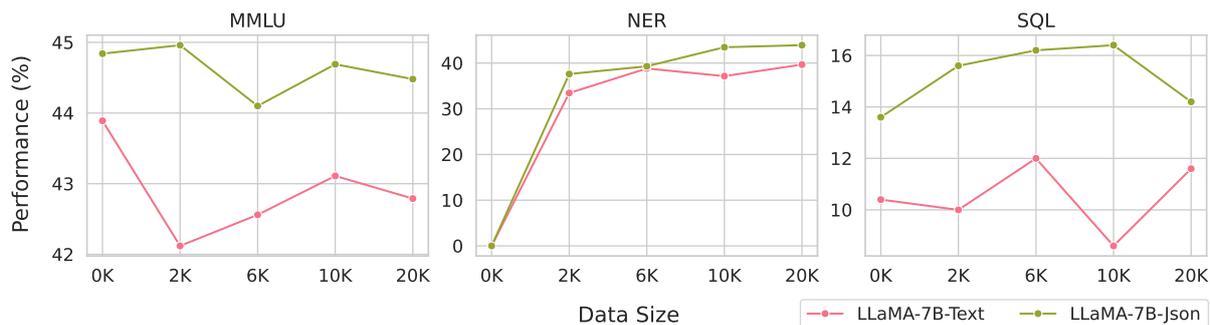


Figure 6: Performance of LLaMA-7B trained using JsonTuning and TextTuning with different numbers of examples of structured tasks.

D Introduction to JSON and Its Utilization

D.1 JSON Data Types and Syntax

JSON data is represented using a combination of the following data types:

- **Object:** An unordered collection of key-value pairs, enclosed in curly braces `{}`. The keys are strings, and the values can be any of the JSON data types.
- **Array:** An ordered list of values, enclosed in square brackets `[]`. The values can be any of the JSON data types.
- **String:** A sequence of Unicode characters, enclosed in double quotes.
- **Number:** A numeric value, which can be an integer or a floating-point number.
- **Boolean:** A value that is either true or false.
- **Null:** A special keyword denoting a null value.

In this paper, we focus on the object, array, and string types, as the number and boolean types can be represented as the string type for simplicity. By combining these simple data types, JSON can represent various structured data. This flexibility allows language models that understand basic data types to potentially generalize to more complex structures.

D.2 JSON Schema

JSON Schema employs a JSON-based format for defining the structure of JSON data, specifying properties like data types, required fields, and permissible values for JSON objects. It uses many keywords to define and validate JSON data. In this

paper, we use the following keywords to construct the control information C :

- **type:** Specifies the data type of a JSON value, such as object, array, and string.
- **description:** Provides explanations and clarifications about the purpose and constraints of a specific element or property.
- **items:** Defines the elements of an array and their data types.
- **properties:** Describes the properties of an object, including their data types and constraints.

We may introduce more keywords to further improve the model’s controllability in the future.

E Named Entity Recognition and Relation Extraction Training Prompts

We create prompts for both the named entity recognition (NER) and relation extraction (RE) tasks, as shown in Figures 7-9. For the RE task, we develop two sets of prompts: one for datasets with entity categories and another for datasets without entity categories. Each prompt comprises an input template and an output template, which are highlighted in blue and orange, respectively.

F Evaluation Prompts and Examples

For MMLU and BBH, we utilize the prompts from the Flan2022 collection designed for question answering². For other evaluation tasks, we create prompts based on their respective task components and definitions. Each prompt includes an input template and an output template, highlighted in blue

²For more details, see <https://github.com/google-research/FLAN/blob/main/flan/v2/templates.py>.

```

Prompt 1: [definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nentities,
{entities}]
Prompt 2: [definition: {definition}\nentity categories: {entity categories}\ntext: {text}\nentities,
{entities}]
Prompt 3: [text: {text}\ndefinition: {definition}\nentity categories: {entity categories}\nentities,
{entities}]
Prompt 4: [text: {text}\nentity categories: {entity categories}\ndefinition: {definition}\nentities,
{entities}]
Prompt 5: [entity categories: {entity categories}\ntext: {text}\ndefinition: {definition}\nentities,
{entities}]
Prompt 6: [entity categories: {entity categories}\ndefinition: {definition}\ntext: {text}\nentities,
{entities}]
Prompt 7: [{definition}\ntext: {text}\nentity categories: {entity categories}\nentities, {entities}]
Prompt 8: [{definition}\nentity categories: {entity categories}\ntext: {text}\nentities, {entities}]
Prompt 9: [text: {text}\nentity categories: {entity categories}\n{definition}\nentities, {entities}]
Prompt 10: [entity categories: {entity categories}\ntext: {text}\n{definition}\nentities, {entities}]

```

Figure 7: NER training prompts.

and orange, respectively. Further details can be found in the subsequent sections.

F.1 Generalization

All tasks employ a single prompt for evaluation, except for the RE task. The RE task utilizes two prompts: one for datasets with entity categories and another for datasets without entity categories. The prompts and examples are presented in Figures 10-18.

F.2 Robustness

For evaluation, we employ 10 prompts for the MMLU benchmark, the NER task, and the NL2SQL task. Prompts are demonstrated in Figure 19, Figure 7, and Figure 20, respectively.

Prompt 1: [definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nrelations: {relations}\nrelational triplets: {relational triplets}]

Prompt 2: [definition: {definition}\nentity categories: {entity categories}\nrelations: {relations}\ntext: {text}\nrelational triplets: {relational triplets}]

Prompt 3: [definition: {definition}\ntext: {text}\nrelations: {relations}\nentity categories: {entity categories}\nrelational triplets: {relational triplets}]

Prompt 4: [definition: {definition}\nrelations: {relations}\nentity categories: {entity categories}\ntext: {text}\nrelational triplets: {relational triplets}]

Prompt 5: [text: {text}\ndefinition: {definition}\nentity categories: {entity categories}\nrelations: {relations}\nrelational triplets: {relational triplets}]

Prompt 6: [text: {text}\nentity categories: {entity categories}\nrelations: {relations}\ndefinition: {definition}\nrelational triplets: {relational triplets}]

Prompt 7: [text: {text}\ndefinition: {definition}\nrelations: {relations}\nentity categories: {entity categories}\nrelational triplets: {relational triplets}]

Prompt 8: [text: {text}\nrelations: {relations}\nentity categories: {entity categories}\ndefinition: {definition}\nrelational triplets: {relational triplets}]

Prompt 9: [entity categories: {entity categories}\nrelations: {relations}\ntext: {text}\ndefinition: {definition}\nrelational triplets: {relational triplets}]

Prompt 10: [relations: {relations}\nentity categories: {entity categories}\ndefinition: {definition}\ntext: {text}\nrelational triplets: {relational triplets}]

Figure 8: RE (with entity categories) training prompts.

Prompt 1: [definition: {definition}\ntext: {text}\nrelations: {relations}\nrelational triplets: {relational triplets}]

Prompt 2: [definition: {definition}\nrelations: {relations}\ntext: {text}\nrelational triplets: {relational triplets}]

Prompt 3: [text: {text}\ndefinition: {definition}\nrelations: {relations}\nrelational triplets: {relational triplets}]

Prompt 4: [text: {text}\nrelations: {relations}\ndefinition: {definition}\nrelational triplets: {relational triplets}]

Prompt 5: [relations: {relations}\ntext: {text}\ndefinition: {definition}\nrelational triplets: {relational triplets}]

Prompt 6: [relations: {relations}\ndefinition: {definition}\ntext: {text}\nrelational triplets: {relational triplets}]

Prompt 7: [{definition}\ntext: {text}\nrelations: {relations}\nrelational triplets: {relational triplets}]

Prompt 8: [{definition}\nrelations: {relations}\ntext: {text}\nrelational triplets: {relational triplets}]

Prompt 9: [text: {text}\nrelations: {relations}\n{definition}\nrelational triplets: {relational triplets}]

Prompt 10: [relations: {relations}\ntext: {text}\n{definition}\nrelational triplets: {relational triplets}]

Figure 9: RE (without entity categories) training prompts.

Task	Dataset	Schema
Training		
NER	ACE2004 (Walker et al., 2006)	7
	ACE2005 (Walker et al., 2006)	7
	broad_twitter_corpus (Derczynski et al., 2016)	3
	CoNLL2003 (Tjong Kim Sang and De Meulder, 2003)	4
	multiNERD (Tedeschi and Navigli, 2022)	16
	Ontonotes (Hovy et al., 2006)	18
	polyglot-NER (Al-Rfou et al., 2015)	3
	tweetNER7 (Ushio et al., 2022)	7
	wikiann (Pan et al., 2017)	3
	wikineural (Tedeschi et al., 2021)	3
	AnatEM (Pyysalo and Ananiadou, 2013)	1
	bc2gm (Kocaman and Talby, 2021)	1
	bc4chemd (Krallinger et al., 2015)	1
	bc5cd (Li et al., 2016)	2
	FabNER (Kumar and Starly, 2021)	12
	FindVehicle (Guan et al., 2023)	21
	HarveyNER (Chen et al., 2022)	4
	ncbi-disease (Dogan et al., 2014)	1
RE	GIDS (Nayak et al., 2021)	4
	kbp37 (Zhang and Wang, 2015)	18
	NYT (Riedel et al., 2010)	24
	NYT11 HRL (Takanobu et al., 2019)	12
	SciERC (Luan et al., 2018)	7
	semeval RE (Hendrickx et al., 2010)	10
Evaluation		
NER	CrossNER_AI (Liu et al., 2021)	14
	CrossNER_literature (Liu et al., 2021)	12
	CrossNER_music (Liu et al., 2021)	13
	CrossNER_politics (Liu et al., 2021)	9
	CrossNER_science (Liu et al., 2021)	17
RE	CoNLL2004 (Roth and Yih, 2004)	5
	FewRel (Han et al., 2018)	5
EE	ACE2005 (Walker et al., 2006)	33(22)
	CASIE (Satyapanich et al., 2020)	5(26)
	PHEE (Sun et al., 2022)	2(16)

Table 7: Information Extraction (IE) datasets utilized for training and evaluation. |Schema| represents the number of entity categories in the named entity recognition (NER) task, the number of relations in the relation extraction (RE) task, and the number of event categories (outside the parenthesis) along with the number of argument categories (inside the parenthesis) in the event extraction (EE) task.

Method	Input	Prediction
Task: Translate a sentence to multiple languages		
Text	Sentence: How are you? Translate the sentence to multiple languages: French and Spanish. Translations:	Comment allez-vous?
Json	<code>{ "input": { "sentence": "How are you?", "instruction": "Sentence: {sentence}. Translate the sentence to two languages: French and Spanish. Translations: {translations}" }, "output control": { "translations": { "type": "object", "properties": { "French": { "type": "string", "description": "The French translation of the sentence." }, "Spanish": { "type": "string", "description": "The Spanish translation of the sentence." } } } }</code>	<code>{ "translations": { "French": "Comment allez-vous?", "Spanish": "C\u00f3mo est\u00e1s?" } }</code>
Task: Joke generation with humor style		
Text	Generate the joke with a specific humor style (e.g., pun, sarcasm):	What do you call a deer with no eyes? No idea.
Json	<code>{ "input": { "instruction": "Generate the joke with a specific humor style (e.g., pun, sarcasm): {humor style} {joke}" }, "output control": { "humor style": { "type": "string" }, "joke": { "type": "string" } } }</code>	<code>{ "humor style": "pun", "joke": "What did the pirate say when he was given a piece of paper?" }</code>
Task: Sentiment analysis with emotion		
Text	Given the sentiment labels: positive, negative, and neutral, and text: I just won a lottery! Life is amazing!, provide the sentiment label and emotion associated with the text. Sentiment label and emotion:	positive
Json	<code>{ "input": { "text": "I just won a lottery! Life is amazing!", "sentiment labels": ["positive", "negative", "neutral"], "instruction": "Given the sentiment labels: {sentiment labels} and text: {text}, provide the sentiment label and emotion associated with the text. Sentiment label and emotion: {sentiment label} {emotion}" }, "output control": { "sentiment label": { "type": "string" }, "emotion": { "type": "string" } } }</code>	<code>{ "sentiment label": "positive", "emotion": "joy" }</code>

Table 8: Case studies focusing on controllability. Each example displays its input along with the model’s prediction.

Prompt: `[{question}\n{options_}\nAnswer: {answer}]`

TextTuning Example:
Input: The following is a multiple choice question about global facts.\nControlling for inflation and PPP-adjustment, about how much did GDP per capita increase from 1950 to 2016 in Japan?
Options:\n(A) by 5 fold\n(B) by 10 fold\n(C) by 15 fold\n(D) by 20 fold\nAnswer:
Output: (C)

JsonTuning Example:
Input: `{ "input": { "question": "The following is a multiple choice question about global facts.\nControlling for inflation and PPP-adjustment, about how much did GDP per capita increase from 1950 to 2016 in Japan?", "options_": "Options:\n(A) by 5 fold\n(B) by 10 fold\n(C) by 15 fold\n(D) by 20 fold", "instruction": "{question}\n{options_}\nAnswer: {answer}" }, "output control": { "answer": { "type": "string" } } }`
Output: `{ "answer": "(C)" }`

Figure 10: MMLU evaluation prompt and examples.

Prompt: [Q: {question}\nA:, {answer}]

TextTuning Example:

Input: Q: $((-1 + 2 + 9 * 5) - (-2 + -4 + -4 * -7)) =$ \nA:

Output: 24

JsonTuning Example:

Input: { "input": { "question": "((-1 + 2 + 9 * 5) - (-2 + -4 + -4 * -7)) =", "instruction": "Q: {question}\nA: {answer}" }, "output control": { "answer": { "type": "string" } } }

Output: { "answer": "24" }

Figure 11: BBH evaluation prompt and examples.

Prompt: [definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nentities:, {entities}]

TextTuning Example:

Input: definition: Given a text and entity categories, your task is to scan the text and identify a list of named entities in it. Each entity contains an entity category and an entity span. An entity span refers to the specific portion of the text that represents an entity. An entity category refers to the category to which an entity belongs.\ntext: He also co-wrote Possible, which has been used as a theme song for the 2005 Southeast Asian Games.\nentity categories: location, event, country, band, person, song, musical artist, music genre, else, album, organization, award, musical instrument\nentities:

Output: [{"song", "Possible"}, {"event", "2005 Southeast Asian Games"}]

JsonTuning Example:

Input: { "input": { "definition": "Given a text and entity categories, your task is to scan the text and identify a list of named entities in it. Each entity contains an entity category and an entity span. An entity span refers to the specific portion of the text that represents an entity. An entity category refers to the category to which an entity belongs.", "text": "He also co-wrote Possible, which has been used as a theme song for the 2005 Southeast Asian Games.", "entity categories": ["location", "event", "country", "band", "person", "song", "musical artist", "music genre", "else", "album", "organization", "award", "musical instrument"], "instruction": "definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nentities: {entities}" }, "output control": { "entities": { "type": "array", "items": { "type": "object", "properties": { "entity category": { "type": "string", "description": "The entity category should be one of the entity categories provided in the input." }, "entity span": { "type": "string", "description": "The entity span should be a continuous span in the text provided in the input." } } } } } }

Output: { "entities": [{ "entity category": "song", "entity span": "Possible" }, { "entity category": "event", "entity span": "2005 Southeast Asian Games" }] }

Figure 12: NER evaluation prompt and examples.

Prompt: [definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nrelations: {relations}\nrelational triplets: {relational triplets}]

TextTuning Example:

Input: definition: Given a text, entity categories, and relations, your goal is to scan the text and identify a list of relational triplets in it. Each relational triplet contains a head entity category, a head entity span, a relation, a tail entity category, and a tail entity span. The head entity is the subject from which the relation originates. The relation represents the specific relation between the head entity and the tail entity. The tail entity is the object which the relation points. An entity span refers to the specific portion of the text that represents an entity. An entity category refers to the category to which an entity belongs.\ntext: In 1822, the 19th president of the United States, Rutherford B. Hayes, was born in Delaware, Ohio. \nentity categories: Organization, Location, People\nrelations: Kill, Work for, Located in, Live in, Organization based in\nrelational triplets:

Output: [[“People”, “Rutherford B. Hayes”, “Live in”, “Location”, “Delaware, Ohio”]]

JsonTuning Example:

Input: { “input”: { “definition”: “Given a text, entity categories, and relations, your goal is to scan the text and identify a list of relational triplets in it. Each relational triplet contains a head entity category, a head entity span, a relation, a tail entity category, and a tail entity span. The head entity is the subject from which the relation originates. The relation represents the specific relation between the head entity and the tail entity. The tail entity is the object which the relation points. An entity span refers to the specific portion of the text that represents an entity. An entity category refers to the category to which an entity belongs.”, “text”: “In 1822, the 19th president of the United States, Rutherford B. Hayes, was born in Delaware, Ohio.”, “entity categories”: [“Organization”, “Location”, “People”], “relations”: [“Kill”, “Work for”, “Located in”, “Live in”, “Organization based in”], “instruction”: “definition: {definition}\ntext: {text}\nentity categories: {entity categories}\nrelations: {relations}\nrelational triplets: {relational triplets}” }, “output control”: { “relational triplets”: { “type”: “array”, “items”: { “type”: “object”, “properties”: { “head entity category”: { “type”: “string”, “description”: “The head entity category should be one of the entity categories provided in the input.” }, “head entity span”: { “type”: “string”, “description”: “The head entity span should be a continuous span in the text provided in the input.” }, “relation”: { “type”: “string”, “description”: “The relation should be one of the relations provided in the input.” }, “tail entity category”: { “type”: “string”, “description”: “The tail entity category should be one of the entity categories provided in the input.” }, “tail entity span”: { “type”: “string”, “description”: “The tail entity span should be a continuous span in the text provided in the input.” } } } } }

Output: { “relational triplets”: [{ “head entity category”: “People”, “head entity span”: “Rutherford B. Hayes”, “relation”: “Live in”, “tail entity category”: “Location”, “tail entity span”: “Delaware, Ohio” }] }

Figure 13: RE (with entity categories) evaluation prompt and examples.

Prompt: [definition: {definition}\ntext: {text}\nrelations: {relations}\nrelational triplets: {relational triplets}]

TextTuning Example:

Input: definition: Given a text and relations, you are required to scan the text and identify a list of relational triplets in it. Each relational triplet contains a head entity span, a relation, and a tail entity span. The head entity is the subject from which the relation originates. The relation represents the specific relation between the head entity and the tail entity. The tail entity is the object which the relation points. An entity span refers to the specific portion of the text that represents an entity.\ntext: The Peasants is a novel written by Nobel Prize-winning Polish author Wadysaw Reymont in four parts between 1904 and 1909.\nrelations: place served by transport hub, winner, field of work, location of formation, occupant\nrelational triplets:

Output: [[“Nobel Prize”, “winner”, “Wadysaw Reymont”]]

JsonTuning Example:

Input: {“input”: { “definition”: “Given a text and relations, you are required to scan the text and identify a list of relational triplets in it. Each relational triplet contains a head entity span, a relation, and a tail entity span. The head entity is the subject from which the relation originates. The relation represents the specific relation between the head entity and the tail entity. The tail entity is the object which the relation points. An entity span refers to the specific portion of the text that represents an entity.”, “text”: “The Peasants is a novel written by Nobel Prize-winning Polish author Wadysaw Reymont in four parts between 1904 and 1909.”, “relations”: [“place served by transport hub”, “winner”, “field of work”, “location of formation”, “occupant”], “instruction”: “{definition}\ntext: {text}\nrelations: {relations}\nrelational triplets: {relational triplets}” }, “output control”: { “relational triplets”: { “type”: “array”, “items”: { “type”: “object”, “properties”: { “head entity span”: { “type”: “string”, “description”: “The head entity span should be a continuous span in the text provided in the input.” }, “relation”: { “type”: “string”, “description”: “The relation should be one of the relations provided in the input.” }, “tail entity span”: { “type”: “string”, “description”: “The tail entity span should be a continuous span in the text provided in the input.” } } } } }

Output: { “relational triplets”: [{ “head entity span”: “Nobel Prize”, “relation”: “winner”, “tail entity span”: “Wadysaw Reymont” }] }

Figure 14: RE (without entity categories) evaluation prompt and examples.

Prompt: [definition: {definition}\ntext: {text}\nevent categories: {event categories}\nargument categories: {argument categories}\nevents:, {events}]

TextTuning Example:

Input: definition: Given a text, event categories, and argument categories, you are expected to scan the text and identify a list of events in it. Each event contains an event category, an event trigger, and a list of arguments. Each argument contains an argument category and an argument span. An event category represents the type of an event. An event trigger is the word or phrase in the text that explicitly denotes the occurrence of an event. Arguments are entities associated with an event and play specific roles or functions in relation to the event. An argument span refers to the specific portion of the text that represents an argument. An argument category refers to the category to which an argument belongs.\ntext: Until Basra, U.S. and British troops had encountered little resistance, even when they seized nearby Umm Qasr, and moved to secure key oil fields.\nevent categories: transfer money, start organization, extradite, meet, appeal, attack, convict, born, execute, transport, release parole, merge organization, sentence, divorce, end position, end organization, transfer ownership, start position, injure, sue, die, trial hearing, marry, nominate, charge indict, elect, declare bankruptcy, phone write, acquit, arrest jail, pardon, demonstrate, fine\nargument categories: instrument, vehicle, agent, seller, place, beneficiary, organization, destination, plaintiff, person, giver, recipient, victim, target, defendant, origin, prosecutor, entity, attacker, artifact, buyer, adjudicator\nevents:",

Output: [[“attack”, “seized”, [[“attacker”, “troops”], [“place”, “Umm Qasr”]]]]

Figure 15: EE evaluation prompt and a TextTuning example.

Prompt: [definition: {definition}\nquestion: {question}\ndatabase schema: {database schema}\nSQL query:, {SQL query}]

TextTuning Example:

Input: definition: Given a question and database schema that consists of table names and column names in the database, the text-to-SQL parsing task aims to translate the natural language question to a sql query that can be executed on the database to produce answers.\nquestion: List the title of all cartoons in alphabetical order.\ndatabase schema: Table: tv_channel; Columns: id, series_name, country, language, content, pixel_aspect_ratio_par, hight_definition_tv, pay_per_view_ppv, package_option. Table: tv_series; Columns: id, episode, air_date, rating, share, 18_49_rating_share, viewers_m, weekly_rank, channel. Table: cartoon; Columns: id, title, directed_by, written_by, original_air_date, production_code, channel\nSQL query:

Output: select title from cartoon order by title

JsonTuning Example:

Input: {"input": {"definition": "Given a 'question' and 'database schema' that consists of table names and column names in the database, the text-to-SQL parsing task aims to translate the natural language question to a sql query that can be executed on the database to produce answers.", "question": "List the title of all cartoons in alphabetical order.", "database schema": [{"table name": "tv_channel", "column names": ["id", "series_name", "country", "language", "content", "pixel_aspect_ratio_par", "hight_definition_tv", "pay_per_view_ppv", "package_option"]}, {"table name": "tv_series", "column names": ["id", "episode", "air_date", "rating", "share", "18_49_rating_share", "viewers_m", "weekly_rank", "channel"]}, {"table name": "cartoon", "column names": ["id", "title", "directed_by", "written_by", "original_air_date", "production_code", "channel"]}]}, {"output control": {"SQL query": {"type": "string"}}}

Output: {"SQL query": "select title from cartoon order by title" }

Figure 17: NL2SQL evaluation prompt and examples.

Prompt: [Q: {question}\nA:, {answer}]

TextTuning Example:

Input: Q: Who created the Superman cartoon character?\nA:

Output: Superman was created by Jerry Siegel and Joe Shuster in 1938.

JsonTuning Example:

Input: {"input": {"question": "Who created the Superman cartoon character?", "instruction": "Q: {question}\nA: {answer}"}, {"output control": {"answer": {"type": "string"}}}

Output: {"answer": "Superman was created by Jerry Siegel and Joe Shuster in 1938."}

Figure 18: AlpacaEval evaluation prompt and examples.

Prompt 1: [{question}\n{options_}\nAnswer:, {answer}]
 Prompt 2: [{question}\n\n{options_}\nAnswer:, {answer}]
 Prompt 3: [{question}\n{options_}, {answer}]
 Prompt 4: [Q: {question}\n\n{options_}\nA:, {answer}]
 Prompt 5: [Answer the following question: {question}\n\n{options_}\nAnswer:, {answer}]
 Prompt 6: [{options_}\n\n{question}\nAnswer:, {answer}]
 Prompt 7: [{options_}\nQ: {question}\nA:, {answer}]
 Prompt 8: [{question}\n\n{options_}\nThe answer is:, {answer}]
 Prompt 9: [{options_}\nGiven those answer options, answer the question: {question}\nAnswer:, {answer}]
 Prompt 10: [Q: {question}\n\n{options_}\nThe answer is:, {answer}]

Figure 19: MMLU robustness evaluation prompts.

Prompt 1: [definition: {definition}\nquestion: {question}\ndatabase schema: {database schema}\nSQL query:, {SQL query}]
 Prompt 2: [definition: {definition}\ndatabase schema: {database schema}\nquestion: {question}\nSQL query:, {SQL query}]
 Prompt 3: [question: {question}\ndefinition: {definition}\ndatabase schema: {database schema}\nSQL query:, {SQL query}]
 Prompt 4: [question: {question}\ndatabase schema: {database schema}\ndefinition: {definition}\nSQL query:, {SQL query}]
 Prompt 5: [database schema: {database schema}\nquestion: {question}\ndefinition: {definition}\nSQL query:, {SQL query}]
 Prompt 6: [database schema: {database schema}\ndefinition: {definition}\nquestion: {question}\nSQL query:, {SQL query}]
 Prompt 7: [{definition}\nquestion: {question}\ndatabase schema: {database schema}\nSQL query:, {SQL query}]
 Prompt 8: [{definition}\ndatabase schema: {database schema}\nquestion: {question}\nSQL query:, {SQL query}]
 Prompt 9: [question: {question}\ndatabase schema: {database schema}\n{definition}\nSQL query:, {SQL query}]
 Prompt 10: [database schema: {database schema}\nquestion: {question}\n{definition}\nSQL query:, {SQL query}]

Figure 20: NL2SQL robustness evaluation prompts.