

# Robustness of Learning from Task Instructions

Jiasheng Gu<sup>†</sup>, Hongyu Zhao<sup>§</sup>, Hanzi Xu<sup>\*</sup>, Liangyu Nie<sup>◇</sup>, Hongyuan Mei<sup>¶</sup>, Wenpeng Yin<sup>‡</sup>

<sup>†</sup>University of Southern California; <sup>§</sup>University of Chicago; <sup>\*</sup>Temple University

<sup>◇</sup>University of Birmingham; <sup>¶</sup>Toyota Technological Institute at Chicago; <sup>‡</sup>Penn State University

gujiashe@usc.edu; hongyuz@uchicago.edu; abdielnie@gmail.com

hanzi.xu@temple.edu; hongyuan@ttic.edu; wenpeng@psu.edu

## Abstract

Traditional supervised learning mostly works on individual tasks and requires training on a large set of task-specific examples. This paradigm seriously hinders the development of task generalization since preparing a task-specific example set is costly. To build a system that can quickly and easily generalize to new tasks, textual instructions have been adopted as an emerging trend of supervision recently. These instructions give the model the definition of the task and allow the model to output the appropriate answer based on the instructions and inputs. However, task instructions are often expressed in different forms, which can be interpreted from two threads (Lou et al., 2023): firstly, some instructions are short sentences and are pretrained language model (PLM) oriented, such as prompts, while other instructions are paragraphs and are human-oriented, such as those annotation guidelines in Amazon MTurk; secondly, different end-users very likely explain the same task with instructions of different textual expressions. A robust system for task generalization should be able to handle any new tasks regardless of the variability of instructions.

However, the system robustness in dealing with a wide range of instructions is still unexplored. This work investigates the system robustness when the instructions of new tasks are (i) manipulated, (ii) paraphrased, or (iii) from different levels of conciseness. To our knowledge, this is the first work that systematically studies how robust a PLM is when it is supervised by instructions with different factors of variability.<sup>1</sup>

## 1 Introduction

One goal of AI is to achieve cross-task generalization, i.e., the model can handle new tasks after being trained on existing tasks. To solve target tasks, the system needs task-specific supervision.

<sup>1</sup>All code and data: <https://github.com/jiashenggu/Tk-Instruct/tree/main>

In the conventional machine learning paradigm, task-specific supervision means a large set of labeled examples particularly for this task. Despite the progress of conventional machine learning, collecting task-specific examples of large volumes is costly and even infeasible if we want the system to cope with new tasks immediately. Therefore, another form of task-specific supervision—task instructions—is attracting increasing attention in the community. Task instructions, commonly a piece of text, describe the task semantics and tell the model how to approach the expected outputs for the inputs. A couple of recent studies (Mishra et al., 2022a; Yin et al., 2022; Wang et al., 2022; Sanh et al., 2022) have shown promising progress in cross-task generalization with only instructions as supervision.

In the real world, task instructions are often presented in different forms, which can be interpreted from two dimensions: on the one hand, as Lou et al. (2023) defined, *instructions have different degrees of conciseness and were created for different audiences*—some instructions are short sentences and are **PLM-oriented**, such as prompts like “translate from English to French”, while other instructions are paragraphs and are **human-oriented**, e.g., those annotation guidelines in Amazon MTurk; on the other hand, the same task is likely to be explained by different end-users with instructions of varying textual expressions. For example, “translate from English to French” can also be expressed as “given an English sentence, translate it into French”. Therefore, a robust system for task generalization is expected to handle any new tasks regardless of the variability of instructions.

However, despite the limited earlier work studying the PLM robustness in terms of prompts (Webson and Pavlick, 2022; Khashabi et al., 2022), it is still unclear how robust a pretrained instruction learning system can handle broader types of instructions, including human-oriented as well as

PLM-oriented ones, and their cross-type generalization. This work focuses on human-oriented instructions and explores the system robustness when the instructions of new tasks are (i) manipulated, e.g., words removal, words replacement, sentence shuffle, etc., (ii) paraphrased with dramatic surface change, or (iii) from different levels of conciseness, e.g., training tasks have human-oriented instructions while testing tasks are equipped with PLM-oriented prompts.

To get the answers for the above three threads of instruction variability, we finetune a T5 (Rafael et al., 2020) model on the *train* of SUPER-NATURALINSTRU (Wang et al., 2022), a dataset with the largest number of NLP tasks and each task is described by an Amazon-Mturk-style instruction, and test it on (i) the *test* of SUPER-NATURALINSTRU, where we modify the original instructions with word-level, sentence-level and instruction-level perturbations, (ii) a newly crowd-sourced dataset, Para-Instructions, with 23 unseen tasks and each has 4 distinct human-oriented instructions in the same format as the *test* of SUPER-NATURALINSTRU, and (iii) a new dataset DIVERSEPROMPT (Honovich et al., 2022), in which each task has prompt-style instructions.

To our knowledge, this is the first work that systematically studies how robust a PLM is when it is supervised by instructions with different factors of variability. Our experiments show that:

- The presence of demonstrations minimizes the impact of instruction perturbations and paraphrasing, as they typically have an unpredictable effect on performance. Demonstrations play a crucial role in ensuring system robustness, even when the instruction provided is a random text fragment, i.e., “random instruction + positive examples” mostly outperforms “correct instruction only”;
- When demonstrations are absent, applying a model that learns from training tasks with human-oriented instructions to unseen tasks with PLM-oriented prompts works poorly;
- So far, we still lack a method to effectively learn instruction to solve tasks in the absence of demonstration.

## 2 Related work

Our work is related to prompts (a special case of instruction learning), cross-task generalization, and

the study of robustness under adversarial attack.

**Prompt as instruction.** Brown et al. (2020a) attempted to unleash the power of PLMs, such as GPT-3 (Brown et al., 2020b), through prompts. Since they discovered prompts, many studies have followed that direction in various ways, with robustness being an aspect that cannot be ignored. Liu et al. (2021) introduced the effectiveness of prompt learning and the wide application prospects. Min et al. (2022) explored how the model learns prompts and what aspects of the prompt contribute to the final task performance. Webson and Pavlick (2022) tried to figure out whether prompt-based models truly understand the meaning of their prompts. Zhao et al. (2021) investigated ways to improve the stability of prompt-based models. Bach et al. (2022) created a platform for the production, distribution, and use of natural language prompts. Despite the large number of papers on prompts, little attention has been paid to studying the robustness of systems under prompt learning. Our work has a broader vision that considers the robustness in terms of PLM-oriented prompts as well as more natural human-oriented instructions.

**Instruction-driven task generalization.** Recent work has worked on using the information in task instructions to allow PLMs to respond to inputs that match the task instructions after they have understood task instructions. Efrat and Levy (2020) tested GPT-2 (Radford et al., 2019) to understand real-world MTurk instructions to annotate some popular datasets, and concluded that GPT-2 works poorly. Mishra et al. (2022b) collected more than 60 distinct NLP tasks with MTurk instructions consisting of items like title, definition, things to avoid, etc., and claimed that BART (Lewis et al., 2020) and GPT-3 benefit from instructions to generalize across tasks. Wang et al. (2022) further extended this dataset to cover over 1.6k cross-lingual tasks, each with a piece of natural language instruction. More work in this thread can be found in (Lou et al., 2023), which provides a comprehensive survey for various instruction learning. The robustness studied in this work is also based on instruction-driven task generalization; however, our emphasis lies on analyzing the variability in system performance as opposed to actively enhancing overall system performance.

**Robustness under adversarial attack.** There is a lot of work on adversarial attacks in NLP systems,

	Definition	In this task, you will be shown an incorrect English sentence. You need to generate a corrected form of the input sentence.
Positive	input	The car’s wheel are loose.
	output	The car’s wheel is loose.
	explanation	The instance of are is replaced by the word is, because the wheel is a singular word. This makes the sentence grammatically correct.
Negative	input	This way is the way to go.
	output	This way may be the way to go.
	explanation	The example does not correct the misuse of the word way. Instead, it should shorten the sentence to: this is the way to go.
Instances	input	For example , nobody is going to ask his personal doctor which he sees when he has a flue if he can also do a heart surgery or transplant organs .
	output	For example , nobody is going to ask his personal doctor , which he sees when he has a flu , if he can also do a heart surgery or transplant organs .

Table 1: A task sample (task1557 jfleg grammar error correction) from SUPER-NATURALINSTRU: each task has a definition, a couple of positive examples, a couple of negative examples, and a large set of instances.

metric	Task Category	Abbr.
Exact Match	Textual Entailment	TE
	Cause Effect Classification	CEC
	Coreference Resolution	CR
	Dialogue Act Recognition	DAR
	Answerability Classification	AC
	Word Analogy	WA
ROUGE-L	Overlap Extraction	OE
	Keyword Tagging	KT
	Question Rewriting	QR
	Title Generation	TG
	Data to Text	DTT
	Grammar Error Correction	GEC

Table 2: 12 task categories in the *test* of SUPER-NATURALINSTRU and their official evaluation metrics.

and adversarial attacks are often used as a way to enhance robustness (Zhang et al., 2020; Goel et al., 2021; Wang et al., 2020; Morris et al., 2020; Raman et al., 2023). Adversarial perturbations are considered to be worst-case scenarios that do not often occur in the real world and represent a very specific kind of noise. To generate valid adversarial examples, most attack methods require access to the structure, internal weights, and hyperparameters of the NLP model. In addition, adversarial perturbations should not be perceived by humans. This is a serious challenge because even small changes in the text may be easily recognized by the user but ignored by the machine. In contrast, some minor changes may cause the machine to react drastically but are not important to humans. Our investigation encompasses not only deliberate perturbations, but also more natural discrepancies that arise between

the training instructions and test instructions.

### 3 Experiment

This work studies the robustness of learning from task instructions from three scenarios: (i) *the system robustness when dealing with instructions that are deliberately perturbed* (Section 3.1); (ii) *the system robustness when a task confronts diverse instructions that are written by different annotators* (Section 3.2); (iii) *the system robustness when it tests on tasks that have instructions with distinct abstractiveness* (Section 3.3). The following subsections present our experiments about the three dimensions of robustness study, respectively.

**The pretrained instruction learning model.** To test the robustness of a pretrained model that learned NLP tasks from instructions, we choose the T5 model (Raffel et al., 2020) with 3B and 11B parameters finetuned on the training set of SUPER-NATURALINSTRU (Wang et al., 2022).

SUPER-NATURALINSTRU consists of 1,616 NLP tasks in total and each task is explained by four items: Definition (a short paragraph describing the task semantics), Positive Examples (a couple of examples with inputs, correct outputs, and explanations), Negative Examples (a couple of examples with inputs, incorrect outputs, and explanations), and Instances (a large set of instances with inputs and correct outputs). One sample is shown in Table 1. As far as we know, SUPER-NATURALINSTRU is the largest dataset containing human-oriented task instructions, and T5 denotes one of the state-of-the-art PLM models (we do not

Original Instruction	In this task, you will be shown an incorrect English sentence. You need to generate a corrected form of the input sentence.	
word-level	Delete words	In this <del>task</del> , you <del>will</del> be <del>shown</del> an incorrect English sentence. You need to generate a corrected <del>form</del> of <del>the</del> input sentence.
	Delete stopwords	<del>In this</del> task, <del>you will be</del> shown <del>an</del> incorrect English sentence. <del>You</del> need <del>to</del> generate <del>a</del> corrected form <del>of the</del> input sentence.
	Insert words	In this , task, , you will be shown an incorrect ( English input sentence. You need to generate a corrected form of the English input sentence.
	Replace words	In this case, you will be shown an incorrect English language. You need to generate a corrected form of the input sentence.
	Shuffle words	a You generate English incorrect be need task, In the will this you sentence. shown to sentence. input an form of corrected
sent-level	Repeat sentences	In this task, you will be shown an incorrect English sentence. You need to generate a corrected form of the input sentence. <u>You need to generate a corrected form of the input sentence.</u>
	Shuffle sentences	You need to generate a corrected form of the input sentence. In this task, you will be shown an incorrect English sentence.
instruction-level shuffle	In this task, you are given a sentence and a question, you would be asked to create the answer which is contained in the sentence provided.	

Table 3: Examples of original instructions, word-level, sentence-level, and instruction-level perturbations.

consider GPT-3 due to the hardware constraints). As the data for training the instruction learning model, the *train* of SUPER-NATURALINSTRU has 1,462 tasks.

### 3.1 Robustness when instructions are perturbed

**Dataset.** This experiment uses the *test* of SUPER-NATURALINSTRU, which has 154 tasks (119 English tasks and 35 tasks in other languages). The robustness of the pretrained T5 model on *train* is tested on this *test* where instructions are perturbed. All *test* tasks are clustered into 12 categories with the official evaluation metrics either EXACTMATCH or ROUGE-L (Lin, 2004). Table 2 lists all the 12 categories and their respect evaluation metrics.

Note that Wang et al. (2022) concatenated Task name, Definition, Positive Examples and Negative Examples in various combinations as a long instruction, referred to as “Instruction with Examples” (INSTRU-W-EXAM). This paper uses the definition and two positive examples as INSTRU-W-EXAM configuration. In order to study the robustness when those examples are available or not, we report INSTRU-W-EXAM as well as INSTRU-WO-EXAM which uses only the Definition as instructions.

**Instruction perturbation.** To imitate how instructions are perturbed in the real world, we design various perturbation methods that try to change the

original instruction by its surface form or even semantics. Table 3 summarizes all the perturbation methods this work uses, including **word-level perturbations**, such as i) delete stop words; ii) randomly delete five to ten words; iii) insert masks at random positions and replace with words predicted by pretrained BERT (Devlin et al., 2019); iv) randomly replace words by the predictions of pretrained BERT, and v) shuffle the words in the instruction, and **sentence-level perturbations**, such as i) repeating a random sentence, and ii) shuffle the sentences in the instruction, and **instruction-level perturbation**, where we try to replace the original task definition with the definition of a randomly chosen task.

Spacy (Schmitt et al., 2019) is used to remove stop words from instructions and to split instructions into separate sentences, and BERT is used to insert and replace words in instructions.

**Results.** Each metric was run three times, and the mean and standard deviation of the three runs were reported. Tables 4&10 presents the performance on English and non-English subsets of SUPER-NATURALINSTRU’s *test*, respectively. Each table summarizes the perturbation results by word-level, sentence-level, and instruction-level for both the INSTRU-WO-EXAM and INSTRU-W-EXAM setups. Our goal is to study i) the effects of different levels of perturbations; ii) the effects of including positive examples in the instructions; iii) whether there are any distinct phenomena between English tasks and



metric		AC	CEC	CR	DAR	WA	TE
EXACTMATCH	original instru.	49.8±0.6	50.0±0.8	36.1±0.5	43.2±0.8	19.7±0.7	40.1±0.8
		63.6±1.2	48.2±3.0	48.1±0.9	57.0±1.2	51.8±0.4	46.6±0.5
	delete stopwords	44.2±0.4	45.0±1.2	32.3±0.7	28.7±0.9	7.60±0.1	31.6±0.5
		63.2±0.6	49.0±0.5	49.1±2.2	49.9±1.7	50.5±1.9	44.5±0.6
	delete [5,10] words	45.3±3.3	40.8±4.3	32.5±0.4	34.0±2.6	18.6±2.6	34.9±3.1
		61.1±2.0	50.7±1.2	49.5±1.3	52.9±4.5	50.8±1.3	44.0±1.4
	insert [5,10] words	44.5±1.9	40.3±4.2	34.0±2.0	38.9±1.2	18.5±1.2	35.1±0.4
		60.4±1.4	48.7±1.3	47.5±0.9	56.6±0.9	52.5±1.6	44.6±0.5
	replace [5,10] words	48.7±1.0	45.9±1.6	31.3±1.9	35.9±1.8	18.2±0.6	33.8±1.7
		60.3±0.7	49.2±1.5	47.7±1.0	50.0±0.9	51.6±1.9	43.9±1.6
shuffle words	20.1±2.7	4.20±2.5	18.1±2.0	7.80±1.6	14.2±2.0	8.60±0.5	
	59.4±1.8	46.5±0.9	46.7±0.8	49.7±2.2	50.5±2.0	41.5±0.4	
repeat sentences	51.1±0.8	46.4±1.0	34.8±1.0	37.6±4.0	16.4±2.3	41.3±1.4	
	61.4±0.5	48.0±0.9	48.4±0.9	53.7±3.3	51.1±1.2	48.5±1.1	
shuffle sentences	49.9±1.2	46.6±1.3	34.0±1.0	42.4±0.1	16.4±1.5	40.0±0.9	
	62.2±1.2	49.6±0.8	48.6±1.4	54.4±0.2	50.1±0.7	45.2±1.7	
shuffle instructions	2.40±2.6	3.00±3.9	4.20±2.2	2.70±3.7	1.80±1.3	1.20±1.5	
	53.3±1.4	45.9±1.1	41.1±1.1	33.8±1.1	38.2±1.0	28.3±1.0	
		DTT	GEC	TG	KT	OE	QR
ROUGE-L	original instru.	31.6±0.4	83.1±1.4	33.8±0.6	56.3±1.7	24.4±2.8	48.2±0.4
		43.3±0.5	85.7±1.1	37.4±0.5	64.8±1.9	32.7±1.4	68.0±0.2
	delete stopwords	27.8±0.4	83.3±0.6	31.5±0.3	49.9±1.1	22.6±0.6	37.2±0.2
		43.8±0.4	85.2±0.5	37.2±0.3	63.6±0.3	34.2±1.0	68.0±0.3
	delete [5,10] words	29.9±1.3	78.2±8.8	31.0±0.8	49.7±4.6	24.0±0.8	41.0±2.3
		42.2±0.5	85.0±2.0	36.8±0.2	64.0±1.0	34.4±1.8	67.5±0.0
	insert [5,10] words	30.9±0.3	81.2±1.2	33.2±0.7	54.2±1.1	25.3±0.5	46.8±0.4
		42.3±0.5	86.6±0.4	37.0±0.2	65.8±1.2	36.7±1.9	67.3±0.3
	replace [5,10] words	31.3±0.9	82.6±0.7	30.9±0.7	53.5±2.8	26.9±1.2	42.9±3.6
		42.5±0.4	84.0±1.5	36.9±0.4	65.6±1.6	36.0±3.3	68.2±0.5
shuffle words	20.7±2.6	71.1±5.2	21.9±0.5	33.9±2.0	22.9±1.0	24.3±1.3	
	44.9±0.4	84.6±0.7	36.3±0.4	62.4±0.2	33.8±3.6	67.7±0.1	
repeat sentences	33.1±0.5	83.5±0.2	31.6±0.6	52.0±3.1	26.9±0.6	52.0±0.4	
	42.9±0.4	85.9±0.3	37.4±0.3	64.8±1.6	36.8±2.5	67.3±0.3	
shuffle sentences	31.1±0.2	83.6±2.3	33.9±0.5	55.1±1.2	25.8±0.6	44.6±2.3	
	41.9±0.3	86.4±0.7	37.1±0.9	64.6±1.2	38.2±1.8	67.6±0.4	
shuffle instructions	9.90±2.2	7.20±8.4	10.1±1.4	18.8±11.3	15.0±6.0	16.7±5.8	
	42.0±0.3	87.2±0.6	27.5±0.4	52.2±0.4	31.4±3.0	66.0±0.4	

Table 4: The INSTRU-WO-EXAM (in black) and INSTRU-W-EXAM (in gray) results of the T5-3B model on the 119 English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations. We use different colors to denote perturbations in **word level**, **sentence level** and **instruction level**.

non-English tasks or not. From the two tables, we observe that:

- **Effects of different levels of perturbations in INSTRU-WO-EXAM:** i) word-level perturbations degrade the performance slightly, among which the “shuffle words” has the worst effect; ii) generally, sentence-level perturbations have less influence than the word-level perturbations: we only notice tiny drop compared with “original instructions”

and no clear difference between “repeat sentences” and “shuffle sentences” is observed; iii) “shuffle instructions” doesn’t work; this is within expectation since a random instruction provides no useful supervision for the target task.

- **When the instructions contain examples (i.e., INSTRU-W-EXAM):** i) the system is pretty robust among all perturbations; ii) “shuffle instructions” (note that the examples are still specific to

metric		AC	CEC	CR	DAR	WA	TE
EXACTMATCH	original instru.	68.00	54.00	57.86	63.14	47.50	64.54
	delete stopwords	54.15	53.14	54.93	52.86	43.00	61.79
	delete [5,10] words	65.46	54.00	54.29	57.29	46.50	61.33
	insert [5,10] words	59.62	54.29	57.29	56.00	47.50	64.12
	replace [5,10] words	64.77	54.43	55.57	56.71	47.75	58.67
	shuffle words	47.23	53.86	51.07	39.86	41.88	56.12
	repeat sentences	67.08	53.86	57.64	64.71	48.88	64.17
	shuffle sentences	66.69	54.14	57.43	66.14	47.25	64.08
	shuffle instructions	43.69	44.14	49.64	38.00	39.50	48.29
		DTT	GEC	TG	KT	OE	QR
ROUGE-L	original instru.	49.17	87.80	39.11	67.78	60.35	66.25
	delete stopwords	49.56	87.19	38.21	65.07	54.82	66.46
	delete [5,10] words	49.26	87.69	38.15	65.94	58.24	66.19
	insert [5,10] words	49.28	87.55	38.62	66.17	57.08	65.98
	replace [5,10] words	49.91	88.77	38.32	66.46	57.68	65.42
	shuffle words	49.39	87.33	36.90	65.89	52.85	67.77
	repeat sentences	49.26	88.81	38.76	67.34	60.19	66.33
	shuffle sentences	49.41	87.80	39.05	66.86	59.17	66.17
	shuffle instructions	48.93	86.32	31.08	62.11	38.08	64.82

Table 5: The INSTRU-W-EXAM results of the T5-11B model on the 119 English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations. We use different colors to denote perturbations in **word level**, **sentence level** and **instruction level**.

Original Instruction	In this task, you will be shown an incorrect English sentence. You need to generate a corrected form of the input sentence.
New Instruction 1	Given an incorrect English sentence. Your task is to correct the input sentence and write out his correct form.
New Instruction 2	You will see an erroneous English sentence in this task. You must create the input sentence in its corrected form.
New Instruction 3	You are given an English sentence that has some small grammatical errors. Your task is to output the given sentence and fix all the errors inside. The output should be a grammatically correct sentence.

Table 6: An example in our constructed dataset where four diverse instructions are prepared for every single task.

the target task) only decreases the performance to a small extent, which indicates that “*correct examples+incorrect instruction*” is better than “*correct instruction without examples*”. In other words, **providing a couple of examples in the instruction rather than merely a short piece of task description is still the key to instruction-driven learning**.

- **Does larger models lead to higher robustness?** We report INSTRU-W-EXAM results of T5-11B in Table 5. We only observe occasional performance drops by word-level perturbations in a few task categories, such as CR, DAR, and OR. In general, only “shuffle instruction” can lead to clear deterioration while other perturbations do not have a noticeable effect.

- **Difference between English and non-**

**English:** To save space, we move the results of non-English *test* tasks to the Table 10 in Appendix B. By comparing Table 4 and Table 10, we do not notice any different phenomena in non-English tasks.

**How to interpret the phenomena caused by instruction perturbations?** Based on the aforementioned observations, we conclude that perturbations, retaining as much task semantic as the original instructions, can yield less influence on the system’s robustness. In other words, two instructions that are more likely to be paraphrased tend to have similar performance. We further validate this argument in the subsequent Section 3.2.

Since INSTRU-W-EXAM generally performs bet-

T5-3B			AC	CEC	CR	DAR	WA	TE
	EM	orig.	69.5	69.2	58.5	48.7	44.5	44.8
		reph.	67.3±1.0	69.3±0.2	56.8±1.6	43.3±2.7	43.2±1.9	52.7±5.8
			DTT	GEC	TG	KT	OE	QR
ROUGE-L	orig.	44.7	85.0	25.5	73.2	34.7	81.0	
	reph.	44.2±0.3	84.3±0.3	24.6±0.2	72.0±0.6	32.7±0.1	81.3±0.2	
T5-11B			AC	CEC	CR	DAR	WA	TE
	EM	orig.	63.5	74.0	62.0	51.5	51.5	67.0
		reph.	61.8±0.6	74.5±1.1	56.7±5.3	50.3±4.1	50.0±0.4	65.0±2.4
			DTT	GEC	TG	KT	OE	QR
	ROUGE-L	orig.	48.8	85.0	26.4	71.5	52.1	75.0
		reph.	48.7±0.5	85.4±0.2	26.3±0.2	71.6±0.5	51.6±1.0	76.9±2.1

Table 7: The results of the T5-3B&11B models on our Para-Instructions dataset.

ter than INSTRU-WO-EXAM, hereafter, we report INSTRU-W-EXAM for the remaining experiments.

### 3.2 Robustness when instructions are rephrased

In addition to the instruction perturbation experiments which may lead to grammatically incorrect instructions, this subsection studies a scenario where multiple valid instructions are available for the same task. This is a real-world challenge since the end users are free to write any open-form instructions to guide the system.

**Data construction.** Unfortunately, there are no existing datasets that contain tasks with diverse human-oriented instructions. To build a dataset, we randomly choose 2 *English*<sup>2</sup> tasks from each of the 12 categories in the *test* of SUPER-NATURALINSTRU; this results in 23 tasks<sup>3</sup> that are uniformly distributed over those categories. We hire three graduate students in the NLP area to construct a dataset of diverse instructions with the following two steps:

- **Writing step:** for each task, each annotator is provided with 10 labeled instances, the task name, and the original task Definition. Then each annotator is required to write a new definition based on her/his understanding of the task. After this step, each task will have 3 new versions of raw instructions.

- **Validation step:** For each raw instruction, we forward it to two checkers who judge if it is valid for the task and if it is different enough from the original Definition. If any issues are found, the

<sup>2</sup>We considered English only since it is hard to find annotators who are familiar with many non-English languages.

<sup>3</sup>The “grammar error correction” has only one task.

original author will be asked to modify the instruction. This process is repeated until both checkers are satisfied with the new instruction. After this step, each task will have 4 versions of high-quality instructions (one original, three new).

Table 6 lists an example task with 4 instructions in our constructed data. Compared with the perturbed instructions demonstrated in Table 3, it is clear that the tasks in our new dataset are equipped with more valid instructions and they are more like paraphrases of each other while conveyed with more distinct textual expressions. Our constructed dataset, “Para-Instructions”, as the first one that contains multiple human-oriented instructions for each task, is released to the community.

**Results.** For all three new instructions, we report their averages and standard deviations. Table 7 compares the performance of the original instruction and the three new instructions by T5-3B and T5-11B models. We notice that the system is considerably robust among all four versions of instructions. **It suggests that providing a paraphrasing instruction can generally keep the performance once demonstrations are available.**

### 3.3 Robustness when instructions have different abstractiveness

Another challenge in instruction-driven task generalization is to encourage the same system to work well on datasets that contain instructions constructed from totally different templates or abstractiveness. This phenomenon often exists between two different datasets. For example, SUPER-NATURALINSTRU has instructions that are mostly a paragraph while DIVERSEPROMPT (Honovich et al., 2022) contains prompts only—mostly short

	Task	T5-3B	T5-11B
EXACTMATCH	sum	0.00±0.0	11.4±4.2
	synonyms	1.00±0.0	2.40±0.9
	rhymes	1.00±0.0	0.00±0.0
	larger animal	5.00±2.5	45.8±7.6
	diff	1.10±0.9	3.30±1.1
	first word letter	1.00±0.0	46.0±9.3
	second word letter	1.30±0.5	19.9±6.5
	singular to plural	0.80±0.4	55.0±1.2
	antonyms	13.6±1.6	33.1±0.6
	word in context	9.00±12.2	31.7±0.0
	letters list	2.00±0.0	1.80±0.6
	sentence similarity	0.00±0.0	0.00±0.0
	sentiment	34.1±8.9	49.0±2.8
	num to verbal	0.00±0.0	0.00±0.0
orthography	1.20±0.4	7.00±2.3	
ROUGE-L	translation en-es	8.50±0.5	23.3±2.8
	translation en-de	9.50±1.0	10.1±5.8
	translation en-fr	12.4±0.8	37.0±0.8
	informal to formal	30.3±0.4	32.8±2.3
	negation	70.7±11.6	71.6±6.1
	cause and effect	50.5±15.1	31.8±9.7
	taxonomy animal	24.0±3.9	26.6±7.0
	common concept	0.00±0.0	1.70±2.2
	active to passive	60.3±2.4	94.4±3.5

Table 8: Train T5-3B&11B on SUPER-NATURALINSTRU and test it on DIVERSEPROMPT (Honovich et al., 2022).

and abstractive sentences.

**Dataset.** To check the robustness of the model finetuned on SUPER-NATURALINSTRU, we test it on the DIVERSEPROMPT dataset. DIVERSEPROMPT has 24 tasks and each task has an average of 7.75 prompts. DIVERSEPROMPT was the only dataset that provides multiple prompts for a task before we wrote this paper, and the prompts are much more concise (average word size 10.75) than the training instructions of SUPER-NATURALINSTRU.

**Results.** Table 8 presents the results on DIVERSEPROMPT by the model finetuned on SUPER-NATURALINSTRU. For this experiment, we observe a totally different story:

- First, the pretrained model on SUPER-NATURALINSTRU does not work well on DIVERSEPROMPT. Regardless of the task types (shown by the evaluation metrics: EXACTMATCH or ROUGE-L), most tasks show unsatisfactory performance. This can result from two

reasons: on the one hand, the instruction format (i.e., prompts in DIVERSEPROMPT) is significantly different from that in SUPER-NATURALINSTRU—DIVERSEPROMPT has prompts that mostly contain a few words while instructions in SUPER-NATURALINSTRU are more specific, informative and are human-oriented; on the other hand, the tasks in DIVERSEPROMPT are very distinct from those in SUPER-NATURALINSTRU, such as “sum”, “diff”, “first word letter”, “common concept”, etc.

- Second, even for some tasks that co-exist in SUPER-NATURALINSTRU and DIVERSEPROMPT, e.g., “word in context”, “translation”, “sentiment”, etc., their performances fall short of expectations. For example, “Text Generation” (TG) in Table 7 gets consistent performance around 24.6 by ROUGE-L, whereas the translation tasks in DIVERSEPROMPT ranging from 8 to 13.

The preceding two items hint at the fragility of the model when dealing with unseen tasks and unfamiliar instruction formats, despite its training on more than 1,000 instruction-laden tasks. So far, we still lack a method to effectively learn instruction to solve tasks in the absence of demonstration. The majority of instruction-tuning datasets commonly furnish only a single instruction for each input-output pair, aiming to enhance the system’s capacity for instruction following through an increased number of training tasks. However, this approach may prove suboptimal in supervising a system that can effectively generalize to unfamiliar tasks with unfamiliar formats of instructions.

- Third, it is worth mentioning that this DIVERSEPROMPT dataset does not provide demonstrations. In this case, we notice mostly *greater standard deviations* in Table 8 than that in Table 7. This indicates that even paraphrasing prompts can not guarantee consistent performance if no demonstrations exist. Then we dive into the performance per prompt for tasks that have larger standard deviations of ROUGE-L, such as “negation”, “cause and effect”, “taxonomy animal” and “active to passive”, in Table 9. We found the following factors that affect prompt performance in text generation tasks: (i) it is better to use a verb that clearly has a single sense. For example, “write” and “extract” are generally better than “output”, “list”, which can be verbs or nouns; (ii) explicitly presenting the action is better than that implicitly. For example, “... write the ...” is better than a question “which of the ... is ...?”.



	ROUGE-L	prompt
negation	63.32	For each input, write a sentence that expresses the exact opposite meaning of the input.
	65.95	Change the fact stated in the sentence to an opposite fact.
	66.50	You will be given a sentence that states a fact (that might be true or not). Try to state the opposite fact.
	68.92	Negate the following sentence:
	74.05	output the negation of the input
	75.39	Negate the given sentence
	75.87	write the negation
	82.83	Write a negated version of the given sentence
cause & effect	16.26	Which of the two events is the cause?
	21.23	Which of the following sentences is the cause?
	21.97	Output the cause (other sentence describes what happened as a result)
	28.71	Output the sentence describing the cause (the other sentence is what happened as a result).
	32.17	Each input consists of two sentences, where one is the cause and the other is the outcome. Write the cause sentence.
	39.03	The input consists of two sentences. One is the cause of the other. Write the cause sentence.
	40.30	Find the cause in the following cause and effect pair
	41.83	The input is a cause and effect. Write the cause.
44.78	The input is a cause and effect, write the cause	
taxonomy animal	12.06	extract animals
	21.46	List which of the following are animals
	22.62	Find the animals in the list
	24.08	List the animals from the given words
	27.86	find the animals in the following list of words
	28.25	write all animals from the list of words
	33.88	write only the animals from the list of words
	33.89	Extract all animals from the list
35.32	Extract all animals from the input list	
active_to_passive	88.18	Write the following sentence in passive language
	88.63	output the passive form of the input
	95.71	For each input, write the passive form of the sentence.
	95.87	Rewrite the sentence in passive form
	96.08	Change the sentence from active to passive.
	96.21	Rewrite the input sentence in passive form
	96.92	Change the wording of the following sentence from active to passive.
	97.80	turn the sentence from active tense to passive tense.

Table 9: ROUGE-L per prompt for four tasks in DIVERSEPROMPT.

This story shows that there is still a long way to go for instruction-driven task generalization, especially when the system deals with less familiar tasks with abstractiveness-distinct instructions.

#### 4 Conclusion

This paper proposed a new issue, the robustness of instruction-tuned models. Experiments show that the instruction-tuned model may not learn the

content of the instructions itself in some tasks, but rather the pattern of positive examples. The robustness of the instruction-tuned model is very important, which determines the stability and reliability of instruction fine-tuning models in real-world applications. Using only instructions to tell models what they should do is a long-term goal of AI, and we hope that future work will be done to improve the robustness of the instruction-tuned model.

## Acknowledgments

The authors would like to thank all the anonymous reviewers for their insightful comments and suggestions.

## Limitations

We summarize the limitations of this work as follows:

- First, it is still challenging to figure out why a particular instruction results in better or worse performance. We tried to explain a part of the observations, but some of the phenomena are still mysterious.
- Technically we can create a much larger Para-Instruction dataset than 23 that provides diverse instructions for NLP tasks. Due to the budget limitation, we only collected 23 tasks.
- Some prior work showed that larger PLMs often bring better performance in dealing with instruction learning. Due to hardware constraints, we are unable to try PLMs that are larger than T5-11B.

## Ethics Statement

We do not anticipate any ethical issues particularly to the topics of this work. Nevertheless, some work presented here extensively uses large-scale pre-trained models with self-attention, which may lead to substantial financial and environmental costs.

## References

Stephen Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-david, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Fries, Maged Alshaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Dragomir Radev, Mike Tian-jian Jiang, and Alexander Rush. 2022. [Prompt-Source: An integrated development environment and repository for natural language prompts](#). In [Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations](#), pages 93–104, Dublin, Ireland. Association for Computational Linguistics.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda

Askeell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020a. Language models are few-shot learners. In [NeurIPS](#).

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askeell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020b. Language models are few-shot learners. In [Proceedings of NeurIPS](#).

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In [NAACL-HLT \(1\)](#), pages 4171–4186. Association for Computational Linguistics.

Avia Efrat and Omer Levy. 2020. The turking test: Can language models understand instructions? [CoRR](#), abs/2010.11982.

Karan Goel, Nazneen Fatema Rajani, Jesse Vig, Zachary Taschdjian, Mohit Bansal, and Christopher Ré. 2021. Robustness gym: Unifying the NLP evaluation landscape. In [Proceedings of NAACL-HLT](#), pages 42–55. Association for Computational Linguistics.

Or Honovich, Uri Shaham, Samuel R. Bowman, and Omer Levy. 2022. Instruction induction: From few examples to natural language task descriptions. [CoRR](#), abs/2205.10782.

Daniel Khashabi, Chitta Baral, Yejin Choi, and Hananeh Hajishirzi. 2022. Reframing instructional prompts to gptk’s language. In [Findings of ACL](#), pages 589–612. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In [Proceedings of ACL](#), pages 7871–7880.

Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In [Text Summarization Branches Out](#), pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of

- prompting methods in natural language processing. CoRR, abs/2107.13586.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019. OpenReview.net.
- Renze Lou, Kai Zhang, and Wenpeng Yin. 2023. [Is prompt all you need? no. A comprehensive and broader view of instruction learning](#). CoRR, abs/2303.10475.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) CoRR, abs/2202.12837.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022a. [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.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022b. [Cross-task generalization via natural language crowdsourcing instructions](#). In Proceedings of ACL, pages 3470–3487. Association for Computational Linguistics.
- John X. Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. [Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in NLP](#). In Proceedings of EMNLP: System Demonstrations, pages 119–126. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#). OpenAI blog.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). J. Mach. Learn. Res., 21:140:1–140:67.
- Mrigank Raman, Pratyush Maini, J. Zico Kolter, Zachary C. Lipton, and Danish Pruthi. 2023. [Model-tuning via prompts makes NLP models adversarially robust](#). CoRR, abs/2303.07320.
- 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 V. 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 Févry, 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 ICLR. OpenReview.net.
- Xavier Schmitt, Sylvain Kubler, Jérémy Robert, Mike Papadakis, and Yves Le Traon. 2019. [A replicable comparison study of NER software: Stanfordnlp, nltk, opennlp, spacy, gate](#). In SNAMS, pages 338–343. IEEE.
- Tianlu Wang, Xuezhi Wang, Yao Qin, Ben Packer, Kang Li, Jilin Chen, Alex Beutel, and Ed H. Chi. 2020. [Catgen: Improving robustness in NLP models via controlled adversarial text generation](#). In Proceedings of EMNLP, pages 5141–5146. Association for Computational Linguistics.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma Doshi, Maitreya Patel, Kuntal Kumar Pal, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Shailaja Keyur Sampat, Savan Doshi, Siddhartha Mishra, Sujan Reddy, Sumanta Patro, Tanay Dixit, Xudong Shen, Chitta Baral, Yejin Choi, Hannaneh Hajishirzi, Noah A. Smith, and Daniel Khashabi. 2022. [Benchmarking generalization via in-context instructions on 1, 600+ language tasks](#). CoRR, abs/2204.07705.
- Albert Webson and Ellie Pavlick. 2022. [Do prompt-based models really understand the meaning of their prompts?](#) In Proceedings of NAACL-HLT, pages 2300–2344. Association for Computational Linguistics.
- Wenpeng Yin, Jia Li, and Caiming Xiong. 2022. [ConTinTin: Continual learning from task instructions](#). In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 3062–3072, Dublin, Ireland. Association for Computational Linguistics.
- Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, and Chenliang Li. 2020. [Adversarial attacks on deep-learning models in natural language processing: A survey](#). ACM Trans. Intell. Syst. Technol., 11(3):24:1–24:41.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. [Calibrate before use: Improving few-shot performance of language models](#). In ICML, volume 139 of Proceedings of Machine Learning Research, pages 12697–12706. PMLR.

	EXACTMATCH			ROUGE-L
	AC	CEC	TE	TG
original	76.3±3.4	43.1±0.7	19.7±1.4	10.3±2.0
	72.0±0.8	54.2±0.6	20.0±1.2	8.80±0.7
delete stopwords	69.0±4.3	20.8±0.6	10.1±0.2	5.70±2.2
	69.7±1.9	54.8±0.2	19.0±1.4	6.30±0.8
delete [5,10] words	56.3±2.8	28.9±2.2	11.2±4.9	9.70±1.6
	73.0±2.9	54.6±0.6	21.0±1.9	8.40±0.1
insert [5,10] words	72.0±2.2	35.5±1.4	17.6±3.4	9.50±0.6
	69.3±4.1	55.7±0.3	19.6±0.6	7.20±1.0
replace [5,10] words	68.3±0.5	39.7±1.5	17.8±3.1	10.5±1.6
	67.0±7.9	54.2±0.3	17.4±0.8	6.80±1.9
shuffle words	3.00±1.6	0.20±0.1	1.40±2.0	8.60±2.3
	70.0±5.7	53.2±0.1	16.0±2.6	8.30±1.7
repeat sentences	73.3±2.6	32.6±2.7	14.6±3.5	9.20±1.7
	74.7±1.2	55.2±0.6	20.9±2.4	9.00±0.7
shuffle sentences	69.7±3.1	34.8±1.7	17.3±0.9	5.80±1.7
	67.3±1.2	54.9±0.5	19.1±1.6	7.00±1.8
shuffle instructions	0.00±0.0	14.9±1.7	0.00±0.0	0.60±0.9
	40.0±2.9	52.6±0.4	5.30±0.7	1.90±1.5

Table 10: The INSTRU-WO-EXAM (in black) and INSTRU-W-EXAM (in gray) results of the T5-3B model on the 35 non-English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations.

	EXACTMATCH			ROUGE-L
	AC	CEC	TE	TG
original	86.00	69.57	32.00	13.27
delete stopwords	78.0	69.17	28.00	14.00
delete [5,10] words	85.0	69.07	28.00	13.00
insert [5,10] words	87.0	69.27	30.00	12.83
replace [5,10] words	85.0	67.47	30.67	13.83
shuffle words	76.0	68.03	29.67	13.17
repeat sentences	87.0	69.53	31.00	13.27
shuffle sentences	85.0	69.97	29.33	13.27
shuffle instructions	60.0	69.07	19.67	4.47

Table 11: The INSTRU-W-EXAM results of the T5-11B model on the 35 non-English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations.

## A Finetuning Details

In this work, we finetune T5 using the HuggingFace library. We use a scheduler with a constant learning rate of  $5e-05$ . All models are optimized by AdamW (Loshchilov and Hutter, 2019) with a batch size of 8. T5-3B is fine-tuned on one NVIDIA RTX A5000 GPU and T5-11B on four NVIDIA A100 GPUs. DeepSpeed<sup>4</sup> is used to parallelize the models, and bfloat16 precision is enabled to save GPU memory.

<sup>4</sup><https://github.com/microsoft/DeepSpeed>

## B T5-3B Results on Non-English Tasks

Table 10 presents the INSTRU-WO-EXAM (in black) and INSTRU-W-EXAM (in gray) results of the T5-3B model on the 35 non-English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations.

## C T5-11B Results on Non-English Tasks

Table 11 presents the INSTRU-W-EXAM results of the T5-11B model on the 35 non-English *test* tasks of SUPER-NATURALINSTRU by different levels of perturbations.



## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
5
- A2. Did you discuss any potential risks of your work?  
1
- A3. Do the abstract and introduction summarize the paper's main claims?  
1
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

3

- B1. Did you cite the creators of artifacts you used?  
*No response.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*No response.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*No response.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*No response.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*No response.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*No response.*

### C Did you run computational experiments?

3

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*No response.*

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*No response.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*No response.*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*No response.*

**D  Did you use human annotators (e.g., crowdworkers) or research with human participants?**

3

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*No response.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*No response.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*No response.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*No response.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*No response.*