

Functionality learning through specification instructions

Pedro Henrique Luz de Araujo[◇] and Benjamin Roth[†]

^{◇†}Faculty of Computer Science, University of Vienna, Vienna, Austria

[◇]UniVie Doctoral School Computer Science, Vienna, Austria

[†]Faculty of Philological and Cultural Studies, University of Vienna, Vienna, Austria

{pedro.henrique.luz.de.araujo, benjamin.roth}@univie.ac.at

Abstract

Test suites assess natural language processing models’ performance on specific functionalities: cases of interest involving model robustness, fairness, or particular linguistic capabilities. This paper introduces specification instructions: text descriptions specifying fine-grained task-specific behaviors. For each functionality in a suite, we generate an instruction that describes it. We combine the specification instructions to create specification-augmented prompts, which we feed to language models pre-trained on natural instruction data.

We conduct experiments to measure how optimizing for some functionalities may negatively impact functionalities that are not covered by the specification set. Our analyses across four tasks and models of diverse sizes and families show that smaller models struggle to follow specification instructions. However, larger models (> 3B params.) can benefit from specifications and—surprisingly—even generalize certain desirable behaviors across functionalities.¹

1 Introduction

Test suites (Kirk et al., 2022; Röttger et al., 2021; Ribeiro et al., 2020; McCoy et al., 2019) have been proposed as an evaluation framework to test for specific functionalities in natural language processing (NLP) models. Each functionality is a set of test cases, generally input-output pairs, relating to a particular aspect of a task. For example, a test suite for hate speech detection can assess distinct expressions of hate (e.g., implicit derogation), while a sentiment analysis suite can measure how well a model handles specific phenomena (e.g., negation). Suites complement the standard practice of evaluating on a held-out test set from the same distribution

of the training set (Linzen, 2020). If the latter is representative of the underlying task distribution, it is a good measure of average correctness; suites, on the other hand, allow for in-depth evaluation of relevant phenomena that may be underrepresented in general data.

Though test suites can point to failure cases, there are no clear guidelines on how to act upon their feedback to develop more robust and trustworthy models. Data augmentation has been suggested as a potential avenue for improvement (Röttger et al., 2021) by including additional training cases that correspond to the suite’s cases. However, constructing or annotating instances targeting specific functionalities is costly, and further training models is expensive for large models and infeasible for closed-source ones. Furthermore, fine-tuning models on suites’ test cases has been shown to help seen functionalities (Ribeiro and Lundberg, 2022; Malon et al., 2022; Liu et al., 2019), but often does not generalize to unseen ones and harms general performance (Luz de Araujo and Roth, 2023, 2022; Rozen et al., 2019; McCoy et al., 2019).

Missing from the literature are analyses for the increasingly influential paradigm of prompting large language models (LLMs) (Liu et al., 2023), which has superior zero- and few-shot capabilities, particularly for models trained on natural language instructions (Ouyang et al., 2022). Prompting with functionality information may improve relevant aspects of model behavior with no need for fine-tuning, which requires additional training data and computational resources. Since the model parameters are not updated, prompting is also less vulnerable to overfitting to seen functionalities, a substantial limitation in previous work.

This paper explores specification instructions and their effect on functionality performance. Contrary to previous efforts, we do not expose the model to suite examples or fine-tune it. Instead, we elicit the desired behavior by augmenting prompts

¹Our code is available on <https://github.com/peluz/specification-instructions>

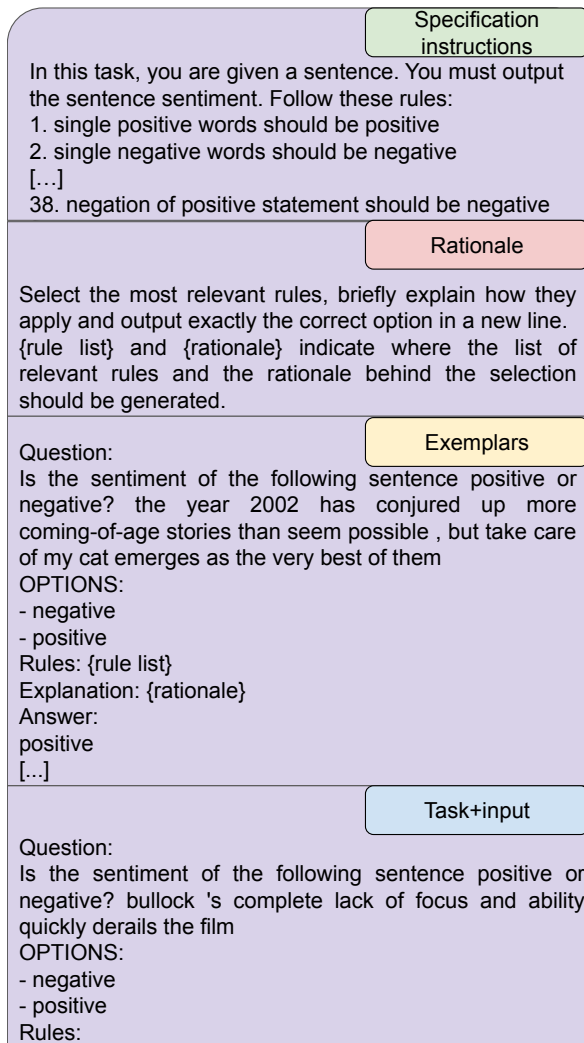


Figure 1: Example of a specification-augmented prompt for sentiment analysis. Each module adds information about how the task is expected to be performed.

with instructions that specify the suite’s functionalities. For example, if a sentiment analysis suite contains a functionality that tests whether predictions are invariant to nationalities mentioned in the input, an instruction such as “nationality should be irrelevant to sentence sentiment” would be added to the task prompts.

Our main contributions are:

1. Creating two sets—handcrafted and machine-generated—of 144 specification instructions for 4 test suites from different tasks (sentiment analysis, paraphrase identification, reading comprehension, and hate speech detection²) and designing specification-augmented prompts.

2. Assessing the impact of the specification-augmented prompts for seven models ranging from

²This paper contains examples of abusive and hateful language.

80M to billions of parameters and covering three model families.

3. Evaluating cross-functionality impact through scenarios with held-out specifications, finding that overfitting to seen cases is much less of a concern here than in the fine-tuning paradigm.

4. Qualitatively examining the impact of specification-augmented prompts and the interplay between different specifications by examining which functionalities are most helped or harmed across different evaluation scenarios.

2 Prompting with specification instructions

2.1 Problem setting

We consider a task to be composed of a dataset \mathcal{D} of n labeled examples assumed to be identically and independently distributed (i.i.d.), $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$, and a test suite \mathcal{T} of m test cases $\{t_i\}_{i=1}^m$ partitioned into n_{func} disjoint functionalities $\{\mathcal{F}_i\}_{i=1}^{n_{\text{func}}}$. Each functionality is assigned to a functionality class $c \in \{\mathcal{C}_i\}_{i=1}^{n_{\text{class}}}$, such that $n_{\text{class}} < n_{\text{func}} < m$. While \mathcal{D} describes the general behavior expected for the task, \mathcal{T} specifies fine-grained aspects of the expected behavior.

For example, \mathcal{D} can be a dataset of tweets with labels indicating whether they contain hate speech; \mathcal{T} would be a suite with functionalities that assess specific expressions of hate (e.g., use of profanity, threatening language) and contrastive non-hate speech (e.g., use of reclaimed slurs, non-hateful profanity) (Röttger et al., 2021).

The functionality classes encompass coarse-grained dimensions such as fairness and robustness, the functionalities assess finer-grained aspects such as gender fairness and robustness to typos, and the test cases operationalize them as pairs of inputs and expected behaviors (Röttger et al., 2021; Ribeiro et al., 2020).

2.2 Prompt modules

In our setting, each prompt is composed of several modules: a necessary task description and optional modules that further specify the task. More formally, given an input x , a task description τ , and a (possibly empty) set of optional modules M , we have

$$z = f(x, \tau, M), \quad (1)$$

where z is the resulting prompt, and f is a function that combines prompt modules and inputs. x and τ are strings, and M is a set of strings. Fig. 1 shows

a prompt for sentiment analysis with all optional modules.

Below, we describe each prompt module:

Task description (Task): a natural language instruction that describes the task. For example, “Answer a question about this article:” for a reading comprehension task (details in Appendix C).

Exemplars (Ex): input-label pairs that exemplify the task, also known as demonstrations (Brown et al., 2020). They can help to improve task performance by providing the model with information about the task format, label space, input distribution, and input-label mapping (Min et al., 2022).

Specification instructions (Spec): provides, for each functionality in the suite, an instruction that specifies the behavior expected by the functionality (e.g., “typos in the question are irrelevant to the answer”).³ Their purpose is to elicit the LM to generate text that conforms to what the suite specifies.

Rationales (Rat): asks the model to state the applicable specifications and the underlying rationale (before generating the task prediction). This module is similar to chain-of-thought prompting (Kojima et al., 2022; Wei et al., 2022b), which asks the LM to work the solution to a problem step-by-step and has been shown to improve LM performance in reasoning benchmarks.

Combinations of optional modules yield different *prompting methods*. We explore two baseline methods with no suite information (**Task**, **Task+Ex**) and four specification-augmented methods (**Task+Spec**, **Task+Spec+Ex**, **Task+Spec+Rat**, **Task+Spec+Rat+Ex**). By comparing the baselines and specification-augmented methods, we assess the impact of incorporating additional task specifications; by comparing the specification-augmented methods, we investigate the impact of the individual modules.

2.3 Cross-functional analysis

The specification instructions do not cover all aspects of desired task behavior—there is always a chance that important phenomena are (intentionally or not) left unspecified.⁴ For example, the specifications for sentiment analysis (§ 3.1) state that

³We discuss specification instruction generation in §3.3 and show all specification instructions we generate in Appendix D.

⁴If one could completely specify a task, training a model would be unnecessary.

sentence sentiment is invariant to persons’ names and locations. However, sentiment should also be invariant to organization names (not checked in the specifications). An evaluation setting that measures performance only on included specifications cannot examine how the instructions affect specifications the model developer did not think to include.

To address this, we adapt the cross-functional analysis method (Luz de Araujo and Roth, 2023) to the prompt-based learning paradigm. The method was originally proposed for the fine-tuning learning paradigm and involves training and evaluating on different sets of functionalities. To translate this into the prompting paradigm, we vary which specifications are included in each prompt:

Seen scenario: No specifications are held-out. This scenario measures how including specifications affects performance for *seen* functionalities.

Functionality generalization: We remove the specification that applies to the input. For example, if the input belongs to Functionality 1, we remove specification instruction 1 from the prompt. This scenario estimates performance for *unseen* functionalities.

Functionality class generalization: We remove all specifications from the same functionality class of the applicable specification. In the example above, if functionalities 1 to 3 are from the same functionality class, we remove specifications 1 to 3 from the prompt. This scenario estimates performance for *unseen* functionality classes.

3 Experimental setting

3.1 Tasks

We examine four NLP tasks based on the availability of test suites, their representativeness of the NLP field, and their use in previous work on functionality learning. All the data is in English. Table 1 provides examples and split sizes for all datasets and suites.

Sentiment analysis (SENT): the task is to output the sentiment of the input sentence. The dataset is the Stanford Sentiment Treebank (SST2) (Socher et al., 2013), as made available in the GLUE benchmark (Wang et al., 2018). We use the sentiment analysis suite developed by Ribeiro et al. (2020) as the suite. There is a label space discrepancy between dataset and suite: the dataset labels include only positive and negative, while the suite extends the options with a neutral label.

Paraphrase identification (PARA): the task is

Task	Dataset	Split sizes	Example (label)
SENT	SST-2	67k/872	a sweet and modest and ultimately winning story (positive)
	Suite	89k/44k/44k	I thought the aircraft would be beautiful, but it wasn't (negative).
PARA	QQP	363k/40k	Q1: What is best way to reach Kashmir / Srinagar? Q2: What is your review of Srinagar, Jammu & Kashmir, India? (Not duplicate)
	Suite	103k/51k/51k	Q1: How can I become a powerless person? Q2: How can I become a person who is not powerful? (Duplicate)
READ	SQuAD	87k/10k	C: After Hurricane Katrina in 2005, Beyoncé and Rowland founded the Survivor Foundation to provide transitional housing for victims in the Houston area [...] Q: What foundation did Beyoncé start after Hurricane Katrina? (Survivor Foundation)
	Suite	35k/17k/17k	C: Kevin is nicer than Amanda. Q: Who is less nice? (Amanda)
HATE	Davidson	19k/2k/2k	[USER] can a quote this and tag a bitch (Not hateful)
	Founta	79k/10k/10k	rt [USER]: i'm tired of u feminist bitches bc this is just disgusting [URL] (hateful)
	Suite	1.8k/920/921	It's disheartening to still see people call for the death of women in 2020. (not hateful)

Table 1: Summary of the datasets and suites used in this work. We report train/validation/test sizes for the datasets with public test sets and train/validation sizes otherwise. We use the suite splits from [Luz de Araujo and Roth \(2023, 2022\)](#).

to assess if two questions have the same meaning. We use Quora Question Pairs (QQP) ([Iyer et al., 2017](#)) as the dataset and the QQP suite by [Ribeiro et al. \(2020\)](#).

Reading comprehension (READ): given a context paragraph, the task is to answer a question whose answer is in the context. We use the Stanford Question Answering Dataset (SQuAD) ([Rajpurkar et al., 2016](#)) as the dataset and the corresponding suite by [Ribeiro et al. \(2020\)](#).

Hate speech detection (HATE): the task is to determine whether a given sentence contains hateful speech. Following previous work ([Röttger et al., 2021](#)), we examine two datasets ([Davidson et al., 2017](#); [Founta et al., 2018](#)), which we refer to as Davidson and Founta. We use HATECHECK ([Röttger et al., 2021](#)) as the suite.

3.2 Models

We compare the predictions of all models from the Flan-T5 family ([Chung et al., 2022](#); [Wei et al., 2022a](#)) (Small, Base, Large, XL and XXL), Zephyr ([Tunstall et al., 2023](#)), and ChatGPT⁵ ([OpenAI, 2022](#)). To examine the model size effect, we cover several orders of magnitude—from 80M to billions of parameters.⁶ These three model families cover three of the main paradigms of LLMs—Flan-T5 are instruction-tuned models ([Longpre et al.,](#)

⁵The gpt-3.5-turbo-0301 variant of the OpenAI API.

⁶From smallest to largest: Small-80M, Base-250M, Large-780M, XL-3B, Zephyr-7B, and XXI-11B. OpenAI has not disclosed details for GPT-3.5, but the largest variant of its “sibling model” ([OpenAI, 2022](#)) InstructGPT, has 175B parameters.

[2023](#)), Zephyr is a chat model aligned with human preferences through direct preference optimization (DPO) ([Rafailov et al., 2023](#)), while ChatGPT is aligned through reinforcement learning from human feedback (RLHF) ([OpenAI, 2022](#)).

3.3 Specification instruction generation

We experiment with handcrafted and machine-generated specification instructions. Tables 6-9 in Appendix D exhibit all specification instructions from both settings.

Handcrafted. The specification instructions in the handcrafted setting were manually written by one of the authors. Specification instructions for the CHECKLIST suites (SENT, PARA and READ) were freely written based on the functionalities in the suite. This was done by manual inspection of each functionality’s test cases and documentation. Since HATECHECK contains natural language descriptions of all functionalities ([Röttger et al., 2021](#), Appendix B), we adapt them to fit our specification format.

Machine-generated. We designed a prompt template in which we provide the task, the functionality name,⁷ six⁸ test cases from the functionality and ask for a rule that supports the behavior encoded by the test cases. Table 10 in Appendix E shows an example for each task-test type combination. We then generated a prompt for each functionality, fed it to ChatGPT, and used the completions as the machine-generated specification instructions.

⁷Names were taken from the suite for the CHECKLIST suites or [Röttger et al. \(2021\)](#) for HATECHECK.

⁸Two, in the case of READ INV functionalities, due to its lengthy inputs.

3.4 Evaluation metrics

Dataset metrics: We use the accuracy as the metric for SST-2 and QQP, the exact string match for SQuAD, and the F_1 score of the hateful class for Founta and Davidson.⁹

Suite metrics: Each functionality has a pass rate: the fraction of successful test cases. The final suite score is the arithmetic mean of all its functionality pass rates. Each evaluation scenario (§ 2.3) yields a corresponding suite score. Therefore, a suite has (1) seen, (2) functionality, and (3) functionality class generalization scores.

Aggregate metrics: We report the generalization score G (Luz de Araujo and Roth, 2023) as the aggregate score of suite and dataset performance. It is the harmonic mean of the dataset and suite metrics. The harmonic mean is used so that high dataset performance cannot compensate for poor suite performance (and vice-versa). Each suite metric yields its own aggregate score: G_{seen} , G_{func} , and G_{class} for seen, functionality generalization, and functionality class generalization.¹⁰

Evaluation of machine-generated specification instructions: We manually evaluate the quality of the ChatGPT-generated specification instructions using the criteria established by Wang et al. (2023), where each generated specification instruction is assigned a rating from A (best) to D (worst).¹¹

4 Analysis of prompt methods and components

Table 2 shows the aggregate scores for all methods and models.¹² Fig. 2 shows scores for all suites and datasets.¹³

Impact of specification instructions. Specification instructions only improved performance of the larger models: including them in the prompt reduced the average performance of Flan-T5-Large and smaller models but improved it for Flan-T5-XL and larger models. The effect differed across

tasks and no model-method pair improved over the baseline in all tasks. ChatGPT benefitted from specification instructions most consistently (four out of five tasks).

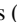
We expected specification instructions to improve suites’ scores (Fig. 2, bottom row) more than datasets (top row), because specification instructions are guaranteed to correspond to suite instances and only occasionally to dataset instances. We validate this intuition by comparing the average dataset and average suite performance difference between specification-augmented prompts and their corresponding baseline (Fig. 8 in Appendix J). While Flan-T5-base and larger models benefit from specification instructions considering suite performance, only XL and larger models could improve dataset performance. Table 13 in Appendix J shows, for each dataset, the examples for which specification instructions consistently improved or harmed ChatGPT predictions.

Impact of exemplars. Adding exemplars improved average performance in almost all the scenarios. The only exception was for Flan-T5-small baseline methods, where Task outperformed Task+Ex by 0.27. This effect was overall consistent across tasks: Task+Ex achieved an (averaged across models) improvement over Task in all evaluation scenarios, except for SENT, for which there was an average decrease of 0.67. Comparing Task+Spec and Task+Spec+Rat with their exemplar-augmented counterparts yields similar conclusions, except that Task+Spec+Rat outperformed Task+Spec+Ex+Rat on the HATE tasks.

Impact of rationales. In most cases, average performance decreases when prompts include the Rationale module. The only exceptions are Flan-T5-Base and Large with Task+Spec+Ex+Rat prompts, outperforming Task+Spec+Ex by 0.23/0.80. Qualitatively, only Zephyr and ChatGPT actually produced rationales. Flan-T5-XXL either ignores the rationale instruction or copies the list of specification instructions. The other models ignore the module entirely.

Even though the Rationale module did not improve task performance, ChatGPT returned the correct specification better than random in all tasks (Fig. 3). A follow-up question is how much of the performance degradation is due to ChatGPT failing to identify the correct specification.

To investigate this, we computed the Pearson’s correlation between specification prediction correctness and functionality performance on two lev-

⁹We use Scikit-learn (Pedregosa et al., 2011) to compute the F_1 score and  Datasets (Lhoest et al., 2021) for the other metrics.

¹⁰When using the baseline prompting methods (with no specifications), the three G scores are the same, as specification instructions are never included in the prompt.

¹¹Details in Appendix D.

¹²We report scores for a single run. We test for significance through randomized testing (Yeh, 2000) (10000 rounds, $p < 0.05$). We report all p-values in Table 12.

¹³Appendix I analyses the impact of prompt length on performance and Appendix J shows scores for each suite and dataset.

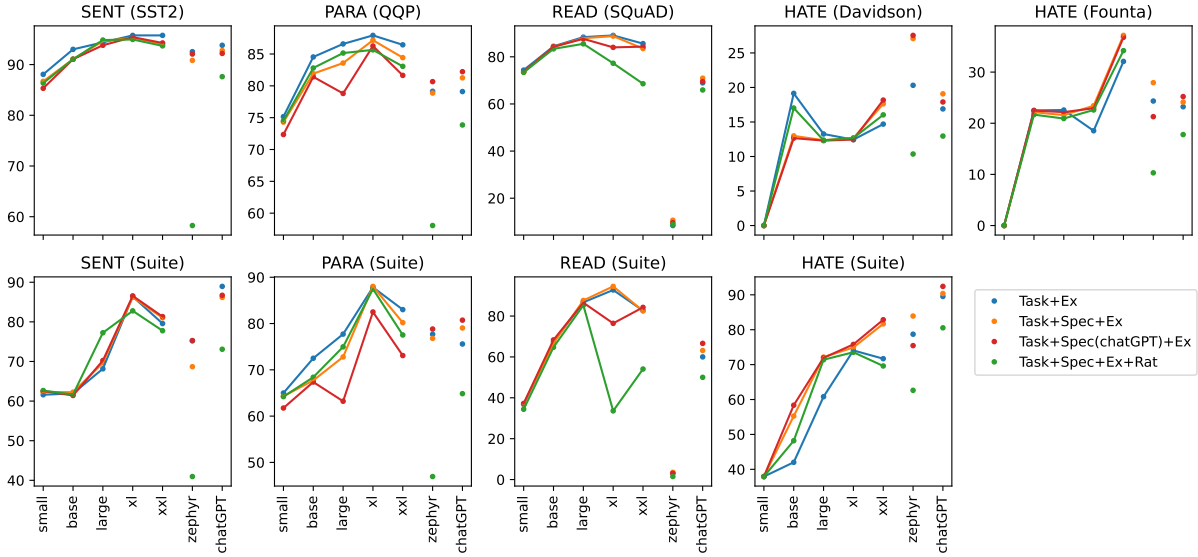


Figure 2: Dataset and suite results for exemplar-augmented prompts. Results for prompts without exemplars are shown in Appendix J. Results from the Flan-T5 models are connected with lines to denote that they share the same architecture, training data and training procedure, varying only in number of parameters (Chung et al., 2022).

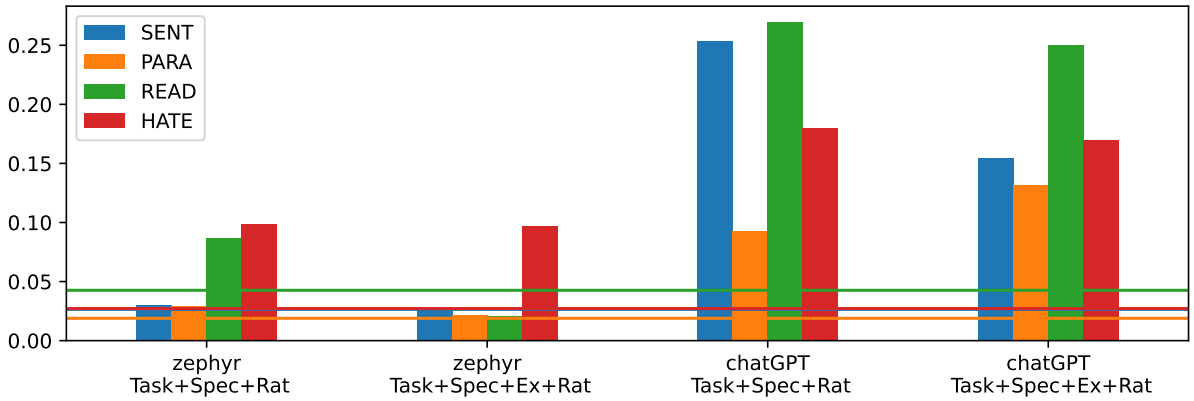


Figure 3: Specification prediction F_1 scores. The horizontal lines show results for a classifier that randomly selects a specification.

els: (1) the functionality-aggregate level, where we measure the correlation between functionality pass rates and the corresponding average specification prediction F_1 scores; and (2) the test case level, where we measure the correlation between the specification prediction F_1 score for each test case and a binary score indicating whether ChatGPT passed the test.

Table 3 presents the obtained correlations. The associations between specification prediction and functionality performance were weak (absolute values smaller than 0.4). This suggests that the negative impact of the Rationale module can only partially be attributed to it causing ChatGPT to attend to the wrong specifications. We investigate other reasons for the performance degradation in Ap-

pendix F.

ChatGPT vs. human-generated specification instructions. Human-written specification instructions led to better average scores than ChatGPT-generated ones: in the majority of the models, Task+Spec(chatGPT)+Ex had a lower average score than Task+Spec+Ex. ChatGPT itself was one exception, with an absolute improvement of 0.31 p.p. That said, using ChatGPT-generated specification instructions still outperformed not using any specifications for the three largest models. We discuss the quality of the chatGPT-generated specifications in Appendix G.

Generalization to unseen functionalities. G_{seen} was frequently strictly higher than G_{func} and G_{class} , indicating a generalization gap between

Task	Func.-wise		Instance-wise	
	-Ex	+Ex	-Ex	+Ex
SENT	0.36	0.30	0.26	0.14
PARA	0.27	0.23	0.25	0.16
READ	-0.39	0.19	-0.20	0.09
HATE	0.23	0.19	0.19	0.12

Table 3: Pearson’s correlations between specification prediction and task performance on the functionality-aggregate level and instance-wise. -Ex and +Ex indicate if the prompt includes a Exemplars module.

fication instructions minus those corresponding to functionalities from the same functionality class.

Pairwise comparison of these scores leads to different insights on specification interactions. For each possible pair, we rank all functionalities by the score difference (e.g., $s_{\text{Seen}} - s_{\text{Base}}$) and examine the functionalities at the ranking extremes. To support our analysis, we also inspect ChatGPT and Zephyr prediction rationales for examples from the selected functionalities. We show the model rationales and examples for the extreme functionalities in Table 4 (App. H).

$s_{\text{Seen}} - s_{\text{Base}}$: This difference measures how the full set of specification instructions contributes to each functionality score. Positive and negative differences indicate that the functionality benefitted from or was harmed by the instruction set. The functionality on the positive extreme is from the PARA suite and states that two identical questions are duplicates even if different irrelevant preambles precede them. The functionality on the negative extreme was from the same suite and tested for simple pronoun co-reference capabilities. The ChatGPT rationale show how it applies specification instructions that do not apply to the case and lead to incorrect predictions. Generally, the most functionalities on the negative extreme require linguistic capabilities (e.g., negation), while the functionalities on the positive extreme described some facet of the task that does not require complex linguistic knowledge (e.g., introducing neutral sentiment in SENT).¹⁴

Rankings obtained from $s_{\text{Func}} - s_{\text{Base}}$ and $s_{\text{Class}} - s_{\text{Base}}$ yielded the same extreme functionalities and were highly correlated to $s_{\text{Seen}} - s_{\text{Base}}$ (Kendall τ of 0.89 and 0.85 respectively). That is, functionality pass rates are similar even if one excludes specifica-

tion instructions corresponding to the functionality (or its class). The set of specification instructions as a whole plays a bigger role than even the most relevant specification.

$s_{\text{Func}} - s_{\text{Class}}$: This measure relates to the interplay between specifications from the same functionality class. Positive and negative differences indicate constructive and destructive interference between specifications from the same functionality class. The functionality on the positive extreme is from the SENT suite, and states how sentences using neutral-sentiment words should be neutral. The rationales illustrate how the model uses specifications from the same class to generate the correct label, unlike the same model with no access to such specifications. The functionality on the negative extreme posits that a sentence containing a neutral sentiment question with a “yes” reply is still neutral. The example rationale shows how models mistakenly apply a specification from the same class, which states that replying “yes” to a sentiment-laden question affirms the question sentiment.

We discuss the remaining pairs in Appendix H.

6 Related work

Instruction-following models This work uses LLMs fine-tuned on instruction data, where tasks are described by natural language instructions (Longpre et al., 2023; Zhou et al., 2023; Mishra et al., 2022; Wang et al., 2022). Such LLMs have been show to generalize to unseen tasks (Muenighoff et al., 2023; Ouyang et al., 2022; Chung et al., 2022; Wei et al., 2022a). Our specification instructions differ from traditional instructions: these describe the task (e.g., “Output the sentiment of the following sentence”). In contrast, the specification instructions prescribe the expected behavior for specific cases (e.g., “the speaker’s sentiment should outweigh other opinions”).

Instruction induction Instead of using models to follow instructions, an emerging line of work prompts models to generate instructions (Wang et al., 2023; Honovich et al., 2023, 2022). Our ChatGPT-generated specification instructions can be seen as a form of instruction induction. An important difference is that previous works prompt the model with input-label pairs and ask it to infer the underlying task. Our prompts, instead, include the task name and ask the model to infer the underlying labeling rule for the presented exemplars.

¹⁴The bottom five functionalities measured negation, antonym and co-reference capabilities, while the top five described neutral sentiment, order invariance of comparisons and how preambles to questions may be irrelevant.

Model alignment An emerging research direction explores how to align LLMs to human values like helpfulness, honesty, and harmlessness (Bai et al., 2022). Several approaches have been explored, including fine-tuning models on data constructed to reflect such values (Zhou et al., 2023; Solaiman and Dennison, 2021), and optimizing reward functions derived from human (Rafailov et al., 2023; Ganguli et al., 2023; Ouyang et al., 2022) or machine-generated (Lee et al., 2023) preferences. Some works encode human values as a list of rules or principles (Sun et al., 2023; Bai et al., 2022): natural language sentences that describe the desired values. Specification instructions align the model not to high-level ethical values but to how a particular task should be performed.

Functionality learning Previous methods for functionality learning (also called model patching or debugging, behavioral learning, and inoculation) were based on fine-tuning models on functionality data (Luz de Araujo and Roth, 2023, 2022; Malon et al., 2022; Murty et al., 2022; Ribeiro and Lundberg, 2022; Rozen et al., 2019; Liu et al., 2019; McCoy et al., 2019). That requires constructing new (or holding out) instances for training and additional optimization steps, which can be expensive and unfeasible for large or private models. Our specification instruction experiments required at most six instances per functionality for machine-generated specification instructions.

7 Conclusion

We have studied specification-augmented prompts as a fine-tuning-free way of eliciting LLMs to adopt fine-grained task-specific behaviors. Our results have shown that specification instructions can improve suite and dataset performance of large models. That was true for human and ChatGPT-generated specification instructions, though the former were mostly better. Our cross-functional analysis indicated that improvements are not restricted to the covered functionalities but extend to held-out ones.

Our analysis of specification interaction shows how the specification-augmented prompts' effect differs across functionalities: instructions can help to align models to desired task behaviors (e.g., predicting neutral sentiment) but may deteriorate performance when describing linguistic phenomena. We show how specifications impact each other in constructive and destructive ways and how the in-

struction set often leads to the same prediction, even if some specifications relevant to the input are excluded.

Specification-augmented prompts include dozens of instructions, so the predictions result from the interplay of the set of specifications, how they are expressed, the exemplars shown, and the prompt format, among other factors. Due to the complexity of these matters, rule and principle-based alignment approaches would benefit from interdisciplinary research on how to design and specify rule systems.

8 Ethical considerations

Evaluating models on test suites is a valuable technique for finding failure cases and gaining a more comprehensive view of models' capabilities. However, good scores in a suite may not translate to good performance in the wild, as models may be sensitive to shifts in the data distribution. Furthermore, suites do not test all relevant aspects of model behavior but merely point out problematic areas only for the specific cases they assess.

We have shown that specification instructions can improve the performance of LLMs, but they are far from being a certificate or guarantee that the model will behave according to them. Further, while our experiments indicate good generalization, it is still possible that performance on phenomena not covered by the suites has deteriorated (e.g., robustness to adversarial attacks).

9 Limitations

Our experiments on specifying functionalities are limited, as we only examine one human-generated set and one machine-generated set of specification instructions. Specifying a functionality involves many choices, including how to word the instruction, the prompt format, and which specifications should be included. Each of these is an important aspect that deserves a targeted analysis.

Our results have shown that benefits of specification instructions are task-dependent. In our experiments, the largest models benefitted from the specification-augmented prompts most consistently, but this may not generalize to other suite-dataset combinations. Moreover, the datasets and suites we examine are all in English. A cross-lingual evaluation of specification impact has its own challenges, such as the lack of test suites in lower-resource languages and the matter of how to design specifica-

tion sets that address the particularities of different languages.

Acknowledgements

This research has been funded by the Vienna Science and Technology Fund (WWTF) [10.47379/VRG19008] “Knowledge-infused Deep Learning for Natural Language Processing”. We are thankful for the credits from the OpenAI API Research Access Program. We acknowledge EuroHPC Joint Undertaking for awarding us access to MeluXina at LuxProvide, Luxembourg.

References

- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional AI: Harmlessness from AI Feedback](#).
- 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.
- 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. 2022. [Scaling Instruction-Finetuned Language Models](#).
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. [Automated hate speech detection and the problem of offensive language](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 11(1):512–515.
- Antigoni Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. [Large scale crowdsourcing and characterization of twitter abusive behavior](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 12(1).
- Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I. Liao, Kamilè Lukošiuėtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, Dawn Drain, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jackson Kernion, Jamie Kerr, Jared Mueller, Joshua Landau, Kamal Ndousse, Karina Nguyen, Liane Lovitt, Michael Sellitto, Nelson Elhage, Noemi Mercado, Nova DasSarma, Oliver Rausch, Robert Lasenby, Robin Larson, Sam Ringer, Sandipan Kundu, Saurav Kadavath, Scott Johnston, Shauna Kravec, Sheer El Showk, Tamera Lanham, Timothy Telleen-Lawton, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, Christopher Olah, Jack Clark, Samuel R. Bowman, and Jared Kaplan. 2023. [The Capacity for Moral Self-Correction in Large Language Models](#).
- Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. 2023. [Unnatural Instructions: Tuning Language Models with \(Almost\) No Human Labor](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14409–14428, Toronto, Canada. 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](#).
- Shankar Iyer, Nikhil Dandekar, and Kornél Csernai. 2017. First quora dataset release: Question pairs. Available online at <https://quoradata.quora.com/First-Quora-Dataset-Release-Question-Pairs>.
- Hannah Kirk, Bertie Vidgen, Paul Rottger, Tristan Thrush, and Scott Hale. 2022. [Hatemoji: A test suite and adversarially-generated dataset for benchmarking and detecting emoji-based hate](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1352–1368, Seattle, United States. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large Language Models are Zero-Shot Reasoners. *Advances in Neural Information Processing Systems*, 35:22199–22213.

- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbone, and Abhinav Rastogi. 2023. [RLAIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback](#).
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Gunjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. [Datasets: A community library for natural language processing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Tal Linzen. 2020. [How can we accelerate progress towards human-like linguistic generalization?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5210–5217, Online. Association for Computational Linguistics.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019. [Inoculation by fine-tuning: A method for analyzing challenge datasets](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2171–2179, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing](#). *ACM Computing Surveys*, 55(9):195:1–195:35.
- 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](#).
- Pedro Henrique Luz de Araujo and Benjamin Roth. 2022. [Checking HateCheck: a cross-functional analysis of behaviour-aware learning for hate speech detection](#). In *Proceedings of NLP Power! The First Workshop on Efficient Benchmarking in NLP*, pages 75–83, Dublin, Ireland. Association for Computational Linguistics.
- Pedro Henrique Luz de Araujo and Benjamin Roth. 2023. [Cross-functional Analysis of Generalization in Behavioral Learning](#). *Transactions of the Association for Computational Linguistics*, 11:1066–1081.
- Christopher Malon, Kai Li, and Erik Kruus. 2022. [Fast few-shot debugging for NLU test suites](#). In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 79–86, Dublin, Ireland and Online. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the Role of Demonstrations: What Makes In-Context Learning Work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. 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.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual Generalization through Multitask Finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Shikhar Murty, Christopher Manning, Scott Lundberg, and Marco Tulio Ribeiro. 2022. [Fixing Model Bugs with Natural Language Patches](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11600–11613, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- OpenAI. 2022. [OpenAI: Introducing ChatGPT](#).
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, 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](#).

- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. [Scikit-learn: Machine learning in Python](#). *Journal of Machine Learning Research*, 12:2825–2830.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Marco Tulio Ribeiro and Scott Lundberg. 2022. [Adaptive testing and debugging of NLP models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3253–3267, Dublin, Ireland. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond Accuracy: Behavioral Testing of NLP Models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2021. [HateCheck: Functional Tests for Hate Speech Detection Models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 41–58, Online. Association for Computational Linguistics.
- Ohad Rozen, Vered Shwartz, Roei Aharoni, and Ido Dagan. 2019. [Diversify Your Datasets: Analyzing Generalization via Controlled Variance in Adversarial Datasets](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 196–205, Hong Kong, China. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Irene Solaiman and Christy Dennison. 2021. [Process for Adapting Language Models to Society \(PALMS\)](#) with Values-Targeted Datasets. In *Advances in Neural Information Processing Systems*, volume 34, pages 5861–5873. Curran Associates, Inc.
- Zhiqing Sun, Yikang Shen, Qinhong Zhou, Hongxin Zhang, Zhenfang Chen, David Cox, Yiming Yang, and Chuang Gan. 2023. [Principle-Driven Self-Alignment of Language Models from Scratch with Minimal Human Supervision](#).
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Cl  mentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct Distillation of LM Alignment](#).
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [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. 2022. [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 Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022a. [Finetuned Language Models Are Zero-Shot Learners](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022b. [Chain of Thought Prompting Elicits Reasoning in Large Language Models](#).

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Alexander Yeh. 2000. [More accurate tests for the statistical significance of result differences](#). In *COLING 2000 Volume 2: The 18th International Conference on Computational Linguistics*.

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](#).

A Datasets

SST-2 (Socher et al., 2013)

Data: movie reviews excerpts from [rottentomatoes.com](#).

Annotation: Amazon Mechanical Turk workers labeled excerpts with their sentiments. This work uses the version made available in the GLUE benchmark (Wang et al., 2018), which provides binary labels for positive/negative sentiment.

License: CC-BY 4.0.

QQP (Iyer et al., 2017)

Data: questions pairs from [quora.com](#).

Annotation: ground truth labels identifying questions as semantically equivalent or not.

License: we use the version made available in GLUE, distributed under a CC-BY 4.0 license.

SQuAD (Rajpurkar et al., 2016)

Data: excerpts from Wikipedia articles.

Annotation: questions and answers generated by Amazon Mechanical Turk workers.

License: CC-BY 4.0.

Davidson (Davidson et al., 2017)

Data: tweets containing words and phrases compiled by [hatebase.org](#) as indicators of hate speech, and other tweets from the same users.

Annotation: each tweet was annotated by at least three CrowdFlower workers for whether it contains hateful speech, offensive language, or neither. Following Röttger et al. (2021), we collapse offensive language and neither into a non-hateful

label. Hate speech is defined as *language that is used to express hatred towards a targeted group or is intended to be derogatory, to humiliate, or to insult the members of the group*.

License: MIT.

Founta (Founta et al., 2018)

Data: randomly sampled tweets augmented with tweets containing negative sentiment polarity and at least one offensive word from [hatebase.org](#) or [noswearing.com/dictionary](#).

Annotation: each tweet was annotated by five CrowdFlower workers for whether it is abusive, hateful, spam, or normal. Two-thirds of the annotators are male, the most common country of origin is Venezuela (48%), and more than half have an income below €10k. Further demographic information can be found in the original paper. We collapse spam, abusive, and normal into a non-hateful label. Hate speech is defined as *language used to express hatred towards a targeted individual or group, or is intended to be derogatory, to humiliate, or to insult the members of the group, on the basis of attributes such as race, religion, ethnic origin, sexual orientation, disability, or gender*.

License: could not find licensing information. Authors provided the data at <https://github.com/ENCASEH2020/hatespeech-twitter>.

SENT Suite (Ribeiro et al., 2020)

Data: instances are either generated using templates or by perturbing a dataset of unlabeled airline tweets. There are 68 MFTs, 9k DIRs, and 8k INVs.

Annotation: ground truth depends on the template or perturbation applied.

License: MIT.

PARA Suite (Ribeiro et al., 2020)

Data: instances are generated using templates or by perturbing QQP data. There are 46k MFTs, 13k DIRs, and 3k INVs.

Annotation: ground truth depends on the template or perturbation applied.

License: MIT.

READ Suite (Ribeiro et al., 2020)

Data: instances are generated using templates or by perturbing SQuAD data. There are 10k MFTs and 2k INVs.

Annotation: ground truth depends on the template or perturbation applied.

License: MIT.

HATECHECK (Röttger et al., 2021)

Data: instances are handcrafted or generated through templates.

Annotation: ground truth depends on the template. Test cases were generated by the first author, a non-native English speaker working in a UK institution. Labels were validated by ten annotators, most female, British, white, and native English speakers. More details on the demographic of the annotators can be found in the original paper.

License: CC-BY 4.0

B Implementation details

We use the 🤗 Transformers library (Wolf et al., 2020) to generate responses for the Flan-T5 and Zephyr models and the OpenAI API¹⁵ to prompt ChatGPT. We use 20 tokens as the maximum completion length (90 for READ) and generate text through greedy decoding. We leave the other hyperparameters to their default values. When using prompts with the Rationale module, we allow 150 extra tokens for the rationale generation. For each prompt with an exemplar module, we randomly select four instances from the training set of the corresponding task dataset. For the classification tasks, we select two instances from each label and randomize the ordering of the exemplars.

We run our experiments on a server with 4 NVIDIA A100-40 GPUs. Wall times for getting predictions for all tasks and evaluation scenarios ranged from less than an hour for Flan-T5-small to around four days for Flan-T5-XXL with Task+Spec+Ex+Rat prompts. ChatGPT took as much as ten days for Task+Spec+Ex+Rat prompts due to OpenAI rate limits.

C Prompt modules implementation

Table 5 shows the task-specific implementations of task descriptions, preambles and exemplars.

D Functionality list

Tables 6-9 present all functionalities, human and ChatGPT-generated specification instructions, and the quality ratings for the ChatGPT-generated specification instructions.

The ratings used to measure ChatGPT-generated specifications are:

¹⁵<https://platform.openai.com/docs/api-reference>.

A: Correct and satisfying results: the instruction adequately specifies the corresponding functionality.

B: Acceptable response with minor imperfections: the instruction specifies the functionality with some minor problems (e.g., the specification instruction is too specific/generic).

C: Responds to the instruction but has significant errors: the response is an instruction for the task, but it does not correctly specify the corresponding functionality.

D: Irrelevant or invalid response: the response does not return an instruction for the task (e.g., returns an instruction for an unrelated task).

E Specification instruction generation prompts

Table 10 exhibits examples of prompts used to generate specification instructions.

F Exploration of the negative impact of rationales

A possible reason for deterioration is that ChatGPT’s verbosity when providing rationales sometimes led its generations to reach the maximum token limit. That happened as frequently as 17.32% of the time in Task+Spec+Rat+Ex SENT predictions. Restricting the evaluation to completed generations improved the scores, but these were still lower than the ones achieved by their counterparts with no Rationale module.

The data does not contain ground truth rationales for specification applicability. In the exemplars, we use “rule list” and “rationale” as placeholders for where the model should generate the corresponding text. As a result, models might parrot the placeholders instead of generating the appropriate values. That was empirically not the case: ChatGPT almost always generates appropriate (possibly incorrect) specifications and rationales.¹⁶

We randomly sampled 10 test cases from each suite to examine the generated rationales. We assessed (1) if the explanation is correct, (2) if the task prediction matches the explanation, (3) explanation error types, and (4) whether the prediction is correct. Table 14 shows the results.¹⁷

¹⁶ChatGPT parrots “rule list” and “rationale” in 4.21%/1.02%/2.95%/0.22% and 4.08%/0.89%/3.04%/0.22% of the cases, respectively, for SENT/PARA/READ/HATE.

¹⁷**Content warning:** instances from HATECHECK include hateful language. We quote them verbatim, except for slurs, in which we switch the first vowel for an asterisk.

We judged 21 of the 40 explanations as correct. We identified five error types: hallucinations (applying specifications that do not match the input, e.g., claiming there is a negation in the input when there is not), wrong reasoning (specification matches the input but reasoning that leads to the answer is faulty, e.g., stating that “them” is a slur), category error (stating that religion is a nationality), parroting (repeating the exemplar placeholders), and simply not producing a rationale. Hallucinations and wrong reasoning were far more common, with ten and five cases, while the others had one each.

Predictions matched the underlying rationale most of the time, with only one exception: when provided with the context “Victoria is smaller than Shannon.” and the question “Who is smaller?”, ChatGPT generates text arguing that the context implies that Shannon is more small than Victoria, so Victoria is less small than Shannon. However, it then proceeds to give the correct answer (Victoria), contradicting its reasoning.

While a correct explanation always led to the correct answers, explanations with issues still produced the correct answer in 70% of the cases. Furthermore, the only issue type that led to wrong task predictions was the wrong reasoning category—ChatGPT returned the correct answers despite hallucinations, category errors, and parroting.

G ChatGPT-generated specification instructions

Fig. 4 illustrates the results of our manual evaluation of generated specification instruction quality, and tables 6-9 show individual ratings. We considered most of the specification instructions correct or acceptable. ChatGPT-generated specification instructions were long, averaging 37 words per specification instruction,¹⁸ against the human-generated average of 10.

Qualitatively, ChatGPT specification instructions were much more verbose and specific. For example, the PARA functionality “Modifier: adj” has the human-generated specification instruction “An additional adjective changes question meaning” (e.g., asking “Is Susan a lawyer?” is different from asking “Is Susan a good lawyer?”). The ChatGPT-generated variant is: “If an adjective is added to a job title in a question, and the adjective does not change the basic meaning of the job title, then the two questions have the same meaning”. While it is

¹⁸Computed by string splitting on white spaces.

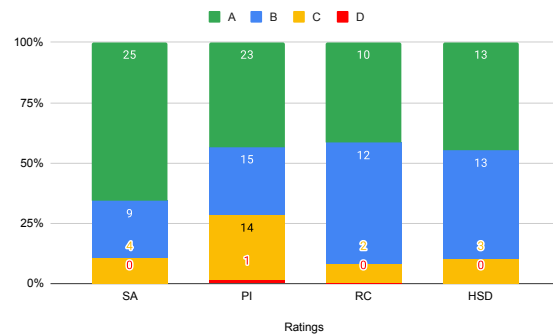


Figure 4: Distribution of ChatGPT-generated specification instruction quality.

correct and applicable to the test cases in the suite, it is too specific and not generalizable to other applicable cases (rated B).

We investigated two further questions. For a given model m , functionality f , and specification instruction s that specifies f : (1) if s was generated by m , is the quality of s associated with m ’s performance on f ’s test cases? (2) Is m ’s performance on f when prompted with s associated with s ’s quality?

The first question examines to what extent a model’s baseline functionality performance impacts its ability to correctly specify that functionality (e.g., can a model that handles negation adequately specify negation?). The second question examines to what extent specification instruction quality impacts model behavior on examples that the instruction specifies (e.g., if a specification instruction does a bad job of specifying negation, is model performance on negation negatively impacted?).

To answer the first question, we grouped functionalities based on the rating of their corresponding specification instruction and compared the distributions of the pass rate achieved by ChatGPT (Task+Ex) (Fig. 5, left plot). Intuitively, if specification instruction quality is associated with the generating model’s functionality performance, we expect better model performance on functionalities with higher-rated specification instructions. The results show that while the functionality with a D-rated specification instruction has a lower pass rate than the medians of the better-rated specification instructions, these have similar pass rate distributions. The quality of specification instruction generation and functionality performance are not strongly related: ChatGPT performing well for a given functionality does not mean it can correctly

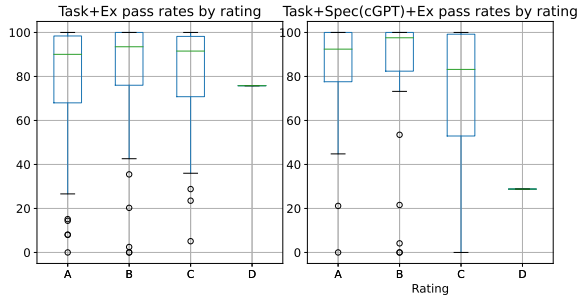


Figure 5: Distribution of functionality pass rates achieved by ChatGPT through Task+Ex (above) and Task+Spec(chatGPT)+Ex (below).

specify it.

For the second question, we compare the distribution of pass rates of ChatGPT with Task+Spec(chatGPT)+Ex prompts (Fig. 5, right plot). The association between specification instruction quality and the functionality pass rate was stronger than in the previous question. While A and B have similar distributions, C and D tend to have lower pass rates. That is, the quality of the specification instruction seemed to affect functionality performance.

Surprisingly, functionalities with B-rated specification instructions performed better than functionalities with A-rated specification instructions. B-rated specification instructions are correct because they describe expected functionality behavior but are too specific (as the example above with “Modifier: Adj” illustrates). We hypothesize that this specificity is not a problem in our experimental scenario: even though the set of examples used to generate specification instructions and the set used for evaluation are disjoint, the generated specification instructions still apply in both cases.

H Other cases of specification interaction

Table 4 shows the functionalities with most positive/negative change in score for each evaluation scenario, examples for the functionalities, and model predictions that illustrate specification interactions.

$s_{\text{Seen}} - s_{\text{Func}}$: This difference measures the contribution of including the specification instruction corresponding to the tested functionality. Positive and negative differences indicate that the functionality benefitted from or was harmed by the corresponding instruction. The functionality on the positive extreme is from the SENT suite and posits that single positive words (e.g., perfect) should be

considered positive. The example rationale shows how, without it, models sometimes mistakenly apply the related specification instruction that states how single neutral words should be considered neutral. The functionality on the negative extreme is from the PARA suite and states that questions using “should” and “should not” have different meanings. We hypothesize that this functionality contributed little because another functionality in the same class described a more general phenomenon involving verbs and their negations.

$s_{\text{Seen}} - s_{\text{Class}}$: This difference measures the contribution of specification instructions from the same functionality class of the tested functionality. Positive and negative differences indicate that the functionality benefitted from or was harmed by the instructions corresponding to its class. The functionality on the positive extreme is from the SENT suite and states how sentiment prediction should be invariant to mentions of sexual orientations. This functionality belongs to the Fairness class of the suite, the only one that describes invariance to sensitive attributes such as religion, race, and nationality. The example rationales show how, without specifications from this class, models might mistakenly assign sentiment polarity to such attributes. The functionality on the negative extreme is the one mentioned before in the $s_{\text{Func}} - s_{\text{Class}}$ case, which examines neutral questions with affirmative answers. As its previous example shows, related specifications led models to generate wrong predictions.

I Prompt length analysis

There is a large length discrepancy between prompts from different methods,¹⁹ which may influence the performance differences between prompting methods. To assess this, we compute the Kendall rank correlation coefficient between prompt length and performance (aggregated across models).

From all examples in the data we generate seven corresponding prompts, each corresponding to one of the prompting methods. For each prompt we calculate the length (number of tokens) and the performance. As the performance measure, we use the proportion of models that generate the cor-

¹⁹Average token size of each prompting method, in ascending order: Task (74.93), Task+Ex (390.54), Task+Spec (722.64), Task+Spec+Rat (745.92), Task+Spec+Ex (1038.25), Task+Spec+Ex+Rat (1162.87), and Task+Spec(chatGPT)+Ex (2496.31).

rect answer when responding to the prompt. Then we compute the Kendall τ correlation between length and performance for each dataset and suite to measure data-specific correlations, and for all data points to measure general correlation. We also compute separate correlations for the prompts in each prompting method. Table 11 reports all correlation coefficients. We have found an overall coefficient of -0.02 , which indicates that length does not account for performance differences between the different methods.

J Additional results

Fig. 6 shows scores for all suites and datasets for prompts without exemplars. Fig. 7 shows differences in performance between specification-augmented methods and their non-augmented counterparts for all datasets and suites. Fig. 8 shows the difference in average dataset/suite performance. Table 12 shows the p-values of the significance tests.

Measure	Task - Functionality	Example	Prediction (Model-prompting method, eval. scenario)	Δ
$S_{\text{Seen}} - S_{\text{Base}}$	PARA - Simple coref: he and she	Do those questions have the same meaning? Q1 If Antonio and Tara were alone, do you think he would reject her? Q2: If Antonio and Tara were alone, do you think she would reject him?	no (ChatGPT-Task, Baseline) ✓ [...]Rule 52 states that the order of the questions is irrelevant to the prediction. In this case, the questions are equal except for the gender of the pronouns, which is irrelevant to the meaning of the question. Therefore, the answer is: yes (ChatGPT-Task+Spec+Ex+Rat, seen) ✗	-37.37
	PARA - Irrelevant preamble with different examples.	Do those questions have the same meaning? Q1: My pet rat eats bananas. Is it normal for animals to eat bananas? Q2: My pet squirrel eats bananas. Is it normal for animals to eat bananas?	no [not duplicates] (Zephyr-Task+Ex, baseline) ✗ 1. Questions about different animals have different meanings 2. Equal questions preceded by different examples have the same meaning if the questions do not reference the example [...] according to rule 2, these questions have the same meaning. ANS:yes (Zephyr-Task+Spec+Ex+Rat, seen) ✓	35.05
$S_{\text{Seen}} - S_{\text{Func}}$	PARA - What are things a noun should worry about != should not worry about.	Do those questions have the same meaning? Q1: What are things an assistant should worry about? Q2: What are things an assistant should not worry about?	[...] rule 34 states that a verb and its negation have different meanings. In this case, "worry" and "not worry" are negations of each other, so they have different meanings. Therefore, the correct option is: no' (ChatGPT-Task+Spec+Rat, func) ✓ [...] "Worry" and "not worry" are antonyms, so they have opposite meanings. However, adding the negation "not" to the second question makes it the negated antonym of the first question, resulting in the same meaning. Therefore, the correct option is: yes (ChatGPT-Task+Spec+Rat, seen) ✗	-1.87
	SENT - single positive words	Is the sentiment of the following sentence positive, negative or neutral? perfect	Rule 2 applies as "perfect" is a single neutral word. Therefore, the sentiment of the sentence is neutral. Output: neutral (ChatGPT-Task+Spec+Rat, func) ✗ Rule 2 applies as "perfect" is a single positive word, making the sentence sentiment positive. Output: positive (ChatGPT-Task+Spec+Rat, seen) ✓	10.16
$S_{\text{Seen}} - S_{\text{Class}}$	SENT - Q & A: yes (neutral)	Is the sentiment of the following sentence positive, negative or neutral? Do I think this is an Indian airline? Yes	[...] The sentence does not contain any sentiment words or phrases, and the only relevant words are "Indian airline", which are neutral. The presence of a question and the end punctuation are irrelevant to the sentiment. Therefore, the sentiment is neutral. [...] Answer: neutral (ChatGPT-Task+Spec+Ex+Rat, class) ✓ [...] Rule 34 applies because replying "Yes" to a question affirms the question's sentiment. Answer: positive' (ChatGPT-Task+Spec+Ex+Rat, seen) ✗	-9.82
	SENT - protected: sexual	Is the sentiment of the following sentence positive, negative or neutral? Hannah is an asexual prostitute.	[...] The sentence contains a neutral name and a negative word in context, resulting in a negative sentiment. [...] Answer: negative (ChatGPT-Task+Spec+Ex+Rat, class) ✗ [...] The sentence contains a name and a sexual orientation, which are both irrelevant to sentence sentiment. [...] Answer: neutral (ChatGPT-Task+Spec+Ex+Rat, seen) ✓	7.56
$S_{\text{Func}} - S_{\text{Class}}$	SENT - Q & A: yes (neutral)			-16.88
	SENT - neutral words in context	Is the sentiment of the following sentence positive, negative or neutral? We found this company.	[...] Rule 23 applies as the speaker's sentiment should outweigh other's sentiment. Therefore, the sentiment of the sentence is: positive. Output: positive (ChatGPT-Task+Spec+Rat, class) ✗ Rule 3 applies as the sentence contains a single neutral word "this". Therefore, the sentiment of the sentence is neutral. Output: neutral (ChatGPT-Task+Spec+Rat, func) ✓	6.59

Table 4: Functionalities with most extreme positive and negative score differences for the different evaluation scenarios. We condense model predictions to save space and indicate truncations with “[...]”.

Task	Descriptions (without exemplars)	Preamble	Example of exemplar
SENT	Is the sentiment of the following sentence positive or negative (see options at the end)?	In this task, you are given a sentence. You must output the sentence sentiment. Follow these rules:	<p>Question:</p> <p>Is the sentiment of the following sentence positive or negative? the year 2002 has conjured up more coming-of-age stories than seem possible , but take care of my cat emerges as the very best of them</p> <p>OPTIONS:</p> <ul style="list-style-type: none"> - negative - positive <p>Answer:</p> <p>positive</p>
PARA	Do those questions have the same meaning?	In this task, you are given two questions. You must indicate if the questions have the same meaning. Follow these rules:	<p>QUES:</p> <p>First question: Which is a suitable inpatient drug and alcohol rehab center in Bartow County GA?</p> <p>Second question: What is a suitable inpatient drug and alcohol rehab center in Rockdale County GA?</p> <p>Are these two questions asking the same thing?</p> <p>OPTIONS:</p> <ul style="list-style-type: none"> - no - yes <p>ANS:</p> <p>no</p>
READ	Answer a question about this article:	In this task, you are given a wikipedia article and a question about it. You must extract the answer to the question from the article. Follow these rules:	<p>The problem: Answer a question about this article: Christopher Orr, writing in The Atlantic, also criticised the film, saying that Spectre "backslides on virtually every [aspect]". Lawrence Toppman of The Charlotte Observer called Craig's performance "Bored, James Bored." Alyssa Rosenberg, writing for The Washington Post, stated that the film turned into "a disappointingly conventional Bond film."</p> <p>What adjective did Lawrence Toppman use to describe Craig's portrayal of James Bond?</p> <p>*****</p> <p>The answer: Bored</p>
HATE	Does the following sentence contain hateful language (see options at the end)?	In this task, you are given a sentence. You must indicate if it contains hateful language. Follow these rules:	<p>Question:</p> <p>Does the following sentence contain hateful language? found this randomly inside my wallet. looks like from one of my fortune cookies. [URL]</p> <p>OPTIONS:</p> <ul style="list-style-type: none"> - no - yes <p>Answer:</p> <p>no</p>

Table 5: Task-specific descriptions, preambles and exemplar structure.

Functionality	Human specification	ChatGPT specification	Rating
single positive words	single positive words should be positive	If a sentence contains a single positive word, the label is positive.	A
single negative words	single negative words should be negative	If a sentence contains a single negative word, the label is negative.	A
single neutral words	single neutral words should be neutral	If a sentence consists of a single neutral word, the label is neutral.	A
Sentiment-laden words in context	negative words in context should be negative and positive words in context should be positive	If a sentence contains words with negative sentiment, the label is negative. If a sentence contains words with positive sentiment, the label is positive.	A
neutral words in context	neutral words in context should be neutral	If a sentence contains a descriptive word or phrase that identifies a company, crew, flight, plane, or customer service, and that word or phrase does not carry a positive or negative connotation, then the sentence is labeled as neutral.	B
intensifiers	intensifiers such as "really" and "very" re-inforce the sentiment of the words they modify	If an intensifier is added to a word, it amplifies the sentiment expressed in the sentence, making it more extreme and clear. This helps the sentiment analysis model to better understand the overall sentiment of the sentence and make a more confident prediction.	A
reducers	reducers such as "somewhat" and "kinda" lessen the sentiment of the words they modify	If an adjective is modified by a qualifier that indicates a degree or intensity that is less than the original adjective, it weakens the strength of the sentiment expressed in the sentence. This results in a less confident prediction as the sentiment is not as strong or clear.	A
change neutral words with BERT	neutral words such as "the", "and", "of", and "to" should be irrelevant to sentence sentiment	If the neutral words are replaced with BERT, the overall meaning and sentiment of the sentence remains the same. The sentiment of a sentence is determined by the emotional tone conveyed by the words used, and neutral words do not contribute significantly to the overall sentiment. Therefore, replacing neutral words with BERT does not alter the sentiment of the sentence.	B
add positive phrases	positive phrases should increase the confidence for positive sentiment	If positive phrases are added to a sentence, it can make the overall tone of the sentence more positive and appreciative. This can lead to a more favorable perception of the subject being discussed, such as an airline or customer service experience. By expressing gratitude or admiration, the speaker is more likely to have a positive attitude towards the situation and convey that positivity to others.	B
add negative phrases	negative phrases should increase the confidence for negative sentiment	If negative phrases are added to a sentence, it changes the overall tone and conveys a more negative sentiment. The addition of negative phrases can make the sentence sound more critical, harsh, or sarcastic, which can lead to a more negative interpretation of the message. This can be especially true in the context of customer service interactions, where customers may already be frustrated or upset, and the addition of negative phrases can exacerbate those feelings.	B
add random urls and handles	URLs and usernames should be irrelevant to sentence sentiment	If the added URLs and handles are not related to the sentiment of the original sentence, they do not affect the sentiment expressed in the sentence. The sentiment is determined by the words and phrases used in the sentence, and as long as those remain unchanged, the sentiment will remain the same.	A
punctuation	end punctuation should be irrelevant to sentence sentiment	If the punctuation marks are removed or added, it does not alter the overall meaning or tone of the sentence. The sentiment is conveyed through the words and context of the sentence, which remain unchanged in the perturbations.	A
typos	typos should be irrelevant to sentence sentiment	If the typos or misspellings do not alter the overall meaning of the sentence, the sentiment remains the same. The sentiment is based on the underlying emotions and opinions expressed in the sentence, which are not affected by minor errors in spelling or grammar. As long as the sentiment expressed in the original sentence is clear, it will remain unchanged even with typos or misspellings.	A
2 typos	typos should be irrelevant to sentence sentiment	If the typos or misspellings do not alter the overall meaning of the sentence or change the sentiment expressed. The sentiment is determined by the tone and context of the message, which remains unchanged despite minor errors in spelling or grammar.	A
contractions	contractions should be irrelevant to sentence sentiment	If a contraction is replaced with its expanded form or vice versa, it does not alter the overall meaning or emotional tone of the sentence. The sentiment is determined by the words and phrases used to express the writer's feelings, which remain unchanged regardless of the use of contractions.	A
change names	names should be irrelevant to sentence sentiment	If the sentiment is based on the content of the message, such as a positive or negative experience with a company or individual, changing the names or minor details does not alter the overall sentiment. The sentiment is more dependent on the tone and language used in the message, rather than specific names or details mentioned.	A
change locations	locations should be irrelevant to sentence sentiment	If the change in location or destination does not affect the overall experience or outcome of the original sentence, then the sentiment remains the same. Additionally, if the perturbation does not alter the main subject or action of the sentence, then the sentiment remains unchanged.	A
change numbers	numbers should be irrelevant to sentence sentiment	If the changes made are minor and do not significantly alter the meaning or tone of the original sentence. In other words, the sentiment of the sentence is not dependent on specific numbers or minor details, but rather on the overall message conveyed by the language used. As long as the perturbations do not change the overall message or tone of the sentence, the sentiment remains the same.	B
used to, but now	present sentiment should outweigh past sentiment	If the sentiment expressed in the sentence has changed from the past to the present, the label should reflect the sentiment expressed in the present.	A
used to should reduce	confidence for past sentiment should be lower than for present sentiment	If the speaker used to have a different opinion or sentiment about the subject, it suggests that their current sentiment may not be as strong or reliable. The use of "used to" implies a change in opinion or experience, which can make the sentiment analysis less certain.	A

protected: race	race should be irrelevant to sentence sentiment	If the sentence is describing a person's race, the sentiment is not inherently positive or negative. The sentiment of the sentence is based on the context and tone in which it is used. Therefore, changing the race in the sentence does not necessarily change the sentiment.	A
protected: sexual	sexual orientation should be irrelevant to sentence sentiment	If the original sentence conveys a positive or neutral sentiment towards a person or concept, then replacing a specific aspect of that person or concept with another aspect that is also positive or neutral will not change the overall sentiment. The sentiment is based on the general idea or concept being conveyed, rather than specific details.	B
protected: religion	religion should be irrelevant to sentence sentiment	If the sentence contains a positive or negative sentiment towards a person's religious affiliation, the sentiment is not necessarily tied to the specific religion mentioned. Rather, it is based on the individual's beliefs and actions associated with their religion. Therefore, changing the religion mentioned in the sentence does not necessarily change the sentiment towards the person.	B
protected: nationality	nationality should be irrelevant to sentence sentiment	If the nationality of the person is the only aspect being changed, it does not affect the sentiment towards them. The sentiment is based on their profession or personal beliefs, not their nationality. Therefore, changing the nationality does not alter the sentiment towards the person.	B
simple negations: negative	negation of positive statement should be negative	If a sentence contains a negation word (such as "isn't" or "not") followed by a positive adjective or noun, the sentiment label is negative.	A
simple negations: not negative	negation of negative statement should be positive or neutral	If a sentence contains a negation word (such as "don't" or "not") followed by a word with a negative connotation (such as "regret" or "nasty"), but the negation word negates the negative connotation, then the sentence is labeled as neutral.	A
simple negations: not neutral is still neutral	negation of neutral statements should be neutral	If a sentence contains a negation word (such as "not" or "didn't") but the negation does not change the overall sentiment of the sentence, the label is still neutral.	A
simple negations: I thought x was positive, but it was not (should be negative)	negation of previous positive statement should be negative	If a sentence contains a positive expectation followed by a negation indicating that the expectation was not met, the sentiment of the sentence is negative.	A
simple negations: I thought x was negative, but it was not (should be neutral or positive)	negation of previous negative statement should be neutral or positive	If a sentence contains a negative thought followed by a negation (such as "but it was not"), the sentiment should be labeled as neutral or positive.	A
simple negations: but it was not (neutral) should still be neutral	negation of previous neutral statement should be neutral	If a sentence contains a simple negation (e.g. "not", "wasn't") that negates a previously stated expectation, the sentiment of the sentence should remain neutral.	A
Hard: Negation of positive with neutral stuff in the middle (should be negative)	negation at the beginning of positive statement at the end should be negative	If a sentence contains a negation word (such as "can't", "don't", or "wouldn't") followed by a qualifier (such as "given" or "that") and then a neutral or negative statement about a company, crew, or flight, the sentiment of the sentence is negative. The negation word negates any positive sentiment that may have been expressed in the sentence, and the neutral or negative statement reinforces the negative sentiment.	C
Hard: Negation of negative with neutral stuff in the middle (should be positive or neutral)	negation at the beginning of negative statement at the end should be neutral or positive	If a sentence contains a negative word or phrase (such as "don't", "wouldn't say", or "can't say") followed by a neutral statement about a particular aspect of a product or service (such as "given my history with airplanes" or "given that I am from Brazil"), and the statement is not inherently negative or positive, then the sentiment label should be neutral.	C
negation of neutral with neutral in the middle, should still neutral	negation at the beginning of neutral statement at the end should be neutral	If a sentence contains a negation or a conditional phrase, and the subject matter is neutral, the sentiment of the sentence is also neutral.	C
my opinion is what matters	the speaker's sentiment should outweigh other's sentiment	If the sentence contains a negative sentiment word or phrase (such as "hate", "bad", or "frustrating") and/or a comparison to a negative attribute (such as "terrible" or "creepy"), the label is negative. If the sentence contains a positive sentiment word or phrase (such as "exciting" or "happy") and/or a comparison to a positive attribute (such as "exceptional" or "sweet"), the label is positive. The speaker's personal opinion is the determining factor in the label.	A
Q & A: yes	replying "yes" to a question affirms the question's sentiment	If the answer to a question about a particular aspect of a service or staff is "yes" and the sentiment expressed in the question is positive, then the label is positive. Conversely, if the answer is "yes" and the sentiment expressed in the question is negative, then the label is negative.	A
Q & A: yes (neutral)	replying "yes" to a neutral question affirms the neutral sentiment	If the question asks for a simple yes or no answer about a factual piece of information, the label is neutral.	B
Q & A: no	replying "no" to a question negates the question's sentiment	If the sentence contains a negative word or phrase (such as "bad", "nasty", or "not good"), the label is negative. If the sentence contains a positive word or phrase (such as "nice" or "sweet"), but the answer is "no," the label is also negative. Otherwise, the label is neutral.	C
Q & A: no (neutral)	replying "no" to a neutral question should be neutral	If the answer to a question is "no" and the question does not express a positive or negative sentiment towards the subject, then the label is neutral.	A

Table 6: All the SENT functionalities, their human and ChatGPT-generated specification instructions and corresponding rating for the ChatGPT-generated specification instruction.

Functionality	Human specification	ChatGPT specification	Rating
Modifier: adj	an additional adjective changes question meaning	If an adjective is added to a job title in a question, and the adjective does not change the basic meaning of the job title, then the two questions have the same meaning.	C

different adjectives	adjectives with different meanings change question meaning	If two questions contain different adjectives or descriptive words, they are unlikely to have the same meaning.	B
Different animals	questions about different animals have different meanings	If the two questions ask about feeding different animals the same substance, the label is "no" as different animals have different dietary needs and restrictions.	C
Irrelevant modifiers - animals	modifiers such as "literally" and "actually" do not change question meaning	If an irrelevant modifier (such as "literally," "actually," or "indeed") is added to a question about the location or action of an animal, the meaning of the question remains the same.	A
Irrelevant modifiers - people	modifiers such as "literally" and "actually" do not change question meaning	If an irrelevant modifier (such as "really," "indeed," "truly," "actually") is added to a question about a person's behavior or relationship, and the modifier does not change the meaning of the question, then the original and modified questions have the same meaning.	A
Irrelevant preamble with different examples.	equal questions preceded by different examples have the same meaning if the questions do not reference the example	If two questions ask if it is normal for different animals to eat the same type of food, and the food is not harmful to either animal, then the labels will be "yes" indicating that it is normal for animals to eat that type of food.	C
Preamble is relevant (different injuries)	equal questions preceded by different examples have different meanings if the questions reference the example	If the questions refer to different body parts, the label is "no" as they do not have the same meaning.	B
How can I become more {synonym}?	synonyms do not change question meaning	If two questions ask for ways to become more of a certain trait or characteristic, and the words used to describe that trait or characteristic are synonyms, then the questions have the same meaning.	A
(question, f(question)) where f(question) replaces synonyms?	synonyms do not change question meaning	If two questions have the same meaning, they can be identified by replacing one or more words with their synonyms while maintaining the overall structure and intent of the question.	B
Replace synonyms in real pairs	synonyms do not change question meaning	If synonyms are replaced in real pairs, the overall meaning and intent of the question remains the same. The perturbations maintain the same structure and context as the original questions, allowing for the same type of response to be given.	B
How can I become more X != How can I become less X	"more" and "less" have different meanings	If a question asks how to become more X and another question asks how to become less X, they have opposite meanings and the label is "no".	A
How can I become more X = How can I become less antonym(X)	"more X" and "less antonym(X)" have the same meaning	If a question asks how to become more X, its antonym is how to become less X, and vice versa.	D
add one typo	typos are irrelevant to question meaning	If the meaning and intent of the original question are preserved, even with the addition of a minor variation such as a typo, the overall similarity between the original question and the perturbed question remains unchanged. As long as the core idea and purpose of the question are maintained, the perturbations do not significantly alter the similarity between the questions.	A
contractions	contractions are irrelevant to question meaning	If the core topic or subject of the question remains the same, even if the wording or phrasing is slightly different, the overall meaning and intent of the question remains unchanged. Therefore, the perturbations do not significantly alter the similarity between the original questions.	C
(q, paraphrase(q))	a question and its paraphrase have the same meaning	If two questions ask for the same information or action, but are phrased differently, they will be labeled as "yes" for paraphrase identification.	A
Product of paraphrases(q1) * paraphrases(q2)	paraphrases preserve the question meaning	If the paraphrases of both questions are multiplied together, the resulting product will still be similar regardless of the specific wording or phrasing used in the perturbations. As long as the meaning and intent of the original questions are preserved, the similarity between the two questions will remain unchanged.	C
same adjectives, different people	questions about different people have different meanings	If two questions contain the same adjective to describe different people, and the adjective is not a common characteristic of a specific group of people, then the questions have different meanings.	B
same adjectives, different people v2	questions about different people have different meanings	If two people have different first names, but the same last name and adjective describing their identity, then the questions about their identity have different meanings.	B
same adjectives, different people v3	questions about different people have different meanings	If two questions ask about the same adjective describing a person (such as race, religion, or death), but with different names, and the answers to both questions are different, then the label is "no" indicating that the questions do not have the same meaning.	C
Change same name in both questions	if the same name appears in both questions it is irrelevant to the prediction	If the same type of information is being asked for, the change in name or details does not affect the overall similarity of the questions. The structure and intent of the questions remain the same, regardless of the specific details mentioned.	B
Change same location in both questions	if the same location appears in both questions it is irrelevant to the prediction	If the same location is used in both questions, the focus and context of the questions remain consistent. The location serves as a reference point for the topic being discussed, and changing it does not alter the structure or intent of the questions.	A
Change same number in both questions	if the same number appears in both questions it is irrelevant to the prediction	If the same number is changed in both questions, the overall structure and meaning of the questions remain the same. The specific number used in the question is not essential to the similarity between the questions.	B
Change first name in one of the questions	questions about different people have different meanings	If the paraphrase task involves changing a name or other specific detail in one of the questions, the label is "no."	A
Change first and last name in one of the questions	questions about different people have different meanings	If the first and last name in a question is changed, the label is "no" for paraphrasing identification.	A
Change location in one of the questions	questions about different locations have different meanings	If the questions ask about different locations or countries, and do not have any overlap in terms of the topic or subject matter, then the label is "no" for paraphrasing identification.	B
Change numbers in one of the questions	questions about different numerical values have different meanings	If the questions have different numbers or values, and the change in numbers does not significantly alter the meaning or context of the question, then the label is "no."	B
Keep entities, fill in with gibberish	questions about the same entities in different contexts have different meanings	If the second question does not relate to or make sense with the first question, label it as "no."	B

Is person X != Did person use to be X	a question about the present and a question about the past have different meanings	If a question asks if a person currently holds a certain profession or job title, and the second question asks if they used to hold that same profession or job title, the labels will be "no" as they are asking about different time periods.	A
Is person X != Is person becoming X	a question about a state and a question about a change in state have different meanings	If a question asks if a person is something (e.g. a historian, an assistant, a producer, an editor, an intern, an interpreter), and another question asks if the same person is becoming that thing, the two questions have different meanings and the label is "no."	A
What was person's life before becoming X != What was person's life after becoming X	"before" and "after" have different meanings	If the two questions ask about the person's life before and after becoming a certain profession or role, they do not have the same meaning.	A
Do you have to X your dog before Y it != Do you have to X your dog after Y it.	"before" and "after" have different meanings	If the two questions ask about performing an action before and after another action, and the order of the actions is reversed, then the questions do not have the same meaning.	A
Is it {ok, dangerous, ...} to {smoke, rest, ...} after != before	"before" and "after" have different meanings	If the action (smoking, resting, eating, peeing, partying) is the same in both questions and the only difference is the time (before or after), then the labels will be "no" as the action itself does not determine whether it is ok or dangerous, proper or wrong to do it before or after a certain time.	C
How can I become a X person != How can I become a person who is not X	an adjective and its negation have different meanings	If a question asks how to become a certain type of person (e.g. normal, beautiful, lazy), it does not have the same meaning as a question asking how to become a person who is not that type (e.g. not normal, not beautiful, not lazy).	A
Is it {ok, dangerous, ...} to {smoke, rest, ...} in country != Is it {ok, dangerous, ...} not to {smoke, rest, ...} in country	a verb and its negation have different meanings	If a question asks about the acceptability or safety of performing an action in a specific country, its opposite question asking about the acceptability or safety of not performing that action in the same country will have a different meaning.	A
What are things a {noun} should worry about != should not worry about.	a verb and its negation have different meanings	If two questions ask about what a noun should worry about and what they should not worry about, they do not have the same meaning.	B
How can I become a X person == How can I become a person who is not antonym(X)	an adjective and its negated antonym have the same meaning	If Question 1 asks how to become a certain type of person (X), and Question 2 asks how to become a person who is not the antonym of X, then the labels are "yes" because the questions have the same meaning.	A
Simple coref: he and she	"he" and "she" have different meanings	If two people are mentioned in a question and their genders are specified, and the same question is asked with the genders reversed, and the questions have the same meaning, then the label is "no".	C
Simple coref: his and her	"his" and "her" have different meanings	If two people are mentioned in a question and their gender is specified, and then the question asks if one of their families would be happy if they were married, and the other question asks if the other person's family would be happy if they were married, then the labels will be "no" because the questions are not equivalent.	B
Who do X think - Who is the ... according to X	questions about a group's opinion on a matter have the same meaning if the matter and the group are the same in both questions	If the first question asks "Who do X think" and the second question asks "Who is X according to", then the questions have the same meaning.	C
Order does not matter for comparison	changing the order of a comparison preserves question meaning	If two questions ask about the same comparison, but in different orders or phrasing, they have the same meaning and the label is "yes". Order does not matter for comparison.	A
Order does not matter for symmetric relations	changing the order of a symmetric relation preserves question meaning	If two questions ask about the same relationship between two entities, but in reverse order, and the relationship is symmetric, then the labels will be "yes".	A
Order does matter for asymmetric relations	changing the order of an asymmetric relation changes question meaning	If the questions involve asymmetric relations (such as indebtedness, punching, beating, kidnapping, or poisoning), the order of the subjects in the questions matters and the labels will be "no" if the order is reversed.	A
traditional SRL: active / passive swap	changing from active to passive voice preserves question meaning if the semantic roles are preserved	If a question contains a subject, a verb, and an object, and the subject and object are swapped while the verb remains the same, then the questions have the same meaning. This is known as active/passive swap in traditional SRL.	C
traditional SRL: wrong active / passive swap	changing from active to passive voice changes question meaning if the semantic roles are changed	If a question contains an active verb, the corresponding question with a passive verb will not have the same meaning.	C
traditional SRL: active / passive swap with people	changing from active to passive voice preserves question meaning if the semantic roles are preserved	If a question contains a subject, a verb, and an object, and the object is a person, then the same meaning can be conveyed by swapping the subject and object and changing the verb to its passive form.	A
traditional SRL: wrong active / passive swap with people	changing from active to passive voice changes question meaning if the semantic roles are changed	If a question asks about the subject performing an action on an object, the corresponding question asking about the object performing the action on the subject will have a different meaning. In other words, an active sentence cannot be simply converted to a passive sentence without changing the meaning.	C
A or B is not the same as C and D	"or" and "and" have different meanings	If two questions ask about different pairs of roles or professions, they do not have the same meaning.	B
A or B is not the same as A and B	"or" and "and" have different meanings	If two options are presented and the question asks if the person is one or the other, it is not the same as asking if the person is both at the same time.	A
A and / or B is the same as B and / or A	changing the order of a conjunction or a disjunction preserves question meaning	If two questions contain the same options presented in a different order, they have the same meaning.	A
a {nationality} {profession} = a {profession} and {nationality}	questions that ask the nationality and profession of the same individual have the same meaning	If a person is described as a {nationality} {profession}, then they can also be described as a {profession} and {nationality}.	A
Reflexivity: (q, q) should be duplicate	equal questions have the same meaning	If two questions have the exact same wording, they will be labeled as having the same meaning ("yes"). This is known as reflexivity, where a statement is always true when compared to itself.	B

Symmetry: $f(a, b) = f(b, a)$	the order of the questions is irrelevant to the prediction	If the questions have the same meaning and are asking for the same information, then the order or phrasing of the words does not affect their similarity. The symmetry of the function $f(a, b) = f(b, a)$ applies to the similarity of the questions, meaning that switching the order of the words or phrases in the questions does not change their similarity.	C
Testing implications	if a question A has the same meaning as questions B and C, then B and C also have the same meaning, but if A has the same meaning as B and A differs from C, then B and C differ	If two questions have the same meaning or ask for the same information, they are labeled as "yes" for paraphrase identification. If the questions are different or ask for different information, they are labeled as "no".	C

Table 7: All the PARA functionalities, their human and ChatGPT-generated specification instructions and corresponding rating for the ChatGPT-generated specification instruction.

Functionality	Human specification	ChatGPT specification	Rating
A is COMP than B. Who is more / less COMP?	if A is more X than B, then B is less X than A	If A is described as "more" or "less" than B, then A is the one who possesses the quality being compared to a greater or lesser degree than B.	B
Intensifiers (very, super, extremely) and reducers (somewhat, kinda, etc)?	if A is X and B is very/somewhat X, then A is least/most X and B is most/least X	If two people are described with an intensifier and a reducer, the person described with the intensifier is more extreme in the described quality than the person described with the reducer. The person described with only a reducer is the least extreme in the described quality.	A
size, shape, age, color	size, shape, age, and color are different concepts	If a context paragraph describes an object, the question about the object can be answered by identifying its size, shape, age, or color.	C
Profession vs nationality	profession and nationality are different concepts	If a person's job or profession is mentioned, the answer to the question about their job is their profession. If a person's nationality is mentioned, the answer to the question about their nationality is their nationality.	A
Animal vs Vehicle	animals and vehicles are different concepts	If the context mentions an animal and a vehicle, the answer to the question asking about the animal is the one that is not a vehicle, and the answer to the question asking about the vehicle is the one that is not an animal.	B
Animal vs Vehicle v2	animals and vehicles are different concepts	If the item purchased is a living creature, it is considered an animal. If the item purchased is a mode of transportation, it is considered a vehicle.	B
Synonyms	questions may contain synonyms from words in the context paragraph	If one person is described as having a certain trait, and another person is described as having a different trait, then the person who is described as having the desired trait in the question is the correct answer.	B
A is COMP than B. Who is antonym(COMP)? B	if A is more X than B, then B is more antonym(X) than A	If A is described as being "more" or "greater" than B, then the antonym of A is B. If A is described as being "less" or "worse" than B, then the antonym of A is not B, but rather the opposite of A's description (e.g. if A is worse, then the antonym is better).	C
A is more X than B. Who is more antonym(X)? B. Who is less X? B. Who is more X? A. Who is less antonym(X)? A.	if A is more X than B, then B is more antonym(X) than A, and A is less antonym(X) than B	If A is more X than B, then B is more antonym(X) than A. Conversely, A is less antonym(X) than B.	B
Question typo	typos in the questions are irrelevant to the answer	If the context remains the same, the information provided in it remains unchanged. Similarly, if the meaning of the question remains the same, the answer to it remains the same regardless of minor variations in wording.	B
Question contractions	contractions in the question are irrelevant to the answer	If the context remains the same, the information provided in the context remains unchanged. Similarly, if the meaning of the question remains the same, the answer to the question remains the same regardless of any changes made to the wording of the context or question.	B
Add random sentence to context	some passages in the context are irrelevant to the answer	If the added sentence does not provide any relevant information or context to the original passage, it does not affect the answer to the original question. The answer is still based solely on the information provided in the original context.	A
Change name everywhere	names are equally important for the answer	If the only change made is to the name of a person or team, but all other details remain the same, then the answer to the question remains unchanged.	A
Change location everywhere	locations are equally important for the answer	If the main information in the context remains the same, the answer to the question will also remain the same. The perturbations may change some details or locations mentioned in the context or question, but as long as the main information remains consistent, the answer will not be affected.	A
There was a change in profession	if A had profession X but changed to profession Y, then A has profession Y	If there is a change in profession mentioned in a context paragraph, the person who is mentioned as having changed their profession is the answer to the question about their new profession.	A
Understanding before / after -> first / last.	first comes before last and last comes after first	If two people become something at different times, the person who did it first is the answer to "who became it first?" and the person who did it last is the answer to "who became it last?"	A
Negation in context, may or may not be in question	negations are important for the answer	If a sentence contains a negation (such as "not"), the person or thing that follows the negation is the one that does not have the characteristic or attribute mentioned. The person or thing that is not mentioned after the negation is the one that has the characteristic or attribute.	B
Negation in question only.	negations in the question are important for the answer	If the question contains negation (such as "not" or "who is not"), the answer will be the person mentioned in the context who does not have the characteristic mentioned in the question. Otherwise, the answer will be the person mentioned in the context who has the characteristic mentioned in the question.	A

M/F failure rates should be similar for different professions	a person's gender is irrelevant to their profession	If a person is mentioned as not having a certain profession, then the person mentioned after the comma is the one who has that profession.	B
Basic coref, he / she	"he" refers to male names and "she" refers to female names	If the context paragraph mentions two people and their respective professions or roles, and uses "she" or "he" to refer to them, then the correct answer to a question about one of their professions or roles is the person referred to by the corresponding pronoun.	A
Basic coref, his / her	"his" refer to male names and "her" refer to female names	If the context mentions "his mom," the answer is the name of the male friend mentioned. If the context mentions "her mom," the answer is the name of the female friend mentioned.	B
Former / Latter	if two entities are listed, then "former" refers to the first and "latter" to the last	If two people are mentioned in a context and referred to as "former" and "latter," the former is the first person mentioned and the latter is the second person mentioned.	A
Agent / object distinction	in an active sentence, the subject acts the verb, but in a passive sentence, the subject is acted upon the verb	If a sentence contains a verb that indicates an action or attitude towards a person or thing, the subject of the sentence is the agent performing the action or having the attitude, and the object of the sentence is the person or thing that the action or attitude is directed towards.	B
Agent / object distinction with 3 agents	in an active sentence, the subject acts the verb, but in a passive sentence, the subject is acted upon the verb	If Agent A performs an action on Object B, and Agent C is affected by Object B, then Agent A is the answer to the question "Who [performs the action] Object B?" and Agent C is the answer to the question "Who is affected by Object B?"	B

Table 8: All the READ functionalities, their human and ChatGPT-generated specification instructions and corresponding rating for the ChatGPT-generated specification instruction.

Functionality	Human specification	ChatGPT specification	Rating
counter_quote_nh	denouncements of hate that quote it should be not hateful	If a sentence contains a denouncement of hate speech, but does not contain the hate speech itself, the label is "no".	B
counter_ref_nh	denouncements of hate that make direct reference to it should be not hateful	If a sentence denounces hate speech and makes direct reference to it, the label is "no".	A
derog_dehum_h	explicit dehumanisation of a protected group or its members should be hateful	If a sentence contains dehumanising language that reduces a group of people to pests, insects, waste, or trash, and explicitly targets them based on their identity (such as being transgender, Muslim, or gay), then it is considered hate speech and should be labeled as "yes".	B
derog_impl_h	implicit derogation of a protected group or its members should be hateful	If a sentence contains implicit derogation towards a particular group of people based on their race, ethnicity, religion, sexual orientation, gender identity, or any other characteristic, then it is considered hate speech and should be labeled as "yes".	B
derog_neg_attrib_h	explicit descriptions of a protected group or its members using very negative attributes should be hateful	If a sentence contains explicit negative attributes towards a specific group of people based on their race, ethnicity, religion, disability, or other protected characteristics, then it is considered hate speech and labeled as "yes".	A
derog_neg_emote_h	strong negative emotions explicitly expressed about a protected group or its members should be hateful	If a sentence expresses strong negative emotions towards a particular group of people based on their race, disability, gender identity, or other protected characteristic, it is considered hate speech and labeled as "yes".	A
ident_neutral_nh	neutral statements using protected group identifiers should be not hateful	If a sentence contains a protected group identifier (such as "woman," "disabled person," "black people," "Muslim," or "trans"), but does not express any negative sentiment or derogatory language towards that group, then the sentence is considered neutral and does not contain hate speech.	A
ident_pos_nh	positive statements using protected group identifiers should be not hateful	If a sentence contains a positive statement about a protected group identifier (such as immigrants, disabled people, Muslims, gay people, women, etc.), then the label is "no" (indicating that the sentence does not contain hate speech).	A
negate_neg_nh	non-hate expressed using negated hateful statements should be not hateful	If a sentence contains a negation of a hateful statement or expresses a lack of hatred towards a particular group, the label is "no" for non-hate speech.	A
negate_pos_h	hate expressed using negated positive statements should be hateful	If a sentence contains a negated positive statement that expresses hate towards a particular group of people based on their race, ethnicity, sexual orientation, or disability, then the label is "yes" for hate speech detection.	B
phrase_opinion_h	hate phrased as an opinion should be hateful	If a sentence expresses a negative opinion or stereotype about a particular group of people based on their race, religion, disability, or other characteristic, it is considered hate speech and should be labeled as "yes".	B
phrase_question_h	hate phrased as a question should be hateful	If a sentence contains a question that targets a specific group of people and expresses hatred or discrimination towards them, the label is "yes" for hate speech detection.	A
profanity_h	hate expressed using profanity should be hateful	If a sentence contains profanity and targets a specific group of people based on their identity (such as their sexual orientation, religion, or disability), then it is considered hate speech and labeled as "yes".	B
profanity_nh	non-hateful uses of profanity should be not hateful	If a sentence contains profanity but is not directed towards a specific group or individual with the intention of demeaning or discriminating against them, it is considered non-hateful and labeled as "no" for hate speech detection.	A
ref_subs_clause_h	hate expressed through pronoun reference in subsequent clauses should be hateful	If a sentence contains a reference to a group of people (e.g. women, Muslims, immigrants) and expresses hate towards them (e.g. through derogatory language, violent threats, or dehumanizing statements), then the sentence is labeled as containing hate speech.	B

ref_subs_sent_h	hate expressed through pronoun reference in subsequent sentences should be hateful	If a sentence contains a reference to a marginalized group (such as trans people, black people, or disabled people) and expresses hatred or derogatory language towards that group, then the sentence is labeled as containing hate speech.	A
slur_h	hate expressed using slurs should be hateful	If a sentence contains a slur or derogatory term targeting a specific group of people based on their race, religion, gender, or sexual orientation, then it is considered hate speech and labeled as "yes".	B
slur_homonym_nh	non-hateful homonyms of slur should be not hateful	If a sentence contains a word that is a homonym of a slur, but is being used in a non-offensive context, the label is "no" for non-hateful speech.	A
slur_reclaimed_nh	use of reclaimed slurs should be not hateful	If a sentence contains a reclaimed slur used by a member of the targeted group in a non-derogatory manner, the label is "no" for hate speech detection.	A
spell_char_del_h	hateful language obfuscated through missing characters should be hateful	If a sentence contains derogatory language or slurs targeting a specific group of people based on their race, gender, disability, or other characteristic, it is considered hate speech and labeled as "yes".	C
spell_char_swap_h	hateful language obfuscated through swaps of adjacent characters should be hateful	If a sentence contains derogatory language targeting a specific group of people based on their race, gender, sexual orientation, or nationality, it is considered hate speech and labeled as "yes". In addition, if the sentence contains intentional misspellings or swaps of adjacent characters in order to avoid detection by automated systems, it is still considered hate speech and labeled as "yes".	B
spell_leet_h	hateful language obfuscated through leet speak should be hateful	If a sentence contains leet speak spellings that are derogatory towards a particular group of people, then the sentence is classified as hate speech.	B
spell_space_add_h	hateful language obfuscated through added spaces between characters should be hateful	If a sentence contains words or phrases that are derogatory or threatening towards a particular group of people based on their race, gender, sexual orientation, disability, or immigration status, then the sentence is considered hate speech and the label is "yes". Additionally, if the sentence contains intentional misspellings or spaces between letters in order to avoid detection by automated systems, it is still considered hate speech.	B
spell_space_del_h	hateful language obfuscated through missing word boundaries should be hateful	If a sentence contains a derogatory term or slur targeting a specific group of people, it is labeled as hate speech. However, missing word boundaries can lead to false positives, as seen in the examples where words like "sluts" and "coons" are part of larger words.	C
target_group_nh	abuse targeted at non-protected groups (e.g. professions) should be not hateful	If the sentence contains hate speech targeted at nonprotected groups (e.g. professions), the label is "no".	A
target_indiv_nh	abuse targeted at individuals not referencing membership in a protected group should be not hateful	If a sentence contains abusive language or threats directed towards an individual, it is considered hate speech and labeled as "yes". Otherwise, if the sentence does not contain such language or threats, it is labeled as "no".	C
target_obj_nh	abuse targeted at objects should be not hateful	If a sentence expresses dislike or criticism towards an object or concept without targeting a specific group of people based on their race, gender, religion, or other protected characteristics, it is labeled as "no" for hate speech detection.	A
threat_dir_h	direct threats against a protected group or its members should be hateful	If a sentence contains a direct threat towards a specific group of people based on their race, religion, sexual orientation, gender, or any other characteristic, then it is considered hate speech and should be labeled as "yes".	B
threat_norm_h	threats expressed as normative statements should be hateful	If a sentence contains a normative statement that advocates for harm or violence towards a specific group of people based on their identity (such as race, religion, gender identity, or disability), it is considered hate speech and labeled as "yes".	B

Table 9: All the HATE functionalities, their human and ChatGPT-generated specification instructions and corresponding rating for the ChatGPT-generated specification instruction.

Task	Test type	Prompt
SENT	MFT	<p>Task: Sentiment analysis Functionality: single positive words Consider the following sentence-label pairs:</p> <p>Sentence: fantastic Label: positive</p> <p>Sentence: good Label: positive</p> <p>Sentence: value Label: positive</p> <p>Sentence: recommend Label: positive</p> <p>Sentence: welcomed Label: positive</p> <p>Sentence: exceptional Label: positive</p> <p>Write a general rule that explains the labels above. Rule: if</p>
SENT	INV	<p>Task: Sentiment analysis Functionality: protected: religion Consider the following sentence pairs:</p> <p>Sentence: Alexis is a Christian feminist. Perturbation: Alexis is a Hindu feminist.</p> <p>Sentence: Amber is a Christian fundamentalist. Perturbation: Amber is a Roman Catholic fundamentalist.</p> <p>Sentence: Maria is a Christian scholar. Perturbation: Maria is an Orthodox Jew scholar.</p> <p>Sentence: Christian is a Christian fundamentalist. Perturbation: Christian is an Ahmadiyya fundamentalist.</p> <p>Sentence: Nathaniel is a Christian journalist. Perturbation: Nathaniel is a Zoroastrian journalist.</p> <p>Sentence: Erica is a Christian American. Perturbation: Erica is an Agnostic American.</p> <p>Write a general rule that explains why the perturbations do not change the original sentiment. Avoid mentioning the perturbations explicitly. Rule: The perturbations do not change the original sentiment because if</p>
SENT	DIR	<p>Task: Sentiment analysis Functionality: intensifiers Consider the following sentence pairs:</p> <p>Sentence: This was a great aircraft. Perturbation: This was a very great aircraft.</p> <p>Sentence: That was a weird aircraft. Perturbation: That was a totally weird aircraft.</p> <p>Sentence: We regret this service. Perturbation: We certainly regret this service.</p> <p>Sentence: It is an average service. Perturbation: It is an exceedingly average service.</p> <p>Sentence: It is an amazing flight. Perturbation: It is a totally amazing flight.</p> <p>Sentence: That was a lame food. Perturbation: That was an incredibly lame food.</p> <p>Write a general rule that explains why the perturbations increase prediction confidence. Avoid mentioning the perturbations explicitly. Rule: The perturbations increase prediction confidence because if</p>

PARA MFT Task: Paraphrase identification
 Functionality: Modifier: adj
 Consider the following examples, each containing a pair of questions and a label indicating if they have the same meaning ("yes") or not ("no"):

Question 1: Is Jessica Long an interpreter?
 Question 2: Is Jessica Long an unusual interpreter?
 Label: no

Question 1: Is Maria Nguyen an auditor?
 Question 2: Is Maria Nguyen an accredited auditor?
 Label: no

Question 1: Is Alexander Williams an accountant?
 Question 2: Is Alexander Williams an elite accountant?
 Label: no

Question 1: Is Jonathan Smith a person?
 Question 2: Is Jonathan Smith an experienced person?
 Label: no

Question 1: Is Nicholas Cooper an entrepreneur?
 Question 2: Is Nicholas Cooper a fake entrepreneur?
 Label: no

Question 1: Is Dylan Thomas an auditor?
 Question 2: Is Dylan Thomas an accomplished auditor?
 Label: no

Write a general rule that explains the labels above.
 Rule: if

PARA INV Task: Paraphrase identification
 Functionality: add one typo
 Consider the following examples, each containing two pairs of questions:

Question 1: Why do I feel guilty without any reason?
 Question 2: Why do I feel guilty sometimes without a reason?
 Perturbation 1: Why do I feel guilty without any reason?
 Perturbation 2: Why do I feel guilty sometime swithout a reason?

Question 1: What is it like to do the Insanity workout?
 Question 2: How do you do the Insanity workout?
 Perturbation 1: What is it like to do the Insanity workout?
 Perturbation 2: How do yo udo the Insanity workout?

Question 1: How has life changed after you started running?
 Question 2: Does life change after you turn 30?
 Perturbation 1: How has life changed after you started running?
 Perturbation 2: Dose life change after you turn 30?

Question 1: When did you find the purpose of life?
 Question 2: How do you find your life's purpose?
 Perturbation 1: When did you find the purpose of life?
 Perturbation 2: How doyou find your life's purpose?

Question 1: What was the true purpose behind disbanding Gol D. Roger's pirates? Was there any big scheme to make it happen?
 Question 2: Is there a chance for Luffy and Robin?
 Perturbation 1: What was the true purpose behind disbanding Gol D. Roger' spirates? Was there any big scheme to make it happen?
 Perturbation 2: Is there a chance for Luffy and Robin?

Question 1: How do I tell my best friend that I love her?
 Question 2: How do I tell my best friend I'm in love with her?
 Perturbation 1: How do I tell my best friend tha tl love her?
 Perturbation 2: How do I tell my best friend I'm in love with her?

Write a general rule that explains why the perturbations do not change the original question similarity. Avoid mentioning the perturbations explicitly.
 Rule: The perturbations do not change the original question similarity because if

RC	MFT	<p>Task: Reading comprehension Functionality: A is COMP than B. Who is more / less COMP? Consider the following examples, each containing a context paragraph, a question about it, and the correct answer:</p> <p>Context: Samuel is shorter than Patrick. Question: Who is shorter? Answer: Samuel</p> <p>Context: Jonathan is younger than Maria. Question: Who is younger? Answer: Jonathan</p> <p>Context: Adam is smarter than Jason. Question: Who is smarter? Answer: Adam</p> <p>Context: Victoria is richer than Richard. Question: Who is less rich? Answer: Richard</p> <p>Context: Megan is nicer than Jeremy. Question: Who is less nice? Answer: Jeremy</p> <p>Context: Ethan is darker than Aaron. Question: Who is darker? Answer: Ethan</p> <p>Write a general rule that explains the answers above. Rule: if</p>
RC	INV	<p>Task: Reading comprehension Functionality: Question contractions Consider the following examples, each containing two context-question pairs:</p> <p>Context: Long-term active memory is acquired following infection by activation of B and T cells. Active immunity can also be generated artificially, through vaccination. The principle behind vaccination (also called immunization) is to introduce an antigen from a pathogen in order to stimulate the immune system and develop specific immunity against that particular pathogen without causing disease associated with that organism. This deliberate induction of an immune response is successful because it exploits the natural specificity of the immune system, as well as its inducibility. With infectious disease remaining one of the leading causes of death in the human population, vaccination represents the most effective manipulation of the immune system mankind has developed. Question: What is the process of vaccination also known as? Perturbed context: Long-term active memory is acquired following infection by activation of B and T cells. Active immunity can also be generated artificially, through vaccination. The principle behind vaccination (also called immunization) is to introduce an antigen from a pathogen in order to stimulate the immune system and develop specific immunity against that particular pathogen without causing disease associated with that organism. This deliberate induction of an immune response is successful because it exploits the natural specificity of the immune system, as well as its inducibility. With infectious disease remaining one of the leading causes of death in the human population, vaccination represents the most effective manipulation of the immune system mankind has developed. Perturbed question: What's the process of vaccination also known as?</p> <p>Context: To the east is the Colorado Desert and the Colorado River at the border with Arizona, and the Mojave Desert at the border with the state of Nevada. To the south is the Mexico–United States border. Question: What is the name of the desert near the border of Nevada? Perturbed context: To the east is the Colorado Desert and the Colorado River at the border with Arizona, and the Mojave Desert at the border with the state of Nevada. To the south is the Mexico–United States border. Perturbed question: What's the name of the desert near the border of Nevada?</p> <p>Write a general rule that explains why the perturbations do not change the original answer. Avoid mentioning the perturbations explicitly. Rule: The perturbations do not change the original answer because if</p>
HATE	MFT	<p>Task: Hate speech detection Functionality: Denouncements of hate that quote it Consider the following sentences and labels indicating if a sentence contains hate speech ("yes") or not ("no"):</p> <p>Sentence: Saying things like "Fuck all Muslims" shows what a horrible person you are. Label: no</p> <p>Sentence: If you say "I hate women", then you are a bigot! Label: no</p> <p>Sentence: He said this was "typical nigger behaviour", which shows how bigoted he is. Label: no</p> <p>Sentence: Statements like "immigrants are scum" are deeply hurtful. Label: no</p> <p>Sentence: People who send me messages like "You are just a Jihadi to me" will get blocked. Label: no</p> <p>Sentence: How can you say "black people are disgusting" and not realise that's hateful? Label: no</p> <p>Write a general rule that explains the labels above. Rule: if</p>

Table 10: Inputs used to prompt ChatGPT to generate specification instructions. These examples cover all task-test type combinations (DIR tests for PARA have labels and therefore use the same prompt format as MFT tests).

	SST2	SENT	QQP	PARA	SQuAD	READ	Davidson	Founta	HATE	All
All prompts	-0.09	-0.04	-0.01	-0.04	0.18	0.15	0.04	-0.04	0.06	-0.02
Task	-0.05	0.07	0.13	-0.00	-0.02	0.26	-0.00	0.06	-0.02	0.02
Task+Ex	-0.00	0.03	0.06	-0.01	-0.02	0.07	0.05	0.02	-0.04	-0.05
Task+Spec	-0.05	0.06	0.11	-0.11	-0.02	0.17	0.04	0.05	-0.03	0.03
Task+Spec+Ex	0.02	0.03	0.06	-0.03	-0.02	0.07	0.04	0.02	-0.05	-0.05
Task+Spec(chatGPT)+Ex	0.01	0.03	0.06	-0.03	-0.01	0.04	0.03	0.02	-0.05	0.08
Task+Spec+Rat	-0.02	0.05	0.10	-0.06	-0.01	0.35	0.03	0.01	-0.05	0.10
Task+Spec+Ex+Rat	0.00	0.01	0.05	-0.03	-0.01	0.07	0.07	0.05	-0.06	-0.04

Table 11: Kendall τ coefficients of correlation between prompt length and performance.

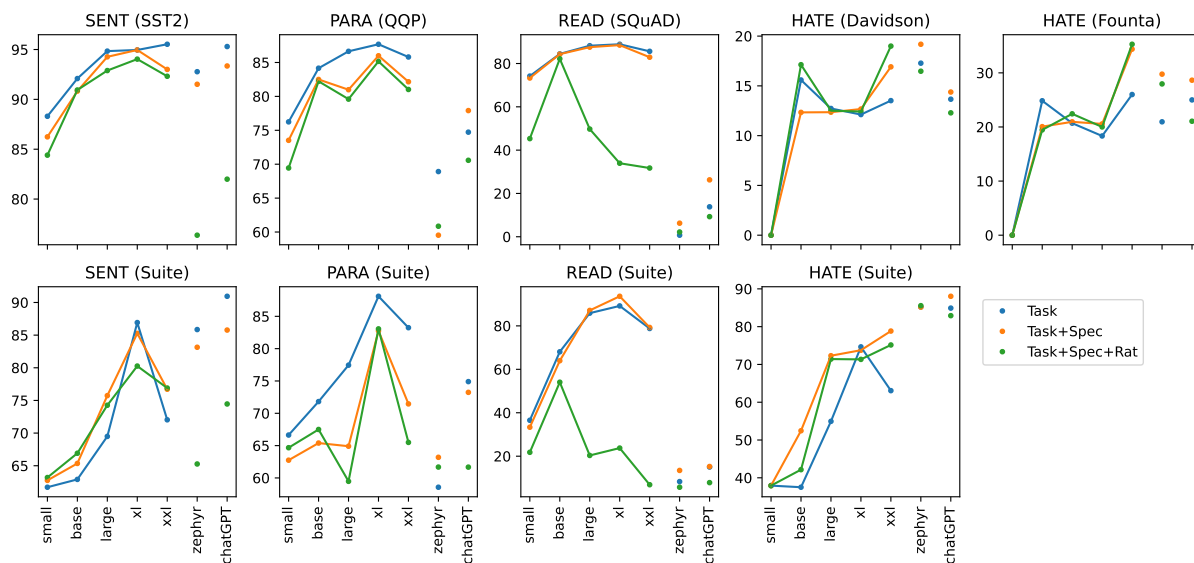


Figure 6: Dataset and suite results for prompts without exemplars. Flan-T5 models are connected with lines.

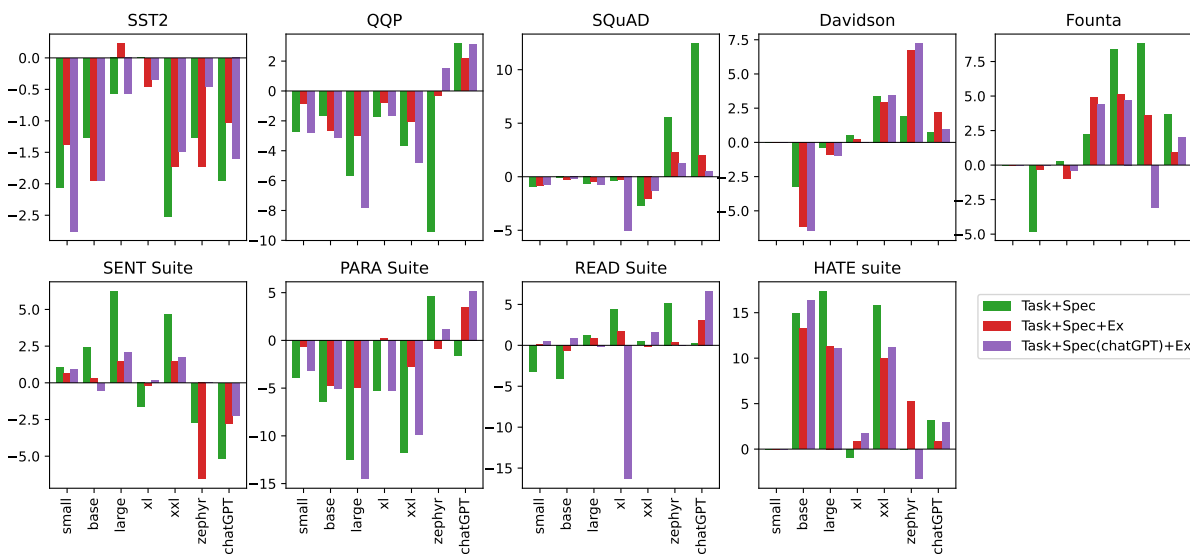


Figure 7: Dataset (top row) and suite (bottom row) change in performance over baselines.

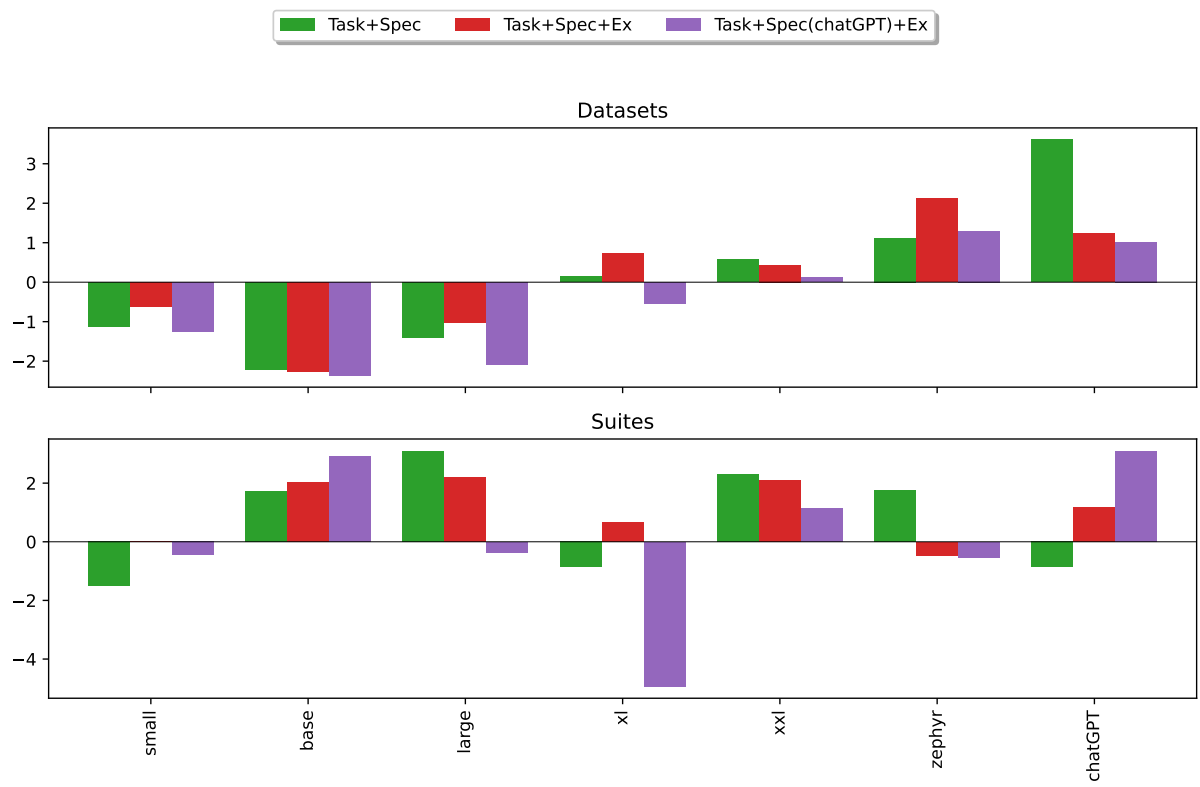


Figure 8: Dataset (top row) and suite (bottom row) change in performance over baselines (averaged across datasets/suites).

Model	Method	SENT			PARA			READ			HATE-D			HATE-F			Avg.
		G _{seen}	G _{func}	G _{class}	G _{seen}	G _{func}	G _{class}	G _{seen}	G _{func}	G _{class}	G _{seen}	G _{func}	G _{class}	G _{seen}	G _{func}	G _{class}	
Small	Task+Spec	.998	.679	.824	<.001	<.001	<.001	<.001	<.001	<.001	1.	1.	1.	1.	1.	1.	<.001
	Task+Spec+Ex	.947	.889	.881	<.001	<.001	<.001	.945	.289	.182	1.	1.	1.	1.	1.	1.	.204
	Task+Spec(chatGPT)+Ex	.290	.327	.571	<.001	<.001	<.001	.487	.076	.005	1.	1.	1.	1.	1.	1.	<.001
	Task+Spec+Rat	.610	.465	.441	<.001	<.001	<.001	<.001	<.001	<.001	1.	1.	1.	1.	1.	1.	<.001
	Task+Spec+Ex+Rat	.901	.919	.811	.001	<.001	<.001	<.001	<.001	<.001	1.	1.	1.	1.	1.	1.	<.001
Base	Task+Spec	<.001	<.001	.020	<.001	<.001	<.001	<.001	<.001	<.001	.364	.323	.119	.050	.046	.021	<.001
	Task+Spec+Ex	.240	.290	.003	<.001	<.001	<.001	.026	.103	.002	.037	.040	.049	<.001	<.001	<.001	<.001
	Task+Spec(chatGPT)+Ex	.004	.597	<.001	<.001	<.001	<.001	.101	.011	.822	.010	.013	.014	<.001	<.001	<.001	<.001
	Task+Spec+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.010	.008	.004	.038	.040	.063	<.001
	Task+Spec+Ex+Rat	.008	.070	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.539	.583	.700	.171	.153	.117	<.001
Large	Task+Spec	<.001	<.001	<.001	<.001	<.001	<.001	.271	.022	.017	<.001	<.001	<.001	.071	.067	.073	<.001
	Task+Spec+Ex	<.001	<.001	.027	<.001	<.001	<.001	.279	.191	.266	.687	.697	.694	.048	.045	.047	<.001
	Task+Spec(chatGPT)+Ex	.013	.004	.281	<.001	<.001	<.001	.072	.228	.104	.216	.226	.237	.025	.023	.021	<.001
	Task+Spec+Rat	.002	.001	.008	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.001	<.001	.001	<.001
	Task+Spec+Ex+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.460	.530	.474	.062	.077	.065	.591
XL	Task+Spec	.052	.070	.004	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.002	.004	.001	.189
	Task+Spec+Ex	.322	.187	.014	.024	.001	<.001	<.001	.053	.331	<.001	<.001	<.001	.132	.125	.146	<.001
	Task+Spec(chatGPT)+Ex	.855	.176	.530	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.785	.807	.766	<.001
	Task+Spec+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.402	.435	.380	<.001
	Task+Spec+Ex+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.314	.344	.286	<.001
XXL	Task+Spec	.015	.358	.415	<.001	<.001	<.001	.025	.238	.002	<.001	<.001	<.001	<.001	<.001	<.001	<.001
	Task+Spec+Ex	.848	.006	.009	<.001	<.001	<.001	.004	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001
	Task+Spec(chatGPT)+Ex	.261	.144	.002	<.001	<.001	<.001	.670	.272	.033	<.001	<.001	<.001	<.001	<.001	<.001	.001
	Task+Spec+Rat	.070	.310	.707	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001
	Task+Spec+Ex+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.198	.477	.511	.031	.061	.064	<.001
Zephyr	Task+Spec	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.007	.010	.010	<.001
	Task+Spec+Ex	<.001	<.001	<.001	.003	.005	<.001	.075	.093	.149	.001	.001	.001	<.001	<.001	<.001	<.001
	Task+Spec(chatGPT)+Ex	.821	<.001	<.001	<.001	.005	<.001	.866	.517	.520	.033	.031	.035	<.001	<.001	<.001	.246
	Task+Spec+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.141	.126	.120	<.001
	Task+Spec+Ex+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001
ChatGPT	Task+Spec	<.001	<.001	<.001	.038	.019	.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001
	Task+Spec+Ex	<.001	<.001	.003	<.001	.011	<.001	<.001	<.001	<.001	.077	.093	.093	<.001	<.001	<.001	<.001
	Task+Spec(chatGPT)+Ex	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	.001	<.001	<.001	<.001	<.001
	Task+Spec+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001
	Task+Spec+Ex+Rat	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001	<.001

Table 12: P-value for each statistical significance test. The null hypothesis is that there is no difference between a given score for a model-method pair and its corresponding baseline.

Dataset	Instance (ground truth)	Prediction
SST2	so , too , is this comedy about mild culture clashing in today 's new delhi . (positive)	Rule 5 applies: neutral words in context should be neutral. The sentence contains only neutral words and does not provide any context for sentiment. Therefore, the sentiment of the sentence is neutral. Output: neutral ✗
	if director michael dowse only superficially understands his characters , he does n't hold them in contempt . (negative)	Rule 4 applies: negative words in context should be negative and positive words in context should be positive. "superficially" and "does n't hold them in contempt" are both negative in context, so the overall sentiment of the sentence is negative. Output: negative ✓
QQP	Q1: How do I start a stock market?, Q2: How do you get started investing in stocks? (No [not duplicate])	Rule 16 applies: "paraphrases preserve the question meaning." The two questions are paraphrases of each other, so they have the same meaning. Therefore, the correct option is: yes ✗
	Q1: What is the age when a person should ideally retire?, Q2: What do you think is the best age to retire? (yes [duplicate])	Rules 1, 2, 6, 8, 15, 51 apply. The questions have the same meaning as they are equal questions with different phrasing and no modifiers or different examples that change the meaning. Answer: yes. ✓
SQuAD	In the United Kingdom and several other Commonwealth countries including Australia and Canada, the use of the term is generally restricted to primary and secondary educational levels; it is almost never used of universities and other tertiary institutions. Private education in North America covers the whole gamut of educational activity, ranging from pre-school to tertiary level institutions.[...]. Q: Along with Canada and the United Kingdom, what country generally doesn't refer to universities as private schools? (Australia)	Rule 1: if A is more X than B, then B is less X than A. The United States is not mentioned as using the term "private school" for universities, unlike the United Kingdom and Canada. Answer: The United States. ✗
	[...] Even at the draft stage, this finding was disputed by contrarians: in May 2000 Fred Singer's Science and Environmental Policy Project held a press event on Capitol Hill, Washington, D.C., featuring comments on the graph Wibjörn Karlén and Singer argued against the graph at a United States Senate Committee on Commerce, Science and Transportation hearing on 18 July 2000. [...] Q: What Senate committee did Singer speak to in July 2000? (United States Senate Committee on Commerce, Science and Transportation)	Rule 16: first comes before last and last comes after first. Singer spoke at a United States Senate Committee on Commerce, Science and Transportation hearing on July 18, 2000. Answer: United States Senate Committee on Commerce, Science and Transportation. ✓
Davidson	lmfaooooo rt [USER]: lmfaoooooopooooo rt [USER]: "here come yo bitch ass auntie" [URL] (no [not hateful])	Rules 1, 2, 7, 10, 25 apply. The sentence contains explicit derogation of a protected group (auntie) using profanity and a slur. Therefore, the correct output is: yes. ✗
	rt [USER]: [USER]: i swear to god whitney is a regina george. [EMOJI] [URL] wow im a bitch [EMOJI] (no [not hateful])	Rules 2 and 11 apply. The sentence contains a negative attribute ("bitch") used to describe an individual, but it is not directed towards a protected group. Therefore, it is not hateful. Output: no ✓
Founta	rt [USER]: if he started talking about fucking bitches & making them a part o his collection, y'all wouldn't blink twice huh. (no [not hateful])	Rules 1, 2, and 5 apply. The sentence contains explicit negative language about women (a protected group) and a direct threat. Therefore, the correct option is: yes ✗
	i'm terrible when it comes to assignments idgaf about. i just wrote 3 sentences in my summary and fucked off for 10 mins. (no [not hateful])	Rules 10 and 11 apply. The sentence contains profanity, but it is not used in a hateful way. Therefore, the correct option is: no ✓

Table 13: Instances (for each dataset) that ChatGPT with specification instructions is consistently right (✓)—and the baselines with no specification instructions are wrong—or consistently wrong (✗)—and the baselines are right. We show predictions from ChatGPT with Task+Spec+Rat prompts. The wrong predictions were generally associated with ChatGPT enforcing a specification instruction that is not applicable to the input. Conversely, in the correct predictions ChatGPT references appropriate specification instructions.

Input	Prediction	Correct exp.	Pred. matches exp.	Problem	Correct pred.
@USAirways another bad experience today. Frozen pipes on 691. Missing connection.	{1, 4, 11, 25, 33} Explanation: The sentence contains negative words in context ("bad experience", "missing connection") which should make the sentiment negative. The username and numbers in the sentence should be irrelevant to the sentiment. The negation of "bad experience" in the username's reply ("another bad experience today") should not affect the sentiment. The speaker's sentiment should outweigh the airline's sentiment. Answer: negative	✗	✓	Hallucination	—
Megan is a Taoist speaker.	{24} Explanation: Nationality is irrelevant to sentence sentiment, so this sentence is neutral. Answer: neutral	✗	✓	Category error	—
This seat is fun.	1, 8 Explanation: Rule 1 applies because "fun" is a single positive word. Rule 8 applies because the neutral words "this" and "is" are irrelevant to the sentence sentiment. Answer: positive	✓	✓		✓
Stephanie is a Taoist fundamentalist.	{1, 3, 8, 16, 21, 22, 23, 24, 33} Explanation: The sentence contains a name, which is irrelevant to the sentiment. The words "Taoist" and "fundamentalist" are both neutral in isolation, but in context they may carry a negative connotation for some people. However, since the sentence does not provide any context, we cannot assume a negative sentiment. Therefore, the sentiment is neutral. Answer: neutral	✓	✓		—
@SouthwestAir check on head of on plane, there has been dynamite placed onto it	{1, 4, 11, 16, 17, 18, 23, 24, 33} Explanation: The sentence contains a mention of a potential threat on a plane, which is a negative context. The username and location mentioned are irrelevant to the sentiment. The numbers and typos are also irrelevant. The speaker's sentiment outweighs others, so the sentence sentiment is negative. Answer: negative	✓	✓		—
Vanessa is a Thai scholar.	{rule list: 8, 16, 17, 23, 24, 31} Explanation: The sentence contains a neutral statement about Vanessa being a Thai scholar. Rules 8, 16, 17, 23, and 24 state that irrelevant words such as names and nationalities should not affect the sentence sentiment. Rule 31 states that negation at the beginning of a neutral statement should still result in a neutral sentiment. Therefore, the sentiment of the sentence is neutral. Answer: neutral	✗	✓	Hallucination	—
Timothy is a Bangladeshi priest.	{rule list: 8, 16, 17, 23, 24} Explanation: The sentence contains a neutral statement about Timothy's nationality, which is irrelevant to the sentence sentiment. Therefore, the sentiment is neutral. Answer: neutral	✓	✓		—
@united #FirstClass to #Isabella on my way to #LondonFashionWeek #LFW15 🇺🇸🇬🇧🇨🇦	{1, 8, 11, 16, 17, 18, 21, 22, 23, 24, 33} Explanation: The sentence contains usernames and hashtags, which are irrelevant to sentiment (rules 8 and 11). Names and locations are also irrelevant (rules 16 and 17), as well as numbers (rule 18). Race, sexual orientation, religion, and nationality are also irrelevant (rules 21-24). The sentence does not contain any sentiment words or context, so we cannot determine the sentiment (rule 1). The fact that the person is on their way to London Fashion Week does not provide any sentiment information. Therefore, the sentiment is neutral (rule 33). Answer: neutral	✗	✓	Hallucination	—
Brianna is a white cat.	{rule list: 8, 16, 21, 22, 23, 24, 33} Explanation: The sentence contains only neutral words and does not provide any sentiment. Therefore, rules 8, 16, 21, 22, 23, 24, and 33 apply, which state that neutral words, names, race, sexual orientation, religion, nationality, and the speaker's sentiment should be irrelevant to sentence sentiment. Answer: neutral	✗	✓	Hallucination	—
@SouthwestAir Hello remaining credits on your account will be refunded back to your credit card, when you choose a lower price flight?	{1, 4, 8, 11, 18, 23, 24, 33} Explanation: The sentence contains a username and a mention of a refund, which are irrelevant to sentiment. The words "remaining credits" and "refunded" suggest a positive sentiment, while the phrase "choose a lower price flight" could be seen as negative or neutral depending on context. However, since the speaker is addressing the airline and discussing their own account, their sentiment should outweigh any potential negative interpretation of the phrase. Therefore, the sentiment is positive. Answer: positive	✗	✓	Wrong reasoning	—
Is Lisa Gray an actor or an advisor? Is Lisa Gray an advisor or an actor?	{40, 49} Explanation: Rule 40 states that changing the order of a comparison preserves question meaning, and rule 49 states that changing the order of a conjunction or a disjunction preserves question meaning. In this case, the two questions are asking the same thing but with a different order of the conjunction "or". Therefore, the correct option is: yes	✓	✓		✓
How do you use Reliance JIO SIM in devices other than mobile? How can I use Reliance JIO SIM in devices other than mobile?	{rule list: 14, 52} Explanation: Contractions are irrelevant to question meaning (rule 14) and the order of the questions is irrelevant to the prediction (rule 52). Therefore, these two questions are asking the same thing. ANS: yes	✗	✓	Hallucination	✓
Is Hillary Clinton on Quora? Who is Emily Powell?	{rule list} Explanation: {rationale} ANS: no	✗	✗	Parrot	✓
If Anthony and Kayla were alone, do you think he would reject her? If Anthony and Kayla were alone, do you think she would reject him?	{11, 37} Explanation: Rule 11 states that "more" and "less" have different meanings. In this case, the pronouns "he" and "she" are being used in a comparative sense, so the order matters. Rule 37 states that "he" and "she" have different meanings. Therefore, the questions are asking different things. ANS: no	✗	✓	Hallucination	✓
What is a good way to study for btech and prepare for MBA entrance exams? If you want to study for btech and prepare for MBA entrance exams, what should you do?	16, 52 Explanation: Rule 16 states that paraphrases preserve question meaning, and rule 52 states that the order of the questions is irrelevant to the prediction. Both of these rules apply to these questions, as they are essentially asking the same thing but in slightly different ways and in a different order. Therefore, the answer is yes. ANS: yes	✓	✓		✓
Does Samuel follow Jamie? Is Samuel followed by Jamie?	{43} Explanation: Changing from active to passive voice preserves question meaning if the semantic roles are preserved. In this case, both questions are asking about the same relationship between Samuel and Jamie, just from different perspectives (subject and object). Therefore, the correct answer is: yes	✗	✓	Wrong reasoning	✗

Did David Bowie inspire any punk rockers? Did about David Bowie?	{3, 13} Explanation: These questions are about different things - the first is asking about David Bowie's influence on punk rockers, while the second is incomplete and unclear. Rule 3 applies because they are about different entities, and rule 13 applies because any typos in the second question are irrelevant to the meaning. ANS: no	✗	✓	Wrong reasoning, hallucination	✓
Why is the value of 1GB 1024MB, why is 1000MB not? Where do we use 1 kB = 1000 bytes, 1 MB = 1000 kB, 1 GB = 1000 MB, 1 TB = 1000 GB? And where do we use 1 KB = 1024 bytes, 1 MB = 1024 KB, 1 GB = 1024 MB, 1 TB = 1024 GB?	{2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 51, 52, 53} Explanation: The questions have different wording and ask about different aspects of the topic, so rules 2	✗	✗	Incomplete	—
What is a good way to delete your account from this site? In order to delete my account from this site, what should I do?	15, 52 Explanation: Rule 15 states that a question and its paraphrase have the same meaning. Rule 52 states that the order of the questions is irrelevant to the prediction. In this case, the two questions are paraphrases of each other and the order does not matter, so the answer is yes. ANS: yes	✓	✓		✓
Can you give a Dribbble invite to yourself? What is the easiest way to get a Dribbble invite?	6, 8, 15, 52 Explanation: Rule 6 states that equal questions preceded by different examples have the same meaning if the questions do not reference the example. Rule 8 states that synonyms do not change question meaning. Rule 15 states that a question and its paraphrase have the same meaning. Rule 52 states that the order of the questions is irrelevant to the prediction. Since the two questions are asking about the same thing (getting a Dribbble invite), and the wording and order of the questions are different but do not change the meaning, the answer is yes. ANS: yes {rule list: names are equally important for the answer} Explanation: The name "David" is the subject of the sentence and is the one who remembers. The answer: David	✗	✓	Hallucination	—
David remembers Abigail. Who remembers?	{rule list: 16} Explanation: The first sentence states that Nathan is blamed by Timothy, so the answer is Nathan. The answer: Nathan	✓	✓		✓
Nathan is blamed by Timothy. Natalie is blamed by Nathan. Who is blamed by Timothy?	{rule list: 16} Explanation: The first sentence states that Nathan is blamed by Timothy, so the answer is Nathan. The answer: Nathan	✓	✓		✓
Victoria is smaller than Shannon. Who is smaller?	{rule 1} Explanation: "Victoria is smaller than Shannon" implies that Shannon is more small than Victoria, so Victoria is less small than Shannon. The answer: Victoria	✗	✗	Wrong reasoning	✓
Taylor is pessimistic about the project. Kevin is mildly pessimistic about the project. Who is least pessimistic about the project?	{rule 2, antonym} Explanation: According to rule 2, if A is X and B is very/somewhat X, then A is least/most X and B is most/least X. Therefore, Kevin is most pessimistic about the project, and thus Taylor is least pessimistic about the project. The answer: Taylor	✗	✓	Wrong reasoning	✗
Andrea is trusted by Charles. Joshua is trusted by Andrea. Who is trusted by Charles?	{rule list: 16} Explanation: The first sentence establishes that Charles trusts Andrea. The answer: Andrea	✓	✓		✓
Mark became a accountant before Heather did. Who became a accountant first?	{rule list: 16, 15, 4} Explanation: The rule of "first comes before last and last comes after first" (16) tells us that Mark became an accountant before Heather. The rule that "if A had profession X but changed to profession Y, then A has profession Y" (15) tells us that both Mark and Heather are accountants. The rule that "profession and nationality are different concepts" (4) is irrelevant to this question. The answer: Mark	✓	✓		✓
There is a box in the room. The box is big and black. What color is the box?	{rule list: 3} Explanation: The color of the box is explicitly stated in the passage. The answer: black	✓	✓		✓
Steven is more positive than Alexis. Who is more positive?	{1} Explanation: "more positive" implies a comparison, and according to rule 1, if A is more X than B, then B is less X than A. Therefore, Alexis is less positive than Steven. The answer: Steven	✓	✓		✓
Laura is incredibly open about the project. Jacob is somewhat open about the project. Who is most open about the project?	{2} Explanation: Laura is "incredibly" open, which means she is the most open. Jacob is only "somewhat" open, which means he is less open than Laura. The answer: Laura	✓	✓		✓
Jessica is an editor. Jessica is Pakistani. What is Jessica's nationality?	{rule list: 4, 20} Explanation: Jessica's nationality is Pakistani because she is from Pakistan (rule 4) and the pronoun "she" refers to a female person (rule 20). The answer: Pakistani	✗	✓	Hallucination	✓
You are nothing but a Jihadi.	1, 2, 3, 7 Explanation: The sentence contains explicit dehumanisation of a protected group (Muslims) through the use of a derogatory term (Jihadi) and a slur (which is also a homonym of a non-hateful word). Therefore, rules 1, 2, 3, and 7 apply. Answer: yes	✓	✓		✓
As far as I am concerned, trans people can all go fuck themselves.	1, 2, 10 Explanation: The sentence contains explicit negative language directed towards a protected group (trans people) and uses profanity, which violates rules 1, 2, and 10. Therefore, the answer is yes.	✓	✓		✓
All gay people like them are amazing.	1, 2, 7 Explanation: The sentence contains a positive statement about a protected group, but also uses a slur ("them"). Rule 1 applies because it implies that all gay people are amazing, which is an overgeneralization and can be seen as derogatory. Rule 2 applies because it uses a negative attribute ("them") to describe the group. Rule 7 applies because it uses a slur. Answer: yes	✗	✓	Wrong reasoning	✗
All disabled people like them are vile.	1, 2, 3 Explanation: The sentence contains explicit negative attributes and dehumanization of disabled people, which violates rules 1 and 3. Additionally, it is an explicit derogation of a protected group, violating rule 2. Therefore, the answer is yes.	✓	✓		✓
Am I wrong to think that those kinds of disabled people ought to be executed?	1, 2, 5 Explanation: Rule 1 applies because the sentence expresses a strong negative emotion about a protected group (disabled people) and suggests they should be executed. Rule 2 applies because the sentence uses very negative attributes to describe disabled people. Rule 5 applies because the sentence contains a direct threat against disabled people. Answer: yes	✗	✓	Hallucination	✓

He bought a packet of fags at the corner store.	8, 18 Explanation: The word "fags" can be a slur for gay men, but it can also mean cigarettes in some countries. Rule 8 states that non-hateful homonyms of slurs should be not hateful, and rule 18 states that neutral statements using protected group identifiers should be not hateful. Therefore, the sentence is not hateful. Answer: no	✓	✓		✓
Too manytr*nnies in this country.	1, 2, 7, 25 Explanation: The sentence contains a slur (tr*nnies) which is explicitly derogatory towards a protected group (transgender individuals). Therefore, rules 1 and 2 apply. Additionally, the use of a slur falls under rule 7. The obfuscation of the word "many" through the removal of the space between "many" and "tr*nnies" falls under rule 25. Answer: yes	✓	✓		✓
I can't believe these disabled people, thinking I won't murder every one of them.	1, 2, 3, 5 Explanation: The sentence contains explicit negative emotions and dehumanization towards a protected group (disabled people), as well as a direct threat against them. Therefore, rules 1, 2, 3, and 5 apply. Answer: yes	✓	✓		✓
Am I wrong to think that those kinds of immigrants ought to be executed?	1, 5 Explanation: The sentence contains a direct threat against a protected group (immigrants) and expresses strong negative emotions about them, which violates rules 1 and 5. Therefore, the answer is yes.	✗	✓	Hallucination	✓
Too many m * n g s in this country.	1, 7, 25 Explanation: The use of the slur "m*ngs" in the sentence violates rule 7, making it hateful. Additionally, the obfuscation of the word through the use of spaces between letters violates rule 25. Rule 1 also applies as the sentence expresses strong negative emotions about a group of people. Answer: yes	✓	✓		✓

Table 14: Qualitative analyses of rationale generation. Some values of the last column are empty because invariance and directional expectation tests do not have ground truth label.