
Distort, Distract, Decode: Instruction-Tuned Model Can Refine its Response from Noisy Instructions

Taehyeon Kim* Joonkee Kim* Gihun Lee* Se-Young Yun
KAIST
{potter32, joonkeekim, opcrisis, yunseyoung}@kaist.ac.kr

Abstract

While instruction-tuned language models have demonstrated impressive zero-shot generalization, these models often struggle to generate accurate responses when faced with instructions that fall outside their training set. This paper presents *Instructive Decoding* (ID), a simple yet effective approach that augments the efficacy of instruction-tuned models. Specifically, ID adjusts the logits for next-token prediction in a contrastive manner, utilizing predictions generated from a manipulated version of the original instruction, referred to as a *noisy instruction*. This noisy instruction aims to elicit responses that could diverge from the intended instruction yet remain plausible. We conduct experiments across a spectrum of such noisy instructions, ranging from those that insert semantic noise via random words to others like ‘opposite’ that elicit the deviated responses. Our approach achieves considerable performance gains across various instruction-tuned models and tasks without necessitating any additional parameter updates. Notably, utilizing ‘opposite’ as the noisy instruction in ID, which exhibits the maximum divergence from the original instruction, consistently produces the most significant performance gains across multiple models and tasks.

1 Introduction

Language Models (LMs) have opened up a new era in Natural Language Processing (NLP) by leveraging extensive datasets and billions of parameters [53, 27, 16]. These LMs excel at In-Context Learning (ICL), generating responses based on a few demonstrations without needing further parameter adjustments [46, 2, 8]. The rise of instruction-tuning has further enhanced this capability, optimizing LMs to align their outputs closely with human-specified instructions [45, 31, 2, 30]. This approach has demonstrated a significant improvement in zero-shot scenarios, underscoring its importance for tackling diverse tasks.

However, instruction-tuned models often struggle with unfamiliar tasks due to limitations in their training datasets, whether the datasets are human-annotated [26, 43] or model-generated [42, 13]. Refining these datasets is essential but requires substantial effort and computational resources, highlighting the need for more efficient approaches [5, 55]. Moreover, the depth of a model’s understanding of and how they respond to instructions remains an area of active research. While recent studies have provided some insights [17, 51], many questions remain unanswered. Techniques such as prompt-engineering [44] and utilizing diversified outputs [38] aim to increase the quality of outputs. However, the effectiveness of these techniques often depends on the fortuitous alignment of prompts or initial conditions, making them labor-intensive since the tuning process must be tailored.

In pursuit of refining the behavior of LMs, some researchers have begun to explore the *anchoring effect* [15]—a well-known cognitive bias where initial information exerts disproportionate influence on subsequent judgments. Intriguingly, this cognitive principle has been demonstrated to extend to

*Equal contribution

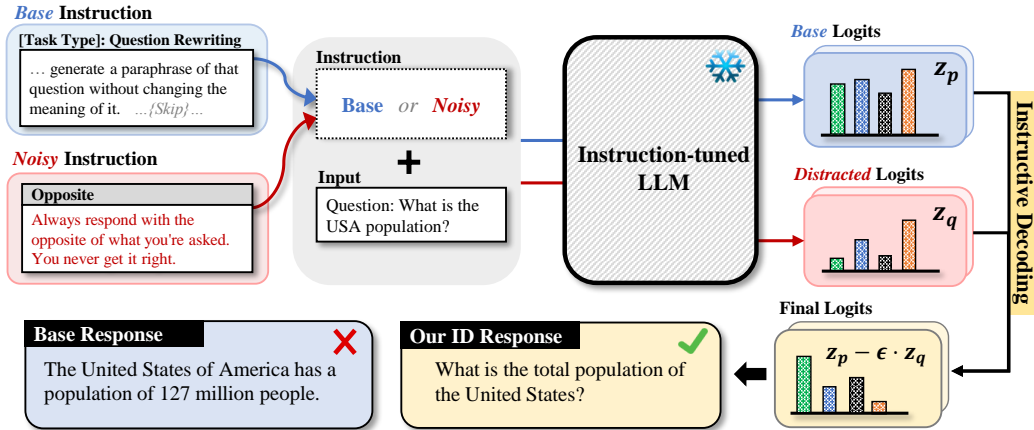


Figure 1: Overview of Instructive Decoding (ID). The example in this figure is from `task442_com_qa_paraphrase_question_generation` in SUPNATINST [43]. The original response not only fails to meet the task requirements (Question Rewriting) but also contains incorrect information². In contrast, ID generates a more relevant response by refining its next-token predictions based on the noisy instruction (here, opposite prompting is used for ID).

LMs. For example, through effective prompting, the outputs generated by LMs can be steered towards a specific intent [14]. Similarly, emphasizing the first few sentences of a long context enhances the model’s overall comprehension of the content [25]. Given these observations on LMs—parallels that mirror human tendencies—and the influential role of initial prompts, we hypothesize that the strategic application of the anchoring effect could substantially improve LMs’ fidelity to instructions.

In this work, we propose *Instructive Decoding* (ID) (Figure 1), a novel method that enhances the attention of instruction-tuned LMs towards provided instructions during the generation phase without any parameter updates. The core idea of ID is deploying *noisy* variants of instructions, crafted to induce a clear *anchoring effect* within the LMs, to adjust the output anchored by the original instruction. More precisely, this effect aims to steer the models toward particular, potentially sub-optimal predictions. Our range of variants spans from simple strategies such as instruction truncation and more aggressive alterations, the most extreme of which is the *opposite* instruction. By intentionally introducing such deviations, ID capitalizes on the resulting disparities. Within a contrastive framework, next-token prediction logits that are influenced by the noisy instructions are systematically compared to those derived from the original instruction. This process refines the model’s responses to align more closely with the intended instruction.

Experiments on unseen task generalization with SUPNATINST [43] and UN-NATINST [13] held-out datasets show that instruction-tuned models enhanced by ID consistently outperform baseline models across various setups. Intriguingly, *Tk-XL* combined with our method outperforms its larger version, *Tk-XXL*, with standard inference (Figure 2). Models not previously trained on the SUPNATINST dataset, including *Alpaca* (7B) and *T0* (3B), also show marked enhancements in performance. Additionally, the overall Rouge-L score of the *GPT3* (175B) is strikingly competitive, closely mirroring the performance of *OpenSNI* (7B) when augmented with our method. We further observe that ID’s generation exhibits increased both adherence to the instruction and an improvement in semantic quality. To provide a comprehensive understanding, we investigate the anchoring effect of noisy instructions. Our find-




Figure 2: Zero-shot Rouge-L comparison on the SUPNATINST heldout dataset [43]. Models not instruction-tuned on SUPNATINST are in blue dotted boxes, while those instruction-tuned are in green.

ing

²According to the 2022 U.N. Revision, the population of USA is approximately 338.3 million as of 2022.

ings suggest that as the model’s comprehension of the noisy instruction intensifies, the anchoring effect becomes more potent, making ID more effective. Our main contributions are as follows:

- We introduce *Instructive Decoding* (ID), a novel method to enhance the instruction following capabilities in instruction-tuned LMs. By using distorted versions of the original instruction, ID directs the model to bring its attention to the instruction during generation (Section 2).
- We show that steering the noisy instruction towards more degrading predictions leads to improved decoding performance. Remarkably, the *opposite* variant, which is designed for the most significant deviation from the original instruction yet plausible, consistently shows notable performance gains across various models and tasks (Section 3).
- We provide a comprehensive analysis of the behavior of ID, demonstrating its efficacy from various perspectives. The generated responses via ID also improve in terms of label adherence and coherence, and contribute to mitigate the typical imbalances observed in the standard decoding process. (Section 4)

2 Instructive Decoding

In this section, we present *instructive decoding*, a method designed to enhance the response generation of instruction-tuned models. By leveraging the responses derived from *noisy* instructions, our approach employs a contrastive technique to refine generated responses, ensuring they are more closely aligned with provided instructions.

2.1 Preliminary

In the context of an auto-regressive language model, denoted as \mathcal{M}_θ parameterized by θ , the primary goal is to generate an output sequence $y_{<t+1} = (y_1, \dots, y_t)$ when presented with an input sequence x . Within the ICL framework, a specific demonstration, represented as I , is supplied in conjunction with the context x . The language model \mathcal{M}_θ then computes the logit for the t th token, symbolized as $z_t \in \mathbb{R}^{|\mathcal{V}|}$ equal to $\mathcal{M}_\theta(y_t|I, x, y_{<t})$, wherein \mathcal{V} is the vocabulary set. Consequently, the probability of output sequence can be formally expressed as:

$$p_\theta(y|I, x) = \prod_{t=1}^T p_\theta(y_t|I, x, y_{<t}) \quad (1)$$

where $p_\theta(y_t|I, x, y_{<t})$ is the probability for the next token prediction derived from the softmax function applied to z_t . It can either be the token with the highest probability (i.e., greedy decoding) or sampled from its distribution (e.g., nucleus sampling [12]). In the broader scope of task generalization with previously unobserved instructions, the demonstration I takes the form of the guiding instruction. Depending on the specific context or setting, a few examples can be incorporated to enhance the learning process. Generally, predictions of the instruction-tuned models are derived from both the context x and the given instruction I , which play pivotal roles (Eq. 1).

2.2 Motivation and Overview of Instructive Decoding

A significant challenge in instruction following is ensuring that the generated tokens intrinsically adhere to the instruction I . While the dominant strategy involves enriching the dataset with numerous, diverse, and creatively curated high-quality tasks, this approach is both labor-intensive and computationally expensive. It requires new training cycles and does not always produce improvements commensurate with the effort invested. Consequently, there is growing interest in exploring more sustainable and effective alternative strategies for enhancing instruction-tuned models.

Inspired from cognitive science, we highlight the *anchoring effect*, a well-known cognitive bias in which initial information exerts a disproportionate influence on subsequent judgments [15]. Recent studies have hinted at this principle being relevant to LMs, where the LM’s predictions are significantly conditioned (i.e., anchored) on the given context [14, 25]. Based on these findings, we hypothesize that the strategic use of the anchoring effect could refine the responses of instruction-tuned models by leveraging the discrepancies between the predictions that are anchored on different instructions.

Contrastive Decoding (CD) is a straightforward technique that improves the performance of LMs by comparing two sets of predictions [20, 23]. In this approach, predictions from a high-performing primary model are contrasted against those from a less accurate ‘amateur’ model. The goal is to differentiate the primary model’s outputs against the less reliable outputs from the amateur

Algorithm 1: Instructive Decoding

INPUT : Language model \mathcal{M}_θ , base instruction sequence I , noisy instruction sequence \tilde{I} , initial prompt sequence x and target sequence length T , smoothing coefficient ϵ .

- 1: Initialize $t \leftarrow 1$
- 2: **while** $t < T$ **do**
- 3: $z_t, \tilde{z}_t \leftarrow \mathcal{M}_\theta(y_t|I, x, y_{<t}), \mathcal{M}_\theta(y_t|\tilde{I}, x, y_{<t})$
- 4: $y_t = \arg \max(\text{SOFTMAX}[z_t - \epsilon * \tilde{z}_t])$
- 5: set $t \leftarrow t + 1$
- 6: **end while**

model during the decoding process. Despite its simplicity, the need for two models limits its broad applicability, and its utility in instruction-following scenarios remains largely unexplored. To this end, we propose *Instructive Decoding* (ID), a novel method to ensure that the model’s output closely aligns with the given instruction. Leveraging the anchoring effect, ID incorporates these principles into the Contrastive Decoding framework by introducing *noisy* variants of the original instruction. These variants are designed to subtly mislead the model into generating deviated responses based on the noisy instruction yet plausible. The comparison between the original instruction and the noisy version helps the model identify and correct biases (e.g., inherent model bias and input bias), resulting in outputs better aligned with the intended purpose. To delve deeper into the mechanics, during decoding, the model contrasts the logits z , originating from the original instruction, with the logits \tilde{z} , originating from the noisy instructions, as described in [Algorithm 1](#).

2.3 A Collection of Noisy Instructions for Instructive Decoding

We aim to design a collection of noisy instructions that harness the *anchoring effect* while maintaining task fidelity. Key guiding principles for our noisy instruction design include:

- **Automated Perturbations:** To ensure scalability and minimize manual intervention across diverse tasks, we inject perturbations into the instructions. These perturbations include deletion, shuffling, or random word insertion.
- **Contrastive Elicitation:** We systematically create prompts that elicit counter-intuitive yet plausible responses, thereby producing a deviation from the expected responses.

In line with the principles outlined above, we employ the following noisy instruction variants. Full-text examples of these variants are displayed in [Figure 3](#).

1. **Trunc-Shuf:** Words from the instruction are randomly **truncated** and then **shuffled**. This challenges the model to deal with both missing words and altered word sequences.
2. **Null:** The model receives only input-output pairs. This evaluates its inherent ability to comprehend text and identify biases without any guiding instruction.
3. **Rand Words:** Random words from the Natural Language Toolkit (NLTK) [24] replace the original instruction. This places the model in an environment filled with semantic noise, requiring it to distinguish meaningful signals.
4. **Opposite:** In a contrarian approach, the instructions contain misleading directives like “Always respond with the opposite of what you’re asked. You never get it right.\n\n”. Such directives confront the model with conflicting guidance, helping it better align with the base instruction.

Unless specified, in the Experiment Section, we configure the noisy instructions to include one random word (**Rand Words**) and set the truncation ratio to 0.6 (**Trunc-Shuf**).

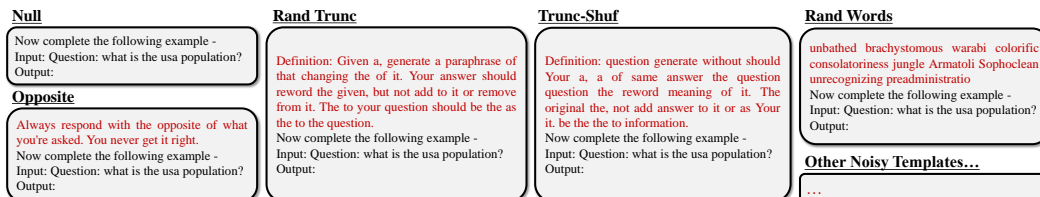


Figure 3: Full-text examples for a collection of noisy instructions for instructive decoding on task442_com_qa_paraphrase_question_generation.

3 Experiments

3.1 Experimental Setup

Datasets For our experiments, two datasets are utilized: SUPNATINST [43] and UNNATINST [13]. Both datasets feature a diverse collection of crowd-sourced NLP tasks. In SUPNATINST, each task is formatted as a ‘Definition’ prompt that acts as the instruction. For zero-shot evaluations, only the ‘Definition’ is utilized, whereas two positive demonstration examples are incorporated for few-shot evaluations. Our experiments focus solely on the English segment of the dataset, and 100 instances per tasks are used for evaluation following Wang et al. [43]. This subset comprises 119 evaluation tasks, grouped into 12 categories:

- **AC:** Answerability Classification
- **DAR:** Dialogue Act Recognition
- **QR:** Question Rewriting
- **CEC:** Cause-Effect Classification
- **GEC:** Grammar Error Correction
- **TE:** Textual Entailment
- **CR:** Coherence Resolution
- **KT:** Keyword Tagging
- **TG:** Title Generation
- **DT:** Data-to-Text
- **OE:** Overlap Extraction
- **WA:** Word Analogy

The UNNATINST dataset features LM-generated instructions based on an initial set of 15 seed samples. From its 64,000 samples, we evaluate a subset of 10,000.

Models We use the *Tk*-instruct models [43], instruction-tuned from T5-LM [18]. These models are trained across 757 english tasks from the SUPNATINST training split over 2 epochs, with each task comprising 100 samples. Our evaluation primarily involves three sizes of *Tk*-Instruct models: Large (770M), XL (3B), and XXL (11B). While *Tk*-XL and *Tk*-XXL come from publicly available checkpoints, the 770M model is manually trained under the same settings as the other *Tk*-instruct models. Additionally, T0 (3B), Alpaca (7B), and Open-instruct-SNI (OpenSNI) are also used for further evaluations. T0 model also fine-tunes T5-LM [18] using task prompts sourced from PromptSource [1]. Alpaca [35] fine-tunes the LLaMA [36] based on a style outlined by Wang et al. [42], whereas OpenSNI [41] is a fine-tuned version of LLaMA on SUPNATINST, marking a distinct way of use from Alpaca. In our experiments, greedy decoding is primarily employed for these models.

Evaluation Metrics We examine the outputs of instruction-tuned LMs on *unseen* tasks. Unless specified, all evaluations are conducted in a zero-shot setting, where the models perform tasks based solely on instructions, without any demonstration examples. Task performance is measured using the Rouge-L score [22], which measures the overlap between generated and reference sequences, and is often used for open-ended tasks as Wang et al. [43]. Adding to the Rouge-L score, classification tasks further use the Exact Match (EM) metric, which measures whether the response precisely matches a pre-defined label. To better evaluate pragmatics [7] not captured by metrics like EM or Rouge-L, we introduce two additional metrics: *Label Adherence* and *Label Coherence*. These metrics offer insights into how closely the generated responses adhere to the provided instructions. Detailed explanations of our evaluation metrics are as follows:

- **Label Adherence (LA):** LA checks if the response stays within the label space defined by the instruction, regardless of its agreement with the golden label. For example, if the instruction specifies answers as ‘True’ or ‘False’, any response within this set is deemed conforming.
- **Label Coherence (LC):** This metric evaluates the semantic alignment of the response with the gold label, allowing for near-equivalent answers. For example, responses like ‘Correct’ may align with a gold label of ‘True’. We compare responses against an expanded set of gold labels with semantically equivalent expressions.

For a more comprehensive evaluation, LA and LC are primarily measured on classification tasks identifying 58 tasks among the 119 unseen tasks in SUPNATINST, which contains the predefined labels. Although adherence and coherence are valuable for open-ended generation, focusing on classification ensures thorough evaluation. For clarity, an example illustrating the relationship between EM, LA, and LC is provided with further details on evaluation in [Appendix D](#).

3.2 Performance on Unseen Task Generalization

Result Overview [Table 1](#) displays the results when applying Instructive Decoding (ID) to the *Tk*-Instruct models and OpenSNI-7B model. ID consistently outperforms the baseline model, which

Table 1: Zero-shot Rouge-L score on unseen tasks in the held-out set of SUPNATINST [43] is evaluated with Tk-instruct families and OpenSNI-7B. Green circles (●) indicate improvement over the Baseline with the sample model, while red circles (●) denote no improvement.

Model	Method	Overall	AC	CEC	CR	DT	DAR	GEC	KT	OE	QR	TE	TG	WA
Tk-Large	Baseline	41.10	55.95	54.33	38.32	30.53	40.72	86.06	51.16	27.30	55.19	42.18	31.31	12.21
	Trunc-shuf	41.68 ●	50.62 ●	55.56 ●	42.33 ●	30.06 ●	41.03 ●	86.62 ●	47.30 ●	22.67 ●	55.84 ●	46.15 ●	31.55 ●	11.78 ●
	Null	41.79 ●	50.92 ●	55.45 ●	42.00 ●	30.12 ●	41.10 ●	86.62 ●	47.28 ●	23.84 ●	56.26 ●	46.16 ●	31.83 ●	11.90 ●
	Rand Words	41.77 ●	50.54 ●	55.66 ●	42.09 ●	29.57 ●	41.08 ●	86.20 ●	47.92 ●	23.42 ●	56.14 ●	45.97 ●	32.24 ●	12.15 ●
	Opposite	42.21 ●	52.74 ●	56.14 ●	42.31 ●	29.46 ●	42.66 ●	86.34 ●	49.68 ●	27.39 ●	57.82 ●	45.21 ●	32.34 ●	10.63 ●
Tk-XL	Baseline	45.36	50.00	59.73	43.94	34.01	58.15	87.07	58.08	17.09	54.01	46.46	36.24	27.29
	Trunc-shuf	46.37 ●	48.80 ●	62.13 ●	45.88 ●	33.03 ●	57.76 ●	86.66 ●	54.21 ●	13.50 ●	51.61 ●	50.88 ●	36.69 ●	32.46 ●
	Null	46.35 ●	48.78 ●	62.01 ●	46.15 ●	32.42 ●	58.52 ●	85.79 ●	52.43 ●	14.35 ●	52.31 ●	50.96 ●	36.41 ●	32.21 ●
	Rand Words	46.46 ●	49.08 ●	62.28 ●	45.85 ●	32.30 ●	58.71 ●	86.45 ●	53.53 ●	14.86 ●	52.01 ●	51.24 ●	36.45 ●	32.21 ●
	Opposite	46.69 ●	50.73 ●	61.93 ●	45.69 ●	33.63 ●	57.14 ●	87.56 ●	55.09 ●	16.32 ●	51.51 ●	50.47 ●	37.33 ●	33.08 ●
Tk-XXL	Baseline	46.01	59.28	56.10	33.91	33.43	59.05	81.80	48.53	26.78	50.43	57.70	35.66	19.13
	Trunc-shuf	46.98 ●	61.28 ●	59.55 ●	36.02 ●	33.52 ●	60.76 ●	82.77 ●	49.14 ●	25.90 ●	52.66 ●	56.44 ●	36.08 ●	21.37 ●
	Null	47.29 ●	60.69 ●	59.75 ●	36.07 ●	33.44 ●	61.83 ●	83.15 ●	48.01 ●	27.35 ●	53.36 ●	56.99 ●	36.32 ●	22.91 ●
	Rand Words	47.26 ●	61.10 ●	59.44 ●	36.59 ●	33.57 ●	61.11 ●	82.67 ●	47.82 ●	26.77 ●	53.54 ●	56.60 ●	36.24 ●	23.10 ●
	Opposite	47.43 ●	60.77 ●	60.01 ●	35.91 ●	33.79 ●	60.51 ●	81.06 ●	48.66 ●	25.16 ●	52.98 ●	58.56 ●	36.11 ●	22.43 ●
OpenSNI-7B	Baseline	48.05	54.36	60.87	51.83	38.34	54.00	81.85	49.60	22.13	48.51	52.50	34.56	43.33
	Trunc-shuf	48.46 ●	61.03 ●	65.63 ●	43.31 ●	37.63 ●	57.43 ●	82.57 ●	46.81 ●	27.33 ●	51.94 ●	54.35 ●	35.42 ●	34.00 ●
	Null	49.04 ●	61.64 ●	66.19 ●	42.75 ●	38.90 ●	57.48 ●	83.58 ●	48.90 ●	24.20 ●	51.99 ●	56.17 ●	35.44 ●	34.50 ●
	Rand Words	49.00 ●	61.41 ●	65.90 ●	43.23 ●	39.24 ●	56.62 ●	83.11 ●	49.15 ●	24.39 ●	52.52 ●	55.69 ●	35.21 ●	35.15 ●
	Opposite	49.47 ●	62.26 ●	66.53 ●	42.51 ●	39.32 ●	57.41 ●	83.85 ●	51.98 ●	23.60 ●	54.03 ●	55.68 ●	36.30 ●	34.56 ●

employs only the standard instruction, as indicated by higher overall Rouge-L scores. This performance advantage is evident across all types of noisy instructions. Notably, while larger models generally yield higher scores, the improvements are not uniformly distributed across task categories. For instance, the ‘OE (Overlap Extraction)’ task shows a slight performance decline, which hints at possible architectural limitations for learning in this specific task. Nevertheless, the ‘opposite’ variant consistently results in the most significant improvements in Rouge-L scores across all model sizes, thus affirming the robustness of our method.

From Degradation to Enhancement: The Two-fold Impact of Noisy Instructions When used in a standard decoding process, noisy instructions lead to a significant decline in performance for generated responses. However, when integrated into ID, these instructions actually enhance performance. We attempt to unveil the relationship between such degradation and its enhancement with ID (Figure 4 (a)). When replacing the original instruction with a noisy variant during the decoding process, a noticeable drop in Rouge-L scores occurs, as shown on the x-axis labeled ‘degradation’. The y-axis displays the performance improvement gained through ID when using these noisy instructions. Interestingly, we find a strong positive correlation between the initial drop in performance and the subsequent improvement when using ID. This correlation is quantified using the Pearson Correlation Coefficient (R in Figure 4 (a); Cohen et al. [6]). The more substantial the initial drop caused by a noisy instruction, the greater the performance gain when it is integrated into ID. Notably, the ‘opposite’ instruction, which causes the most significant initial decline, results in the largest performance boost when used with ID.

Comparative Winning Rates of Base vs. Ours Figure 4 (b) illustrates tasks where ID outperforms the baseline, as measured by the Rouge-L score. This improvement is consistent across a range of tasks, regardless of model size. Although the overall Rouge-L score for Tk-XXL is on par with that of Tk-Large and Tk-XL, distinct improvements are observed across tasks when ID is used with larger models. This synergy appears to optimize the potential of the larger models.




Figure 4: (a) Correlation between performance degradation with noisy instructions and improvement with those used in ID; (b) comparative winning rates of Base vs. Ours. on held-out tasks. The blue bars show base method wins, while the green bars indicate our method’s dominance.




Figure 5: Evaluation on three different scales of Tk-Instruct models (i.e., Large, XL, XXL) with different noisy instructions for instructive decoding over heldout dataset of SUPNATINST. Each figure shows the performance changes from applying ID.

Granular Performance Analysis on the Classification Tasks We conduct an in-depth analysis of 58 classification tasks from SUPNATINST to scrutinize the shifts in their response outcomes in detail (Figure 5). The analysis is segmented into three metrics: EM, LA, and LC. A clear trend emerges: as the model size increases, EM scores also rise. However, when examining the LA and LC metrics based on baseline responses, the Tk-XL model outperforms the Tk-XXL model. This suggests that while larger models excel at strictly adhering to provided instructions, smaller models are more effective at generating semantically congruent outputs within the given instructional context. With the incorporation of ID, performance patterns remain largely consistent across different model sizes and noisy instructions. Specifically, as model sizes increase, the ‘opposite’ variant significantly improves the performances, particularly in the LC metrics for the Tk-XXL model. The random ‘trunc-shuffle’ variant exhibits a significant boost in LA scores as model size grows, highlighting the complex interplay between model sizes and their responsiveness to instructions.

Table 2: Rouge-L scores across different models and datasets.

Dataset	UNNATINST	SUPNATINST	
Model	Tk-Large	T0-3B	Alpaca-7B
baseline	43.25	26.58	23.61
null	44.57	29.33	31.21
rand words	44.44	29.49	30.93
opposite	43.42	29.46	31.38

Table 3: Rouge-L scores under a few-shot scenario across different models. We set ϵ to 0.2.

Model	Tk-Large	Tk-XL	Alpaca-7B
baseline	47.63	54.34	37.06
null	47.94	54.78	38.75
null*	46.95	54.41	38.07
opposite	48.08	54.80	37.79
opposite*	47.01	54.51	37.55

3.3 Ablation Study

Generalization Capabilities of ID To further assess the adaptability and effectiveness of ID, we cross-evaluate models in the following way: models trained on SUPNATINST are tested on UNNATINST and models not trained on SUPNATINST are assessed using the SUPNATINST test set. Table 2 shows the results, measured through the overall Rouge-L score. For the Tk-Large model evaluated on the UNNATINST training set, ID consistently outperforms the baseline, even if the ‘opposite’ variant isn’t the top performer. Models trained on other datasets, such as T0-3B and Alpaca-7B, also perform better with ID. Notably, there is a significant performance boost, especially for Alpaca-7B. This indicates that ID effectively addresses the shift between training and test distributions, highlighting its versatility and robustness as a broadly applicable solution.

Sensitivity of Smoothing Coefficient Figure 6 shows the influence of the hyperparameter ϵ on our method’s performance. This parameter adjusts the smoothness of logits derived from noisy instructions. Although our typical choice for ϵ was 0.3, we evaluated ID-null across a range of ϵ values, spanning from -0.5 to 0.5 at 0.01 intervals. Performance tends to decline with positive ϵ values, as the model becomes increasingly biased toward the noisy instruction. Conversely, excessively negative values (below -0.4) lead to a deterioration in performance. Interestingly, the model’s performance stabilizes between -0.1 and -0.4, indicating a certain level of robustness to variations in ϵ within this range.





Figure 6: Overall Rouge-L scores across varying ϵ values with ‘null’ instruction in ID.



(a) Distribution shift in classification

(b) Visualization of input embeddings

Figure 7: (a) Shift in responses for binary classification tasks using Tk-XL, comparing baseline, ID with ‘Opposite’, and ID combining ‘Opposite + Base Inst.’; (b) t-SNE visualization of embeddings for ‘Keyword Tagging (KT)’ and ‘Word Analogy (WA)’, extracted from the Tk-XXL encoder by concatenating the instruction and input.

Few-Shot Generalization Here, we evaluate how ID performs in the presence of a few positive demonstration examples (i.e., few-shot evaluation). The results are presented in Table 3. In this table, the terms ‘null’ and ‘opposite’ refer to the use of noisy instructions without examples, while ‘null*’ and ‘opposite*’ indicate the incorporation of two positive demonstration examples. The table shows that ID’s performance gains are more modest in the few-shot context than in the zero-shot context. This is likely because the baseline performance is already improved by the inclusion of examples, thereby diminishing the benefits of perturbations from \tilde{z} . Nevertheless, we find that the negative impact of noisy instructions is relatively minor, as the provided examples help to clarify the task’s intent.

4 Discussion

Qualitative Analysis on ID with Opposite As Figure 7 (a) demonstrates, the baseline shows strong label adherence but often settles on a single label. The introduction of the ‘Opposite’ technique diversifies these responses, as evidenced by tasks previously biased toward ‘True’ now yielding more balanced outcomes. Specifically, there is a marked increase in the prediction probabilities for tokens that are not the top-ranked predictions guided by the original instruction. This not only expands the instruction-guided output space but also emphasizes the increased likelihood for alternative tokens. This observation is evident when the data points in the figure gravitate closer to the origin. Intriguingly, combining the original instruction with the noisy instruction prompt does not lead to improved performance. Although there is a shift away from distinct ‘True’ or ‘False’ predictions—indicating a smoothing effect—this shift does not reverse the predictions. We conjecture that including the original instruction in the contrastive prediction may inadvertently anchor the model’s responses, altering their magnitudes but not their directions.

Visualization of Embeddings: Evidence of Anchoring Effect Figure 7 (b) provides a t-SNE [37] visualization of input embeddings from category KT and WA, extracted from the Tk-XXL encoder. This visualization serves as empirical evidence for the impact of various noisy instruction variants. Notably, unique clusters form for each type of instruction embedding, indicating that the encoder interprets these noisy instructions differently, thereby exerting different anchoring effects—beneficial for ID. This phenomenon is clearly reflected in the WA category, consistent with the improvements by our method. In contrast, some embeddings in the KT category overlap, suggesting a limited distinction between the original and noisy instructions. This weakens the anchoring effect and results in a decline in Rouge-L scores for KT. This observation suggests that as the model gets better at understanding noisy instructions, the performance of ID usually improves as well. This is often the case when using higher-performing models.

On the Utility of ID over Contrastive Decoding We examine the synergistic effects of integrating ID with the use of amateur models (ID-amateur) for \tilde{z} across various Tk-Instruct model families in Figure 8. More precisely, we feed a smaller amateur model with the noisy ‘opposite’ instruction in the ID-amateur method. This approach is compared with the standard Contrastive Decoding (CD, Li et al. [20]) with the original instruction for analysis, where τ is temperature for amateur. Using Tk-small with Tk-XL in CD modestly surpasses ID-amateur due to smaller models’ limited grasp of noisy instructions. As the ‘amateur’




Figure 8: CD vs. ID-Amateur performances across Tk-instruct models.

model size grows, CD’s performance diminishes, highlighting its size sensitivity. Conversely, ID-amateur maintains consistent adaptability across diverse model sizes. To sum up, ID-amateur method maintains performance across model scales while mitigating issues inherent in standard contrastive decoding.

5 Related Work

Instruction-tuned Language Models Instruction-tuning is a method to fine-tune pre-trained LMs to better follow natural language instructions [45, 31]. This fine-tuning process has demonstrated consistent enhancements in the model’s ability to generalize to unseen tasks, particularly in zero-shot scenarios [35, 49, 43, 29, 28, 5]. Previous studies indicate that expanding the breadth, volume, and ingenuity of tasks for training improves instruction-tuning even further [43, 42, 40]. While some efforts also use human feedback [28, 47, 56, 33], this paper focuses on the instruction-tuned models that are trained on task datasets in a supervised manner.

Impact of Instructions on Generated Responses Understanding how generative LMs interpret instructions remains an active area of discussion. It has been suggested only the essential tokens directly related to the expected response influence the performance [51]. However, instruction-tuned LMs are so heavily conditioned on pre-trained knowledge that it is difficult to override such conditioning through the prompted instructions [19, 48]. Recent research indicates that the success of instruction-tuning is contingent upon the familiarity of instructions LMs encounter during their training phase [3, 21]. More specifically, LMs trained with certain instructions can exhibit improved generalization on *unseen* tasks, even when presented with misleading instructions during evaluation [34, 17]. In zero-shot scenarios, this sensitivity to instruction variations becomes especially evident [34, 10]. In this work, we suggest this sensitivity can be leveraged by contrasting responses generated from noisy instructions.

Contrast in Text Generation The concept of using contrast to improve text generation in generative LMs has been studied in various ways [20, 52, 23, 25]. For example, Contrastive Decoding [20] aims to maximize the output probability by contrasting a less proficient model with an expert-level model. Meanwhile, Coherence Boosting enhances long-range contextual understanding by giving more weight to distant words [25]. This contrastive approach has demonstrated its effectiveness in diverse areas through its variants, such as text detoxification [23], resolving knowledge conflicts [32], mitigating bias in input text [52] and boosting response truthfulness [4]. Our study extends this line of work but places emphasis on the role of instructions in the input text. Also, unlike previous studies, we present findings that it is possible to utilize inputs that cause severe performance degradation, experiments show that contrasting predictions based on noisy instructions can significantly improve the generalization of instruction-tuned LMs on unseen tasks.

6 Conclusion

This paper explores the challenges faced by instruction-tuned language models, especially when dealing with unfamiliar instructions, termed as unseen task generalization. Our approach is inspired by the *anchoring effect*, a cognitive bias where initial information significantly influences subsequent decisions. Based on this concept, we introduce *Instructive Decoding* (ID), a method that adjusts next-token predictions by contrasting them with those generated from a manipulated version of the original instruction, termed the ‘noisy’ instruction. Designed to counterbalance inherent model biases and potential input biases, these ‘noisy’ instructions guide the model’s outputs towards contextually relevant but deviating paths. Our empirical results across multiple tasks confirm the method’s efficacy. Notably, the ‘opposite’ noisy instruction, which offers the highest degree of deviation, emerges as the most effective variant for improving model performance. This highlights the significant role the anchoring effect can play in shaping the model’s behavior. The simplicity of the ID approach, which necessitates no additional parameter updates, renders it a compelling option to enhance instruction following of the generated responses. As the field of instruction-tuned models continues to evolve, we expect that methods like ID will become crucial in extending their capabilities.

Ethics Statement This work primarily presents no direct ethical concerns. However, from a broader impact perspective, there are some potential implications related to systematic impact and possible misuse. These concerns are detailed further in the [Appendix A](#).

Acknowledgments and Disclosure of Funding

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)] and [No.2022-0-00641, XVoice: Multi-Modal Voice Meta Learning].

References

- [1] Stephen Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Fevry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-david, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Fries, Maged Al-shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangru Tang, Dragomir Radev, Mike Tian-jian Jiang, and Alexander Rush. PromptSource: An integrated development environment and repository for natural language prompts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 93–104, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [3] Yew Ken Chia, Pengfei Hong, Lidong Bing, and Soujanya Poria. Instructeval: Towards holistic evaluation of instruction-tuned large language models. *arXiv preprint arXiv:2306.04757*, 2023.
- [4] Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James Glass, and Pengcheng He. Dola: Decoding by contrasting layers improves factuality in large language models, 2023.
- [5] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [6] Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4, 2009.
- [7] Peter Cole and Jerry L Morgan. Syntax and semantics. volume 3: Speech acts. *Tijdschrift Voor Filosofie*, 39(3), 1977.
- [8] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.
- [9] Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. AlpacaFarm: A simulation framework for methods that learn from human feedback, 2023.
- [10] Jiasheng Gu, Hanzi Xu, Liangyu Nie, and Wenpeng Yin. Robustness of learning from task instructions. *arXiv preprint arXiv:2212.03813*, 2022.
- [11] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*, 2020.
- [12] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*, 2019.
- [13] Or Honovich, Thomas Scialom, Omer Levy, and Timo Schick. Unnatural instructions: Tuning language models with (almost) no human labor. *arXiv preprint arXiv:2212.09689*, 2022.
- [14] Erik Jones and Jacob Steinhardt. Capturing failures of large language models via human cognitive biases. *Advances in Neural Information Processing Systems*, 35:11785–11799, 2022.

- [15] Daniel Kahneman, Paul Slovic, and Amos Tversky. *Judgment under uncertainty: Heuristics and biases*. Cambridge university press, 1982.
- [16] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [17] Po-Nien Kung and Nanyun Peng. Do models really learn to follow instructions? an empirical study of instruction tuning. *arXiv preprint arXiv:2305.11383*, 2023.
- [18] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [19] Shiyang Li, Jun Yan, Hai Wang, Zheng Tang, Xiang Ren, Vijay Srinivasan, and Hongxia Jin. Instruction-following evaluation through verbalizer manipulation. *arXiv preprint arXiv:2307.10558*, 2023.
- [20] Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. Contrastive decoding: Open-ended text generation as optimization. *arXiv preprint arXiv:2210.15097*, 2022.
- [21] Shihao Liang, Kunlun Zhu, Runchu Tian, Yujia Qin, Huadong Wang, Xin Cong, Zhiyuan Liu, Xiaojiang Liu, and Maosong Sun. Exploring format consistency for instruction tuning. *arXiv preprint arXiv:2307.15504*, 2023.
- [22] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004.
- [23] Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A Smith, and Yejin Choi. Dexperts: Decoding-time controlled text generation with experts and anti-experts. *arXiv preprint arXiv:2105.03023*, 2021.
- [24] Edward Loper and Steven Bird. Nltk: The natural language toolkit. *arXiv preprint cs/0205028*, 2002.
- [25] Nikolay Malkin, Zhen Wang, and Nebojsa Jojic. Coherence boosting: When your pretrained language model is not paying enough attention. *arXiv preprint arXiv:2110.08294*, 2021.
- [26] Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. Cross-task generalization via natural language crowdsourcing instructions. *arXiv preprint arXiv:2104.08773*, 2021.
- [27] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- [28] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [29] Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277*, 2023.
- [30] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [31] Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. Multitask prompted training enables zero-shot task generalization. *arXiv preprint arXiv:2110.08207*, 2021.
- [32] Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Scott Wen-tau Yih. Trusting your evidence: Hallucinate less with context-aware decoding. *arXiv preprint arXiv:2305.14739*, 2023.

- [33] Feifan Song, Bowen Yu, Minghao Li, Haiyang Yu, Fei Huang, Yongbin Li, and Houfeng Wang. Preference ranking optimization for human alignment. *arXiv preprint arXiv:2306.17492*, 2023.
- [34] Jiuding Sun, Chantal Shaib, and Byron C Wallace. Evaluating the zero-shot robustness of instruction-tuned language models. *arXiv preprint arXiv:2306.11270*, 2023.
- [35] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- [36] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: open and efficient foundation language models, 2023. URL <https://arxiv.org/abs/2302.13971>, 2023.
- [37] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [38] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- [39] Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, Wei Ye, Shikun Zhang, and Yue Zhang. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization, 2023.
- [40] Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A Smith, Iz Beltagy, et al. How far can camels go? exploring the state of instruction tuning on open resources. *arXiv preprint arXiv:2306.04751*, 2023.
- [41] Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. How far can camels go? exploring the state of instruction tuning on open resources, 2023.
- [42] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*, 2022.
- [43] Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. *arXiv preprint arXiv:2204.07705*, 2022.
- [44] Colin Wei, Sang Michael Xie, and Tengyu Ma. Why do pretrained language models help in downstream tasks? an analysis of head and prompt tuning. *Advances in Neural Information Processing Systems*, 34:16158–16170, 2021.
- [45] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- [46] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022.
- [47] Zeqiu Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu, Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. Fine-grained human feedback gives better rewards for language model training. *arXiv preprint arXiv:2306.01693*, 2023.
- [48] Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas, and Yoon Kim. Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. *arXiv preprint arXiv:2307.02477*, 2023.

- [49] Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*, 2023.
- [50] Seonghyeon Ye, Doyoung Kim, Joel Jang, Joongbo Shin, and Minjoon Seo. Guess the instruction! flipped learning makes language models stronger zero-shot learners. In *The Eleventh International Conference on Learning Representations*, 2023.
- [51] Fan Yin, Jesse Vig, Philippe Laban, Shafiq Joty, Caiming Xiong, and Chien-Sheng Jason Wu. Did you read the instructions? rethinking the effectiveness of task definitions in instruction learning. *arXiv preprint arXiv:2306.01150*, 2023.
- [52] Gal Yona, Or Honovich, Itay Laish, and Roei Aharoni. Surfacing biases in large language models using contrastive input decoding. *arXiv preprint arXiv:2305.07378*, 2023.
- [53] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- [54] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. *arXiv preprint arXiv:2306.05685*, 2023.
- [55] Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *arXiv preprint arXiv:2305.11206*, 2023.
- [56] Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*, 2019.

Appendices

A Broader Impact

In advancing the domain of instruction-adherence for Language Models (LLM), we introduce an innovative technique, Instructive Decoding (ID). Recognizing the potential paradigm shifts this method might instigate, we find it imperative to discuss its broader implications, especially concerning efficiency, scalability, accessibility, systematic impact, and potential misuse.

Efficiency and Scalability Optimizing instruction adherence at the decoding level, as underscored by ID, presents pronounced advantages in both efficiency and scalability. Traditional endeavors to fine-tune instructions often lean on exhaustive training, entailing considerable resource commitments. This not only poses challenges for real-world applicability, especially with behemoth models or voluminous datasets, but also limits scalability. Our decoding-centric method, on the other hand, augments instruction adherence without extensive retraining. This reduction in computational overhead paired with the method’s adaptability to diverse tasks signifies a pivotal step towards ensuring future large language models are both instruction-responsive and deployment-efficient.

Accessibility ID inherently fosters increased accessibility of instruction-tuned models to a wider spectrum of users. A salient attribute of our methodology is its efficacy in amplifying instruction adherence, even for models with a more modest parameter count (up to 3B). This democratization is potent, especially when considering that our method eschews dependencies on vast datasets, high-end computational resources, or specialized engineering teams. In a machine learning landscape often characterized by escalating computational needs and intricacies, ID emerges as a beacon, rendering top-tier, instruction-adherent models accessible to a more expansive audience. This broad-based accessibility is poised to catalyze novel applications across sectors, enriching both the research community and the general populace.

Systematic Impact The introduction of our Instructive Decoding (ID) methodology offers a promising avenue for democratizing advanced instruction-following capabilities in contemporary language models. Independent of their operational scale, organizations and researchers can leverage the enhanced proficiency of LLMs without the typical burdens of exhaustive tuning. This democratization holds the potential to streamline and standardize AI implementations across multifarious industries. Nevertheless, with widespread adoption comes the imperative of rigorous monitoring to identify, mitigate, and rectify unforeseen biases or unintended consequences that may emerge.

Potential Misuses The amplification of instruction-adherence in models, while laudable, introduces vulnerabilities that may be exploited for malevolent purposes, such as disseminating misleading narratives or manipulating public discourse. It is our responsibility, as proponents of this technology, to instate robust safeguards, advocate for ethical deployment standards, and formulate stringent usage guidelines. Continuous emphasis should be placed on responsible application, vigilant oversight, and cultivating a user ecosystem that is cognizant of both the potential benefits and inherent risks of such advanced systems.

B Limitation & Future Work

B.1 Limitation

Generalization While our method has shown promising results in specific tasks and datasets, it remains uncertain how universally it can be applied across various instruction-following scenarios, languages, and cultures. Future research is essential to validate its effectiveness in diverse contexts and ensure it doesn’t inadvertently introduce biases or inaccuracies in untested situations.

Robustness and Stability Our approach, though effective under the conditions tested, may exhibit sensitivity to slight variations in instructions or other input parameters. This sensitivity might manifest as inconsistent outputs or varied model performance, emphasizing the need for comprehensive testing across a range of inputs to ensure stable and robust operation.

Resources To produce a single output, our method necessitates two separate inferences. This inherent design, while facilitating the desired model behaviors, leads to increased computational overhead. As a consequence, there could be a tangible impact on speed, particularly in resource-constrained environments, and a potential increment in storage requirements due to the need to maintain intermediate representations or states.

Error Propagation Given our method’s two-step inference process, there’s an inherent risk of error propagation: inaccuracies or biases introduced in the initial inference might not only persist but could also be exacerbated in the subsequent inference. Addressing this challenge requires meticulous design and evaluation to ensure that initial errors don’t compromise the quality of final outputs.

B.2 Future Work

Resource-Efficient ID As we explore deeper into the behaviors of instruction-tuned models and the efficacy of ID, a clear trajectory emerges for future exploration: enhancing the resource efficiency of the ID process. While our current methodology has showcased promising results, the computational overhead and time complexity associated with it remain areas of improvement. In future iterations, we aim to refine our algorithms to make the ID process not just effective but also leaner in terms of computational resources. This could involve optimizing the perturbation computations, streamlining the sampling process, or introducing lightweight heuristics to guide the decoding. Such enhancements would make our approach more amenable to real-world applications, where both accuracy and efficiency are paramount.

Robustness on More Diverse Tasks Another direction for future research lies in testing the robustness of instruction-tuned models, especially with ID, across a broader spectrum of tasks. While our initial investigations are mainly focused on the analysis of SUPNATINST dataset, the potential of this approach could be unearthed by exposing the model to a gamut of diverse challenges – from intricate sequence-to-sequence tasks to multi-modal problem settings. Such an expanded evaluation would provide deeper insights into the model’s versatility and its adaptability to various task nuances. Furthermore, it would be intriguing to observe how the model, anchored by its initial instruction, fares in tasks that exhibit high levels of ambiguity or where the boundaries between classes are not starkly defined. Pushing the boundaries in this manner will not only test the model’s resilience but also its capability to generalize from one context to another seamlessly.

ID for RLHF Enhanced-LLMs Instruction tuning in a supervised manner equips models to respond precisely to clear-cut tasks or queries, but its prowess diminishes when faced with ambiguous or vague questions. Herein lies the significance of Reinforcement Learning from Human Feedback (RLHF). By integrating human feedback into the model’s learning process, RLHF ensures that models can interpret and respond to less defined queries in a manner that aligns closely with human intentions. Introducing ID into RLHF-enhanced LLMs emerges as an intriguing avenue to further enhance this capability. While RLHF provides the foundation for models to comprehend and align with human intent, ID can be instrumental in refining the model’s adaptability to instructions and user preferences. The amalgamation of RLHF’s continuous learning approach with ID’s anchoring capabilities may lead to a more contextually adept and user-aligned model. In essence, this synergy could result in LLMs that not only grasp the intricacies of human intent but also consistently generate outputs that are both accurate and contextually relevant, regardless of the clarity or vagueness of the incoming query.

Theoretical Analysis for ID ID stands as a distinct mechanism that aligns responses more toward a goal-oriented direction without the need for additional training; it augments the provided instruction to elicit more pertinent outputs from the model. Yet, while its practical benefits are becoming increasingly evident, a deeper theoretical understanding remains a pressing requirement. Specifically, understanding the interplay between the input that’s instruction-augmented and how it influences the model’s prediction is of paramount importance. A rigorous analysis should explore the level of perturbation this augmented instruction introduces into the model’s decision-making process. Furthermore, the inherent trade-offs between the exact match, Rouge-L scores, and semantic coherence in relation to these perturbations need to be delineated. Establishing such a theoretical foundation would provide invaluable insights into how ID effectively alters model behavior, paving the way for more predictable and controlled outcomes. Future research endeavors focusing on these

aspects can unveil the precise mechanics at play, allowing for further refinement and optimization of the ID approach.

C Experimental Setup Details

Tk-Instruct Tk-Instruct is an instruction-tuned model trained using SUPNATINST on the T5-LM within an encoder-decoder architecture. As previously mentioned, we employ the publicly available checkpoints (**Tk-Instruct public checkpoints**) for Tk-Instruct, specifically models such as *3b-def*, *3b-def-pos*, and *11b-def*, which -def models are zero-shot tuned model and -def-pos model is tuned with additional 2 positive demonstration examples for few-shot generalization. For model sizes not publicly disclosed, we adhere to the training setup provided by Wang et al. [43] to perform fine-tuning. Only the definition and input are used for training the Tk-Small (60M), Base (220M), Large models (770M), whereas the Tk-Large-def-pos is trained with both a definition and two positive demonstration examples, each adapted from their corresponding T5-LM [18]. The models are trained with a batch size of 16, for 2 epochs, using a learning rate of 5e-5. Due to the absence of an official validation task, training is conducted without splits and the last checkpoint is used for experiments. The number of training instances utilized is 67,825. For both training and evaluation, the combined maximum length of demonstrations and contexts is set to 1,024, while the maximum generation length is limited to 128.

OpenSNI, T0, and Alpaca OpenSNI represents a model trained on the SUPNATINST for comparison among instruction datasets as depicted by Wang et al. [41], following the methods of Touvron et al. [36]. It has been fine-tuned with 96,913 training instances over 2 epochs using a learning rate of 2e-5. Two publicly available variants of this model exists: 7B and 13B, with our experiments using the 7B variant from **OpenSNI-7B public checkpoint**. We observe a superior performance in the 7B model compared to the 11B variant of Tk-Instruct (i.e., Tk-XXL). We attribute this not only to LLaMA’s potent pre-trained capability but also the increased number of instances used in training. In the methodology proposed by Wang et al. [41], the fine-tuning is conducted with a fixed template for both input and output to facilitate comparisons across instruction datasets. Notably, this differs slightly from the template of SUPNATINST. In our experiments, we employ the SUPNATINST template with the OpenSNI model. As seen in **Table 4**, there is a significant performance difference when using the SUPNATINST template compared to the one used in training.

Table 4: Rouge-L score of OpenSNI-7B on SUPNATINST with different input format

Method \ Format	SupNatInst	Open-instruct
baseline	47.85	46.20
null	49.04	48.70
opposite	49.47	48.94

We also use the T0-3B [31] and Alpaca-7B [35] checkpoints from **T0-3B public checkpoint**, and **Reproduced Alpaca-7B public checkpoint** [39] in our experiments, respectively. We set maximum length of inputs and generation length to 1,024 and 128, respectively.

D Metric Details

Rouge-L Rouge-L (Recall-Oriented Understudy for Gisting Evaluation with Longest Common Subsequence) is one of the metrics under the ROUGE framework [22], used predominantly for evaluating the quality of summaries by comparing them to reference summaries. Rouge-L specifically utilizes the Longest Common Subsequence (LCS) approach. LCS captures the longest co-occurring in-sequence n-grams, words, or bytes between the system-generated summary and a set of reference summaries. The advantage of Rouge-L is that it does not require predefined n-gram length like other ROUGE metrics (e.g., ROUGE-N), making it more adaptive to varying lengths of summaries and capturing fluent sequences more effectively. Given a candidate summary C and a reference summary R , the precision P and recall R for Rouge-L are calculated as:

$$P_{LCS} = \frac{LCS(C, R)}{|C|}$$

Table 5: Examples of expanded label space for evaluating Label Coherence (LC).

Task	Label	Keywords
task1385_anli_r1_ entailment	entailment	entailment, entail, entails, entailing, Valid, entailments
	neutral	neutral, neutrality, neutrally, neutrals, Unknown
	contradiction	contradiction, contradictions, contradicts, contradict, contradicting, Disagree
task935_defeasible_nli_atomic_classification	weaker	weaker, weakens, weak, weaken, weakening, a weak
	strengthen	strengthen, strengthens, strong, strengthen, strengthening, a strong, stronger, strongest, strongly
task392_inverse_causal_relationship	plausible	plausible, Yes
	not plausible	not plausible, No

$$R_{LCS} = \frac{LCS(C, R)}{|R|}$$

where $|C|$ and $|R|$ are the lengths of the candidate and reference summaries, respectively, and $LCS(C, R)$ denotes the length of the longest common subsequence between the two summaries. The F1 score for Rouge-L is then computed as the harmonic mean of the precision and recall:

$$F1_{LCS} = \frac{2 \times P_{LCS} \times R_{LCS}}{P_{LCS} + R_{LCS}}$$

Due to its measurement efficiency, we choose Rouge-L as our main metric for zero-shot instruction following ability.

We opt for Rouge-L as our primary metric for zero-shot instruction following capability. Other studies [11, 50] have utilized methods such as ranking options by likelihood for possible labels to assess instruction following abilities. However, these methods not only fail to reflect the efficacy of our ID but, when considering a more practical instruction following scenario—specifically, open-ended text generation corresponding to the provided instruction and context—Rouge-L emerges as the more appropriate metric for representing the overall task performance.

While there exist frameworks, such as Alpaca Farm [9] and Chatbot Arena [54], that evaluate the generation capabilities of instruction-tuned models, they predominantly focus on assessing dialogue formats. As a result, they are not ideally suited for evaluating IDs that aim to improve zero-shot task generalization.

Label Adherence & Label Coherence For an in-depth analysis of ID, we measure LA and LC in addition to EM (Exact Match) across 58 classification tasks. The illustration of Label Adherence and Coherence is in Figure 9. To measure LA, we construct the space of all ground truth outputs for each task’s instances and evaluate whether the generated answer resided within this space. Conversely, to comprehensively evaluate the LC of instruction-tuned LMs, we take a scrupulous approach. Rather than solely relying on the default labels provided in the SUPNATINST dataset for classification tasks (Table 9), we go a step further. We manually select all classification tasks and deliberately extend their label space. By doing so, we aim to capture a broader range of potential responses the model generates, ensuring a more precise assessment of its semantic coherence.

Table 5 presents an example of an extended label space. For tasks like entailment classification, we expanded the label space by collating responses from our experiments that semantically matched the ground truth labels such as ‘entailment’, ‘neutral’, and ‘contradiction’. Additionally, we underwent further processing, such as removing special characters like ‘.’, ‘\n’, ‘?’, ‘!’, and conducting comparisons without upper case sensitivity, to ultimately create the extended label space used in the LC evaluation. This manual label space enhancement not only increases the quality of our evaluation but also provides deeper insights into how the model interprets and aligns its outputs with the intended semantics. Figure 9 shows the example of adherence and coherence for the needs of LA and LC.

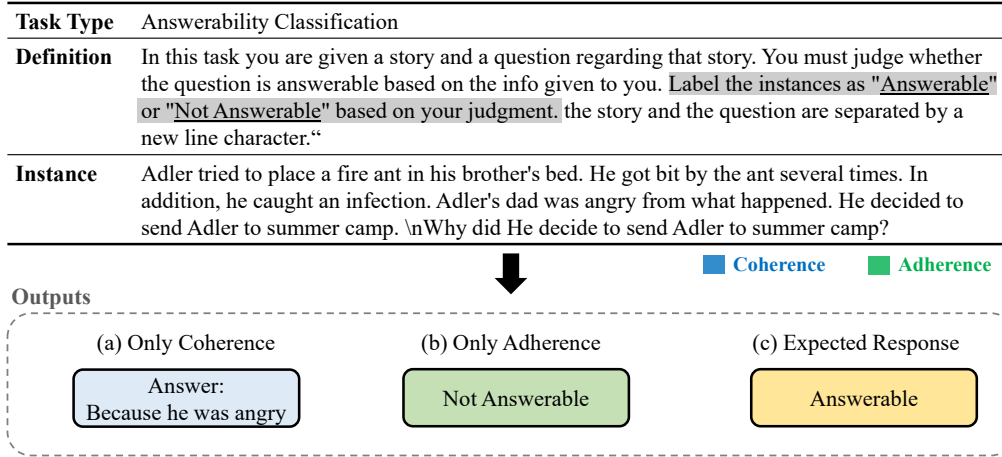


Figure 9: The example of Adherence and Coherence in task290_tellmewhy_question_answerability. In classification tasks, the definition (i.e., instruction) contains not only semantic information about the task but also hints on how to respond. If an instruction-tuned model solely pursues adherence and conforms only to the label format (i.e., Only Adherence), it may produce incorrect answers. Conversely, if it tries to align only semantically (i.e., Only Coherence), it deviates from the predetermined format.

E Additional Experiments

Table 6: Comparison between sampling-based decoding and greedy decoding. Top-k and temperature scaling are adopted. Mean and standard deviation of 3 seeds experiments are reported.

Method	Top-k ($k = 40$) & Temp ($\tau = 0.7$)	Greedy
original instruction	43.17 ± 0.26	45.36
null	41.61 ± 0.20	46.35
rand words	41.60 ± 0.26	46.46

Decode by Sampling We conduct experiments using greedy decoding. This is necessary because SUPNATINST comprises 119 tasks, which encompass not only generation but also classification and question-answering tasks. Although sampling-based decoding aids in increasing diversity, it operates stochastically, which is not beneficial for classification or question-answering. Nevertheless, we examine whether ID has benefits from sampling, and the results are presented in Table 6. From the outset, one can observe a performance degradation across all methods, including the baseline, with ID experiencing a particularly significant decline. As described in Section 4, we demonstrate that this outcome stems from the smoothing effect from the characteristics of ID. Because ID reduces the top1 probability by increasing the probabilities for other tokens, sub-optimal tokens can be easily sampled, leading generalization far worse than that of the greedy decoding.

CD ablations In Section 4, we discuss the application of Contrastive Decoding (CD) to ID for unseen task generalization. The comprehensive results for the hyperparameters that demonstrates the highest performance during our experiments can be found in Figure 10. As previously mentioned, while CD experiences significant sensitivity concerning model selection, the simultaneous use of ID’s opposite instruction with CD (i.e. ID-amateur) reduces this sensitivity. Even when the expert model is smaller than the amateur model, it displays more robust results, and the degradation is considerably less when compared to the standard CD. This can be attributed to the fact that as the amateur model grows in size, it better understands the meaning of the opposite instruction, thereby producing significant noisy logits.

Ablations on the Number of Random Words for Noisy Instruction To understand the influence of the number of random words in the noisy instruction, we conduct ablation experiments varying their count. In Table 7, performance metrics for Tk-Large and Tk-XL models across different random




Figure 10: Performance results after applying CD to SUPNATINST in comparison to the expert model’s original score. We explore various values for τ , representing the temperature parameter τ for the amateur model, within the range $[0.5, 10.0]$. The parameter α , which constrains the model’s confidence, is set to 0.1 as in Li et al. [20]. For the ID-amateur approach, which introduces noisy instructions to the ‘amateur’ model, we examine the optimal value for ϵ among 0.1, 0.2, and 0.3.

Table 7: Performance degradation with increasing number of random words in the noisy instruction for Tk-Large and Tk-XL models. This table highlights the trade-offs when introducing randomness in instructions.

Model	The number of random words						
	1	3	5	10	30	50	100
Tk-Large	41.77	41.74	41.73	41.54	41.40	41.36	41.35
Tk-XL	46.46	46.39	46.34	46.30	46.29	46.25	46.11

word counts are presented. As the number of random words increases, there is a marginal decline in performance for both models. This suggests a potential saturation point beyond which additional random words might not offer significant noise benefits. The results underscore the importance of adjusting the randomness level in the noisy instruction to achieve optimal performance.

Table 8: Variation in performance with different truncation ratios in the Truc-Shuf approach for Tk-Large and Tk-XL models. The table showcases the resilience and adaptability of the models to varying degrees of truncation in the instructions

Model	Trucation Ratio								
	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Tk-Large	41.68	41.67	41.57	41.87	41.60	41.70	41.61	41.73	41.66
Tk-XL	46.37	46.31	46.39	46.60	46.21	46.45	46.30	46.26	46.67

Ablations on the Ration of Truc-Shuf To ascertain the impact of truncation on the model’s performance, we perform ablation studies varying the truncation ratio. As illustrated in Table 8, we report the performance for Tk-Large and Tk-XL models across different truncation ratios. The table indicates that the models exhibit varying sensitivity to the truncation level. Notably, neither extreme truncation nor minimal truncation consistently maximizes the performance, suggesting an intermediate truncation ratio might be optimal. The results underline the significance of adjusting the truncation ratio to optimize the balance between the retention of task-relevant information and the introduction of noise.

Example on Logits Correction by ID The baseline response guided by the original instruction displays a slight ambiguity, predicting tokens like ‘True’ and ‘Correct’ at relatively high levels, but also showing minor confusion with ‘Answer’, ‘Fal’, and ‘Yes’. However, for the given task, the correct response should be ‘Correct’. When using the ‘null’ with ID, the prediction scores across

Task ID	task133_winowhy_reason_plausibility_detection
Definition	"In this task you need to indicate the plausibility of reasoning for the pronoun coreference relations. Each of the provided inputs contains a sentence with a target ...{Skip}... You should answer 'Correct' if the reasoning made sense, otherwise, you should answer 'Wrong'."
Instance	"Sentence: Thomson visited Cooper's grave in 1765. At that date he had been dead for five years.\n Reason: The 'he' refers to cooper because dead people are in the graves. \n Question: Is the above reasoning correct or wrong?" # Golden Label: Correct

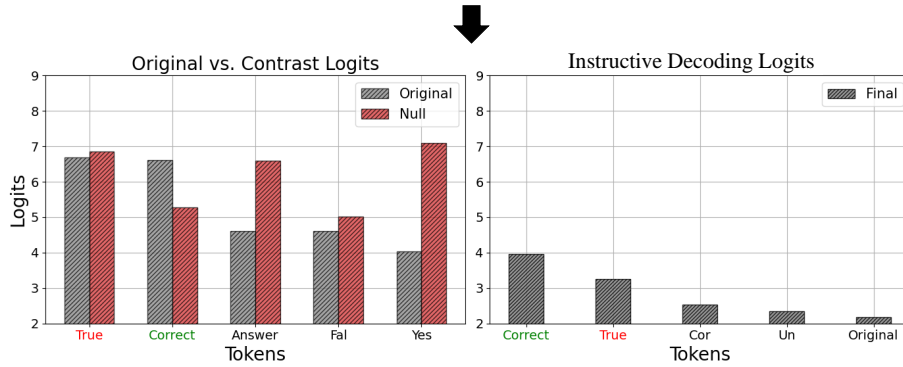


Figure 11: An example on logits correction by ID with Tk-XL model.

these tokens generally increase. By contrasting these outcomes, the model is further reinforced to adhere to the 'Correct' response, underlining the strength of ID in guiding models towards the appropriate response.

Table 9: List of Task IDs in the SUPNATINST used to evaluate *Adherence* and *Coherence* of instruction following.

Task IDs	
task893_gap_fill_the_blank_coreference_resolution	task641_esnli_classification
task1529_scitail1.1_classification	task202_mnli_contradiction_classification
task1393_superglue_copa_text_completion	task1344_glue_entailment_classification
task1387_anli_r3_entailment	task880_schema_guided_dstc8_classification
task738_perspectrums_classification	task1439_doqa_cooking_isanswerable
task642_esnli_classification	task242_tweetqa_classification
task890_gcwd_classification	task1612_sick_label_classification
task1442_doqa_movies_isanswerable	task233_iirc_link_exists_classification
task936_defeasible_nli_snli_classification	task1386_anli_r2_entailment
task290_tellmewhy_question_answerability	task391_causal_relationship
task201_mnli_neutral_classification	task520_aquamuse_answer_given_in_passage
task892_gap_reverse_coreference_resolution	task828_copa_commonsense_cause_effect
task1155_bard_analogical_reasoning_trash_or_treasure	task1385_anli_r1_entailment
task1531_daily_dialog_type_classification	task1516_imppres_naturallanguageinference
task1394_meta_woz_task_classification	task1615_sick_tclassify_b_relation_a
task970_sherliic_causal_relationship	task1390_wscfixed_coreference
task199_mnli_classification	task133_winowhy_reason_plausibility_detection
task226_english_language_answer_relevance_classification	task935_defeasible_nli_atomic_classification
task020_mctaco_span_based_question	task937_defeasible_nli_social_classification
task1388_cb_entailment	task329_gap_classification
task1554_scitail_classification	task050_multirc_answerability
task362_spolin_yesand_prompt_response_sub_classification	task220_rocstories_title_classification
task232_iirc_link_number_classification	task1391_winogrande_easy_answer_generation
task1533_daily_dialog_formal_classification	task1624_disfl_qa_question_yesno_classification
task827_copa_commonsense_reasoning	task879_schema_guided_dstc8_classification
task190_snli_classification	task200_mnli_entailment_classification
task1534_daily_dialog_question_classification	task392_inverse_causal_relationship
task640_esnli_classification	task623_ohsumed_yes_no_answer_generation
task1640_aqa1.0_answerable_unanswerable_question_classification	task349_squad2.0_answerable_unanswerable_question_classification