

Deconstructing In-Context Learning: Understanding Prompts via Corruption

Namrata Shivagunde, Vladislav Lialin, Sherin Muckatira, Anna Rumshisky

University of Massachusetts Lowell
{nshivagu, vlialin, smuckati, arum}@cs.uml.edu

Abstract

The ability of large language models (LLMs) to "learn in context" based on the provided prompt has led to an explosive growth in their use, culminating in the proliferation of AI assistants such as ChatGPT, Claude, and Bard. These AI assistants are known to be robust to minor prompt modifications, mostly due to alignment techniques that use human feedback. In contrast, the underlying pre-trained LLMs they use as a backbone are known to be brittle in this respect. Building high-quality backbone models remains a core challenge, and a common approach to assessing their quality is to conduct few-shot evaluation. Such evaluation is notorious for being highly sensitive to minor prompt modifications, as well as the choice of specific in-context examples. Prior work has examined how modifying different elements of the prompt can affect model performance. However, these earlier studies tended to concentrate on a limited number of specific prompt attributes and often produced contradictory results. Additionally, previous research either focused on models with fewer than 15 billion parameters or exclusively examined black-box models like GPT-3 or PaLM, making replication challenging. In the present study, we decompose the entire prompt into four components: task description, demonstration inputs, labels, and inline instructions provided for each demonstration. We investigate the effects of structural and semantic corruptions of these elements on model performance. We study models ranging from 1.5B to 70B in size, using ten datasets covering classification and generation tasks. We find that repeating text within the prompt boosts model performance, and bigger models (≥ 30 B) are more sensitive to the semantics of the prompt. Finally, we observe that adding task and inline instructions to the demonstrations enhances model performance even when the instructions are semantically corrupted. The code is available at this [URL](#).

Keywords: ICL, prompting, prompt components, prompt corruption, zero-shot evaluation

1. Introduction

The ability of language models to respond to prompts and learn in context has led to an explosive growth in their use, culminating in the proliferation of AI assistants such as ChatGPT (OpenAI, 2023), Claude (Anthropic, 2023), and Bard (Google AI, 2023), which use large pre-trained language models as the backbone. AI assistants built on top of backbone models are robust to prompt variation, in large part due to alignment techniques involving learning from human feedback (Ouyang et al., 2022). However, the underlying backbone models are notoriously brittle in this respect, and their performance often varies widely with slight prompt modifications. Building a high-quality backbone model remains a core challenge, and one of the more common ways to gauge their quality is to conduct in-context evaluation, which suffers from high sensitivity to prompt variation. Despite this sensitivity, models have shown remarkable resilience to corruption in certain parts of the prompt. Recently proposed explanations for in-context learning, such as implicit gradient descent (Dai et al., 2022; von Oswald et al., 2022), fail to account for this resiliency.

A number of previous studies have examined the impact of prompts on model performance across different tasks (Brown et al., 2020; Radford et al.,

2019; Lu et al., 2021; Lialin et al., 2022; Talmor et al., 2020; Webson and Pavlick, 2021; Lampinen et al., 2022; Reynolds and McDonell, 2021; Min et al., 2022; Zhao et al., 2021; Raman et al., 2022; Kim et al., 2022). However, the results have sometimes been contradictory. In particular, the studies of individual prompt components have been plagued by inconsistency. For instance, Webson and Pavlick (2021) found that meaningless instructions don't have a significant effect on model performance. On the other hand, evidence from Mishra et al. (2021b) and Reynolds and McDonell (2021) suggested that meaningful prompts are crucial for zero-shot performance. Similarly, while Min et al. (2022) and Wei et al. (2023b) demonstrated that label semantics aren't necessary for zero-shot performance, both Kim et al. (2022) and Webson and Pavlick (2021) argued otherwise.

Additionally, the majority of prior research has focused either on smaller models with < 15 B parameters or black-box LLMs like GPT-3 (Brown et al., 2020), InstructGPT (Ouyang et al., 2022), and PaLM (Chowdhery et al., 2022), and therefore don't offer a complete understanding of the significance of different prompt components across model sizes.

In the present study, we decompose the input prompt into four components: task instructions, inline instructions, and demonstrations that consist

of input/target label pairs (see Figure 1) and investigate the effect of structural and semantic corruptions to these prompt components across ten models, ranging from 1.5B to 70B. We evaluate them on ten datasets, covering both classification and generation tasks. Building on techniques from model interpretability research, we also examine the average per component attention of two of the models to determine which components contribute more to model output. Our results show that:

1. Including **repeated text in the prompt boosts model performance**.
2. Addition of both **task and inline instructions improves model performance**, even when these instructions are random words.
3. **Larger models exhibit higher sensitivity to prompt semantics** and pay more attention to the semantically relevant prompt components.

2. Related work

Several prior studies have investigated in-context learning (ICL) capabilities of large language models (Brown et al., 2020; Radford et al., 2019; Lu et al., 2021; Lialin et al., 2022; Talmor et al., 2020; Webson and Pavlick, 2021; Lampinen et al., 2022; Reynolds and McDonell, 2021; Min et al., 2022; Zhao et al., 2021; Raman et al., 2022; Wei et al., 2023b; Madaan and Yazdanbakhsh, 2022). However, when it comes to the impact of different parts of the prompt on model performance, the conclusions have often been inconsistent. For example, Webson and Pavlick (2021) suggest that relevant and irrelevant instructions in the prompt yield similar model performance, whereas Mishra et al. (2021b) and Reynolds and McDonell (2021) argued the opposite. The latter studies showed that detailed and task-relevant prompts that closely resemble natural human language give better model performance. Similarly, Kim et al. (2022) studied the importance of ground-truth labels for in-context learning and found that ground-truth labels were important for ICL, contradicting the results from Min et al. (2022).

While prior work has not provided a comprehensive analysis of the impact of different prompt components on model performance, a few studies have selectively examined specific elements of the prompt. For example, Lampinen et al. (2022) looked into adding explanations to the demonstration and found that adding task explanations can significantly improve model performance. Min et al. (2022) examined different aspects of in-context demonstrations and found that input-label mapping did not significantly affect model accuracy. Webson and Pavlick (2021) and Gu et al. (2021) studied

instructions and labels and suggested that the labels were more important than instructions. Wei et al. (2023b) investigated the effect of semantic priors associated with the labels during pre-training, relative to the input-label mapping provided in the prompt, showing that the ability to override semantic priors with the prompt is an emergent ability.

Prior work has also examined different prompting strategies, as well as additional fine-tuning to improve in-context performance. For instance, Xu et al. (2023) introduced re-reading prompting strategy where they repeat the question in the prompt and found that this strategy improves performance for ChatGPT and GPT-3. Wei et al. (2023a) proposed “symbol tuning”, fine-tuning models with arbitrary labels, and observed performance improvements on unseen ICL tasks. In a different approach, Gonen et al. (2022) proposed constructing the prompt with lower perplexity for better performance.

Few studies (Dai et al., 2022; von Oswald et al., 2022) also linked attention computation performed during in-context learning to model updates performed with gradient descent. However, it is unclear how this mechanism would account for some aspects of in-context learning, such as the success of zero-shot prompting.

3. Experiment setup

3.1. Prompts

Prompt components We use the term “**prompt**” to refer to the complete input text provided to the model. A prompt consists of four main components: a task instruction, demonstration input, demonstration label, and brief inline instructions accompanying each demonstration (see Figure 1). Two newlines are used as a separator after the task instruction and after each demonstration. We refer to a prompt with all components, including a test instance and its inline instruction, as a **baseline prompt**. Our experiments are conducted in a zero-shot and 4-shot setting. Figure 1 shows the baseline prompt for Twitter Emotion classification dataset. Baseline prompts for all the datasets are provided in the Appendix A.

Prompt design We leverage the task instructions and demonstrations provided by Wang et al. (2022c) for each dataset, as they have been reviewed and refined through multiple iterations. We use inline instructions from PromptSource (Bach et al., 2022). To ensure coherence and simplicity, we made a few changes to the task and inline instructions, following the recommendations of Gonen et al. (2022).

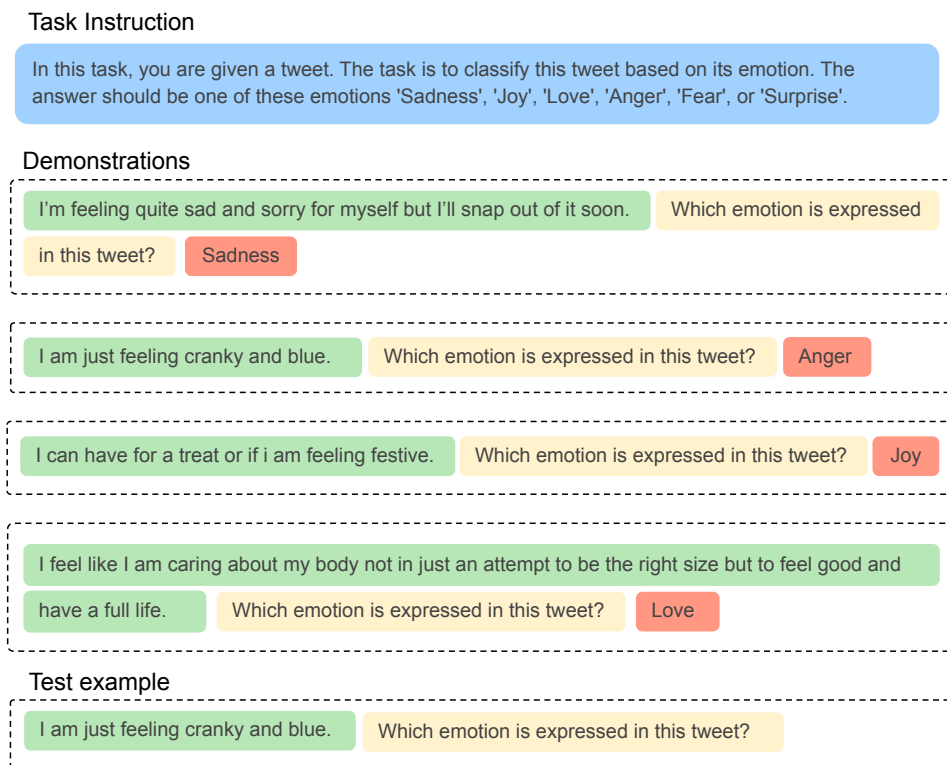


Figure 1: Prompt Components of Twitter Emotion Classification baseline prompt. Demonstration includes input, inline instruction, label. Two newlines are added as separators after task instruction and each demonstration. Prompts taken verbatim from Super-NaturalInstructions and PromptSource.

Prompt corruptions We perform two types of prompt corruptions: **structural corruption** and **semantic corruption**. In structural corruption, we add or remove the prompt components, depending on the setup. We start with the test instance and add components one by one to analyze their effect on model performance. To assess the impact of repeated text in the prompt, we systematically eliminate inline instructions from the baseline prompt. We remove the inline instruction from one demonstration, then two, and continue this process until we have removed the inline instructions from all four demonstrations. These corruptions are referred to as **repeated text corruptions**. We keep the inline instruction which follows the test instance as is.

In semantic corruption, we disrupt the semantics of prompt components. Task and inline instructions are corrupted with random words drawn from the `english_words`¹ set. With a 100% corruption rate, we refer to this corruption as the **random words** corruption. The random word instructions retain the same number of tokens as the original (meaningful) instructions. Labels are perturbed by assigning incorrect labels to the demonstrations. These incorrect labels are drawn from the same label space. This corruption is referred to as the

wrong label corruption and is only applied to classification tasks. In the **random words label** corruption, we replace original labels with random words, similar to the instruction random words corruption, but we use the original labels to assess the model's predictions. To noise the demonstration inputs, we replace them with random sentences sampled from Common Crawl. We refer to this as **Out-Of-Distribution (OOD)** input corruption.

3.2. Models, datasets and metrics

Models To cover a broad range of models, we conducted experiments with ten models ranging in size from 1.5B to 70B. The models are GPT2-xl (Radford et al., 2019), GPT-J-6B (Wang and Komatsuzaki, 2021), Pythia-12B (Biderman et al., 2023), OPT-30B, OPT-30B-IML-MAX², OPT-66B (Zhang et al., 2022a), Vicuna-33B (Chiang et al., 2023), Llama-7B, Llama-2-70B and Llama-2-70B-chat (Touvron et al., 2023). This provides a wide range of model sizes, and types and also doesn't focus on a single model family, making the results more generalizable. Our study included pre-trained language models, as well as instruction-tuned and aligned models. We refer to models as "aligned"

¹<https://pypi.org/project/english-words/>

²Instruction tuned variant of OPT.

when they undergo additional training through reinforcement learning from human feedback (RLHF) (Ouyang et al., 2022).

Datasets The evaluation was conducted on ten datasets from Super-NaturalInstructions (Wang et al., 2022c). Datasets include eight classification tasks: RTE, Medical Question Pair, Financial Phrasebank, Twitter Emotion classification, CoLA, AgNews, COPA, Com2sense, and two generation tasks: TriviaQA and Mathdataset answer generation (Wang et al., 2022c). Following Wang et al. (2022c), we used 100 randomly sampled balanced test instances for each of the 10 datasets. Data statistics are shown in Table 3 in the Appendix.

Evaluation method and metrics For evaluation, we used Exact Match for classification tasks, and Rouge-L for generative tasks. Following Wang et al. (2022c), we strip the model response at the first full stop symbol. We used the jackknife variance estimate method to calculate the mean of model performance. The mean performance, averaged across tasks, is reported for each model in Tables 1 and 2. For all figures, we plot the mean as the “average score”. For generation tasks, we used the greedy decoding strategy, setting the top-p and temperature values to 1 and limiting the maximum number of new tokens to 10.

Attention computation To understand the significance of different prompt components, we computed the average attention norm per prompt component for GPT-J-6B (Wang and Komatsuzaki, 2021) and OPT-30B (Zhang et al., 2022a). Following Kobayashi et al. (2020), we compute the L2 norm of the sum of the attention-weighted value vector $\|\sum \alpha V(x)\|$, where α is the attention weight, x is the input vector, and $V(x) = W_O(W_V x)$ is the value projection of x , followed output transformation W_O . Specifically, we used the last token of the prompt as the query token and extracted attention norms for the other tokens. For each token, we averaged these norms across all layers. We then averaged the resulting scores over all tokens corresponding to a given prompt component. This average is reported in Figures 6, 13 and 12.

For GPT-J-6B, each plot shows the average attention norm over 100 samples (10 samples per dataset). For OPT-30B, due to computing costs, we focused on datasets with shorter baseline prompts: CoLA, Twitter Emotion classification, and TriviaQA. For each attention plot, we included the average attention norm from 30 samples, selecting only those where the model predictions were correct.

4. Results

Tables 1 and 2 show results for each corruption across 10 datasets. Here’s a breakdown of the prompt configurations used:

- **Test instance:** The input prompt containing only the test instance.
- **+task instr.:** Test instance with the task instruction added.
- **+inline instr.:** Test instance with an inline instruction added instead of the task instruction.
- **+both instr.:** Test instance with both task and inline instructions added.
- **+demos.:** Test instance plus four demonstrations (no instructions included).
- **+task instr. +demos.:** Test instance with task instructions and four demonstrations.
- **+inline instr. +demos.:** Test instance with four demonstrations (each containing an inline instruction) and no task instruction.
- **Baseline:** Includes all components (task instruction, inline instruction, demonstrations, and test instance).
- **Baseline -inputs:** Baseline prompt with demonstration inputs removed.
- **Baseline -labels:** Baseline prompt with demonstration labels removed.
- **Rw both instr.:** Random word corruption applied to both task and inline instructions.
- **Rw labels:** Random word corruption applied to labels.
- **OOD inputs:** Out-of-distribution input corruption.
- **Inline instr. in [n] demos.:** Meaningful inline instructions added to "n" demonstrations.
- **Rw inline instr. in [n] demos.:** Random word inline instructions added to "n" demonstrations.

Adding task and inline instructions boosts the performance even when the instructions are random words. Our experiments with four models (see Figure 2) highlight that the addition of demonstrations to the test instance has the most impact, producing a gain of 25-35% across all models. Accuracy is improved further by adding task instructions and inline instructions (Figure 2). The models gain between 5-18% accuracy when meaningful instructions are added. Interestingly, this gain is between 1-12% when instructions are just random words, except for Llama-70B. Figure 3 shows a similar effect for all ten models.

Inline instructions are more important than task instructions. From Figure 3, we see that the performance gained by inline instruction is 2.5-12.5% across models whereas task instruction helps models by only 1-7.5%. This pattern is observed for models of all sizes, except OPT-66B, where the gain obtained by the inline instructions is close to

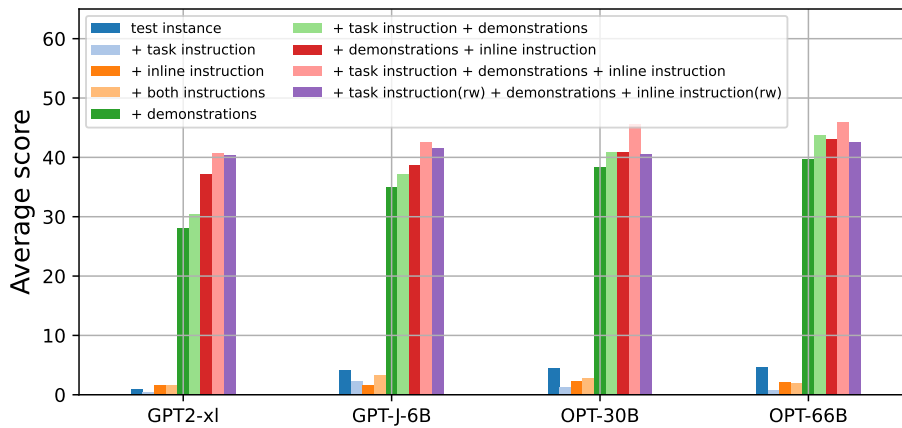


Figure 2: Demonstrations improve the average score, adding task and inline instruction improves it further, even when instructions are just random words. The Y-axis represents the average score across all datasets. The use of random words is indicated with “rw”.

Structural Corruptions	GPT2-xl	GPT-J-6B	OPT-30B	OPT-66B
Test instance	0.9	4.2	4.4	4.7
+ task instr.	0.4	2.3	1.2	0.7
+ inline instr.	1.6	1.6	2.2	2.1
+ both instr.	1.6	3.3	2.7	2.0
+ demo.	28.1	34.9	38.3	39.6
+ task instr. + demo.	30.4	37.1	40.9	43.7
+ inline instr + demo.	37.1	38.7	40.8	43.0
Baseline	<u>40.7</u>	<u>42.6</u>	<u>45.6</u>	<u>45.9</u>
Baseline - labels	0.1	0.2	0.5	0.9
Baseline - inputs	27.1	17.0	22.0	22.3

Table 1: Model performance averaged across all datasets. The highest performance is in bold, *baseline prompt* performance is underlined.

that of the task description (Figure 3). This suggests models benefit from the brief repetitive text more than from a detailed task instruction.

Repeated text boosts performance. We further investigated the effects of repeating inline instructions. In Figure 4, we plot the results for the baseline prompt (which includes all components) and the results obtained when eliminating inline instructions from demonstrations one by one. Note that we always keep the inline instruction that occurs after the test instance. We see a huge drop in performance when removing the inline instruction from all demonstrations. The drop is 20-35% across all models, except OPT-30B-IML, which shows a drop of 8.8%. Interestingly, we observed a similar effect for prompts in which inline instructions were replaced with random words, producing a performance drop of 40-51% (cf. Figure 5). This suggests that the mere presence of repetitive text in the prompt, whether relevant or irrelevant, can boost model performance.

However, how often we introduce these repeti-

tions in the prompt also matters. Table 2 shows that models like OPT-30B and Llama-70B can achieve better performance with only two or three meaningful inline instructions in the prompt. The attention plot in Figure 11 shows that if we introduce inline instruction in four demonstrations rather than in one, the attention to the input segment of each demonstration dropped from 2% to 1.8% and the attention to each of the labels decreased by around 0.4%. We see a similar pattern for repeated text corruptions in prompts with random word instructions (see Figure 15 in Appendix).

Labels must be drawn from the label space, but need not be correct. When we perturb labels with the **wrong label** corruption, the performance drops just by 0-6% across all models (except for Llama-2-70B, where the drop is 18.6%). However, when we apply **random words label** corruption, the accuracy drops to almost zero for all models (except OPT-30B-IML-MAX) (cf. Figure 9). Complete removal of the labels from the prompt has a similar effect (cf. Figure 7).




Figure 3: Adding relevant or meaningless instruction to the prompt improves model performance. The components are added to the test instance. For example ‘+ demonstrations’ means test instance + demonstration. The Y-axis represents the average score across all datasets. Random words are indicated with “rw”.

Corruptions	GPT2 xl	GPT-J 6B	LLama 7B	Pythia 12B	OPT 30B	OPT-30B IML-MAX	Vicuna 33B	OPT 66B	LLama-2 70B	LLama-2 70B-chat
Structural										
+ demos.	28.1	34.9	40.8	33.7	38.3	49.0	48.5	38.6	49.1	48.8
+ task instr. + demos.	30.4	37.1	44.2	37.2	40.9	55.1	49.7	43.7	55.9	56.5
+ inline instr. + demos.	37.1	38.7	46.0	41.1	40.8	64.6	59.0	43.0	61.7	58.3
Baseline	<u>40.7</u>	<u>42.6</u>	<u>47.4</u>	<u>38.9</u>	<u>45.6</u>	<u>67.7</u>	<u>61.7</u>	<u>45.9</u>	<u>64.5</u>	<u>63.4</u>
Semantic										
Rw both instr.	40.4	41.5	44.5	39.3	40.5	49.8	50.5	42.6	48.2	52.7
Rw labels	3.7	1.8	1.4	1.4	3.4	46.5	3.8	2.7	1.2	7.9
OOD inputs	41.7	40.7	43.9	38.1	44.6	67.6	57.1	40.7	50.5	57.4
Repeated Text										
Inline instr. in 3 demos.	43.2	43.2	48.2	39.3	45.8	65.8	61.1	43.1	64.6	63.5
Inline instr. in 2 demos.	40.9	43.1	44.5	41.6	43.7	63.9	59.3	43.5	62.7	62.6
Inline in instr. 1 demos.	40.6	43.1	42.7	39.8	45.7	63.2	58.8	41.9	62.0	61.3
Inline in instr. 0 demos.	13.3	22.6	17.9	14.8	21.4	59.0	30.5	20.3	35.1	29.2
Rw inline instr. in 3 demos.	35.8	38.8	43.1	38.8	35.6	50.9	51.9	39.7	48.2	56.1
Rw inline instr. in 2 demos.	35.9	36.1	41.1	36.0	35.4	45.4	46.6	40.6	49.2	49.0
Rw inline instr. in 1 demos.	33.0	34.9	36.7	31.4	22.4	35.0	38.7	32.8	37.7	41.5
Rw inline instr. in 0 demos.	0.6	0.2	1.3	0.4	0.2	0.2	1.1	0.3	0.7	1.6

Table 2: Model performance averaged across all datasets. Structural corruption is when components are added to the test instance. Repeated text corruptions are performed on baseline prompt which includes inline instruction in all four demonstrations. Random words text is represented by “Rw”. The top two performances for each model are in bold, and baseline prompt performance is underlined.

Bigger models are more sensitive to the semantics of the prompt. We divide the models in the study into smaller (<15B) and bigger (≥ 30B) models. In Figure 3, smaller models show the performance gain between 5-12% with both relevant and irrelevant instructions, whereas bigger models gain more with meaningful instructions (7-18%) and just 1-4% with random word instructions. When we perturb demonstration inputs with OOD sentences (see Figure 10), smaller models’ accuracy drops by 1-4%. In bigger models, this performance decrease is larger (1-6%), with Llama-2-70b showing a huge drop of 18%, which suggests that bigger models are more sensitive to prompt semantics.

Bigger models pay more attention to relevant components. In Figure 6, we plot the average attention per component for GPT-J-6B and OPT-30B baseline prompts. In line with earlier observations about vertical attention patterns (Kovaleva et al., 2019), we find that the models allocate the highest attention weight to separators, after which the most attended to component is labels. Inline instructions are next, followed by demonstration inputs and task instructions. Compared to smaller models, larger models seem to allocate more attention to relevant components and less to separators when generating the target label. For example, OPT-30B allocates 5.4% less attention to separators compared




Figure 4: Repeated text boosts performance. Inline instruction in four demos is the baseline prompt. Inline instruction which occurs after the test instance is kept as is.

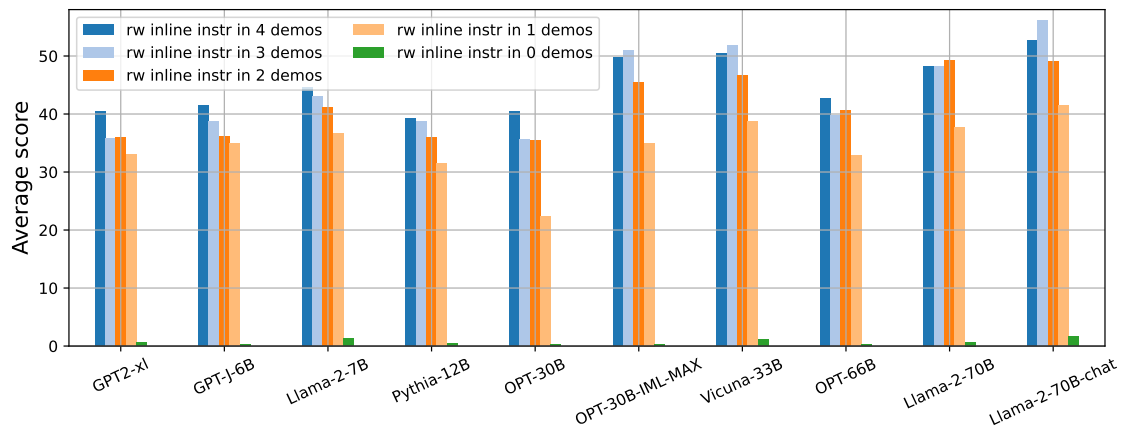


Figure 5: Repeated text boosts performance even when the text is irrelevant; “rw” refers to random words. The prompts include all components but the instructions are replaced with random words.

to GPT-J-6B, and instead increases attention to inline instructions by 3.5% and to demonstration inputs by 2.3%. To compare how models behave when text is corrupted semantically, we plotted attention for prompts with meaningful versus irrelevant instructions (see Figures 12 and 13). We see that GPT-J-6B shifts its attention from separators and labels to the demonstration input and the random inline instructions. OPT-30B does the opposite: it reduces its attention to random words

text and shifts it to the separator.

Results are similar in classification and generation tasks, with few exceptions. As can be seen in Table 4, repeating inline instructions has a big impact on both classification and generation tasks regardless of model size. The first repetition has the most pronounced effect, and this is true even when the inline instructions are random. To see this effect, compare the rows “Inline instr. in 0 demos”




Figure 6: Average attention per component for GPT-J-6B and OPT-30B baseline prompts. ‘D’ stands for Demonstration and “sep” for the new line separator.




Figure 7: Label from label space is important. Complete removal of labels drops the performance to almost zero.

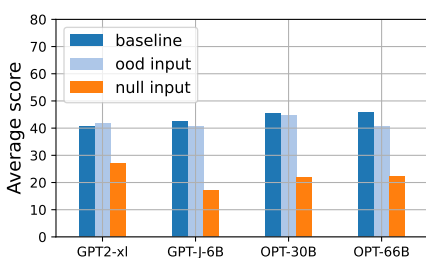


Figure 8: Semantics of the demonstration input is not important. Complete removal of labels drops the performance.

and “Inline instr. in 1 demos”, as well as the rows “Rw Inline instr. in 0 demos” and “Rw Inline instr. in 1 demos” in Table 4. We also observe that adding relevant or random word instructions to demonstrations improves GPT2-xl performance by an average of 12.5% for classification and 3.8% for generation tasks. For the larger model LLama-2-70B, relevant instructions lead to gains of 18.5% for classification and 3.3% for generation tasks. Adding random word instructions yields LLama-2-70B, a marginal improvement of 0.1% in classification tasks and a drop of 1.7% for generation tasks. To see this effect, compare the rows “+demos.”, “Baseline” and “Rw both instr.” in Table 4. The performance drop from meaningful to random word instructions is more pronounced (20.2% drop) in LLama-2-70B in classification tasks (compare rows “Baseline” and “Rw both instr.” in Table 4) suggesting that larger models pay more attention to the meaning of the instruction. In generation tasks, both model sizes exhibit a comparable drop in performance.

Results are consistent across most of the datasets. Tables 5 and 6 show results for each dataset individually for both GPT2-xl and LLama-2-70B. The first repetition of relevant or irrelevant inline instructions in the prompt significantly boosts performance across all datasets for both model sizes. Adding relevant instructions proves beneficial for both GPT2-xl and LLama-2-70B on all 10

datasets. At the same time, adding random word instruction benefits GPT2-xl on 8 out of 10 datasets, but LLama-2-70B only on 6 out of 10 datasets. Corrupting labels with random words impacts GPT-2-xl on 8 out of 10 datasets and LLama-2-70B on all datasets.

5. Conclusion

This study investigated the importance of different components of a prompt for large language models. We systematically corrupted prompts in different ways across 10 models ranging from 1.5 billion to 70 billion parameters and evaluated their performance on 10 diverse datasets. We also examined how much attention the models allocate to different prompt components. One interesting finding was that adding any inline instructions to the prompt, even just random words, actually helps models perform better. We also showed that repeated text improves model performance drastically and that larger models are substantially more sensitive to prompt semantics. We hope our study will pave the way for more refined and effective prompting strategies in future applications.

6. Limitations

Our study was focused on exploring various types of corruption across a diverse range of datasets and model sizes. It involved a large number of experiments and certain prompt elements were held constant, such as demonstrations and instructions. Altering these components might introduce variations in the results, and this aspect should be taken into consideration for further research. Additionally, we limited the attention analysis to datasets with shorter prompts due to the computational intensity and cost associated with computing attention norms. An additional limitation arises from the use of the same prompt template across all model types. This uniformity may lead to some performance discrepancies in instruction-tuned models.

7. Ethic Statements

Our goal with this study is to enrich the understanding of prompting and contribute to the responsible utilization of large language models. We believe that attention analysis can offer meaningful insights to the research community, facilitating the development of more robust language models. It’s noteworthy that all the models and datasets employed in our research are open source, and we meticulously reported all experiment details in the paper to support the transparency and accessibility of the research.




Figure 9: Using labels from the correct label space is crucial for model performance.




Figure 10: Semantics of the demonstration input is not important.




Figure 11: Average OPT-30B attention per component for repeated text corruptions. “Inline” refers to the presence of the number of inline instructions in the baseline prompt. A solid black box represents omitted components.




Figure 12: Average attention per component for OPT-30B: Baseline prompt versus prompt when both task and inline instructions are replaced by random words.




Figure 13: Average attention per component for GPT-J-6B. Baseline prompt versus prompt when both task and inline instructions are replaced by random words.

8. Acknowledgement

This work was funded in part by an Amazon Alexa AI research award to Anna Rumshisky. We would like to express our gratitude to Anton Kovalev for helping with the tables in the paper.

9. References

- Anthropic. 2023. [Model card and evaluations for claude models](#).
- Stephen H. Bach, Victor Sanh, Zheng Xin Yong, Albert Webson, Colin Raffel, Nihal V. Nayak, Abheesht Sharma, Taewoon Kim, M Saiful Bari, Thibault Févry, Zaid Alyafeai, Manan Dey, Andrea Santilli, Zhiqing Sun, Srulik Ben-David, Canwen Xu, Gunjan Chhablani, Han Wang, Jason Alan Fries, Maged S. Al-Shaibani, Shanya Sharma, Urmish Thakker, Khalid Almubarak, Xiangu Tang, Mike Tian-Jian Jiang, and Alexander M. Rush. 2022. [Promptsources: An integrated development environment and repository for natural language prompts](#). *ArXiv*, abs/2202.01279.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Stella Rose Biderman, Hailey Schoelkopf, Quentin G. Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron, Lintang Sutawika, and Oskar van der Wal. 2023. [Pythia: A suite for analyzing large language models across training and scaling](#). *ArXiv*, abs/2304.01373.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Zhifang Sui, and Furu Wei. 2022. Why can gpt learn in-context? language models secretly perform gradient descent as meta optimizers. *arXiv preprint arXiv:2212.10559*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Hila Gonen, Srinu Iyer, Terra Blevins, Noah A Smith, and Luke Zettlemoyer. 2022. Demystifying prompts in language models via perplexity estimation. *arXiv preprint arXiv:2212.04037*.
- Google AI. 2023. [Bard Google AI](#).
- Yuxian Gu, Xu Han, Zhiyuan Liu, and Minlie Huang. 2021. Ppt: Pre-trained prompt tuning for few-shot learning. *arXiv preprint arXiv:2109.04332*.
- Daniel Khoshdel, Shane Lyu, Sewon Min, Lianhui Qin, Kyle Richardson, Sameer Singh, Sean Welleck, Hannaneh Hajishirzi, Tushar Khot, Ashish Sabharwal, et al. 2021. Prompt waywardness: The curious case of discretized interpretation of continuous prompts. *arXiv preprint arXiv:2112.08348*.
- Junyeob Kim, Hyeon Gwon Kim, Hyunsoo Cho, Hwiyeol Jo, Sang-Woo Lee, Sang goo Lee, Kang Min Yoo, and Taek Kim. 2022. [Ground-truth labels matter: A deeper look into input-label demonstrations](#). *ArXiv*, abs/2205.12685.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. Attention is not only a weight: Analyzing transformers with vector norms. *arXiv preprint arXiv:2004.10102*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of bert. *arXiv preprint arXiv:1908.08593*.
- Andrew K. Lampinen, Ishita Dasgupta, Stephanie C. Y. Chan, Kory Matthewson, Michael Henry Tessler, Antonia Creswell, James L. McClelland, Jane X. Wang, and Felix Hill. 2022. [Can language models learn from explanations in context?](#)
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu

- Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Vladislav Lialin, Kevin Zhao, Namrata Shivagunde, and Anna Rumshisky. 2022. Life after bert: What do other muppets understand about language? *arXiv preprint arXiv:2205.10696*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021a. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021b. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*.
- Aman Madaan and Amir Yazdanbakhsh. 2022. [Text and patterns: For effective chain of thought, it takes two to tango](#). *ArXiv*, abs/2209.07686.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? *arXiv preprint arXiv:2202.12837*.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2021a. Reframing instructional prompts to gptk’s language. *arXiv preprint arXiv:2109.07830*.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2021b. Natural instructions: Benchmarking generalization to new tasks from natural language instructions.
- OpenAI. 2023. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- 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 E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. [Training language models to follow instructions with human feedback](#). *ArXiv*, abs/2203.02155.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020a. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020b. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Karthik Raman, Iftekhar Naim, Jiecao Chen, Kazuma Hashimoto, Kiran Yalasangi, and Krishna Srinivasan. 2022. Transforming sequence tagging into a seq2seq task. *arXiv preprint arXiv:2203.08378*.
- Laria Reynolds and Kyle McDonell. 2021. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak,

Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Stella Biderman, Leo Gao, Tali Bers, Thomas Wolf, and Alexander M. Rush. 2021. [Multitask prompted training enables zero-shot task generalization](#).

Teven Le Scao, Angela Fan, Christopher Akiki, Elizabeth-Jane Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagn'e, Alexandra Sasha Luccioni, Francois Yvon, Matthias Gall'e, Jonathan Tow, Alexander M. Rush, Stella Rose Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurenccon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa Etxabe, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris C. Emezue, Christopher Klamm, Colin Leong, Daniel Alexander van Strien, David Ifeoluwa Adelani, Dragomir R. Radev, Eduardo G. Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady ElSahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jorg Frohberg, Josephine L. Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro von Werra, Leon Weber, Long Phan, Loubna Ben Allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, Mar'ia Grandury, Mario vSavsko, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad Ali Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto L'opez, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, S. Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin

Heinzerling, Chenglei Si, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Alshaibani, Matteo Manica, Nihal V. Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Févry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiang Tang, Zheng Xin Yong, Zhiqing Sun, Shaked Brody, Y Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre Francois Lavall'ee, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwaa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aur'elie N'ev'eol, Charles Lovering, Daniel H Garrette, Deepak R. Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najaoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, S. Osher Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdenek Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ananda Santa Rosa Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Olusola Ajibade, Bharat Kumar Saxena, Carlos Muñoz Ferrandis, Danish Contractor, David M. Lansky, Davis David, Douwe Kiela, Duong Anh Nguyen, Edward Tan, Emily Baylor, Ezinwanne Ozoani, Fatim T Mirza, Frankline Ononiwu, Habib Rezanejad, H.A. Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jan Passmore, Joshua Seltzer, Julio Bonis Sanz, Karen Fort, Livia Macedo Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, M. K. K. Ghauri, Mykola Burynok, Nafis

- Abrar, Nazneen Rajani, Nour Elkott, Nourhan Fahmy, Olanrewaju Modupe Samuel, Ran An, R. P. Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas L. Wang, Sourav Roy, Sylvain Viguier, Thanh-Cong Le, Tobi Oyebade, Trieu Nguyen Hai Le, Yoyo Yang, Zachary Kyle Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Kumar Singh, Benjamin Beilharz, Bo Wang, Caio Matheus Fonseca de Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourier, Daniel Le'on Perin'an, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully A. Burns, Helena U. Vrabec, Iman I.B. Bello, Isha Dash, Ji Soo Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthi Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, María Andrea Castillo, Marianna Nezhurina, Mario Sanger, Matthias Samwald, Michael Cullan, Michael Weinberg, M Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patricia Haller, R. Chandrasekhar, R. Eisenberg, Robert Martin, Rodrigo L. Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aaroonsiri, Srishti Kumar, Stefan Schweter, Sushil Pratap Bharati, T. A. Laud, Th'eo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yashasvi Bajaj, Y. Venkatraman, Yifan Xu, Ying Xu, Yun chao Xu, Zhee Xiao Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2022. Bloom: A 176b-parameter open-access multilingual language model. *ArXiv*, abs/2211.05100.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- Timo Schick and Hinrich Schütze. 2022. True few-shot learning with prompts—a real-world perspective. *Transactions of the Association for Computational Linguistics*, 10:716–731.
- Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Auto-prompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*.
- Yusheng Su, Xiaozhi Wang, Yujia Qin, Chi-Min Chan, Yankai Lin, Zhiyuan Liu, Peng Li, Juanzi Li, Lei Hou, Maosong Sun, et al. 2021. On transferability of prompt tuning for natural language understanding. *arXiv preprint arXiv:2111.06719*.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2020. olmpics-on what language model pre-training captures. *Transactions of the Association for Computational Linguistics*, 8:743–758.
- Zhixing Tan, Xiangwen Zhang, Shuo Wang, and Yang Liu. 2021. Msp: Multi-stage prompting for making pre-trained language models better translators. *arXiv preprint arXiv:2110.06609*.
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*, abs/2307.09288.
- Johannes von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2022. Transformers learn in-context by gradient descent. *arXiv preprint arXiv:2212.07677*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. Superglue: A stickier benchmark for general-purpose language understanding systems. *CoRR*, abs/1905.00537.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Ben Wang and Aran Komatsuzaki. 2021. Gpt-j-6b: A 6 billion parameter autoregressive language model.

- Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2022a. [Towards understanding chain-of-thought prompting: An empirical study of what matters](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Han Wang, Canwen Xu, and Julian McAuley. 2022b. Automatic multi-label prompting: Simple and interpretable few-shot classification. *arXiv preprint arXiv:2204.06305*.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, et al. 2022c. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. URL <https://arxiv.org/abs/2204.07705>.
- Albert Webson and Ellie Pavlick. 2021. Do prompt-based models really understand the meaning of their prompts? *arXiv preprint arXiv:2109.01247*.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021a. Finetuned language models are zero-shot learners. *ArXiv*, abs/2109.01652.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021b. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Jerry Wei, Le Hou, Andrew Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, et al. 2023a. Symbol tuning improves in-context learning in language models. *arXiv preprint arXiv:2305.08298*.
- Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. 2023b. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*.
- Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf El-nashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. *arXiv preprint arXiv:2302.11382*.
- Xiaohan Xu, Chongyang Tao, Tao Shen, Can Xu, Hongbo Xu, Guodong Long, and Jian-guang Lou. 2023. Re-reading improves reasoning in language models. *arXiv preprint arXiv:2309.06275*.
- Sen Yang, Yunchen Zhang, Leyang Cui, and Yue Zhang. 2022. Do prompts solve nlp tasks using natural language? *arXiv preprint arXiv:2203.00902*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022a. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Yue Zhang, Hongliang Fei, Dingcheng Li, and Ping Li. 2022b. Promptgen: Automatically generate prompts using generative models. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 30–37.
- Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022c. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.
- Chunting Zhou, Junxian He, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022a. Prompt consistency for zero-shot task generalization. *arXiv preprint arXiv:2205.00049*.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022b. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*.

A. Components of Prompts for all datasets

We show baseline prompts for all datasets. We use 4-shot setting and each prompt consists of four components: **Task instruction**, **inline instruction**, demonstration input and **label**. [Test instance] will vary. Each dataset consists of 100 samples and is balanced. Data statistics for each of the datasets is shown in 3.

Medical Question Pair (Classification task)

In this task you are given a medical question pair. Your task is to classify this question pair into two categories 1) 'Similar' if the given two questions have the same connotation or meaning 2) 'Dissimilar' if the given two questions have a different connotation or meaning.

Question1: After how many hour from drinking an antibiotic can I drink alcohol?
Question2: I have a party tonight and I took my last dose of Azithromycin this morning. Can I have a few drinks? **Are these two questions similar or dissimilar?** **Similar.**

Question1: After how many hour from drinking an antibiotic can I drink alcohol?
Question2: I vomited this morning and I am not sure if it is the side effect of my antibiotic or the alcohol I took last night... **Are these two questions similar or dissimilar?** **Dissimilar.**

Question1: Can coarctation of the aorta cause poor growth in height? Question2: I am 4' 8". My mom said that I have a birth defect (coarctation of aorta). Are the two related? **Are these two questions similar or dissimilar?** **Similar.**

Question1: Aspirin allergy - is it worth getting a bracelet? Question2: How much Aspirin can I take for my headache without causing any side effects? **Are these two questions similar or dissimilar?** **Dissimilar.**

[Test instance.] **Are these two questions similar or dissimilar?**

Twitter Emotion Classification (Classification task)

In this task, you are given a tweet. The task is to classify this tweet based on its emotion. The answer should be one of these emotions 'Sadness', 'Joy', 'Love', 'Anger', 'Fear', or 'Surprise'.

Im feeling quite sad and sorry for myself but ill snap out of it soon. **Which emotion is expressed in this tweet?** **Sadness.**

I am just feeling cranky and blue. **Which emotion is expressed in this tweet?** **Anger.**

I can have for a treat or if i am feeling festive. **Which emotion is expressed in this tweet?** **Joy.**

I feel like im caring about my body not in just an attempt to be the right size but to feel good and have a full life. **Which emotion is expressed in this tweet?** **Love.**

[Test instance.] **Which emotion is expressed in this tweet?**

CoLA (Classification task)

You will be given a sentence. If the sentence is grammatically correct and meaningful, then answer with 'Yes', otherwise 'No'.

Our friends won't buy this analysis, let alone the next one we propose. **Is this sentence meaningful and grammatically correct?** **Yes.**

One more pseudo generalization and I'm giving up. **Is this sentence meaningful and grammatically correct?** **Yes.**

They drank the pub. **Is this sentence meaningful and grammatically correct?** **No.**

Day by day the facts are getting murkier. **Is this sentence meaningful and grammatically correct?** **Yes.**

[Test instance.] **Is this sentence meaningful and grammatically correct?**

Dataset	Statistics
RTE	Yes(50), No(50)
Medical Question Pair	Similar(50), Similar(50)
Financial Phrasebank	Neural(33), Negative(33), Positive(34)
Twitter Emotion classification	Sadness(17), Joy(17), Love(17), Anger(17), Fear(16), Surprise(16).
CoLA	Yes(50), No(50)
AgNews	World(25), Sports(25), Business(25), Sci/Tech(25)
COPA	Cause(50), Effect(50)
Com2sense	Yes(50), No(50)
TriviaQA	-
Mathdataset	-

Table 3: Datasets statistics: labels and total number of samples per label in brackets.

Com2sense (Classification task)

You will be given a piece of text either about an everyday event, or a general statement. If the event seems a plausible event, or the general statement makes sense to you then answer the question as 'Yes', otherwise 'No'.

The glass fell of a three-story building, so it broke into pieces. Does this statement make sense to you? Yes.

Marry was going out to work, so she asked her sixteen-year-old daughter to take care of her five-year-old son. Does this statement make sense to you? Yes.

Johnathan didn't have a hammer, so he used a cotton pad to drive the nail into the wood. Does this statement make sense to you? No.

Suraya's best friend is getting married soon, so she will likely choose to go on a trip instead of helping her friend organize the ceremony. Does this statement make sense to you? No.

[Test instance.] Does this statement make sense to you?

RTE (Classification Task)

In this task, you are given two sentences. Indicate if the first sentence clearly entails the second sentence (i.e., one can conclude the 2nd sentence by reading the 1st one). Indicate your answer with 'Yes' if the first sentence entails the second sentence, otherwise answer with 'No'.

Sentence 1: No Weapons of Mass Destruction Found in Iraq Yet. Sentence 2: Weapons of Mass Destruction Found in Iraq. Does Sentence 1 entail Sentence 2? No.

Sentence 1: A place of sorrow, after Pope John Paul II died, became a place of celebration, as Roman Catholic faithful gathered in downtown Chicago to mark the installation of new Pope Benedict XVI. Sentence 2: Pope Benedict XVI is the new leader of the Roman Catholic Church. Does Sentence 1 entail Sentence 2? Yes.

Sentence 1: Herceptin was already approved to treat the sickest breast cancer patients, and the company said, Monday, it will discuss with federal regulators the possibility of prescribing the drug for more breast cancer patients. Sentence 2: Herceptin can be used to treat breast cancer. Does Sentence 1 entail Sentence 2? Yes.

Sentence 1: Nearly 4 million children who have at least one parent who entered the U.S. illegally were born in the United States and are U.S. citizens as a result, according to the study conducted by the Pew Hispanic Center. That's about three quarters of the estimated 5.5 million children of illegal immigrants inside the United States, according to the study. About 1.8 million children of undocumented immigrants live in poverty, the study found. Sentence 2: Three quarters of U.S. illegal immigrants have children. Does Sentence 1 entail Sentence 2? No.

[Test instance.] Does Sentence 1 entail Sentence 2?

Financial Phrasebank (Classification task)

Classify the given a piece of financial news into three classes: positive, negative, and neutral. Output must be 'Positive', 'Negative', or 'Neutral'.

According to Gran , the company has no plans to move all production to Russia , although that is where the company is growing. Is the sentiment of the sentence 'Negative', 'Neutral', or 'Positive'? **Neutral**.

Technopolis plans to develop in stages an area of no less than 100,000 square meters in order to host companies working in computer technologies and telecommunications , the statement said. Is the sentiment of the sentence 'Negative', 'Neutral', or 'Positive'? **Neutral**.

The international electronic industry company Elcoteq has laid off tens of employees from its Tallinn facility ; contrary to earlier layoffs the company contracted the ranks of its office workers , the daily Postimees reported. Is the sentiment of the sentence 'Negative', 'Neutral', or 'Positive'? **Negative**.

With the new production plant the company would increase its capacity to meet the expected increase in demand and would improve the use of raw materials and therefore increase the production profitability. Is the sentiment of the sentence 'Negative', 'Neutral', or 'Positive'? **Positive**.

[Test instance.] Is the sentiment of the sentence 'Negative', 'Neutral', or 'Positive'?

Mathdataset Answer Generation (Generation task)

Given a simple high-school level math question, you are required to solve it and provide the final answer. The final answer is always a single number. These questions can range from a variety of topics like simple arithmetic, solving equations, converting a quantity from one unit to another, finding remainders/GCD/LCM, finding probabilities etc. Each question has only one correct answer. This answer can be a positive or negative integer, a fraction or a decimal number. If the answer is a negative number use the hyphen (e.g. -42) symbol for the minus sign. For decimal numbers, do not add extra zeros after the decimal point. For fractional numbers, separate the numerator and denominator using a forward slash (e.g. 3/25).

Let $y = -74 - -79$. Solve $0 = -y * q - 13 + 3$ for q . The answer to this math problem is **-2**.

Work out $29.8 + -0.18$. The answer to this math problem is **29.62**.

How many nanometers are there in 610.1077 millimeters The answer to this math problem is **610107700**.

Four letters picked without replacement from bboobleoeewobw. What is prob of picking 3 o and 1 e? The answer to this math problem is **12/1001**.

[Test instance.] The answer to this math problem is

AGNews (Classification task)

In this task, you are given a news article. Your task is to classify the article to one out of the four topics 'World', 'Sports', 'Business', 'Sci/Tech'. If you are not sure about the topic, choose the closest option. Note that URLs in the text have been replaced with [Link].

Comets, Asteroids and Planets around a Nearby Star (SPACE.com) SPACE.com - A nearby star thought to harbor comets and asteroids now appears to be home to planets, too. The presumed worlds are smaller than Jupiter and could be as tiny as Pluto, new observations suggest. **What label best describes this news article?** Sci/Tech.

Oil and Economy Cloud Stocks' Outlook NEW YORK (Reuters) - Soaring crude prices plus worries about the economy and the outlook for earnings are expected to hang over the stock market next week during the depth of the summer doldrums. **What label best describes this news article?** Business.

Russian FM meets with Katsav Russian Foreign Minister Sergey Lavrov met Monday with Israeli 39;s President Moshe Katsav as part of his first tour of the region to discuss, among other topics, a collaboration between the two countries in combating terrorism. **What label best describes this news article?** World.

Murtagh a stickler for success North-eastern field hockey coach Cheryl Murtagh doesn't want the glare of the spotlight that shines on her to detract from a team that has been the America East champion for the past three years and has been to the NCAA tournament 13 times. **What label best describes this news article?** Sports.

[Test instance.] **What label best describes this news article?**

COPA (Classification task)

In this task you are given two statements. You must judge whether the second sentence is the cause or effect of the first sentence. The two sentences are separated by a newline character and the answer can be 'Cause' or 'Effect'.

The women met for coffee. They wanted to catch up with each other. **Is the second sentence cause or effect of the first sentence?** Cause.

The physician misdiagnosed the patient. The patient filed a malpractice lawsuit against the physician. **Is the second sentence cause or effect of the first sentence?** Effect.

The guests of the party hid behind the couch. It was a surprise party. **Is the second sentence cause or effect of the first sentence?** Cause.

My friend was recovering from surgery. I brought her a card and flowers. **Is the second sentence cause or effect of the first sentence?** Effect.

[Test instance.] **Is the second sentence cause or effect of the first sentence?**

TriviaQA (Generation task)

You are given a general knowledge question based on Wikipedia and Web content. Write an answer to this question.

Who was the man behind The Chipmunks? **The answer to this question is David Seville.**

What star sign is Jamie Lee Curtis? **The answer to this question is Scorpio.**

Which Lloyd Webber musical premiered in the US on 10th December 1993? **The answer to this question is Sunset Boulevard.**

The Euro is divided into how many cents? **The answer to this question is 100.**

[Test instance.] **The answer to this question is**

B. More results

Table 4 shows results for classification and generation tasks whereas Tables 5 and 6 show results for individual tasks for GPT-2-xl (smaller model) and Llama-2-70B (bigger model) respectively.

C. Attention plots

Figure 14 shows repeated text corruptions for GPT-J-6B. Figure 15 shows repeated text corruptions for OPT-30B with random word instructions.

Corruptions	GPT2-xl Classification	GPT2-xl Generation	LLama-2-70B Classification	LLama-2-70B Generation
Semantic Corruptions				
+ demos.	28.1	9.8	53.8	30.7
+ task instr. + demos.	30.4	11.2	61.6	33.0
+ inline instr. + demos.	38.3	9.3	69.0	32.5
Baseline	<u>40.7</u>	15.5	<u>72.3</u>	<u>33.4</u>
Semantic Corruption				
Rw both instr.	40.4	11.6	52.1	30.8
Rw labels	3.7	6.3	1.1	1.7
OOD inputs	41.7	13.3	57.6	22.1
Repeated text				
Inline instr. in 3 demos	45.3	15.1	72.4	33.5
Inline instr. in 2 demos	41.4	14.7	70.4	31.9
Inline instr. in 1 demos	41.4	15.3	69.6	31.7
Inline instr. in 0 demos	17.6	9.1	38.1	22.9
Rw Inline instr. in 3 demos	41.1	10.2	53.5	27.0
Rw Inline instr. in 2 demos	41.6	9.1	54.0	29.8
Rw Inline instr. in 1 demos	40.2	3.0	41.5	22.5
Rw Inline instr. in 0 demos	0.9	0.0	0.9	0.2

Table 4: Model performance for classification and generation tasks. The highest performance is in bold, baseline prompt performance is underlined.

Corruption	RTE	MQP	FPH	TE	CoLA	AGN	COPA	C2S	TQ	MATH
Structural										
+ demos.	46.0	46.0	19.0	6.0	49.0	33.0	20.0	42.0	18.6	0.9
+ task instr.	45.0	47.0	25.0	19.0	48.0	39.0	22.0	37.0	20.2	2.1
+ inline instr.	50.0	50.0	35.0	26.0	50.0	61.0	54.0	50.0	15.1	3.6
Baseline	53.0	60.0	34.0	23.0	51.0	58.0	50.0	47.0	17.0	14.0
Semantics										
Rw both instr.	53.0	50.0	58.0	25.0	44.0	54.0	50.0	47.0	15.6	7.7
Rw labels	0.0	0.0	0.0	23.0	0.0	1.0	0.0	0.0	11.9	0.7
OOD inputs	51.0	52.0	46.0	28.0	50.0	63.0	52.0	48.0	12.7	13.9
Repeated text										
Inline instr. in 3 demos	52.0	63.0	50.0	22.0	50.0	41.0	50.0	51.0	18.0	12.2
Inline instr. in 2 demos	54.0	70.0	35.0	21.0	50.0	32.0	50.0	50.0	14.9	14.6
Inline instr. in 1 demos	50.0	50.0	33.0	31.0	50.0	38.0	50.0	50.0	15.1	15.5
Inline instr. in 0 demos	1.0	43.0	0.0	1.0	47.0	16.0	0.0	0.0	9.1	9.3
Rw Inline instr. in 3 demos	50.0	50.0	40.0	20.0	42.0	37.0	50.0	49.0	15.4	5.0
Rw Inline instr. in 2 demos	50.0	50.0	33.0	20.0	50.0	35.0	53.0	50.0	11.5	6.8
Rw Inline instr. in 1 demos	50.0	50.0	33.0	16.0	50.0	25.0	50.0	50.0	0.0	6.0
Rw Inline instr. in 0 demos	0.0	0.0	0.0	0.0	4.0	0.0	0.0	2.0	0.0	0.0

Table 5: **Model performance for each dataset for GPT2-xl.** Datasets are RTE, Medical Question Pair (MQP), Financial Phrasebank (FPH), Twitter Emotion classification(TE), CoLA, AgNews (AGN), COPA, Com2sense (C2S), and two generation tasks: TriviaQA (TQ) and Mathdataset answer generation(MATH)

Corruption	RTE	MQP	FPH	TE	CoLA	AGN	COPA	C2S	TQ	MATH
Structural										
+ demos.	61.0	60.0	60.0	22.0	49.0	58.0	67.0	53.0	41.6	19.8
+ task instr.	67.0	66.0	69.0	22.0	72.0	61.0	64.0	72.0	41.8	24.2
+ inline instr.	70.0	67.0	78.0	42.0	72.0	83.0	75.0	65.0	42.4	22.2
Baseline	84.0	77.0	80.0	34.0	81.0	86.0	64.0	72.0	42.5	24.3
Semantics										
Rw both instr.	66.0	53.0	34.0	27.0	54.0	74.0	56.0	56.0	42.9	18.7
Rw labels	3.0	0.0	0.0	1.0	0.0	0.0	0.0	5.0	1.7	1.7
OOD inputs	70.0	61.0	71.0	20.0	77.0	35.0	54.0	73.0	35.0	9.2
Repeated text										
Inline instr. in 3 demos	76.0	82.0	84.0	44.0	74.0	85.0	63.0	71.0	45.8	21.2
Inline instr. in 2 demos	72.0	80.0	78.0	41.0	72.0	86.0	65.0	69.0	40.4	23.3
Inline instr. in 1 demos	70.0	81.0	82.0	37.0	76.0	84.0	60.0	67.0	43.2	20.3
Inline instr. in 0 demos	43.0	52.0	34.0	24.0	71.0	1.0	25.0	55.0	25.3	20.5
Rw Inline instr. in 3 demos	63.0	55.0	40.0	35.0	49.0	79.0	58.0	49.0	39.5	14.5
Rw Inline instr. in 2 demos	58.0	70.0	33.0	34.0	50.0	77.0	59.0	51.0	39.9	19.7
Rw Inline instr. in 1 demos	50.0	58.0	34.0	18.0	48.0	25.0	50.0	49.0	39.1	6.0
Rw Inline instr. in 0 demos	5.0	0.0	0.0	0.0	1.0	1.0	0.0	0.0	0.4	0.0

Table 6: **Model performance for each dataset for LLama-2-70B.** Datasets are RTE, Medical Question Pair (MQP), Financial Phrasebank (FPH), Twitter Emotion classification(TE), CoLA, AgNews (AGN), COPA, Com2sense (C2S), and two generation tasks: TriviaQA (TQ) and Mathdataset answer generation(MATH)




Figure 14: Average GPT-J-6B attention per component for repeated text corruptions. “Inline” refers to the presence of the number of inline instructions in the baseline prompt. Fully black box represents missing components.




Figure 15: Random words instructions: Average OPT-30B attention per component for repeated text corruptions. “Inline” refers to the presence of the number of inline instructions in the baseline prompt. A solid black box represents omitted components.