

Customizing Language Model Responses with Contrastive In-Context Learning

Xiang Gao, Kamalika Das

Intuit AI Research, 2700 Coast Avenue, Mountain View, CA 94043
{xiang_gao, kamalika_das}@intuit.com

Abstract

Large language models (LLMs) are becoming increasingly important for machine learning applications. However, it can be challenging to align LLMs with our intent, particularly when we want to generate content that is preferable over others or when we want the LLM to respond in a certain style or tone that is hard to describe. To address this challenge, we propose an approach that uses contrastive examples to better describe our intent. This involves providing positive examples that illustrate the true intent, along with negative examples that show what characteristics we want LLMs to avoid. The negative examples can be retrieved from labeled data, written by a human, or generated by the LLM itself. Before generating an answer, we ask the model to analyze the examples to teach itself what to avoid. This reasoning step provides the model with the appropriate articulation of the user’s need and guides it towards generating a better answer. We tested our approach on both synthesized and real-world datasets, including StackExchange and Reddit, and found that it significantly improves performance compared to standard few-shot prompting.

Introduction

In recent years, large language models like GPT, Llama, and PaLM series have made significant progress in natural language processing (Bommasani et al. 2021; Brown et al. 2020; Touvron et al. 2023), enabling them to generate coherent and contextually relevant responses. However, despite their impressive capabilities, these models can still struggle to align with our intent, particularly when it comes to generating content that is preferable over others or when we want the LLM to respond in a certain style or tone that is hard to describe. As a result, there has been a growing interest in developing techniques that can help us better steer the output of LLMs towards our desired goals. In this work, we propose an approach that leverages contrastive examples to better describe our intent, which can significantly improve the performance of LLMs in generating desirable responses.

In light of the challenges faced by current LLMs in aligning with user intent, it is crucial to develop novel techniques that can effectively guide these models towards generating more desirable responses. Prior research has demonstrated

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

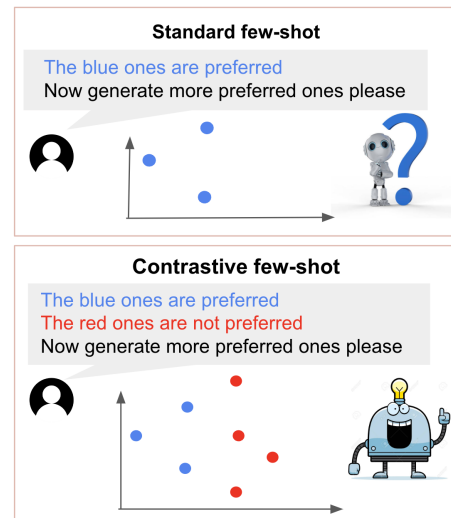


Figure 1: Contrastive examples provide positive examples that illustrate the true intent, along with negative examples that show what characteristics we want LLMs to avoid.

the benefits of few-shot learning (Brown et al. 2020), fine-tuning with smaller models (Gao, Fisch, and Chen 2020), selective annotation frameworks (Su et al. 2022), and visual language modeling (Alayrac et al. 2022) for enhancing LLM performance. However, these approaches do not explicitly address the challenge of guiding LLMs to generate content that adheres to specific preferences, styles, or tones. Additionally, although contrastive learning techniques have shown promise in areas such as image representation (Radford et al. 2021), dialogue response ranking (Gao et al. 2020), and self-supervised learning (Meng et al. 2021), their application to content generation in LLMs remains underexplored. Furthermore, while recent work on prompt optimization (Honovich et al. 2022; Zhou et al. 2022; Sun et al. 2023) has highlighted the importance of effective prompts in steering LLMs, there remains a need for more robust methods that can better capture user intent through diverse and contrastive examples. In this paper, we propose a novel approach that addresses these gaps by leveraging contrastive examples, including positive and negative instances, to more

accurately define user intent and guide LLMs in generating responses that are better aligned with desired outcomes. By incorporating this contrastive reasoning step, our method aims to overcome the limitations of existing techniques and substantially enhance the performance of LLMs in generating preferable content.

Our proposed approach involves providing the LLM with both positive and negative examples to better understand our intent (Figure 1). The positive examples showcase the desired outcomes, while the negative examples highlight what characteristics the LLM should avoid. By analyzing both types of examples before generating an answer, the model can reason about our intent and make a more informed decision about what to generate. It may be challenging for LLMs to understand negative instruction, but we observed that using negative examples helps aligning our preference with LLMs. Moreover, the negative examples can be retrieved from labeled data, written by a human or generated by the model itself, making it a flexible and scalable technique. Our experiments show that this approach can significantly improve the performance of LLMs in generating desirable responses, making them more useful for a wide range of natural language processing applications.

In summary, our proposed approach of using contrastive examples to better describe our intent significantly improves the performance of LLMs in generating desirable responses. This approach can help address the challenge of aligning LLMs with our intended goals and make them more useful for a wide range of natural language processing applications. The key contributions of this work are:

- i) Providing a novel approach that leverages contrastive examples to improve the performance of LLMs in generating desirable responses.
- ii) Demonstrating the effectiveness of this approach on both synthesized and real-world datasets, including StackExchange and Reddit.
- iii) Highlighting the potential of previously discarded negative examples to make LLMs more useful for a wide range of applications by better aligning them with our intended goals.

Related Work

Recent advancements in natural language processing have led to the development of large language models (LLMs) that are capable of few-shot learning, where they can learn new tasks with only a small number of annotated examples. Significant works in this area include Brown et al. (2020), who demonstrated the impressive few-shot performance of GPT-3, Gao, Fisch, and Chen (2020), who proposed LM-BFF for fine-tuning smaller language models, Su et al. (2022), who presented a selective annotation framework, and Alayrac et al. (2022), who introduced Flamingo, a family of visual language models. These works highlight the importance of few-shot learning in LLMs and the need for efficient annotation and prompt generation techniques.

The use of contrastive learning for LLMs has also gained attention, with Radford et al. (2021) proposing Contrastive Language-Image Pre-training (CLIP) for learning image

representations, Gao et al. (2020) leveraging social media feedback data for training a dialogue response ranking model, and Meng et al. (2021) presenting COCO-LM, a self-supervised learning framework that pretrains LLMs by correcting and contrasting corrupted text sequences. These works emphasize the importance of providing diverse and contrastive examples to LLMs and suggest that contrastive learning is a promising direction for LLM research.

Additionally, recent works have focused on improving the performance of LLMs by optimizing the prompts used to steer them towards a desired task or outcome. Honovich et al. (2022) introduce the instruction induction challenge, Zhou et al. (2022) propose Automatic Prompt Engineer (APE) for instruction generation and selection, and Sun et al. (2023) present AutoHint, a framework for automatic prompt engineering and optimization. To apply these approaches for preference alignment, however, is challenging when the characteristics are hard to describe in instruction or measured by automated metrics.

Our work builds upon these previous efforts by proposing an approach that uses contrastive examples to better describe our intent, which we tested on both synthesized and real-world datasets. This approach significantly improves performance compared to standard few-shot prompting, emphasizing the potential for LLMs to become more human-like in their ability to generate natural language instructions.

Contrastive In-Context Learning

Our method comprises two main components: (1) obtaining paired positive and negative examples and (2) forming the prompt. See Figure 2 for an illustration.

Obtaining Paired Contrastive Examples

There are several ways to obtain contrastive examples.

Using Labeled Feedback In some tasks, there exist multiple natural outputs for a single input, and feedback for these outputs is available. For example, Reddit and StackExchange posts typically receive multiple responses, and users provide feedback through upvotes and downvotes. Similarly, in business applications like email and copywriting, multiple versions might be generated given the same constraints, and feedback (e.g., click-through rate) can be obtained through A/B testing or other experiments. For these tasks, we can use the response with the highest feedback as the “positive” example and the one with the lowest feedback as the “negative” example. It is important to note that a “negative” example here does not necessarily mean it is incorrect or unacceptable, but rather less preferred given the specific audience and scenario. It is also important to note, that the highest versus lowest votes is an indication of popularity which may or may not represent personal preference, but we are using this popularity (or lack thereof) signal as a surrogate for general preference of a particular response over others, ignoring confounding factors such as time of response etc.

Using LLM-Generated Responses In cases where labeled feedback is unavailable or negative examples do not capture the characteristics we want LLMs to avoid, we propose a second method. We let the target LLM generate a

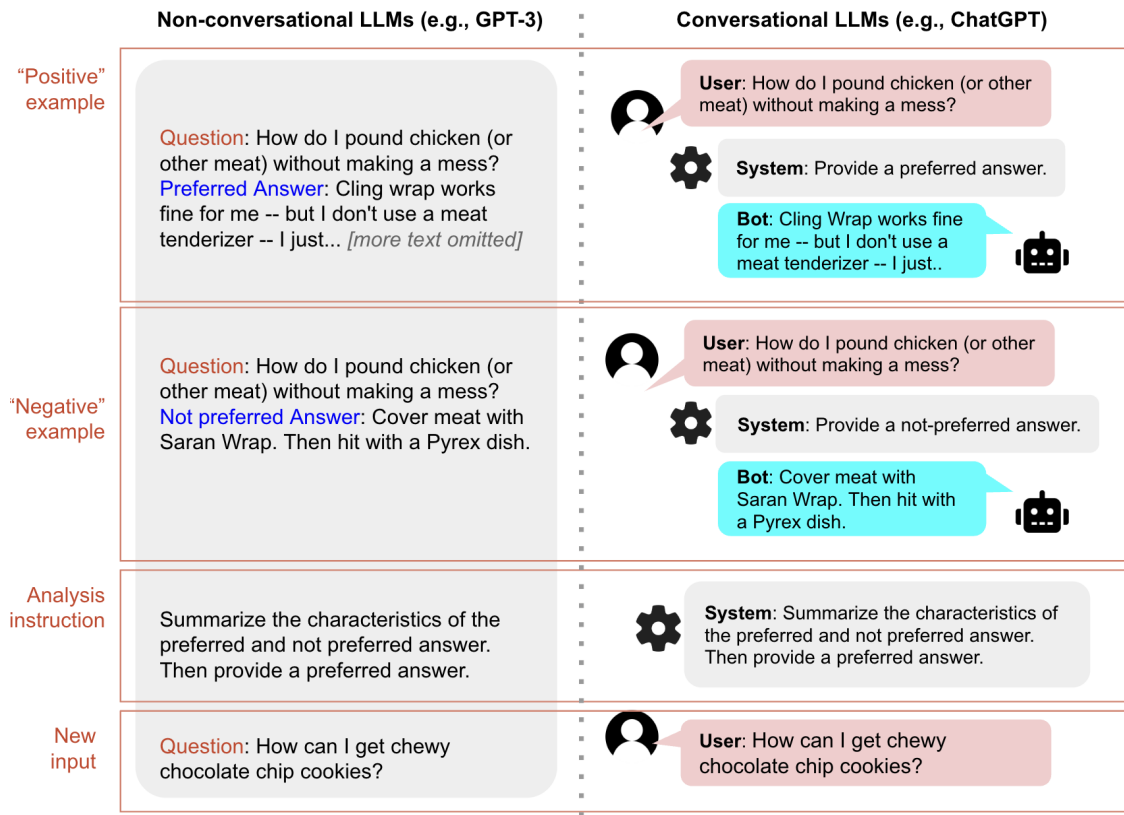


Figure 2: Contrastive in-context learning prompting utilize contrastive few-shot examples, which consists of a “positive” example and a “negative” example for each demonstration input, and a prompt to elicit analysis of the characteristics of the positive/negative examples, before generating an output for the new input. This strategy can be applied to both conversational and non-conversational LLMs. For non-conversational LLMs, the labels and instructions are provided as system messages (the gear symbols).

response and use it as the negative example. LLMs generated responses often appears mechanical, lacking emotion, and impersonal. Pairing this with a highly preferred labeled response as the positive example guides the LLM to generate a response more aligned with the positive example.

Using Automated Evaluator For some tasks, it is possible to define automated evaluation rules or apply classifiers to measure the quality. The evaluation score can be used to select “positive” and “negative” examples. We demonstrate this with a words-constrained generation task, with accuracy determined by whether all given words were included (scored as 1) or not (scored as 0) in the generated sentences.

Forming the Prompt

Once we have obtained paired positive and negative examples, we consider two prompting strategies:

Contrastive Examples as Few-Shot Examples In this strategy, we provide the contrastive example pairs as few-shot examples. Labels are included in the prompt, such as “preferred answer” for the positive example and “less preferred answer” for the negative example, to indicate the role

of each part in the prompt. The LLM is then asked to generate a “preferred answer” for the new input.

Reasoning and Analysis In addition to the first strategy, we ask the LLM to analyze the reasons for preference and the characteristics of the examples before generating a “preferred answer” for the new input. This reasoning step is inspired by the Chain-of-Thought prompting (Wei et al. 2022) and is designed to summarize the characteristics from the contrastive examples, allowing the LLM to automatically generate instructions for itself. Coupled with contrastive examples, this guidance helps the LLM to better align with the preferred intent.

Experiments

In this section, we describe the experiments conducted to evaluate the effectiveness of our proposed approach using contrastive examples. We focus on text generation tasks where user preferences play a significant role, rather than tasks with objective right or wrong answers. We first outline the datasets used in our experiments, which include both real-world and synthetic datasets.

Datasets

We consider the following datasets to investigate the impact of user preferences on text generation tasks:

Human Preference Datasets

User preference is expressed via different forms of feedback publicly available at a large scale on many platforms.

StackExchange People post questions on StackExchange, and other users provide answers, with upvotes and downvotes used to rate the responses. Our focus is on subjective preferences, so we created a dataset using data from `cooking.stackexchange.com`, which includes how-to type cooking-related questions. Cooking often does not have right or wrong answers, and different people have different preferences or ways to make a dish they consider good.

Reddit Users post questions or topics on Reddit to trigger discussions, and other users can comment (comments can have nested comments). In our study, we only considered the comments of the posts as they share a similar context. Redditors can provide upvotes or downvotes as feedback for each post and comment. We experimented a subreddit called ‘NoStupidQuestions’. ‘NoStupidQuestions’ encourages people to ask any question without fear of being judged by conventional social norms. The subjects cover a wide range of topics, including objective subjects and social or personal experiences. These characteristics make this subreddit ideal for our study, as personal (or group collective) preference is an essential factor, rather than objective right or wrong.

For both datasets, We filtered the posts to include those with multiple answers (to have answers with varying “quality” or level of preference). For the remaining posts, we used the top-rated answer as the “positive” example, and bottom-rated answer as the “negative” example. We randomly selected 500 samples for evaluation.

Synthetic Stylistic Datasets

Human preferences play a crucial role in the aforementioned real datasets. However, it is challenging to interpret and analyze the preferences, as it is often difficult even for humans to articulate what exactly makes an answer highly voted. Therefore, we also considered a few simplified synthetic datasets. Using the `cooking.stackexchange.com` dataset, we leveraged ChatGPT to generate the following datasets: i) Funny vs. Serious, ii) Concise vs. Detailed, and iii) British vs. American

We then used the generated responses as contrastive few-shot examples. It is essential to note that by “positive” and “negative”, we do not mean right or wrong, and the choice of “positive” is arbitrary in this case for synthesized datasets. The synthetic datasets aim to illustrate the ability of the proposed method to guide LLMs to generate output in a given direction.

Constrained Generation Dataset

The datasets mentioned above focus on controlling

implicit stylistic aspects of text generation, whether matching a specific attribute such as level of conciseness (as in the synthetic datasets) or an overall holistic style (as in the human preference datasets). In addition to these datasets, we further experiment with a dataset focused on the control of lexical content, constrained generation. This also allows us to test the ability of our method to generalize to tasks with explicit constraints.

Following Zhou et al. (2023), the language model is prompted to craft sentences using specific seed words as constraints. We randomly selected 500 questions from the StackExchange dataset and randomly chose 5 words from each question to serve as lexical constraints. We use the wrong results generated by the LLM as the negative examples.

The model’s performance is measured by the success rate, defined as the percentage of sentences generated that contain all 5 given seed words. Success requires an exact uncased match between the generated sentence and constraint words.

Prompt Settings

In our experiments, we consider two types of large language models (LLMs): non-conversational LLMs, such as GPT-3, and conversational LLMs, including ChatGPT (GPT-3.5-turbo) and GPT-4. We evaluate these LLMs under two different settings: zero-shot and few-shot.

Zero-Shot For non-conversational LLMs, we use a prompt consisting of the post with the prefix “Question:” followed by the second line “Answer:”. For conversational LLMs, we use a system prompt stating, “You are a good StackExchange/Reddit user,” followed by the user prompt containing the post. In both cases, we input the post title. We exclude the post body, as it often includes edits after the original poster has read the replies, potentially leaking preference information.

Few-Shot For each query, we randomly select k labeled examples as few-shot examples¹. For non-conversational LLMs, we include the obtained few-shot examples in the prompt, using the “Question:” and “Answer:” prefixes. For conversational LLMs, we use the few-shot examples to create a conversation history between the user (post) and the bot (top-rated reply). We experimented with $k = 1$ to $k = 4$ and stops when performance does not increase significantly as k increases.

Contrastive Examples Based Prompts

For the few-shot setting, we compare the standard approach, which only includes positive examples, with three settings involving contrastive examples:

Contrastive Examples Only For each few-shot example, we provide a positive and a negative example. For non-conversational LLMs, we use the prefixes “top-rated” and

¹We experimented with retrieving few-shot examples based on query embedding similarity but didn’t observe significant performance difference compared to randomly selected few-shot examples

“low-rated” to indicate positive and negative examples, respectively. For conversational LLMs, we use system messages (“provide a top-rated answer” and “provide a low-rated answer”) to inform the LLMs of the preference.

Contrastive Instruction Only Given the contrastive examples (with labels defined in method 1), we ask LLMs to automatically generate an instruction based on the preference revealed by the positive and negative examples. A sample prompt could be, “Summarize the characteristics of the preferred and not preferred answer.” We then use the generated analysis, followed by “generate a top-rated answer for input,” as the instruction for the test input. The contrastive examples are only used to generate the analysis and are not included in the final prompt.

Contrastive Examples + Instruction This method combines the previous two approaches. We first provide the contrastive examples, followed by the instruction asking LLMs to perform an analysis. Then, the LLM generates a new response. When generating the final answer, both the contrastive examples and the generated analysis are available.

These three settings involving contrastive examples allow us to conduct an ablation study to examine the contributions of specific contrastive examples and more general instructions.

Evaluation Method

In order to evaluate the performance of our approach, we consider two distinct methods: reference-based and reference-free evaluation.

Reference-Based Evaluation The reference-based evaluation method measures the similarity between the generated responses and the top-rated reference answer. We employ two metrics, **BERT Score** (Zhang et al. 2019) and **Emb. Similarity**. The latter is the cosine similarity of the sentence embeddings obtained using the Sentence-BERT model (Reimers and Gurevych 2019).

Reference-Free Evaluation An answer that deviates from the reference could still be preferred by many readers. Therefore, we also consider two reference-free evaluation methods. i) We leverage **DialogRPT** (Gao et al. 2020), a pre-trained dialog response ranking model. This model is trained on Reddit data and shows a high correlation with human upvotes (Gao et al. 2020), so we apply it to the Reddit dataset only. ii) In addition, we employ LLM as an evaluator following Liu et al. (2023). We compute a **GPT Score** by prompting GPT4 to score the generated results using the positive and negative examples as in-context few shot examples. The LLM generated score and human evaluation labels exhibit positive correlation (StackExchange: 0.65, Reddit: 0.58).

Results

We evaluated the performance of the baselines and contrastive in-context prompting strategies on the synthetic and two real datasets. The results are shown in Table 1 and Table 2. The “contrastive-combined” approach (which combines the contrastive examples and the derived instruction)




Figure 3: The dependence of the performance on prompt length. For standard few-shot, we experimented with $k = 1, 2, 3, 4$, shown by the four black dots from left to right in each subfigure. For contrastive prompts, empty circles denotes prompts using generated answers as negative examples, and filled symbols for prompts using low-rated human answers as negative examples.

achieved the best performance for most cases. Using similar number of tokens in prompt, this contrastive in-context learning approach performs significantly better than standard few-shot approach, as shown in Figure 3. For Reddit dataset, we observe the most obvious improvement with ChatGPT model. Standard two-shot prompt only performs better than zero-shot approach for about 64 % test cases, while the “contrastive-combined” approach wins 76 % test cases. For StackExchange, standard few-shot approaches do not show obvious improvement compared to zero-shot, while “contrastive-combined” improve BERT Score by approximately 0.01. For synthetic datasets, the improvement of is more obvious, with BERT Score increased by 0.02 to 0.03, and much higher than the improvement achieved with standard few-shot prompts. As an ablation study, we also evaluated the results for “contrastive - examples only” and “contrastive - instruction only”. Both approaches improve the performance compared to zero shot, and the “contrastive - examples only” performs better than “contrastive - instruction only” for many cases. Combining them together further improves the results, as discussed above.

The Impact of Contrastive Examples As mentioned above, using contrastive in-context learning by introducing negative examples improve the performance compared to standard few-shot approaches. We further investigate different ways to obtain negative examples. The first method, “human”, uses human-written answers, where low-rated replies serve as “negative” examples. The second method, “generated”, uses zero-shot generated replies as “negative” examples. Interestingly, we observed that the second method performs on par with, and sometimes even better than, the first method (see Tables 1 and 2). This finding demonstrates that the proposed contrastive method is not limited to cases

Dataset Method	Funny vs. Serious				Concise vs. Detailed				British vs. American			
	Emb. Similarity		BERT Score		Emb. Similarity		BERT Score		Emb. Similarity		BERT Score	
GPT-3												
Zero-shot	0.606		0.858		0.849		0.903		0.785		0.883	
Few-shot, $k=1$	0.611		0.863		0.887		0.916		0.813		0.892	
Few-shot, $k=2$	0.629		0.864		0.886		0.920		0.808		0.887	
Contrastive, $k=1$	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.
- Examples only	0.614	0.620	0.857	0.859	0.817	0.865	0.895	0.910	0.827	0.807	0.891	0.885
- Instruction only	0.645	0.595	0.865	0.861	0.834	0.845	0.895	0.902	0.824	0.784	0.888	0.879
- Combined	0.655	0.617	0.864	0.863	0.869	0.897	0.910	0.922	0.826	0.838	0.896	0.891
ChatGPT												
Zero-shot	0.604		0.841		0.826		0.878		0.783		0.867	
Few-shot, $k=1$	0.613		0.842		0.829		0.880		0.807		0.870	
Few-shot, $k=2$	0.610		0.846		0.834		0.884		0.806		0.874	
Contrastive, $k=1$	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.
- Examples only	0.608	0.599	0.847	0.844	0.850	0.856	0.881	0.891	0.825	0.818	0.882	0.880
- Instruction only	0.617	0.617	0.849	0.850	0.831	0.835	0.882	0.888	0.819	0.821	0.876	0.876
- Combined	0.638	0.630	0.862	0.862	0.873	0.883	0.914	0.915	0.830	0.827	0.891	0.891
GPT-4												
Zero-shot	0.596		0.838		0.803		0.873		0.786		0.864	
Few-shot, $k=1$	0.607		0.843		0.805		0.877		0.792		0.867	
Few-shot, $k=2$	0.624		0.846		0.809		0.880		0.807		0.868	
Contrastive, $k=1$	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.	human	gen.
- Examples only	0.609	0.611	0.844	0.843	0.828	0.855	0.888	0.893	0.817	0.821	0.870	0.871
- Instruction only	0.624	0.612	0.843	0.841	0.853	0.852	0.893	0.895	0.821	0.820	0.868	0.870
- Combined	0.616	0.630	0.855	0.857	0.875	0.853	0.919	0.917	0.813	0.824	0.886	0.884

Table 1: Results on the synthetic dataset. k is the number of demonstration inputs. The “negative” examples are obtained via two methods, “human” and “gen.” (generated by the LLM itself).

where developers have to provide human-written pairs of positive and negative examples. Instead, it requires the same inputs as standard few-shot settings, which only need a few positive examples since we can generate the negative examples.

This observation also supports our assumption that negative examples obtained from human-written data may not capture all characteristics we want LLMs to avoid. For instance, in the Reddit dataset, a human-written reply may receive downvotes because the content is rude, disrespectful, or violates some rules of the subreddit, such as self-promotion. However, recent LLMs like ChatGPT are generally trained to follow social norms and sometimes even apologize too frequently. As a result, it is highly unlikely for these LLMs to generate offensive language. Providing human-written negative examples, therefore, may not offer much helpful signal in such cases. In contrast, providing generated responses as negative examples can supply the signals of certain characteristics we want the LLM to avoid. For example, although LLM-generated responses are typically fluent and relevant to the question, they often lack emotion, details, examples, or elements that trigger readers’ personal feelings. These characteristics, however, are essential factors that make users prefer certain answers. By providing zero-shot generated responses, we can guide LLMs to move away from the machine-generated style and toward a more human-preferred style.

Distilling Contrastive Examples as an Instruction Although the primary focus of this work is not to automat-

ically generate instructions for LLMs, we were interested in understanding what instructions LLMs could derive from the contrastive examples. Therefore, we asked the LLMs to summarize the characteristics of the provided “positive” and “negative” examples. We then used this analysis in the prompt, instead of the actual contrastive examples, to offer the LLMs insight into user preferences. This compress the contrastive examples into a shorter instruction, and reduce prompt length and cost.

Our experimental results showed that, with this analysis, LLMs performed better than in the zero-shot setting (see Tables 1 and 2). However, the improvement was not significantly better than the standard few-shot setting. This could be partially explained by the fact that the preference was not consistent across the examples chosen based on upvotes. There is room for enhancement in the way the preference is summarized through the use of automated prompt generation approaches (Sun et al. 2023). By refining the way LLMs derive instructions from contrastive examples, we can potentially achieve more significant improvements in performance compared to standard few-shot prompting.

Combining the Complementary Parts We observed that when combining the contrastive examples and the instruction (analysis of the characteristics of “positive” and “negative” examples automatically generated based on the contrastive examples) together, an even better performance was achieved (see Tables 1 and 2). This improved performance can be attributed to the fact that the analysis and the contrastive examples provide complementary information to

Dataset Method	StackExchange				Reddit				Constrained Gen. Success rate
	BERT Score		GPT Score		DialogRPT score		GPT Score		
GPT-3									
Zero-shot	0.840		0.550		0.605		0.533		0.750
Few-shot, $k=1$	0.841		0.555		0.596		0.531		0.760
Few-shot, $k=2$	0.841		0.561		0.610		0.540		0.772
Contrastive, $k=1$	human	generated	human	generated	human	generated	human	generated	0.801
Examples only	0.846	0.845	0.605	0.600	0.630	0.632	0.634	0.638	
Instruction only	0.845	0.846	0.605	0.601	0.606	0.610	0.621	0.630	
Combined	0.847	0.846	0.608	0.606	0.631	0.632	0.635	0.640	
p-value	(0.047)	(0.049)	(0.018)	(0.023)	(0.030)	(0.030)	(0.020)	(0.019)	
ChatGPT									
Zero-shot	0.839		0.550		0.602		0.531		0.780
Few-shot, $k=1$	0.839		0.561		0.636		0.578		0.800
Few-shot, $k=2$	0.839		0.570		0.645		0.583		0.810
Contrastive, $k=1$	human	generated	human	generated	human	generated d	human	generated	0.866
Examples only	0.844	0.842	0.592	0.589	0.654	0.663	0.654	0.659	
Instruction only	0.844	0.843	0.598	0.597	0.640	0.645	0.640	0.642	
Combined	0.845	0.846	0.601	0.603	0.656	0.663	0.690	0.701	
(p-value)	(0.048)	(0.044)	(0.033)	(0.031)	(0.048)	(0.041)	(0.023)	(0.021)	
GPT-4									
Zero-shot	0.839		0.575		0.657		0.580		0.850
Few-shot, $k=1$	0.840		0.578		0.655		0.585		0.870
Few-shot, $k=2$	0.841		0.582		0.656		0.590		0.900
Contrastive, $k=1$	human	generated	human	generated	human	generated	human	generated	0.923
Examples only	0.844	0.842	0.595	0.590	0.658	0.665	0.660	0.671	
Instruction only	0.842	0.842	0.590	0.591	0.658	0.660	0.661	0.672	
Combined	0.847	0.845	0.610	0.607	0.660	0.665	0.665	0.670	
(p-value)	(0.047)	(0.049)	(0.023)	(0.025)	(0.043)	(0.040)	(0.034)	(0.032)	

Table 2: Results on real-world datasets. k is the number of demonstration inputs. The “negative” examples are obtained via two methods, “human” and “generated”. p-value indicate the statistical significance of the comparison between the “Combined” contrastive prompting and the standard few-shot ($k = 2$) prompting.

guide the LLM about user preferences.

The generated instruction exhibits better generalization but may lack the necessary clarity and specificity. This is because the instructions can be vague or provide only general descriptions, which can be challenging for LLMs to interpret. In contrast, the actual contrastive examples offer more detailed information by providing demonstrations of the desired and undesired characteristics.

In summary, our approach of using contrastive examples along with an analysis of the characteristics of positive and negative examples has proven to be a valuable method for aligning LLMs with user preference. This approach not only enhances performance but also provides LLMs with a better understanding of the user’s preferences, resulting in more accurate and satisfactory responses.

Prompt Token Efficiency In real-world and industrial applications, the number of prompt tokens plays a crucial role as it is directly related to latency and monetary cost. Ideally, a good prompt should be short yet lead to high-quality output. Therefore, we measure the token efficiency to understand the effectiveness of our approach. As shown in Figure 3, the standard few-shot strategies do exhibit an increase in performance as we increase the number of examples. However, the performance, as measured by BERT score, is still lower than the contrastive in-context learning

strategies using a similar number of tokens. Furthermore, the contrastive instruction strategy, which summarizes the contrastive examples as an instruction, utilizes fewer tokens but achieves better performance than few-shot examples. This indicates that, given the same budget of prompt tokens, contrastive in-context learning strategies outperform the standard few-shot approaches, thereby demonstrating higher efficiency.

Conclusion

This research introduced an approach to enhance the alignment of large language models (LLMs) with user preference using contrastive examples. By integrating positive examples showcasing desired outputs and negative ones emphasizing unwanted LLM traits, we tested our methodology on both synthesized and real-world datasets, including StackExchange and Reddit.

The results confirmed the superiority of contrastive examples over the standard few-shot prompting, particularly in terms of performance and prompt token efficiency. Notably, negative examples generated from zero-shot outputs were as effective as those from human-written data. Future endeavors might delve into refining LLM instructions based on these examples and innovating automatic prompt generation techniques.

References

- Alayrac, J.-B.; Donahue, J.; Luc, P.; Miech, A.; Barr, I.; Hasson, Y.; Lenc, K.; Mensch, A.; Millican, K.; Reynolds, M.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35: 23716–23736.
- Bommasani, R.; Hudson, D. A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M. S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Gao, T.; Fisch, A.; and Chen, D. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Gao, X.; Zhang, Y.; Galley, M.; Brockett, C.; and Dolan, B. 2020. Dialogue response ranking training with large-scale human feedback data. *arXiv preprint arXiv:2009.06978*.
- Honovich, O.; Shaham, U.; Bowman, S. R.; and Levy, O. 2022. Instruction induction: From few examples to natural language task descriptions. *arXiv preprint arXiv:2205.10782*.
- Liu, Y.; Iter, D.; Xu, Y.; Wang, S.; Xu, R.; and Zhu, C. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Meng, Y.; Xiong, C.; Bajaj, P.; Bennett, P.; Han, J.; Song, X.; et al. 2021. Coco-lm: Correcting and contrasting text sequences for language model pretraining. *Advances in Neural Information Processing Systems*, 34: 23102–23114.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Reimers, N.; and Gurevych, I. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Su, H.; Kasai, J.; Wu, C. H.; Shi, W.; Wang, T.; Xin, J.; Zhang, R.; Ostendorf, M.; Zettlemoyer, L.; Smith, N. A.; et al. 2022. Selective annotation makes language models better few-shot learners. *arXiv preprint arXiv:2209.01975*.
- Sun, H.; Li, X.; Xu, Y.; Homma, Y.; Cao, Q.; Wu, M.; Jiao, J.; and Charles, D. 2023. AutoHint: Automatic Prompt Optimization with Hint Generation. *arXiv preprint arXiv:2307.07415*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35: 24824–24837.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Zhou, Y.; Muresanu, A. I.; Han, Z.; Paster, K.; Pitis, S.; Chan, H.; and Ba, J. 2022. Large language models are human-level prompt engineers. *arXiv preprint arXiv:2211.01910*.