

Demonstrations Are All You Need: Advancing Offensive Content Paraphrasing using In-Context Learning

Anonymous Submission

Abstract

Paraphrasing of offensive content is a better alternative to content removal and helps improve civility in a communication environment. Supervised paraphrasers; however, rely heavily on large quantities of labelled data to help preserve meaning and intent. They also often retain a large portion of the offensiveness of the original content, which raises questions on their overall usability. In this paper we aim to assist practitioners in developing usable paraphrasers by exploring In-Context Learning (ICL) with large language models (LLMs), i.e., using a limited number of input-label demonstration pairs to guide the model in generating desired outputs for specific queries. Our study focuses on key factors such as – number and order of demonstrations, exclusion of prompt instruction, and reduction in measured toxicity. We perform principled evaluation on three datasets, including our proposed Context-Aware Polite Paraphrase dataset, comprising of dialogue-style rude utterances, polite paraphrases, and additional dialogue context. We evaluate our approach using four closed source and one open source LLM. Our results reveal that ICL is comparable to supervised methods in generation quality, while being qualitatively better by 25% on human evaluation and attaining lower toxicity by 76%. Also, ICL-based paraphrasers only show a slight reduction in performance even with just 10% training data.

1 Introduction

Disclaimer: Figures and examples in this work may feature offensive language.

Timely moderation helps curb the spread of hateful content on social-media platforms and prevents the harmful effects it has on a user’s psychological well-being (Waldron, 2012; Ye et al., 2023). Unfortunately, the sheer volume of content generated on these platforms makes it infeasible to enforce a scalable human moderation process (Hassan et al., 2022; Dosono and Semaan, 2019). AI-based mod-




Figure 1: Influence of number and order of demonstrations, and instruction, on BLEU Score performance and measured Toxicity. Comparison is done between BART and three ICL-based approaches. Numbers on the x -axis represent number of demonstrations used in the ICL framework. Note, measured Toxicity for BART in ParaDetox is 82, exceeding the set y -axis limit.

eration systems can help with this problem. However, current systems often remove or flag offensive content, which can reduce user participation and diversity in online discussions (Xiang et al., 2012; Warner and Hirschberg, 2012; Kwok and Wang, 2013; Wang et al., 2014; Burnap and Williams, 2015; Nobata et al., 2016; Davidson et al., 2017; Founta et al., 2019; Jhaver et al., 2019; Ye et al., 2023). A better alternative is to paraphrase offensive content to make it less offensive. Paraphrasing offensive content; however, is nontrivial since the paraphrased output should not only be inoffensive but also retain the original meaning and intent.

Prior works (Atwell et al., 2022; Logacheva et al., 2022) have proposed using supervised generative models (Vaswani et al., 2017) like BART (Lewis et al., 2019), to paraphrase offensive content. However, these methods require sufficient labelled training data, which makes it harder to adapt them to novel settings. Moreover, these models are

optimized to perform well on certain automated metrics (Papineni et al., 2002; Zhang et al., 2019; Lin, 2004; Vedantam et al., 2015) at the expense of possibly retaining a portion of the original toxicity, thereby making us question their overall usability for the targeted task (see Figure 1).

The emergence of few-shot *In-Context Learning* (ICL) has revolutionized the field by complementing the generalization capabilities of *Large Language Models* (LLMs) to quickly and accurately adapt to new tasks. It does this by using a small amount of labeled data, known as *demonstrations* or *demos* or *examples* (Brown et al., 2020). As shown in Figure 1, ICL shows BLEU score performance that is similar to BART, but significantly reduces the measured toxicity (Hanu and Unitary team, 2020). Through detailed, principled experiments we explore the viability of ICL for paraphrasing offensive content, which to the best of our knowledge has not been done before. Our key contributions and findings are summarized below.

1. Influence of the following factors on generation quality, as summarized in Figure 1.
 - (a) *Number of Demonstrations*: Performance improves by increasing number of demos but eventually saturates.
 - (b) *Selection and Order of Demonstrations*: Systematically selecting and ordering demos is better than its random counterpart. It is more effective to select demos that are semantically similar to the query and curate them in a decreasing/increasing order of similarity.
 - (c) *Exclusion of Prompt Instruction in Prompt*: ICL without the main instruction only slightly affects performance but at the expense of toxicity. Thus we need both demonstrations and instructions to simultaneously preserve performance and lower toxicity.
 - (d) *Robustness to Training Data Size*: Carefully ordering demos shows robustness to available training data size, with only small decrease in generation performance when 10% of training data is only made available.
2. We tested the capabilities of OpenAI’s *text-davinci-003*, *gpt-3.5-turbo-0613*, *gpt-3.5-turbo-instruct*, *gpt-3.5-turbo-1106* models and the open-source *Vicuna-13b* model (Chiang et al., 2023). ICL generated paraphrases are comparable to SOTA supervised approaches in performance, but on average show 76% less toxicity and are 25% better us-

ing a manual qualitative assessment, and thus have superior overall usability. We also show that our demonstration curation approach is simpler and faster than other more sophisticated methods that offer only marginal performance improvement at the expense of significant time delays.

3. Current paraphraser are less effective at mitigating offensiveness like rudeness in conversations. They are trained using datasets that focus on social-media content, and hence aren’t directly applicable to dialogue-based environments. To this end we release a new *Context-Aware Polite Paraphrase (CAPP)* dataset, a dialogue-style corpus of rude utterances and corresponding polite paraphrases, with samples accompanied by additional context in the form of prior two turns from the dialogue. We conduct experiments to show the importance and benefit of incorporating context to improve paraphraser performance.

Paper Outline: Section 2 describes ICL in our experimental setting; details about selecting and ordering the demos; and finally our proposed CAPP dataset in detail. Section 3 contains detailed experimental results. Section 4 discusses related work. Section 5 concludes the paper.

2 Method

2.1 In-Context Learning

Prompts used for ICL contain three parts – (1) an instruction I that defines the task; (2) a set of n demonstrations from the training corpus, $D = (x_i, y_i)_{i=1}^n$, where (x_i, y_i) denotes the offensive, in-offensive sentence pair; and (3) the offensive test query sample x_q . Consider the following prompt example with $n = 2$ demonstrations, where the final sentence represents the query for which we want to generate the paraphrase.

Instruction: Paraphrase the following sentence to be more polite.

Sentence: What’s wrong with you?

Paraphrase: Are you feeling alright?

Sentence: Get out of the way.

Paraphrase: Can you please step aside?

Sentence: What’s the matter with you?

Paraphrase:

The impact of each part on the BLEU score and toxicity is briefly illustrated in Figure 1. For instance, prompts with only instruction show the low-

est BLEU scores, followed by prompts with only demos, while prompts that include both have the best BLEU scores. In terms of Toxicity, prompts with just instruction show the least Toxicity, followed by prompts that include both demos and instruction, while prompts that only include demos exhibit a higher toxicity. The order of demos is also crucial, and we discuss this next.

2.2 Selection and Ordering of Demonstrations

Here we describe our approach to select and order the demonstrations. We first compute normalized vector embeddings for each training sample x_i and the query x_q , denoted as e_i and e_q respectively. Next, the cosine similarity scores between e_q and each e_i are used to select n demonstrations. We explored following two variations for selecting the demonstrations – (1) *Least Similar*, (2) *Most Similar*, i.e., select n demos with the lowest and highest cosine similarity scores respectively. These are compared to randomly selecting n demos, that are arranged in no particular order. We further investigated if arranging the n selected demos in either ascending or descending order based on their measured cosine similarity, had any impact on the overall performance. Using BLEU and toxicity, Figure 1 compares *Random* selection to the *Most Similar (Descending order)* approach, with the latter being better on both fronts. Our findings are described in detail in Section 3.2.

2.3 Context-Aware Polite Paraphrase (CAPP) Dataset

Existing datasets (Atwell et al., 2022; Logacheva et al., 2022) contain comments flagged for toxicity and provide non-toxic paraphrases that maintain the core meaning in a neutral manner. However, they are not directly suitable to address rudeness in speech, as speech is often directed at specific participants, while social media posts have a broader audience, resulting in different styles and tones. Additionally, most social media posts can be remedied by removing explicit insults, but rude speech requires additional modifications to make it more polite. For instance, we should not just eliminate offensive language and direct insults in a rude utterance, but also transform an accusation of ignorance into an inquiry about knowledge.

To address the aforementioned differences, we constructed a dialogue-style rude speech dataset by leveraging the OpenSubtitles corpus (Lison and Tiedemann, 2016). Our approach involved

a two-step process to extract target rude utterances. Initially, we fine-tuned a DistilBERT-base model (Sanh et al., 2019) using the Stanford Politeness corpus (Danescu-Niculescu-Mizil et al., 2013) and a subset of manually labeled OpenSubtitles samples to train a three-class model capable of predicting polite, neutral, or rude sentences. We then use this fine-tuned model to annotate a larger portion of the OpenSubtitles corpus, bootstrapping additional training data for our final rudeness detection model. A separate portion of the OpenSubtitles dataset was selected and labeled as rude, polite, or neutral, resulting in an intermediate set containing rude samples without polite paraphrases. Detailed information about the training/evaluation of the rudeness detector is provided in Appendix A.

We captured context, if available, in the form of prior two turns preceding up to the original utterance. We used the gpt-3.5-turbo-0613 model to generate polite paraphrases, and qualitatively checked the impact of incorporating context on the quality of generated paraphrases. More information about this can be found in Appendix A.

3 Experiments and Discussion

We realized ICL using OpenAI’s text-davinci-003, gpt-3.5-turbo-0613 models and their latest stand-ins, and the open-source Vicuna-13b model. We performed evaluation on the APPDIA (Atwell et al., 2022), ParaDetox (Logacheva et al., 2022), CAPP datasets, with the corresponding (#training, #test) being (1584, 199), (11927, 670), (7939, 1120) respectively. APPDIA contains offensive Reddit comments and their inoffensive paraphrases. The ParaDetox corpus consists of toxic and non-toxic sentence pairs, obtained by filtering the ParaNMT corpus (Wieting and Gimpel, 2017). We used the sentence transformer (*all-mpnet-base-v2*) (Reimers and Gurevych, 2019) to generate the normalized embeddings described in Section 2.2. We evaluated generation quality using automated evaluation metrics such as BLEU (Papineni et al., 2002), BERT-F1 (Zhang et al., 2019), ROUGE (Lin, 2004) and CIDEr (Vedantam et al., 2015). For toxicity we used the implementation by (Hanu and Unitary team, 2020). The exact prompt instruction used in all experiments is provided in Appendix C.

3.1 Number of Demonstrations

Figure 2 shows the plot between number of demonstrations versus BLEU (refer to Appendix

D.1, Figure 8 for other metrics). We set the number of demonstrations to [0, 1, 10, 20, 30, 40] for text-davinci-003, gpt-3.5-turbo-0613, and [0, 1, 2, 4, 6, 8, 10] for Vicuna-13b. We used the proposed *Most Similar (Descending Order)* approach to select and order the demos. We observe that BLEU improves rapidly until 10 demos for the OpenAI models and 4 demos for the Vicuna-13b model across all datasets. Further increasing the demos only results in slight improvement, as each additional demo is semantically less similar to the query and thereby less important than the demonstrations selected before (Liu et al., 2021).

We notice in the case for the gpt-3.5-turbo-0613 model on CAPP dataset that BLEU without any demos is better than with 40 demos. One possibility is that the main instruction used here was less effective in the ICL paradigm, and that a different instruction could have increased the BLEU score, as seen in Section 3.3. However, we believe this happens because the polite paraphrases for CAPP were generated using gpt-3.5-turbo-0613. This hints at the possibility of ICL not necessarily improving paraphrasing performance of LLMs, which in turn were used to generate the dataset. We have similar observations in the following sections as well.

3.2 Selection and Order of Demonstrations

We now discuss the effect of selection and ordering the demos in the prompt on BLEU. We set the number of demonstrations to 10 and explore different ordering mechanisms described in Section 2.2. In Figure 3, we observe that the *Random* strategy sometimes achieves better BLEU than the *Least Similar* strategy. However, in most cases *Most Similar* shows better performance than both *Random* and *Least Similar*. This intuitively makes sense since *Most Similar* represents samples from the training corpus that are most semantically similar to the query (Liu et al., 2021). This enables the LLM to generate a paraphrase that is also similar to the Gold-Standard paraphrase of the query.

Next, the order in which the demos are arranged also has an impact on BLEU score. We find that curating the demos in decreasing order of similarity often results in better BLEU than arranging them in increasing order of similarity. We also observe similar trends with other automated evaluation metrics. Note, the above observations do not apply to the gpt-3.5-turbo-0613 model on the CAPP dataset.

While the approaches in Section 2.2 might be

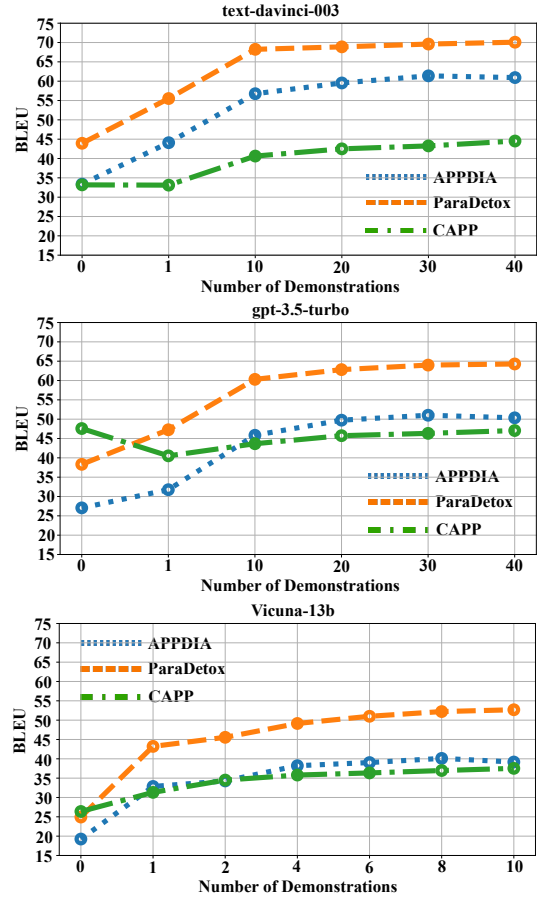


Figure 2: BLEU as a function of number of demos. Noticeable improvement in BLEU is observed in the beginning, with performance saturating after a certain number of demos.

simple and effective, they need not necessarily bring out the best possible performance. Sophisticated methods to select and order demonstrations have been proposed and have shown better performance in other applications (Ye et al., 2022; Zhang et al., 2022; Lu et al., 2021). However, we find that these methods only offer marginal improvement in paraphrasing performance, while taking significantly longer times to process each query test sample. For example, Appendix D.2, Table 5 shows the different performance metrics and compute times taken by (Ye et al., 2022)’s MMR-based approach and compares it to our proposed approach. While the MMR-based approach might be quantitatively better, it is several orders of magnitude slower, which can impact its usability due to potential scalability issues. Compute time limitations for other methods is further discussed in Appendix D.2.

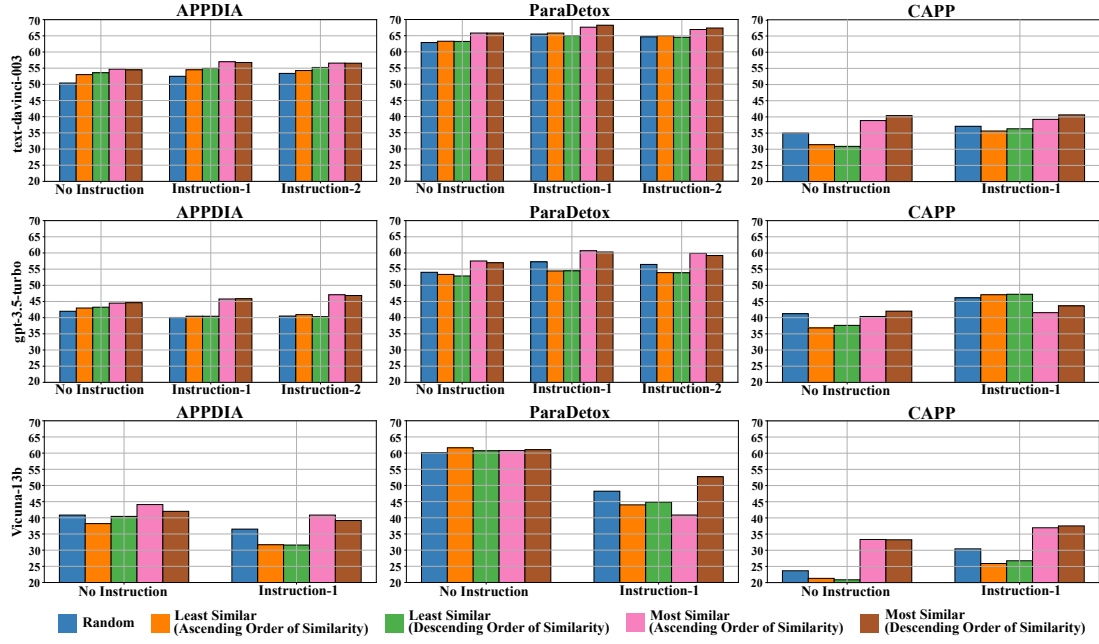


Figure 3: BLEU as a function of order of demonstrations and type of instruction used in the prompt design. Demonstrations that are semantically more similar to the query sample show better performance than less semantically similar and randomly selected samples. Also, prompts that only include demonstrations (*i.e.*, no instruction) show a BLEU score that is comparable to prompts that include instruction and demonstrations.




Figure 4: BLEU score and measured toxicity performance with different instructions but with the same set of demos. Instructions can either complement or work against the selected demos and accordingly affect the BLEU score. No instruction setting shows comparable BLEU to instructions with demos but result in paraphrases with higher toxicity.

3.3 Significance of Instruction

We now investigate the effect of removing the instruction in the prompt. Figure 3 and 4 show prompts that only incorporate demonstrations (*no instruction*) show BLEU scores that are on par with prompts that include both instructions and demon-

strations. This is interesting since it is a common practice with recent instruction-tuned LLMs to always include the instruction even if no demonstrations are provided. Our results suggest that if it is difficult to determine effective instructions for the target paraphrasing task, with ICL one can simply use a few systematically selected demonstrations to get paraphrases that have high generation quality.

In Figure 4; for text-davinci-003 on APPDIA, we observe that the *no instruction* setting retains a significant amount of the original content’s toxicity, thereby making its usability questionable. Similar observations were made with other models (refer to Appendix D.3, Figure 9). Order of demos also plays an important role in the no instruction setting, with the *Most Similar* showing much lower toxicity than both *Random* and *Least Similar* strategies. Cases that include both instruction and demos, the measured toxicity is less impacted by the order of demos, indicating that the main instruction serves as a toxicity regularizer. We want to also highlight that creating a good instruction for paraphrasing tasks is non-trivial. Despite using good demos, a bad instruction can negatively impact the quality of the generated paraphrase. In Figure 4, we see that certain instructions can result in lower BLEU than prompts that have no instruction. Similarly, in Figure 3, the Vicuna-13b model shows better

Dataset	Method	BLEU \uparrow	BERT-F1 \uparrow	ROUGE \uparrow	CIDEr \uparrow	Toxicity \downarrow	Quality \uparrow
APPDIA	<i>Offensive Test Set</i>	-	-	-	-	75.60	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	14.37	3.68 \pm 0.93
	BART (Atwell et al., 2022)	65.0	68.1	65.6	4.77	25.91	3.42 \pm 1.08
	T5 (Atwell et al., 2022)	65.3	69.2	66.5	4.75	20.15	-
	DialoGPT (Atwell et al., 2022)	42.3	46.7	38.0	1.11	14.51	3.52 \pm 0.93
	PDTB+RST (Atwell et al., 2022)	46.2	50.7	42.5	1.54	16.39	-
	text-davinci-003 (10 Demos)	56.8	63.6	57.6	3.70	11.64	3.98\pm1.05
	text-davinci-003 (40 Demos)	60.9	66.7	62.9	4.29	12.67	3.77\pm1.08
	gpt-3.5-turbo-0613 (10 Demos)	45.8	53.3	41.6	2.12	7.00	4.24\pm0.91
	gpt-3.5-turbo-0613 (40 Demos)	50.4	58.2	47.6	2.67	10.08	4.11\pm1.00
	Vicuna-13b (4 Demos)	38.2	46.8	34.9	1.41	12.07	3.87\pm1.00
	Vicuna-13b (10 Demos)	40.3	48.0	37.6	1.79	18.44	3.91\pm1.07
ParaDetox	<i>Offensive Test Set</i>	-	-	-	-	88.64	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	6.56	3.77 \pm 0.97
	BART (Logacheva et al., 2022)	77.3	76.2	69.8	4.94	82.00	2.82 \pm 0.75
	text-davinci-003 (10 Demos)	68.2	67.7	58.9	3.67	6.50	4.34\pm0.91
	text-davinci-003 (40 Demos)	70.1	69.3	60.4	3.95	6.21	4.22\pm0.96
	gpt-3.5-turbo-0613 (10 Demos)	60.3	62.0	50.5	2.72	5.71	3.90\pm1.01
	gpt-3.5-turbo-0613 (40 Demos)	64.3	65.1	54.1	3.08	6.20	3.92\pm1.02
	Vicuna-13b (4 Demos)	49.3	54.1	41.1	1.78	7.23	4.00\pm0.99
	Vicuna-13b (10 Demos)	52.8	56.7	43.7	2.05	9.98	4.53\pm0.84
CAPP	<i>Offensive Test Set</i>	-	-	-	-	25.87	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	0.94	4.38 \pm 0.83
	BART	38.5	48.3	36.3	1.86	3.54	3.78 \pm 0.87
	T5	39.4	50.2	37.9	1.92	2.63	3.84 \pm 0.87
	text-davinci-003 (10 Demos)	40.6	49.6	36.1	1.73	1.04	4.03\pm0.96
	text-davinci-003 (40 Demos)	44.5	53.2	40.7	2.10	1.09	4.10\pm0.94
	gpt-3.5-turbo-0613 (10 Demos)	43.7	51.9	39.6	2.00	0.82	4.44\pm0.81
	gpt-3.5-turbo-0613 (40 Demos)	47.1	55.0	43.0	2.33	0.72	4.58\pm0.76
	Vicuna-13b (4 Demos)	35.8	42.4	31.3	1.34	1.04	4.36\pm0.78
	Vicuna-13b (10 Demos)	37.5	35.9	33.6	1.55	1.02	4.21\pm0.88

Table 1: Quantitative and qualitative assessment of different LLMs using the ICL paradigm and comparison against different baseline supervised approaches. Toxicity of the offensive test set and inoffensive ground-truth paraphrases is also provided. Differences in the reported Mean \pm Std Quality scores between each ICL-based approach and the different baselines is significantly different (*i.e.*, P -value < 0.05).

BLEU with just the curated demonstrations on the APPDIA and ParaDetox datasets.

3.4 Comparison with Supervised Approaches

We compare our ICL-based approach to prior state-of-the-art supervised baselines. For APPDIA we use BART, T5, DialoGPT, and PDTB+RST methods as done in (Atwell et al., 2022); for ParaDetox we use BART as done in (Logacheva et al., 2022); and for CAPP we fine-tuned BART-base and T5-base on the training set. We obtained the generated paraphrases for APPDIA and ParaDetox by their authors. We used the default hyperparameters defined in the Transformers Seq2SeqTrainer for fine-tuning on CAPP. The comparison between our ICL-based approaches and prior baselines is shown in Table 1 (refer to Appendix D.4, Table 9 for latest OpenAI models). Note, the objective of any paraphraser should be to score high on generation quality and have a low Toxicity in the generated paraphrases. For APPDIA and ParaDetox,

the BART and T5 models perform better than our ICL-based approach on generation quality. However, the paraphrases generated by these baselines seem to retain a significant amount of the original toxicity. To better understand this issue, we use the Toxicity for the *Inoffensive Gold-Standard* in each dataset as a point of reference. Ideally, a paraphraser should generate paraphrases whose average Toxicity is no greater than this reference. We observe that all baseline methods except DialoGPT show a higher Toxicity, while the ICL-based methods exhibit Toxicity that is lower or on par with that of the Gold-Standard. Our approach offers a better trade-off between generation quality and Toxicity.

Figure 5 illustrates the Toxicity measured for the different ICL-based methods by varying the number of demonstrations. The *Most Similar (Descending Order)* strategy was used to select and organize the demos. Figure 5 also displays the measured Toxicity of the Offensive Test Set, Gold-Standard and the different baseline approaches. Note that,

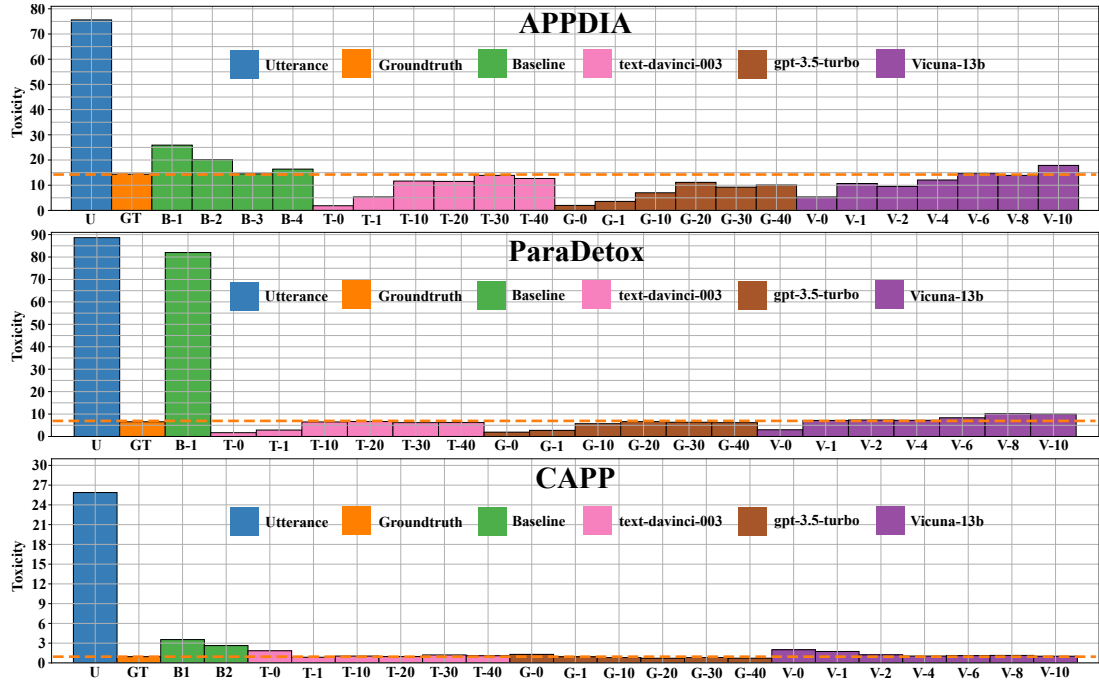


Figure 5: Average Toxicity measured using the *Detoxify* (Hanu and Unitary team, 2020). The orange dotted line serves as a reference for the Gold-Standard’s Toxicity. U, GT, B-#, T-#, G-#, V-# along the x -axis refer to Utterance, Gold-Standard, Baseline methods, text-davinci-003, gpt-3.5-turbo, Vicuna-13b respectively. # in T-#, G-#, V-# indicate number of demonstrations used. Note, T-0, G-0, V-0 only contain an instruction in the prompt.

B1, B2, B3, B4 refer to BART, T5, DialoGPT and PDTB+RST models respectively. For APPDIA and ParaDetox we observe that LLMs without any demonstrations show much lower Toxicity than when any demonstration is used. The reverse is observed for the CAPP dataset. The absence of demos causes LLMs to fallback on their own task definition which results in paraphrases with Toxicity significantly different from that of the Gold-Standard. However, LLMs without any demos show a dramatically lower Toxicity score, but also exhibit a lower BLEU score as seen before in Figure 2. A balance between the main instruction and demos can ensure generation of paraphrases that reduce offensiveness and score high using different automated metrics.

3.5 Additional Dialogue Context Helps

We show preliminary results of using the prior two utterances as additional context in the our ICL-based method. Similar to the example in Section 2.1, we prepend the context for both the demo and the query. Figure 6 shows BLEU score as we add/remove context and vary the number of demonstrations. We clearly see performance improvement by incorporating dialogue context using the text-davinci-003 and gpt-3.5-turbo-0613 models. We were unable to create a prompt for Vicuna-13b that successfully uses additional context in the

ICL framework. We will develop such prompts for Vicuna-13b in future work.

3.6 Robustness to Reduced Training Data

Here we study the impact of available training data on the performance of our best performing strategy i.e. *Most Similar (Descending Order)*. We observe only a minimal fall in BLEU up to 10% of the training data as shown for text-davinci-003 in Figure 7 (refer to Appendix D.5, Figure 10 for other models). Further reducing training data results in noticeable drop in BLEU. We also find that reducing training dataset below 10% results in BLEU score that is similar to *Random* demo selection and arrangement strategy but with access to 100% of the training data. That result shows that our ICL-based method can work with limited training data and thus can be adapted quickly to novel settings.

3.7 Manual Qualitative Assessment

We also perform quality assessment of the Gold-standard and generated paraphrases using a human annotator. We select a subset of 150, 200, and 200 samples from the test-set of APPDIA, ParaDetox and CAPP respectively. For ParaDetox and CAPP, we use all the baseline methods listed in Table 1. For APPDIA we only use BART and DialoGPT models for comparison. We use the scoring guide-




Figure 6: Comparison of BLEU between including and excluding prior context in the form of prior two utterances for both OpenAI models.

lines described in Appendix B, Tables 2 and 3, and the information about the paraphraser model was not made available to the annotator. Table 1 shows that the three ICL-based LLM models received a higher average score that is significantly different (*i.e.*, P -value < 0.05) than the corresponding baseline methods. We also note that Vicuna-13b’s qualitative score was comparable to and in some cases better than the OpenAI models, despite having scored lower on metrics measuring generation quality. This shows that open-source LLM paraphraser are comparable to closed-source LLMs as per human assessment and might be a viable alternative. Refer to Appendix D.6 for analysis between manual evaluation score and toxicity metric.

4 Related Work

Our paper explores the potential use of LLMs with ICL for paraphrasing systems. There has been significant interest in better understanding the capabilities of ICL, but for other applications (Min et al., 2021a; Zhao et al., 2021; Razeghi et al., 2022; Xie et al., 2021; Lampinen et al., 2022; Mishra et al., 2021; Chen et al., 2021; Min et al., 2021b; Chen et al., 2023). (Lu et al., 2021) showed that order of demos has a significant impact on model performance. (Liu et al., 2021) showed that retrieving demonstrations that are semantically similar to the query can be a more effective approach to control the variability in performance. (Rubin et al., 2021) learned an encoding scheme to retrieve better demos for ICL. Other works also explored the influence of number of demos in different settings (Garg et al., 2022; Min et al., 2022; Wei et al., 2023). (Zhou et al., 2022) evaluated the importance of each part in the prompt has towards the final performance. In this paper we study the impact of various components on the final performance, while ensuring that the toxicity of the outputs is within tolerable levels. This enables us to propose a few-shot solution to offensive content paraphras-




Figure 7: BLEU for the “Most Similar (Descending Order)” approach as a function of percentage of training data available, and comparison with Randomly selected demonstrations using 100% of the training data.

ing. Most prior works (Atwell et al., 2022; Logacheva et al., 2022) have modeled paraphrasing as a sequence-to-sequence problem and trained models such as T5, BART on human annotated data. Despite good generation results, these models tend towards higher toxicity and are difficult to adapt to new applications without collecting more data. Our solution addresses those challenges successfully, with only a fraction of the original training set.

5 Conclusion

In this paper, we focus on developing usable offensive content paraphrasing systems by leveraging generalization capabilities of LLMs and quickly adapting them to new tasks using ICL. A paraphraser should generate qualitatively good paraphrases that preserve the original content’s meaning, while also minimizing toxicity. Focusing only on one of these aspects compromises overall usability. Compared to supervised approaches that require lot of training data and often produce undesired yet coherent paraphrases, our ICL-based framework is generally comparable on various evaluation metrics like BLEU, but is qualitatively better and helps significantly reduce toxicity in the generated paraphrases. Through systematic experiments we tested the capabilities and limitations of ICL-based offensive paraphraser. Other key highlights of using our ICL framework include: (1) Selection and arrangement of demos significantly impacts quality of paraphrases; (2) Measured toxicity is lowest when only the instruction is used and highest when only demos are used. Combining both instruction and demos helps ensure quality and usability of generated paraphrases; (3) Robust to limited data, *i.e.*, with just 10% training data we only see a slight decrease in overall performance, thereby enabling us to easily scale and deploy.

Limitations

Here, we list the limitations identified in this paper:

1. We found that ICL fails on datasets that were prepared using the same LLMs used in the ICL framework. Since we used the gpt-3.5-turbo-0613 model to create polite paraphrases for our CAPP dataset, we were unable to see the same observations in the results that we did using the other models and datasets. This could have been avoided by creating manually annotated polite paraphrases. However, manual annotation is a laborious process and isn't scalable. Hence, a manual qualitative assessment was done on a small subset of the final CAPP dataset to ensure usability of the generated paraphrases.
2. Prompt engineering for the Vicuna-13b model with the ICL framework is nontrivial. We found it difficult to create main instructions in the prompt that result in the Vicuna-13b model to behave in a desired way. Also, unlike the two OpenAI models, the number of demonstrations that can be effectively passed into Vicuna-13b is quite limited. In some cases we were able to concatenate more than 10 demos to the prompt but it often resulted in generating incomprehensible outputs.
3. The *No Instruction* prompt explored in the paper resulted in paraphrases that are comparable to prompts that include both instruction and demos, on several automated evaluation metrics. However, we notice that the *No Instruction* setting also retains a significant amount of toxicity from the original content. We propose that in situations where it is difficult to decide on a good main instruction, one could simply use a few carefully curated and ordered demos like the "Most Similar (Descending Order)" approach to generate paraphrases and check if it is within the desired toxicity levels.
4. Our experimental results indicate that there is no single prompt that works in all situations. One must carefully balance the main instruction and the set of demos from the training corpus to get desired paraphrase outputs.
5. We showed preliminary results showcasing the benefit of incorporating additional context in the form of prior two utterances in

the ICL framework. We believe there can be better ways to incorporate this contextual information and further improve performance of LLMs.

6. The closed-source OpenAI models are more powerful, faster and expensive to use. Despite open-source models like Vicuna-13b coming close to OpenAI models on other tasks, they still have a long way to go for offensive content paraphrasing.

Ethics Statement

We have to take great care with our collection of offensive content to protect privacy. We have to ensure judicious use of the collected data to protect the vulnerable against such speech. We recognize that our models cannot entirely eliminate offensive content from a given text. Additionally, we acknowledge that utilizing pretrained models may introduce biases in specific situations, as studies have revealed that pretrained models can be influenced by biases present in the data used for their initial training. We have to continue research on making sure that the LLM's do not hallucinate and end up injecting toxicity since we don't know what they have been trained on.

There is a danger of this kind of technology being used in reverse, i.e., take harmless content and paraphrase to inject toxicity.

We realize that ethics is an ongoing challenge. We are engaged with the Fairness, Accountability and Transparency community and are learning to address key ethics issues on an ongoing basis.

References

- Katherine Atwell, Sabit Hassan, and Malihe Alikhani. 2022. Appdia: A discourse-aware transformer-based style transfer model for offensive social media conversations. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6063–6074.
- 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.
- Pete Burnap and Matthew L Williams. 2015. Cyber hate speech on twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & internet*, 7(2):223–242.

627	Derek Chen, Celine Lee, Yunan Lu, Domenic Rosati, and Zhou Yu. 2023. Mixture of soft prompts for controllable data generation. <i>arXiv preprint arXiv:2303.01580</i> .	682
628		683
629		684
630		685
631	Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He. 2021. Meta-learning via language model in-context tuning. <i>arXiv preprint arXiv:2110.07814</i> .	686
632		687
633		
634	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.	688
635		689
636		690
637		691
638		692
639		693
640	Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. <i>arXiv preprint arXiv:1306.6078</i> .	694
641		695
642		696
643		
644		
645	Thomas Davidson, Dana Warmley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In <i>Proceedings of the international AAAI conference on web and social media</i> , volume 11, pages 512–515.	697
646		698
647		699
648		
649		
650	Bryan Dosono and Bryan Semaan. 2019. Moderation practices as emotional labor in sustaining online communities: The case of aapi identity work on reddit. In <i>Proceedings of the 2019 CHI conference on human factors in computing systems</i> , pages 1–13.	700
651		701
652		702
653		703
654		
655	Antigoni Maria Founta, Despoina Chatzakou, Nicolas Kourtellis, Jeremy Blackburn, Athena Vakali, and Ilias Leontiadis. 2019. A unified deep learning architecture for abuse detection. In <i>Proceedings of the 10th ACM conference on web science</i> , pages 105–114.	704
656		705
657		706
658		707
659		708
660		709
661	Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. 2022. What can transformers learn in-context? a case study of simple function classes. <i>Advances in Neural Information Processing Systems</i> , 35:30583–30598.	710
662		711
663		712
664		713
665		714
666	Laura Hanu and Unitary team. 2020. Detoxify. Github. https://github.com/unitaryai/detoxify .	715
667		
668	Sabit Hassan, Katherine J Atwell, and Malihe Alikhani. 2022. Studying the effect of moderator biases on the diversity of online discussions: A computational cross-linguistic study. In <i>Proceedings of the Annual Meeting of the Cognitive Science Society</i> , volume 44.	716
669		717
670		718
671		719
672		
673	Shagun Jhaver, Darren Scott Appling, Eric Gilbert, and Amy Bruckman. 2019. Did you suspect the post would be removed? understanding user reactions to content removals on reddit. <i>Proceedings of the ACM on human-computer interaction</i> , 3(CSCW):1–33.	720
674		721
675		722
676		
677		
678	Irene Kwok and Yuzhou Wang. 2013. Locate the hate: Detecting tweets against blacks. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 27, pages 1621–1622.	723
679		724
680		725
681		726
		727
	Andrew K Lampinen, Ishita Dasgupta, Stephanie CY 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? <i>arXiv preprint arXiv:2204.02329</i> .	728
		729
		730
		731
	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. <i>arXiv preprint arXiv:1910.13461</i> .	732
		733
		734
		735
		736
	Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In <i>Text summarization branches out</i> , pages 74–81.	
	Pierre Lison and Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles.	
	Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? <i>arXiv preprint arXiv:2101.06804</i> .	
	Varvara Logacheva, Daryna Dementieva, Sergey Ustyantsev, Daniil Moskovskiy, David Dale, Irina Krotova, Nikita Semenov, and Alexander Panchenko. 2022. Paradox: Detoxification with parallel data. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 6804–6818.	
	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. <i>arXiv preprint arXiv:2104.08786</i> .	
	Sewon Min, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2021a. Noisy channel language model prompting for few-shot text classification. <i>arXiv preprint arXiv:2108.04106</i> .	
	Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2021b. Metaicl: Learning to learn in context. <i>arXiv preprint arXiv:2110.15943</i> .	
	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? <i>arXiv preprint arXiv:2202.12837</i> .	
	Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2021. Reframing instructional prompts to gptk’s language. <i>arXiv preprint arXiv:2109.07830</i> .	
	Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In <i>Proceedings of the 25th international conference on world wide web</i> , pages 145–153.	

737	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-	topical feature discovery over a large scale twitter cor-	790
738	Jing Zhu. 2002. Bleu: a method for automatic evalu-	pus. In <i>Proceedings of the 21st ACM international</i>	791
739	ation of machine translation. In <i>Proceedings of the</i>	<i>conference on Information and knowledge manage-</i>	792
740	<i>40th annual meeting of the Association for Computa-</i>	<i>ment</i> , pages 1980–1984.	793
741	<i>tional Linguistics</i> , pages 311–318.		
742	Yasaman Razeghi, Robert L Logan IV, Matt Gardner,	Sang Michael Xie, Aditi Raghunathan, Percy Liang, and	794
743	and Sameer Singh. 2022. Impact of pretraining term	Tengyu Ma. 2021. An explanation of in-context learn-	795
744	frequencies on few-shot reasoning. <i>arXiv preprint</i>	ing as implicit bayesian inference. <i>arXiv preprint</i>	796
745	<i>arXiv:2202.07206</i> .	<i>arXiv:2111.02080</i> .	797
746	Nils Reimers and Iryna Gurevych. 2019. Sentence-bert:	Meng Ye, Karan Sikka, Katherine Atwell, Sabit Hassan,	798
747	Sentence embeddings using siamese bert-networks .	Ajay Divakaran, and Malihe Alikhani. 2023. Multi-	799
748	In <i>Proceedings of the 2019 Conference on Empirical</i>	lingual content moderation: A case study on reddit.	800
749	<i>Methods in Natural Language Processing</i> . Associa-	<i>arXiv preprint arXiv:2302.09618</i> .	801
750	tion for Computational Linguistics.		
751	Ohad Rubin, Jonathan Herzig, and Jonathan Berant.	Xi Ye, Srinivasan Iyer, Asli Celikyilmaz, Ves Stoy-	802
752	2021. Learning to retrieve prompts for in-context	anov, Greg Durrett, and Ramakanth Pasunuru. 2022.	803
753	learning. <i>arXiv preprint arXiv:2112.08633</i> .	Complementary explanations for effective in-context	804
754	Victor Sanh, Lysandre Debut, Julien Chaumond, and	learning. <i>arXiv preprint arXiv:2211.13892</i> .	805
755	Thomas Wolf. 2019. Distilbert, a distilled version	Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q	806
756	of bert: smaller, faster, cheaper and lighter. <i>arXiv</i>	Weinberger, and Yoav Artzi. 2019. Bertscore: Eval-	807
757	<i>preprint arXiv:1910.01108</i> .	uating text generation with bert. <i>arXiv preprint</i>	808
758	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob	<i>arXiv:1904.09675</i> .	809
759	Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz	Yiming Zhang, Shi Feng, and Chenhao Tan. 2022. Ac-	810
760	Kaiser, and Illia Polosukhin. 2017. Attention is all	tive example selection for in-context learning. <i>arXiv</i>	811
761	you need. <i>Advances in neural information processing</i>	<i>preprint arXiv:2211.04486</i> .	812
762	<i>systems</i> , 30.	Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and	813
763	Ramakrishna Vedantam, C Lawrence Zitnick, and Devi	Sameer Singh. 2021. Calibrate before use: Improv-	814
764	Parikh. 2015. Cider: Consensus-based image de-	ing few-shot performance of language models. In <i>In-</i>	815
765	scription evaluation. In <i>Proceedings of the IEEE</i>	<i>ternational Conference on Machine Learning</i> , pages	816
766	<i>conference on computer vision and pattern recogni-</i>	12697–12706. PMLR.	817
767	<i>tion</i> , pages 4566–4575.	Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han,	818
768	Jeremy Waldron. 2012. <i>The harm in hate speech</i> . Har-	Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy	819
769	vard University Press.	Ba. 2022. Large language models are human-level	820
770	Wenbo Wang, Lu Chen, Krishnaprasad Thirunarayan,	prompt engineers. <i>arXiv preprint arXiv:2211.01910</i> .	821
771	and Amit P Sheth. 2014. Cursing in english on twitter.		
772	In <i>Proceedings of the 17th ACM conference on</i>		
773	<i>Computer supported cooperative work & social com-</i>		
774	<i>puting</i> , pages 415–425.		
775	William Warner and Julia Hirschberg. 2012. Detecting		
776	hate speech on the world wide web. In <i>Proceedings</i>		
777	<i>of the second workshop on language in social media</i> ,		
778	pages 19–26.		
779	Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert		
780	Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu,		
781	Da Huang, Denny Zhou, et al. 2023. Larger language		
782	models do in-context learning differently. <i>arXiv</i>		
783	<i>preprint arXiv:2303.03846</i> .		
784	John Wieting and Kevin Gimpel. 2017. Paranzmt-50m:		
785	Pushing the limits of paraphrastic sentence embed-		
786	dings with millions of machine translations. <i>arXiv</i>		
787	<i>preprint arXiv:1711.05732</i> .		
788	Guang Xiang, Bin Fan, Ling Wang, Jason Hong, and		
789	Carolyn Rose. 2012. Detecting offensive tweets via		

A Context-Aware Polite Paraphrase (CAPP) Dataset

Rudeness Detection Model: The initial rudeness detection model was trained on a randomly-selected subset of the OpenSubtitles corpus (Lison and Tiedemann, 2016) that was manually annotated by an in-house human annotator, as well as the Stanford Politeness corpus (Danescu-Niculescu-Mizil et al., 2013). Between these two data sources, this model was trained on 10997 gold-quality samples. This intermediate model achieved an overall accuracy score of 59.40 across the three classes (rude, polite, and neutral). The accuracy for specifically rude utterances in the test data was 71.84 – higher than that of other classes. The final rudeness detection model, trained on a combination of the aforementioned gold data, and a larger subset (250379 samples) of the OpenSubtitles corpus annotated using the intermediate model, achieved an overall accuracy of 57.20 across the three classes, with an accuracy score of 71.84 on rude samples – still higher than the accuracy on other classes. The disparity in rudeness detector performance between classes was deemed acceptable due to its high performance on the rude class. Accurate identification of rude utterances was prioritized in the interest of generating a sufficient number and variety of rude utterances for paraphraser training – with an acceptance of the potential for noise in the form of some non-rude utterances in the training data.

Qualitative Analysis: We employed three different prompts for generating three versions of polite paraphrases – (1) *Context-Free*: No context included in the prompt; (2) *Context-Infused*: Prompt includes context and significantly influences the generated paraphrase; (3) *Context-Aware*: Prompt includes context, with the paraphrase being less impacted by it. 500 rude utterances and their corresponding polite paraphrases were selected for evaluation. An in-house annotator assessed the quality of the paraphrases using the scoring guidelines in Table 2 and was not informed about the type of prompt used. The annotator was compensated monetarily. The annotator identifies as Female (pronouns she/her) and is 28 years old. Table 4 shows the final evaluation scores. All scores are significantly different, with the Context-Aware prompt achieves a score comparable to the Context-Free prompt while still incorporating context like the Context-Infused prompt. Context-Aware combines the benefits of both, and was hence used in the CAPP dataset.

Score	Description
5	Perfect meaning-preserving polite paraphrase.
4	Paraphrase that is polite but somewhat distinct in meaning.
3	Meaning-preserving paraphrase that could be more polite.
2	Paraphrase that is very different in meaning and somewhat more polite than the original.
1	Paraphrase that is very different in meaning and not more polite than the original.

Table 2: Description of the scoring guidelines used for evaluating the CAPP dataset in Section 2.3. The same guidelines were again used in Section 3.7 to evaluate paraphrases generated by the different approaches for the CAPP dataset.

Score	Description
5	Perfect meaning-preserving inoffensive paraphrase.
4	Paraphrase that is inoffensive but somewhat distinct in meaning.
3	Meaning-preserving paraphrase that could be less offensive.
2	Paraphrase that is very different in meaning and somewhat less offensive than the original.
1	Paraphrase that is very different in meaning and not less offensive than the original.

Table 3: Description of the scoring guidelines used in Section 3.7 for evaluating the generated paraphrases for the APPDIA (Atwell et al., 2022) and ParaDetox (Logacheva et al., 2022) datasets.

Prompt	Manual Evaluation Score [↑]
Context-Free	4.214 ± 1.047
Context-Infused	3.324 ± 0.839
Context-Aware	4.096 ± 1.093

Table 4: Human evaluation scores of 500 polite paraphrases generated using different prompts. A higher score indicates a qualitatively better approach.

B Scoring Guidelines

Table 2 shows the scoring guidelines used to evaluate the quality of the CAPP dataset. Table 3 shows the scoring guidelines used to evaluate the quality of the generated paraphrase from the different

paraphrasing methods.

C Instructions Used in Experiments

The exact prompt instruction used for each experiment in Section 3 is shown below. Except Section 3.2 and 3.3, we use only **Instruction-1** for the following Subsections 3.1, 3.4, 3.5, 3.6 and 3.7.

1. Instruction prompts used for text-davinci-003 and gpt-3.5-turbo models on the APPDIA dataset:
 - (a) **Instruction-1:** Instruction: Rewrite the following sentence to be inoffensive
 - (b) **Instruction-2:** Instruction: Rewrite the following sentence to be inoffensive and ensure that the paraphrased sentence is similar to the original sentence
 - (c) **Instruction-3:** Instruction: Generate a paraphrased sentence that expresses a more balanced viewpoint
 - (d) **Instruction-4:** Instruction: Rewrite the following sentence with a more objective tone
2. Instruction prompts used for text-davinci-003 and gpt-3.5-turbo models on the ParaDetox dataset:
 - (a) **Instruction-1:** Instruction: Rewrite the following sentence to be inoffensive
 - (b) **Instruction-2:** Instruction: Rewrite the following sentence to be not toxic
3. Instruction prompts used for text-davinci-003, gpt-3.5-turbo and Vicuna-13b models on the CAPP dataset:
 - (a) **Instruction-1:** The following sentence can contain rude text. Please provide one paraphrased sentence that is semantically similar to it and is polite.
4. Instruction prompts used for Vicuna-13b model on APPDIA and ParaDetox datasets:
 - (a) **Instruction-1:** The following sentence contains offensive text. Please provide one paraphrased sentence that is semantically similar to it and is inoffensive.
5. While incorporating additional context in the form of previous two turns in the dialogue, the instruction prompt used for text-davinci-003 and gpt-3.5-turbo models on the CAPP dataset:

- (a) **Instruction-1:** Paraphrase only the below Sentence to be polite and semantically similar to the Sentence. Use the context as as reference but do not include any part of it in the final paraphrase.

D Experiments

D.1 Number of Demonstrations

Figure 8 provides additional details, supporting information provided in Section 3.1.

D.2 Selection and Order of Demonstrations

Table 5 compares the performance results and compute times between the MMR approach (Ye et al., 2022) and the proposed selection and ordering approach. Note, for the proposed approach, the performance metrics reported correspond to the *Most Similar (Descending Order)* method, however, the compute times can represent any of the proposed demonstration selection and ordering approaches described in Section 3.2. As suggested by the authors, the parameter λ was set to 0.5 for the MMR approach. For all three approaches we used the gpt-3.5-turbo-1106 model.

Distance Compute Time refers to average time taken to perform similarity measurements between every query sample and all the reference training samples. While the original paper suggests using BERTScore to perform the similarity measurement, we also explore cosine similarity measurement of the embeddings extracted using the sentence transformer models. Note, the Distance Compute Time for the cosine similarity variation of the MMR approach will be the same as our proposed approach. *Top-10 Index Retrieval Compute Time* refers to average time taken to select and order 10 demonstrations from the training set.

As seen in Table 5, the MMR-based approaches show only marginal improvement with respect to the different quantitative metrics across the three datasets. With respect to the proposed approach, the MMR-BERT approach is three orders of magnitude slower on the Distance Compute Time, while both MMR approaches are two orders of magnitude slower on the Top-10 Index Retrieval Compute Time.

While there are several demonstration selection and ordering approaches (Ye et al., 2022; Zhang et al., 2022; Lu et al., 2021) that can help push the performance ceiling, one must also make sure if these approaches can be easily implemented and

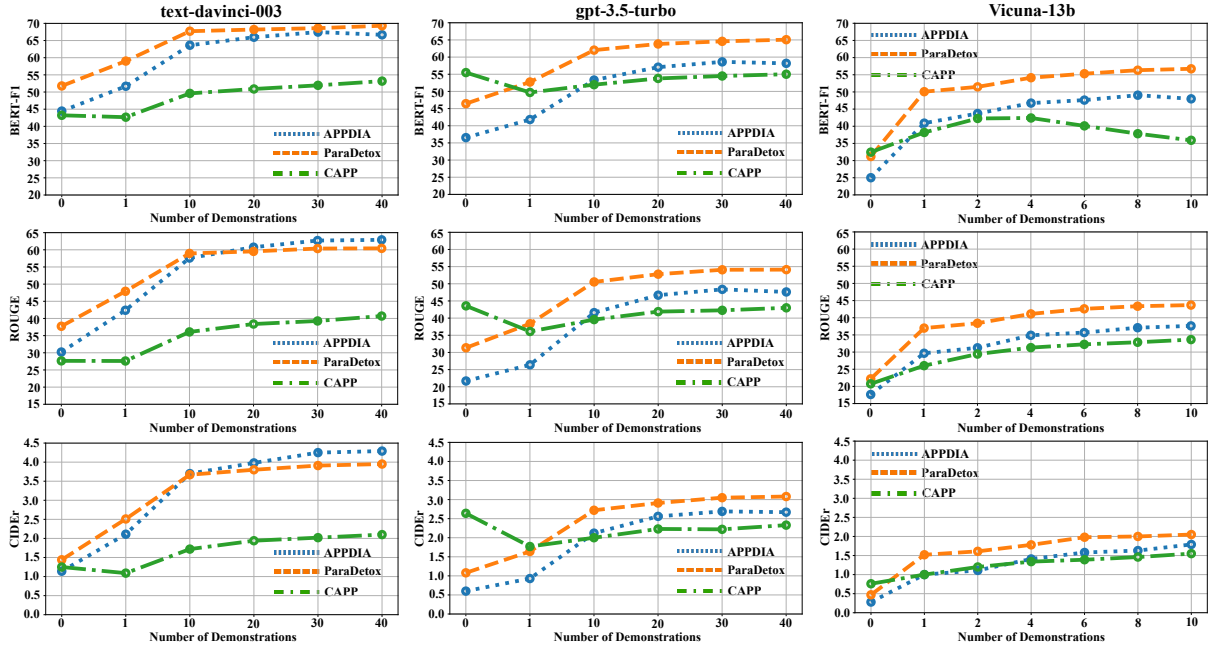


Figure 8: Performance using different evaluation metrics as a function of number of demonstrations used in the prompt. Noticeable improvement in score performance is observed in the beginning, with performance saturating after a certain number of demos.

scaled up in a real-time, real-world application setting. For example, (Zhang et al., 2022) propose a reinforcement learning algorithm to identify generalizable policies to select demonstrations but find that this approach offers diminishing returns on larger, more sophisticated LLMs. The approach described by (Lu et al., 2021) to overcome few-shot prompt order sensitivity is better suited for multi-class classification tasks, since the entropy-based statistics framework discussed to identify performant prompts is not directly applicable to text generation tasks like paraphrasing.

D.3 Significance of Instruction

Figure 9 provides additional details, supporting information provided in Section 3.3.

D.4 Comparison with Supervised Approaches

Table 9 shows an extended version of Table 1. Highlighted rows describe performance results for gpt-3.5-turbo-instruct and gpt-3.5-turbo-1106, and how it differs from their respective predecessor models text-davinci-003 and gpt-3.5-turbo-0613. Despite performance differences across the older and newer models, the newer models continue to exhibit the same patterns and behaviors across the different empirical experiments just as their predecessors.

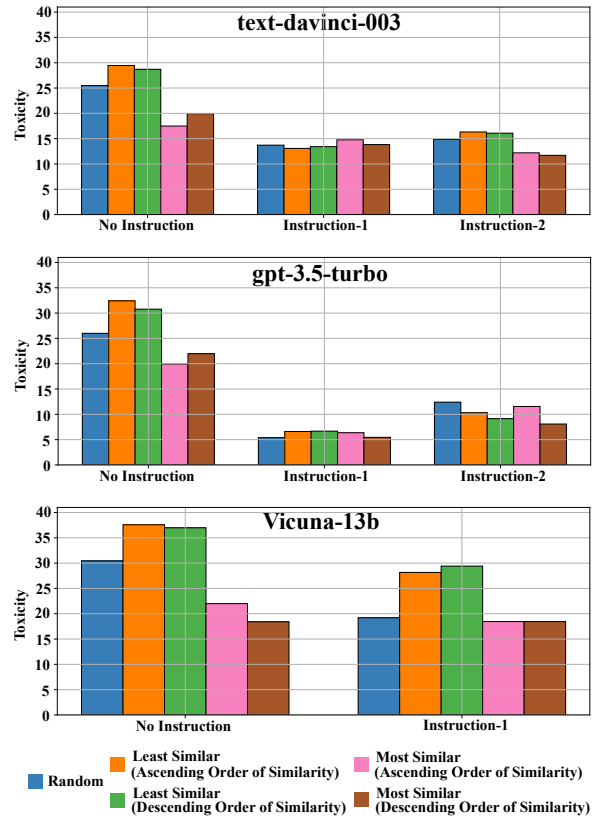


Figure 9: Measured toxicity performance of different models on the APPDIA dataset, with different instructions but with the same set of demos. No instruction setting results in paraphrases with higher toxicity.

Dataset	Method	BLEU \uparrow	BERT-F1 \uparrow	ROUGE \uparrow	CIDEr \uparrow	Toxicity \downarrow	Distance Compute Time \downarrow	Top-10 Index Retrieval Compute Time \downarrow
APPDIA	MMR-BERT (Ye et al., 2022)	58.5	65.1	59.3	3.90	17.54	4.7538 ± 0.0898	$(0.0791) \pm (1.77 \times 10^{-3})$
	MMR-Embedding	57.9	63.8	57.7	3.79	14.11	0.0025 ± 0.0004	$(0.0786) \pm (1.04 \times 10^{-3})$
	Proposed	57.1	64.0	57.2	3.72	14.43	0.0025 ± 0.0004	$(0.0005) \pm (6.49 \times 10^{-6})$
ParaDetox	MMR-BERT (Ye et al., 2022)	68.6	68.0	58.8	3.67	8.47	-	-
	MMR-Embedding	67.6	67.3	57.7	3.52	8.91	-	-
	Proposed	67.6	67.3	57.2	3.45	8.46	-	-
CAPP	MMR-BERT (Ye et al., 2022)	45.5	54.4	41.8	2.20	1.05	-	-
	MMR-Embedding	44.0	52.4	39.9	2.10	1.34	-	-
	Proposed	44.0	53.2	40.5	2.17	1.22	-	-

Table 5: Comparison of the proposed demonstration selection and ordering approach to MMR (Ye et al., 2022). Here, the proposed approach refers to the *Most Similar (Descending Order)* approach outlined in Section 3.2. While MMR provides marginal performance gains on all three datasets, it is several orders of magnitude slower than the proposed approach. Note, *Distance Compute Time* refers to the time taken to compute the distance or similarity between each query sample and all reference samples in the training set; *Top-10 Index Retrieval Compute Time* refers to time taken to select and order 10 demonstrations from the training set.

D.5 Robustness to Reduced Training Data

Figure 10 provides additional details, supporting information provided in Section 3.6.

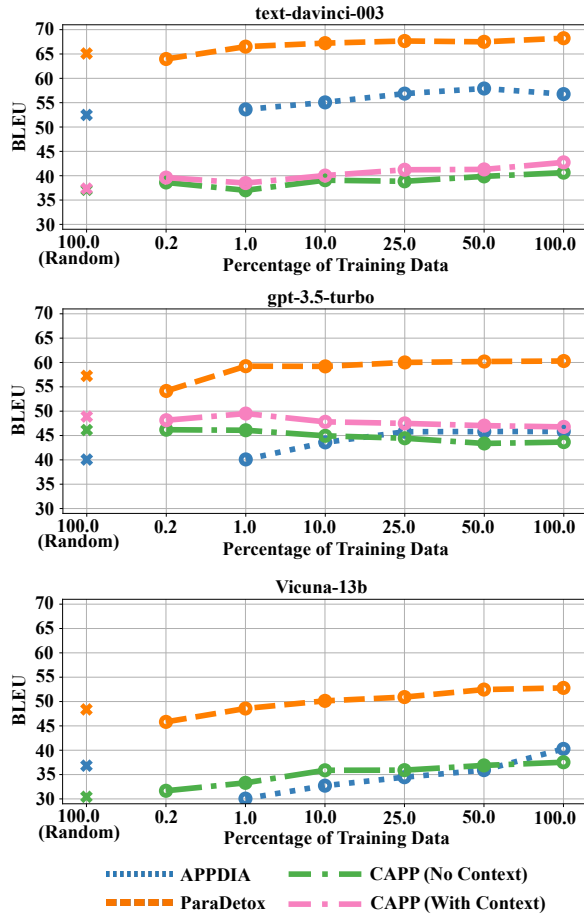


Figure 10: BLEU for the “*Most Similar (Descending Order)*” approach as a function of percentage of training data available and comparison to Random demo selection with access to 100% of the training data.

D.6 Correlation between Manual Evaluation Score and Automated Toxicity Metric

The manual evaluation score considers not just Toxicity minimization in the generated paraphrase but also considers meaning preservation. Directly comparing it to the automated toxicity score is not possible because there is a semantic mismatch between the two metrics. Instead, we first compute the difference in the toxicity measured between the (offensive) utterance and the (inoffensive) paraphrase. Next, we compute the Pearson correlation between this difference in toxicity score to the manual evaluation score. This difference captures the comparisons made by the annotator while coming up with the manual scoring as shown in Table 3. Therefore, if there is correlation between this automatically computed difference and the manual scores, then convergent validity is assured.

Tables 6, 7, 8 show the computed Pearson correlation coefficient between the manual evaluation score and automated toxicity score under two different settings. Type-1 represents the Pearson correlation coefficient between the manual quality score and automated toxicity score of the paraphrased output; and Type-2 represents the Pearson correlation coefficient between the manual quality score and difference in measured toxicity between original utterance and paraphrased output. Note, the dynamic range of the toxicity captured in the CAPP dataset is low because it mostly consists of rude speech that contains little to no foul language.

Dataset	Method	Type-1	Type-2
APPDIA	Gold Standard	-0.28	0.07
	BART	-0.41	0.36
	DialoGPT	-0.40	0.28
	text-davinci-003 (10 demos)	-0.45	0.25
	text-davinci-003 (40 demos)	-0.43	0.27
	gpt-3.5-turbo (10 demos)	-0.39	0.26
	gpt-3.5-turbo (40 demos)	-0.35	0.21
	Vicuna-13b (4 demos)	-0.42	0.24
	Vicuna-13b (10 demos)	-0.41	0.17

Table 6: Computed Pearson correlation coefficient on the APPDIA dataset between manual quality score and automated toxicity score of paraphrased output, denoted as Type-1; manual quality score and difference in measured toxicity between original utterance and paraphrased output, denoted as Type-2.

Dataset	Method	Type-1	Type-2
ParaDetox	Gold Standard	-0.15	0.19
	BART	-0.60	0.56
	text-davinci-003 (10 demos)	-0.22	0.20
	text-davinci-003 (40 demos)	-0.22	0.28
	gpt-3.5-turbo (10 demos)	-0.24	0.25
	gpt-3.5-turbo (40 demos)	-0.26	0.24
	Vicuna-13b (4 demos)	-0.24	0.26
	Vicuna-13b (10 demos)	-0.49	0.41

Table 7: Computed Pearson correlation coefficient on the ParaDetox dataset between manual quality score and automated toxicity score of paraphrased output, denoted as Type-1; manual quality score and difference in measured toxicity between original utterance and paraphrased output, denoted as Type-2.

Dataset	Method	Type-1	Type-2
CAPP	Gold Standard	-0.139	-0.049
	BART	-0.189	-0.094
	T5	-0.188	-0.005
	text-davinci-003 (10 demos)	-0.155	0.022
	text-davinci-003 (40 demos)	-0.218	-0.016
	gpt-3.5-turbo (10 demos)	-0.197	-0.077
	gpt-3.5-turbo (40 demos)	-0.265	-0.129
	Vicuna-13b (4 demos)	-0.096	-0.188
	Vicuna-13b (10 demos)	-0.176	-0.058

Table 8: Computed Pearson correlation coefficient on the CAPP dataset between manual quality score and automated toxicity score of paraphrased output, denoted as Type-1; manual quality score and difference in measured toxicity between original utterance and paraphrased output, denoted as Type-2.

Dataset	Method	BLEU↑	BERT-F1↑	ROUGE↑	CIDEr↑	Toxicity↓	Quality↑
APPDIA	<i>Offensive Test Set</i>	-	-	-	-	75.60	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	14.37	3.68±0.93
	BART (Atwell et al., 2022)	65.0	68.1	65.6	4.77	25.91	3.42±1.08
	T5 (Atwell et al., 2022)	65.3	69.2	66.5	4.75	20.15	-
	DialoGPT (Atwell et al., 2022)	42.3	46.7	38.0	1.11	14.51	3.52±0.93
	PDTB+RST (Atwell et al., 2022)	46.2	50.7	42.5	1.54	16.39	-
	text-davinci-003 (10 Demos)	56.8	63.6	57.6	3.70	11.64	3.98±1.05
	text-davinci-003 (40 Demos)	60.9	66.7	62.9	4.29	12.67	3.77±1.08
	gpt-3.5-turbo-instruct (10 Demos)	51.9	58.9	50.5	2.81	15.86	-
	gpt-3.5-turbo-instruct (40 Demos)	56.2	62.9	55.8	3.42	21.37	-
	gpt-3.5-turbo-0613 (10 Demos)	45.8	53.3	41.6	2.12	7.00	4.24±0.91
	gpt-3.5-turbo-0613 (40 Demos)	50.4	58.2	47.6	2.67	10.08	4.11±1.00
	gpt-3.5-turbo-1106 (10 Demos)	57.1	64.0	57.20	3.72	14.42	-
	gpt-3.5-turbo-1106 (40 Demos)	61.2	66.4	61.9	4.23	19.55	-
ParaDetox	<i>Offensive Test Set</i>	-	-	-	-	88.64	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	6.56	3.77±0.97
	BART (Logacheva et al., 2022)	77.3	76.2	69.8	4.94	82.00	2.82±0.75
	text-davinci-003 (10 Demos)	68.2	67.7	58.9	3.67	6.50	4.34±0.91
	text-davinci-003 (40 Demos)	70.1	69.3	60.4	3.95	6.21	4.22±0.96
	gpt-3.5-turbo-instruct (10 Demos)	65.2	66.4	55.6	3.17	9.97	-
	gpt-3.5-turbo-instruct (40 Demos)	69.1	68.1	58.9	3.68	12.3	-
	gpt-3.5-turbo-0613 (10 Demos)	60.3	62.0	50.5	2.72	5.71	3.90±1.01
	gpt-3.5-turbo-0613 (40 Demos)	64.3	65.1	54.1	3.08	6.20	3.92±1.02
	gpt-3.5-turbo-1106 (10 Demos)	67.6	67.3	57.2	3.45	8.46	-
	gpt-3.5-turbo-1106 (40 Demos)	70.1	69.3	59.6	3.79	9.64	-
	Vicuna-13b (4 Demos)	49.3	54.1	41.1	1.78	7.23	4.00±0.99
	Vicuna-13b (10 Demos)	52.8	56.7	43.7	2.05	9.98	4.53±0.84
CAPP	<i>Offensive Test Set</i>	-	-	-	-	25.87	-
	<i>Inoffensive Gold-Standard</i>	-	-	-	-	0.94	4.38±0.83
	BART	38.5	48.3	36.3	1.86	3.54	3.78±0.87
	T5	39.4	50.2	37.9	1.92	2.63	3.84±0.87
	text-davinci-003 (10 Demos)	40.6	49.6	36.1	1.73	1.04	4.03±0.96
	text-davinci-003 (40 Demos)	44.5	53.2	40.7	2.10	1.09	4.10±0.94
	gpt-3.5-turbo-instruct (10 Demos)	44.9	53.6	41.1	2.18	1.23	-
	gpt-3.5-turbo-instruct (40 Demos)	48.0	56.4	44.5	2.53	1.49	-
	gpt-3.5-turbo-0613 (10 Demos)	43.7	51.9	39.6	2.00	0.82	4.44±0.81
	gpt-3.5-turbo-0613 (40 Demos)	47.1	55.0	43.0	2.33	0.72	4.58±0.76
	gpt-3.5-turbo-1106 (10 Demos)	43.9	53.2	40.5	2.17	1.22	-
	gpt-3.5-turbo-1106 (40 Demos)	45.8	54.8	42.5	2.29	1.30	-
	Vicuna-13b (4 Demos)	35.8	42.4	31.3	1.34	1.04	4.36±0.78
	Vicuna-13b (10 Demos)	37.5	35.9	33.6	1.55	1.02	4.21±0.88

Table 9: Quantitative and qualitative assessment of different LLMs using the ICL paradigm and comparison against different baseline supervised approaches. Toxicity of the offensive test set and inoffensive ground-truth paraphrases is also provided. Differences in the reported Mean±Std Quality scores between each ICL-based approach and the different baselines is significantly different (*i.e.*, P -value < 0.05).