

Leveraging LLMs for Dialogue Quality Measurement

Jinghan Jia¹, Abi Komma², Timothy Leffel², Xujun Peng²
Ajay Nagesh², Tamer Soliman², Aram Galstyan², Anoop Kumar²

¹ Computer Science & Engineering, Michigan State University

² Amazon AGI Foundations

jiajingh@msu.edu, {kommaak, leffelt, penxujun, nagesajg,
tsoliman, argalsty}@amazon.com, anoopkum@gmail.com

Abstract

In task-oriented conversational AI evaluation, unsupervised methods poorly correlate with human judgments, and supervised approaches lack generalization. Recent advances in large language models (LLMs) show robust zero-shot and few-shot capabilities across NLP tasks. This paper explores using LLMs for automated dialogue quality evaluation, experimenting with various configurations on public and proprietary datasets. Manipulating factors such as model size, in-context examples, and selection techniques, we examine “chain-of-thought” (CoT) reasoning and label extraction procedures. Our results show that (1) larger models yield more accurate dialogue labels; (2) algorithmic selection of in-context examples outperforms random selection; (3) CoT reasoning where an LLM is asked to provide justifications before outputting final labels improves performance; and (4) fine-tuned LLMs outperform out-of-the-box ones. Our results indicate that LLMs that are suitably fine-tuned and have sufficient reasoning capabilities can be leveraged for automated dialogue evaluation.

1 Introduction

Evaluating conversational system performance in NLP is challenging. Automating effective evaluation is crucial for enhancing dialogue systems. However, automatic metrics like BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) fall short in accurately measuring perceived quality due to complex mappings (Liu et al., 2016). Advanced methods (USR (Mehri and Eskenazi, 2020b), FED (Mehri and Eskenazi, 2020a), DialogRPT (Gao et al., 2020)) address this but often require extensive training data and human references, making them costly and limited in generalization to new datasets.

Recent advancements in Large Language Models (LLMs) (Bubeck et al., 2023) have demonstrated robust zero- or few-shot capabilities and

reasoning skills across a range of tasks (Brown et al., 2020). Consequently, researchers have begun to explore the application of LLMs to classification problems such as dialogue evaluation (Lin and Chen, 2023; Huynh et al., 2023).

Recent studies (Lin and Chen, 2023; Huynh et al., 2023) demonstrate that LLMs excel on diverse datasets. However, uncertainties persist regarding the impact of factors like model size and in-context examples on open-sourced LLM performance. This paper aims to clarify these influences. Additionally, there is a lack of comprehensive literature on deploying LLMs for dialogue evaluation. To fill this gap, we propose two common evaluation strategies, providing a comparative analysis of their pros and cons.

In this paper, we systematically study different aspects of LLM-based dialogue evaluation by conducting extensive experiments on two benchmark datasets, one publicly available and the other proprietary, Amazon-internal datasets. We initially explore the connection between different attributes such as model size and in-context examples, and their impact on dialogue evaluation performance. Additionally, we present a dialogue evaluation that leverages “chain-of-thought” (CoT) reasoning abilities of LLMs (Wei et al., 2022; Wang et al., 2023).

Our experiments demonstrate that larger model sizes and instruction tuning generally helps with zero shot dialogue evaluation. Furthermore, in few-shot scenario, we find that algorithmic selection of in-context examples yields better results than random selection. Next, we demonstrate that supervised fine-tuning can substantially improve the performance of LLMs on dialogue evaluation task. Finally, we explore and validate a CoT-based evaluation framework which is capable of returning not only dialogue labels but comprehensive explanations and justifications, thereby offering a more coherent and holistic evaluation. Remarkably, our results indicate that CoT-based evaluation is more

accurate when the LLM is prompted to first analyze the dialogue and then produce labels. Combined, our findings confirm the applicability and effectiveness of LLM-based automated dialogue evaluation.

2 Related work

Dialogue evaluation. Evaluating dialog systems poses challenges like accounting for multiple interlocutors, contextual dynamics, and the one-to-many relationship, as highlighted by Zhang et al. (2021) and Zhao et al. (2017). Metrics such as USR and FED address these challenges, showing strong correlation with human evaluation standards. Utilizing models like RoBERTa and DialoGPT, notable for their smaller yet effective versions fine-tuned for specific dialog tasks, these metrics excel in capturing nuanced dialog attributes. Other evaluation metrics such as GRADE and DEB (Huang et al., 2020; Mehri and Eskenazi, 2020a) attempt to measure text coherence, response diversity, engagement, and common sense. With the exponential growth in parameter count of contemporary LLMs and their promising generalization capabilities in NLP tasks, it is anticipated that these model-centric evaluative metrics will undergo further enhancements.

Large language models for evaluations. Recent studies explore LLMs in dialogue evaluation. GPTScore (Fu et al., 2023) uses models such as GPT-3 (Brown et al., 2020), assigning higher probabilities to superior-quality content and employing diverse prompts for holistic evaluation. Similarly, Huynh et al. (2023) investigate using ChatGPT and InstructGPT (Ouyang et al., 2022) for reference-independent text quality assessment, contrasting various LLM methodologies, including explicit scoring, leveraging model confidence for implicit score allocation, and direct pairwise text comparison.

The G-EVAL framework (Liu et al., 2023) is a notable advancement, synergistically integrating LLMs with the chain-of-thought (CoT) paradigm and a form-filling strategy. Notably, using GPT-4 (Bubeck et al., 2023) as its foundational model, G-EVAL shows strong correlation with human evaluations in summarization tasks.

Parameter-efficient fine-tuning (PEFT). As base language models grow in size (Touvron et al., 2023; Zhang et al., 2022), researchers frequently turn to parameter-efficient fine-tuning techniques to tailor models for specific downstream tasks. These fine-

tuning approaches typically fall into three main categories:

(1) **Prefix-Tuning:** This method inserts special tokens among input tokens with trainable embeddings for the task at hand (Li and Liang, 2021).

(2) **Adapter Tuning:** This approach inserts adapter layers between self-attention and MLP modules, providing nuanced control over the model’s behavior without altering the core architecture (Houlsby et al., 2019; Zhang et al., 2023).

(3) **Low-Rank Adaptation:** This technique uses trainable low-rank decomposition matrices in each network layer, simplifying the model for efficient fine-tuning (Hu et al., 2021). It shows promise in adapting large generative models for specific applications (Cuenca and Paul, 2023; Zhang et al., 2023).

These strategies reflect ongoing efforts to make large-scale models more adaptable and efficient, leveraging their vast capacities while mitigating computational and practical challenges.

In-context learning. In-context learning is a prompting technique for LLMs in which example input-output pairs for some task are injected into the prompt before the target input is presented. The idea is that seeing correct examples of the task will help the model to provide a correct target output.

Selecting in-context examples is crucial for effectively prompting LLMs, enabling them to pivot to new tasks without extensive fine-tuning. Examples play a pivotal role in guiding LLMs’ predictive capabilities, with research exploring methods such as semantic proximity evaluation (Liu et al., 2021) and retrieval mechanisms such as BM25 (Robertson et al., 2009), used independently or in an initial training phase for a selector retriever.

These selection approaches excel in few-shot NLP tasks. For instance, in Su et al. (2022), a bifurcated framework effectively annotates and selects in-context samples from unlabeled repositories, achieving impressive performance across various tasks. Similarly, Liu et al. (2021) suggested that choosing examples with congruent sentence embeddings optimizes GPT-3’s efficacy. Despite positive outcomes, there is a need for deeper explorations to discover more general in-context example retrieval methodologies.

3 Methodology

Dialogue evaluation with logits. LLMs like GPT (Radford et al., 2018) with a decoder-only architec-




Figure 1: Schematic overview of LLM dialogue evaluation methods. Left: Pipeline using logits method for generating scores from LLMs. Right: Pipeline employing generation method to produce ratings from LLMs.

ture are autoregressive. They generate sequences one element at a time, each conditioned on the preceding ones. The probability of a token sequence $x = (x_1, x_2, \dots, x_T)$ is modeled as the product of conditional probabilities for each token given its history.

$$p(x) = \prod_{t=1}^T p(x_t | x_{<t}) \quad (1)$$

where $x_{<t} = (x_1, \dots, x_{t-1})$ is the history before x_t , and $p(x_t | x_{<t})$ is typically modeled by a softmax over the vocabulary. LLMs can be prompted to provide a score (example prompt in section A.1), using the returned probabilities to generate ratings.

Building on methods from Huynh et al. (2023), we use these properties to select the top- K ratings r_1, r_2, \dots, r_K , based on their corresponding log probabilities p_1, \dots, p_K . We then perform a weighted sum of these ratings, as illustrated in the left panel of Figure 1. The weights can be calculated using Equation (2):

$$w_i = \frac{p_i}{\sum_{j=1}^K p_j} \quad (2)$$

The final rating can be represented by Equation (3):

$$r = \sum_{i=1}^K r_i * w_i \quad (3)$$

Dialogue evaluation with generation. In addition to the mentioned method, Lin and Chen (2023) proposed a novel framework where LLMs are prompted to directly generate responses for dialogue evaluation. Ratings for the dialogue can then be extracted from the produced LLM responses. See prompts in section A.2.

4 Experiment setup

Model. In this study, we utilize models from the Llama family (Touvron et al., 2023) and the Falcon series (Almazrouei et al., 2023). We also incorporate the instruction-tuning variants of these models as proposed in the Alpaca study (Taori et al., 2023). Temperature is fixed at 0.7 during generation.

Dataset. We experiment on two datasets: the publicly available USS (Sun et al., 2021) dataset includes SGD, MultiWOZ, ReDial, and CCPE subsets, with a 1-5 quality score scale. We randomly allocate 10% as test data, using the rest for training (supervised fine-tuning or in-context learning). We also evaluate our methods on two versions of an Amazon-internal dialogue-quality dataset, which has a human-annotated quality rating on a scale of 1-5 (similar in format to the data described in Komma et al. (2023)). The rating distribution is shown in Figure 2, with the smaller training set’s rating distribution resembling the test set more closely than the larger one.

We binarize the datasets, considering scores of three and below as “defect” (unsatisfactory) and scores of four or five as “non-defect” (satisfactory). This simplified scheme enables us to use standard binary classification metrics to compare the effectiveness of different methods.




Figure 2: Score distribution in train and test splits from the Amazon-internal dataset.

In-context example selection methods. Here we use three methods to select in-context examples. The first is to simply source examples randomly from the training dataset. The second employs a probabilistic algorithm for information retrieval (IR) (Robertson et al., 2009), selecting similar examples based on a defined similarity metric. The third relies on BERT (Devlin et al., 2018), extracting representations from the input dialogue and identifying similar examples using cosine similarity computations. Together, these three methods offer diverse strategies for in-context example selection. We conducted three runs for the random selection method and report the mean value.

Table 1: Performance comparison across model with different sizes on Amazon-internal datasets. Spearman and Pearson correlation values presented, with the best results highlighted in bold.

Models	Alpaca-7b	Alpaca-13b	Llama-7b	Llama-13b	Llama-30b	Falcon-7b-instruct	Falcon-40b-instruct
Spearman correlation	0.47	0.48	0.00	0.01	0.01	-0.02	0.41
Pearson correlation	0.47	0.48	0.03	0.02	0.03	-0.02	0.35

Supervised fine-tuning method. In this study, we also explore supervised fine-tuning to adapt LLMs to the task of dialogue quality evaluation. To manage cost, we adopt the LoRA setting (Hu et al., 2021), fine-tuning a relatively small number of parameters compared to full-rank fine tuning. LoRA fine-tuning efficiently enhances the model’s effectiveness at the target task while addressing computational challenges.

5 Results

We now describe the results of our experiments designed to analyze various aspects of LLM-based dialogue evaluation.

5.1 Larger models help zero-shot dialogue evaluation.

Table 1 shows the relationship between model size and zero-shot ability in dialogue quality evaluation, comparing Spearman and Pearson correlation values with human annotation for different model configurations across the Alpaca, Llama, and Falcon series.

Table 1 suggests a positive relationship between model size and zero-shot ability in dialogue quality evaluation. Notably, the Falcon series shows a significant improvement (Spearman: -0.02 to 0.41 for Falcon-7b versus Falcon-40b). Alpaca 13b sees marginal improvements compared to 7b, but the Llama series exhibits no significant improvement as size increases, with weak correlations across sizes (possibly due to the original Llama’s poor instruction understanding).

This conclusion underscores the potential benefits of employing larger models in dialogue systems, particularly for applications that require zero-shot flexibility.

5.2 Instruction-tuning helps zero-shot dialogue evaluation.

Table 1 also provides insight into the impact of instruction fine-tuning on models’ ability to do zero-shot dialogue quality evaluation. The Alpaca series underwent instruction fine-tuning based on the

foundational Llama models, an important modification to enhance task-specific performance. Our empirical results illustrate the superiority of the Alpaca models in the realm of zero-shot dialogue quality evaluation, as evidenced by consistently higher Spearman and Pearson correlation coefficients. This observation underscores the significance of instruction tuning as a critical technological approach for augmenting dialogue evaluation performance. For example, the improvement from Alpaca-13b compared to the Llama-13b models is 0.47 on the Spearman correlation shown in Table 1¹. We hypothesize that instruction tuning may serve to refine LLMs comprehension of instructions or prompts, thereby optimizing their ability to execute the specified tasks with greater accuracy. The results suggest that such tuning may be instrumental in facilitating a more nuanced understanding of dialogues, opening avenues for further research and development in this domain.

5.3 In-context examples enhance the performance of dialogue evaluation.

In this section, we evaluate the influence of in-context examples on the base model Falcon-7b-instruct, utilizing three distinct in-context example selection methods: Random, BM25, and BERT. For the BM25 and BERT-based ICL approaches, we selected either 1 or 4 of the most semantically similar samples from the training set. The selection was based on the semantic similarity between the examples, as determined by the BM25 and BERT models. In contrast, for the random-selection ICL approach, we randomly picked either 1 or 4 examples from the training set to use as the in-context examples. These experiments are conducted on the open-sourced dataset USS shown in Table 2.

First, we observe that in general, the few-shot performance is better than the zero-shot performance. A closer examination reveals that the best performance across various evaluation metrics is primarily concentrated within the few-shot settings in the first three datasets. If we take a closer look at the MWOZ dataset, all highlighted numbers are in the few-shots settings. This observation substantiates the notion that providing in-context examples can indeed enhance the performance in dialogue evaluation. However, it is worth noting that an excessive provision of examples does not necessar-

¹Llama-7b having a Spearman correlation value of 0 is actually rounded value of a tiny correlation of < 0.005.

Table 2: Comparison of different in-context example selection methods and zero-shot ability for dialogue quality evaluation across different datasets using Falcon-7b-instruct Model on open-source dataset.

Dataset	Model	Defect Rate	Defect Class			Non-Defect Class			Weighted Average			Macro Average			Spearman	Pearson
			Precision	Recall	F1-Score	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score		
CCPE	Zero-shot	62%	0.71	0.39	0.5	0.42	0.74	0.54	0.6	0.52	0.51	0.57	0.56	0.52	0.21	0.27
	Random-1	62%	0.75	0.19	0.31	0.4	0.89	0.56	0.62	0.46	0.4	0.58	0.54	0.43	0.44	0.25
	BM25-1	62%	0.79	0.35	0.49	0.44	0.84	0.58	0.66	0.54	0.52	0.62	0.6	0.54	0.4	0.31
	BERT-1	62%	0.73	0.26	0.38	0.41	0.84	0.55	0.61	0.48	0.45	0.57	0.55	0.47	0.18	0.19
	Random-4	62%	0.71	0.55	0.62	0.46	0.63	0.53	0.61	0.58	0.59	0.58	0.59	0.58	0.16	0.17
	BM25-4	62%	0.73	0.77	0.75	0.59	0.53	0.56	0.67	0.68	0.68	0.66	0.65	0.65	0.23	0.2
	BERT-4	62%	0.68	0.68	0.68	0.47	0.47	0.47	0.6	0.6	0.6	0.58	0.58	0.58	-0.07	-0.13
MWOZ	Zero-shot	55%	0.5	0.29	0.37	0.43	0.64	0.51	0.47	0.45	0.43	0.46	0.47	0.44	-0.05	0.0
	Random-1	55%	0.55	0.11	0.18	0.45	0.89	0.6	0.5	0.46	0.37	0.5	0.5	0.39	0.09	0.1
	BM25-1	55%	0.53	0.15	0.23	0.45	0.84	0.58	0.49	0.46	0.39	0.49	0.49	0.41	0.18	0.06
	BERT-1	55%	0.63	0.18	0.28	0.46	0.87	0.6	0.55	0.49	0.43	0.54	0.52	0.44	0.13	0.0
	Random-4	55%	0.53	0.35	0.42	0.44	0.62	0.51	0.49	0.47	0.46	0.48	0.48	0.47	0.09	0.08
	BM25-4	55%	0.61	0.6	0.61	0.52	0.53	0.53	0.57	0.57	0.57	0.57	0.57	0.57	0.08	0.04
	BERT-4	55%	0.56	0.58	0.57	0.47	0.44	0.45	0.52	0.52	0.52	0.51	0.51	0.51	0.06	0.05
Redial	Zero-shot	44%	0.49	0.55	0.52	0.61	0.55	0.58	0.56	0.55	0.55	0.55	0.55	0.55	0.09	0.11
	Random-1	44%	0.5	0.2	0.29	0.57	0.84	0.68	0.54	0.56	0.51	0.54	0.52	0.49	0.18	0.24
	BM25-1	44%	0.48	0.27	0.35	0.57	0.77	0.66	0.53	0.55	0.52	0.53	0.52	0.5	0.2	0.17
	BERT-1	44%	0.58	0.32	0.41	0.61	0.82	0.7	0.6	0.6	0.57	0.59	0.57	0.55	0.24	0.29
	Random-4	44%	0.5	0.41	0.45	0.59	0.68	0.63	0.55	0.56	0.55	0.55	0.54	0.54	0.12	0.11
	BM25-4	44%	0.4	0.45	0.43	0.52	0.46	0.49	0.47	0.46	0.46	0.46	0.46	0.46	0.1	0.15
	BERT-4	44%	0.44	0.55	0.49	0.57	0.46	0.51	0.51	0.5	0.5	0.5	0.5	0.5	0.05	0.05
SGD	Zero-shot	48%	0.64	0.63	0.63	0.66	0.67	0.67	0.65	0.65	0.65	0.65	0.65	0.65	0.34	0.32
	Random-1	48%	0.40	0.08	0.14	0.51	0.88	0.65	0.46	0.50	0.40	0.46	0.48	0.39	0.20	0.09
	BM25-1	48%	0.39	0.15	0.21	0.50	0.79	0.61	0.45	0.48	0.42	0.44	0.47	0.41	0.29	0.15
	BERT-1	48%	0.50	0.15	0.23	0.52	0.87	0.65	0.51	0.52	0.45	0.51	0.51	0.44	0.33	0.26
	Random-4	48%	0.44	0.33	0.38	0.50	0.62	0.55	0.47	0.48	0.47	0.47	0.47	0.47	-0.01	-0.02
	BM25-4	48%	0.52	0.58	0.55	0.57	0.50	0.53	0.54	0.54	0.54	0.54	0.54	0.54	0.06	0.08
	BERT-4	48%	0.37	0.40	0.38	0.41	0.38	0.40	0.39	0.39	0.39	0.39	0.39	0.39	-0.12	-0.12

ily lead to further improvement. We hypothesize that this limitation may stem from the capacity constraints of LLMs, which can struggle to process overly lengthy inputs, occasionally resulting in performance degradation. This pattern is particularly evident in the Redial dataset, where the performance of the 4-shot approach does not surpass the results obtained from the 1-shot experiment. In addition to these findings, Table 2 reveals that zero-shot outperforms the few-shot settings in the SGD dataset. This is likely attributable to the capacity constraints of LLMs; the dialogue lengths should not be excessively long. Notably, the dialogues in the SGD dataset have more turns (26.7 turns per dialogue) compared to the other three datasets from Sun et al. (2021). Such findings further emphasize the nuanced relationship between the number of in-context examples and the resulting performance, highlighting the importance of careful selection in few-shot learning.

Second, when comparing the performance across different in-context selection methods, we find that algorithmic selection methods result in notable performance improvements over random selection for in-context examples. For instance, the BM25 and BERT methods consistently perform best across all datasets. Upon closer examination of the CCPE dataset in Table 2, we observe that 85% of the highest values across all metrics are derived from the algorithm’s selected method. What is more, the optimal choice for selecting different in-context example methods varies based on the dataset. In the

first two datasets, the BM25 method excels, while in the third, the BERT method stands out.

In conclusion, our results suggest that in-context examples can significantly enhance the quality of dialogue evaluation. Our findings also underscore the importance of employing algorithmic methods for selecting these examples, as the right selection strategy can lead to meaningful performance gains. By revealing these patterns, our study contributes to a deeper understanding of how few-shot learning can be best utilized in dialogue systems.

Table 3: Summary of performance metrics for supervised finetuning models on Amazon-internal datasets across different training datasets and various model architectures.

Models	Spearman	Pearson	Precision	Recall	F1	F1-micro
Alpaca-7b	0.47	0.47	0.52	0.72	0.60	0.69
Alpaca-7b-sft-small	0.61	0.61	0.96	0.59	0.73	0.65
Alpaca-7b-sft-large	0.64	0.66	0.93	0.58	0.72	0.33
Llama-7b	0.00	0.03	0.67	1.00	0.80	0.33
Llama-7b-sft-small	0.58	0.60	0.92	0.62	0.74	0.58
Llama-7b-sft-large	0.64	0.66	0.93	0.58	0.72	0.33
Llama-13b	0.01	0.02	0.67	1.00	0.80	0.36
Llama-13b-sft-small	0.64	0.65	0.94	0.61	0.74	0.46
Llama-13b-sft-large	0.64	0.65	0.97	0.54	0.69	0.40
Falcon-40b-instruct	0.41	0.36	1.00	0.01	0.02	0.67
Falcon-40b-instruct-sft-small	0.60	0.61	0.96	0.52	0.67	0.65
Falcon-40b-instruct-sft-large	0.61	0.63	0.93	0.53	0.68	0.69

5.4 Supervised fine-tuning improves dialogue evaluation quality

Here we examine the influence of supervised fine-tuning (SFT) on the performance of LLMs for the dialogue evaluation task. Specifically, we fine-tuned the models using both Likert-scale and binary label data. We leverage two internal datasets

Table 4: Comparison of CoT methods on the internal dataset over Falcon-7b-instruct. “Rating-first” refers to generating the score first, followed by the reasons, while “Analysis-first” involves generating an analysis first, then determining the scores.

Model	Defect Rate	Defect Class			Non-Defect Class			Weighted Average			Macro Average			Spearman	Pearson
		Precision	Recall	F1-Score	Precision	Recall	F1-score	Precision	Recall	F1-score	Precision	Recall	F1-score		
Rating-first	59%	0.59	1.00	0.74	0.67	0.01	0.02	0.62	0.59	0.45	0.63	0.50	0.38	0.10	0.07
Analysis-first	59%	0.67	0.86	0.75	0.66	0.39	0.49	0.67	0.67	0.65	0.67	0.63	0.62	0.23	0.28

used for training dialogue quality estimation models. The first is a small version of the dataset, resulting in models denoted as “xxx-small,” while the second is a larger version, leading to models labeled as “xxx-large.”

We report on several classification metrics, including precision, recall, F1-score, and F1-micro. Our findings from Table 3 reveal a consistent improvement in the original model’s performance after SFT, especially in the Spearman and Pearson correlations with human annotation. For instance, the model “Falcon-40b-instruct-sft-large” exhibits a 48% relative improvement compared to the original “Falcon-40b-instruct.” This indicates that SFT enhances the alignment of the model’s scoring with human evaluation.

Analysis of F1-micro shows that SFT generally improves performance in comparison to the original model. When comparing the Llama-7b-sft-small to the original Llama-7b, there is a 75% relative improvement. A closer examination also suggests that utilizing a larger dataset can boost overall performance. For example, the highest correlations in the first two columns are consistently associated with models trained on the larger dataset. Interestingly, models trained on smaller datasets occasionally exhibit superior F1-micro scores, as observed for the Llama series. For example, when comparing the Llama-13b-sft-small to the original Llama-13b-sft-large, there is a 15% relative improvement. We hypothesize that this may occur when the score distribution between the small dataset and the test dataset aligns more closely shown in Figure 2.

It is important to note that Llama-7b and Llama-13b predict defects for all test samples, which results in a recall of 1.0 and a precision of 0.67—this obviously does not indicate optimal performance. To better understand the relationship between model predictions and the ground truth, we should consider the Spearman and Pearson correlations over likert scores, which provide more insight into the linear relationship between the predictions and the human labels.

In conclusion, our findings show that SFT can substantially enhance the performance of LLMS on dialogue evaluation. This study underscores the

value of fine-tuning and dataset selection in achieving more accurate and human-aligned evaluations in the context of dialogue systems.

5.5 Chain-of-thought for generation of scores and reasons

In this section, we shift our focus to the “generation and chain-of-thoughts” approach (Wei et al., 2022), which not only generates scores but also provides natural language *reasons* for selecting those scores. We explore two distinct paradigms to accomplish this task. The first paradigm, **Analysis-first**, involves prompting the model to generate an analysis first and then derive ratings based on that analysis. The second paradigm, **Rating-first**, prompts the model to generate a rating first and then elucidate the reasons for choosing that particular score (see Section A.2 for respective prompts).

Interestingly, our findings suggest that the first paradigm—Analysis-first—provides more aligned scores and reasons, as shown in Table 4. We observe consistent improvement compared to the Rating-first approach. For example, Analysis-first methods outperform Rating-first in 85% of all evaluation metrics. Upon conducting a failure analysis, we discovered that for Rating-first, the scores do not always align with the subsequent reasons. However, in the Analysis-first paradigm, there is consistent alignment between ratings and scores. We attribute the observed metric improvement to this consistency.

The implications of these findings may extend to various applications where the alignment between scores and reasoning is essential. Further exploration of these paradigms and their potential advantages and limitations may provide valuable insights into the optimal utilization of LLMs for complex tasks such as rating and explanation generation.

6 Conclusion

This paper explores the application of LLMs to evaluation of task-oriented dialogue systems. Key findings include: the impact of pretrained model size; the importance of instruction fine-tuning; the effectiveness of in-context examples; consistent performance improvement through supervised fine-

tuning; and that the chain-of-thought paradigm is most effective with the Analysis-first approach.

7 Limitations

Our experiments provide valuable insights, but there are limitations. We focus on open-sourced models, excluding closed ones like ChatGPT and Claude. Evaluation primarily centers on user satisfaction, lacking metrics for interestingness and coherence. Performance is influenced by prompt designs, and suboptimal prompts may lead to decline.

8 Ethics Statement

We acknowledge ethical concerns in using LLMs in our evaluation. Firstly, LLMs may carry biases that could impact dialogue evaluation negatively. Secondly, our focus on user satisfaction might overlook issues like toxic responses, leading to inadequate evaluations. Lastly, concerns arise about unintentional release of private information during reason and rating generation. To address these concerns, researchers should exercise caution in using LLMs for reasons and ratings, ensuring accuracy and fairness in interpretation.

References

- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- 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.
- Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. 2023. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*.
- Pedro Cuenca and Sayak Paul. 2023. Using lora for efficient stable diffusion fine-tuning.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jinlan Fu, See-Kiong Ng, Zhengbao Jiang, and Pengfei Liu. 2023. Gptscore: Evaluate as you desire. *arXiv preprint arXiv:2302.04166*.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020. Dialogue response ranking training with large-scale human feedback data. *arXiv preprint arXiv:2009.06978*.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Lishan Huang, Zheng Ye, Jinghui Qin, Liang Lin, and Xiaodan Liang. 2020. Grade: Automatic graph-enhanced coherence metric for evaluating open-domain dialogue systems. *arXiv preprint arXiv:2010.03994*.
- Jessica Huynh, Cathy Jiao, Prakhar Gupta, Shikib Mehri, Payal Bajaj, Vishrav Chaudhary, and Maxine Eskenazi. 2023. Understanding the effectiveness of very large language models on dialog evaluation. *arXiv preprint arXiv:2301.12004*.
- Abishek Komma, Nagesh Panyam Chandrasekarasastri, Timothy Leffel, Anuj Goyal, Angeliki Metallinou, Spyros Matsoukas, and Aram Galstyan. 2023. Toward more accurate and generalizable evaluation metrics for task-oriented dialogs. In *The 61st Annual Meeting Of The Association For Computational Linguistics (Industry Track)*.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yen-Ting Lin and Yun-Nung Chen. 2023. Llm-eval: Unified multi-dimensional automatic evaluation for open-domain conversations with large language models. *arXiv preprint arXiv:2305.13711*.
- Chia-Wei Liu, Ryan Lowe, Iulian V Serban, Michael Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How not to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. *arXiv preprint arXiv:1603.08023*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.

- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruo Chen Xu, and Chenguang Zhu. 2023. Gpteval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Shikib Mehri and Maxine Eskenazi. 2020a. Unsupervised evaluation of interactive dialog with dialogpt. *arXiv preprint arXiv:2006.12719*.
- Shikib Mehri and Maxine Eskenazi. 2020b. **USR: An unsupervised and reference free evaluation metric for dialog generation**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 681–707, Online. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Hongjin Su, Jungo Kasai, Chen Henry Wu, Weijia Shi, Tianlu Wang, Jiayi Xin, Rui Zhang, Mari Ostendorf, Luke Zettlemoyer, Noah A Smith, et al. 2022. Selective annotation makes language models better few-shot learners. *arXiv preprint arXiv:2209.01975*.
- Weiwei Sun, Shuo Zhang, Krisztian Balog, Zhaochun Ren, Pengjie Ren, Zhumin Chen, and Maarten de Rijke. 2021. Simulating user satisfaction for the evaluation of task-oriented dialogue systems. In *Proceedings of the 44rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*. ACM.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen, You Wu, Luke Zettlemoyer, and Huan Sun. 2023. Towards understanding chain-of-thought prompting: An empirical study of what matters. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2717–2739.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. **Chain-of-thought prompting elicits reasoning in large language models**. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Zhenyu Zhang, Tao Guo, and Meng Chen. 2021. Dialoguebert: A self-supervised learning based dialogue pre-training encoder. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 3647–3651.
- Tiancheng Zhao, Ran Zhao, and Maxine Eskenazi. 2017. Learning discourse-level diversity for neural dialog models using conditional variational autoencoders. *arXiv preprint arXiv:1703.10960*.

A Prompts

In this section, we elaborate on the prompts and instructions we used in the logits-based and generation-based evaluation methods.

A.1 Prompts for logits method

The prompts we used in the logits method are formulated in the following manner:

Instruction: Could you please evaluate the subsequent dialogue by assigning a score from the given set [1,2,3,4,5]? A score of 1 implies dissatisfaction, while a 5 signifies high satisfaction.

A.2 Prompts for generation method

Rating-first In this section, we outline the prompts utilized in the Rating-First generation method. This process begins with an instruction that directs the Large Language Models (LLMs) to provide both reasons and ratings for a given dialogue. To enhance the LLMs' comprehension, we also supply evaluation criteria standards. Finally, we detail the evaluation steps necessary to complete the entire procedure for dialogue evaluation. The prompt is shown as follows:

Instruction: Could you please evaluate the subsequent dialogue, providing a score from the set [1,2,3,4,5], and give an explanation for choosing that score?

To help you better evaluate, here is the evaluation Criteria:

A score of 1 means very dissatisfied, where the user repeatedly has to stop or cancel bad responses and repeat their request again;

A score of 2 means dissatisfied, where None of the user goals are achieved, the user expresses negative feedback, steps towards a user goal succeeds but the goal fails;

A score of 3 means normal, where At least one of a user goals succeed, and no negative feedback

A score of 4 means satisfied, where the majority of turns succeeded or moved the user closer to their goal

A score of 5 means very satisfied, all turns either succeeded or moved the user closer to their goal(s), and the user expressed no dissatisfaction, and goal was achieved without unnecessary steps

Steps to conduct the evaluation are:

1. Read the dialog, and the response carefully
2. Rate the response on a scale of 1-5 for satisfaction level from user, according to the criteria above
3. Provide a brief explanation for your rating, referring to specific aspects of the response and the dialog.

Analysis-first In this section, we describe the prompts used in the Analysis-First generation method. Similar to the Rating-First approach, we

begin by providing instructions for the LLMs to carry out dialogue evaluation. We then present specific aspects for the LLMs to analyze. To facilitate better understanding, we also supply the evaluation criteria. Finally, we detail the specific steps for the LLMs to follow.

Instruction: Could you please evaluate the subsequent dialogue overall quality, first analysis the dialogue from the following aspects:

1. User goal.
2. User feedback.
3. System response.
4. System feedback.

Based the above analysis provide a user satisfactory score from the set [1,2,3,4,5]

To help you better evaluate, here is the evaluation Criteria:

A score of 1 means very dissatisfied, where the user repeatedly has to stop or cancel bad responses and repeat their request again;

A score of 2 means dissatisfied, where None of the user goals are achieved, the user expresses negative feedback, steps towards a user goal succeeds but the goal fails;

A score of 3 means normal, where At least one of a user goals succeed, and no negative feedback

A score of 4 means satisfied, where the majority of turns succeeded or moved the user closer to their goal

A score of 5 means very satisfied, all turns either succeeded or moved the user closer to their goal(s), and the user expressed no dissatisfaction, and goal was achieved without unnecessary steps

Steps to conduct the evaluation are:

1. Read the dialog, and the response carefully
2. Give some brief analysis from the aspects mentioned before
3. Rate the response on a scale of 1-5 for satisfaction level from user, according to the criteria above and the analysis.