

Fostering Natural Conversation in Large Language Models with NICO: a Natural Interactive CONversation dataset

Anonymous ACL submission

Abstract

Benefiting from diverse instruction datasets, contemporary Large Language Models (LLMs) perform effectively as AI assistants in collaborating with humans. However, LLMs still struggle to generate natural and colloquial responses in real-world applications such as chatbots and psychological counseling that require more human-like interactions. To address these limitations, we introduce NICO, a Natural Interactive CONversation dataset in Chinese. We first use GPT-4-turbo to generate dialogue drafts and make them cover 20 daily-life topics and 5 types of social interactions. Then, we hire workers to revise these dialogues to ensure that they are free of grammatical errors and unnatural utterances. We define two dialogue-level natural conversation tasks and two sentence-level tasks for identifying and rewriting unnatural sentences. Multiple open-source and closed-source LLMs are tested and analyzed in detail. The experimental results highlight the challenge of the tasks and demonstrate how NICO can help foster the natural dialogue capabilities of LLMs. The dataset will be released.

1 Introduction

The development of large language models (LLMs) has received much attention over the past few years, demonstrating their potent potential across numerous fields (Zhao et al., 2023). In particular, in the field of dialogue, these models have been widely used in areas such as chatbots (Zheng et al., 2023) and virtual assistants (Liu et al., 2023).

Many researchers have sought to enhance the performance of LLMs by designing instruction datasets. Existing datasets for fine-tuning LLMs can be categorized into those that improve instruction-following abilities (He et al., 2024), planning abilities (Dagan et al., 2023), reasoning abilities (Wang et al., 2023), etc. These datasets have indeed improved the performance of LLMs on benchmarks such as MMLU (Hendrycks et al.,

2020) and MMLU-Pro (Wang et al., 2024). However, when engaging in dialogue with humans, the responses generated by LLMs are often unnatural and insufficiently colloquial (Guo et al., 2023; Yi et al., 2024). LLMs can only act as AI assistants rather than human beings. It has led to the failure of the LLMs for applications such as psychological counseling and accompanying children. Jones and Bergen (2024) showed that the success rate of the best-performing GPT-4 in acting as a human was only 50%. The main reasons for the failure included unnatural and too informal responses, weird tones, and so on.

To address these issues, we are committed to constructing a multi-turn daily dialogue dataset in Chinese. We aim to evaluate the ability of LLMs to simulate natural human conversations in daily life, and improve their abilities by fine-tuning with the constructed dataset.

We ensure that the dialogues in this dataset are related to daily life and conform to human speaking habits. We first define 20 broad topics, which are basically involved and discussed in human daily life. For each topic, we construct 200 specific concepts or items. We also draw inspiration from the five types of interactions in sociology. To make the constructed dialogs resemble real human conversations, we ensure that the dataset covers all interaction types. Then, we utilize GPT-4-turbo to generate 4,000 dialogs, which contains 22.1 utterances per dialog on average. We name the constructed dataset NICO. We also perform a human evaluation, and the results show that NICO is better than the current representative datasets in fluency, coherence and naturalness.

Subsequently, we define four tasks, two of which are used to test the daily conversation capabilities of LLMs. The other two are used to test the capabilities to identify and rewrite unnatural sentences. We have selected several representative LLMs for testing and also conducted supervised fine-tuning.

The results indicate that mimicking human-like daily conversation poses a challenge for LLMs. The dialogs we constructed can enhance the performance of LLMs. At the sentence level, LLMs can rewrite unnatural sentences but are almost incapable of identifying them. Further analysis of the results reveals that LLMs perform much worse in conflict-type dialogues than in other types of dialogues, with limited improvement from fine-tuning.

In summary, enabling LLMs to mimic human-like fluent and natural daily conversations remains a challenge. The contributions of our paper are as follows:

- We construct a multi-turn natural interactive conversation dataset in Chinese named NICO, covering 20 daily topics and 5 interaction types. Human annotation ensures that NICO is free of grammatical errors and unnatural utterances.
- We define four tasks, including two dialogue-level tasks to test the natural dialog abilities of LLMs and two sentence-level tasks of identifying and rewriting unnatural sentences. We test several representative LLMs on the four tasks and report their results.
- We conduct a detailed analysis of the results, pointing out the deficiencies of existing LLMs in conducting natural conversations and identifying unnatural sentences. We validate that NICO can help LLMs to generate responses that are more human-like rather than resembling AI assistants.

2 Related Work

In recent years, open-domain dialog systems have witnessed tremendous progress as the number of model parameters keeps growing. Representative dialog systems include ChatGPT (Achiam et al., 2023), Gemini (Reid et al., 2024), Mixtral-Instruct (Jiang et al., 2024), LLama3-Instruct (Meta, 2024), and so on (Cai et al., 2024; Gan et al., 2023; Wei et al., 2023; Yang et al., 2023a).

Many scholars have been working on constructing dialog data to help dialog systems evolve. Some datasets focus on human emotions. DailyDialog (Li et al., 2017) is an English conversation dataset containing various daily conversation scenarios. This dataset especially emphasizes the communication intention and emotion information of the conversations. Persona-chat (Zhang et al., 2018a)

is an English conversation dataset that focuses on daily communication. It makes conversations more like human-to-human communication by establishing background information about different users. Other datasets focus on strategy and behavior. The Ubuntu Dialogue Corpus (Lowe et al., 2015) is a multi-turn dialog dataset in English. It is collected from technical support conversations on the Ubuntu forums. MultiWOZ (Budzianowski et al., 2018) is also a large-scale English conversation dataset. Each of its conversations is labeled with a sequence of conversation states and the respective system dialog behavior.

Apart from the above English conversation datasets, many researchers working on proposing high-quality Chinese conversation datasets and developing Chinese conversation systems. Some Chinese conversation datasets such as Douban Conversation Corpus (Wu et al., 2017), LCCC (Wang et al., 2020) are mined and cleaned from Chinese social media. Some datasets such as E-commerce Dialogue Corpus (Zhang et al., 2018b), JDDC corpus (Chen et al., 2020) are constructed from e-commerce service stuff and customer conversation data. There are also datasets such as DuConv (Wu et al., 2019), Mutual (Cui et al., 2020), KdConv (Zhou et al., 2020), Corecode (Shi et al., 2024) annotated using knowledge bases including news and exams.

Despite the above efforts, existing dialog systems are still considered to be inconsistent with human speaking habits (Jones and Bergen, 2024). Natural dialogue datasets, especially those in Chinese, are still scarce. High-quality Chinese dialog datasets are essential to advance the progress of Chinese dialog systems.

3 The NICO Dataset

We have built a Natural Interactive CONversation dataset named NICO based on daily topics and the social interactions. In this section, we will detail how we build and annotate the dialogues, as well as ensure the quality of the data.

3.1 Data Construction

We aim to construct dialogues that are as relevant as possible to real life. We start by defining 20 broad topics, including culture, technology, and so on. These topics are fundamentally ones that humans frequently engage with and discuss in their daily lives. The complete broad topics are shown

in Table 1. Then, we give prompts for GPT-4-turbo (Achiam et al., 2023) to generate 200 specific items or concepts under each broad concept, e.g., sweaters, artificial intelligence, etc. We carefully check and modify these generations to ensure their alignment with the broad topics.

Topics	动物(Animal), 艺术和手工艺品(Art and Craft Supplies), 图书和文具(Books and Stationery), 衣服(Clothing), 文化(Culture), 教育和学习(Education and Learning), 金融产品和服务(Financial Products and Services), 饮食(Food and Drinking), 杂货和家庭必需品(Groceries and Household Essentials), 健康和美容(Health and Beauty), 历史和事件(History and Events), 音乐和乐器(Music and Instruments), 植物(Plant), 不动产(Real Estate), 体育运动(Sports), 科技(Technology), 玩具和游戏(Toys and Games), 交通运输(Transportation), 旅行(Travel), 工作和职业(Working and Occupation)
--------	---

Table 1: The 20 defined broad topics.

In sociological terms, the process of communication between members of a society is known as social interaction. It describes the way people behave when they cross paths with others. Social interaction is consistent with the definition of behavioral intelligence proposed by Li et al. (2024), and has significant ramifications for human-AI interactions. We argue that high-quality daily dialog datasets need to cover all types of social interactions. There are many types of interactions between humans, and in the paper we accommodate the widely accepted types of social interactions proposed by Nisbet (1970).

Five modes or types of social interaction are probably the most basic and universal ones: exchange, cooperation, conformity, coercion, and conflict.

— (Nisbet, 1970), p.63

We construct dialogs based on the five interaction types and give the definition of each interaction type as follows:

- **Exchange** refers to the process in which an individual or a group takes action with the explicit purpose of receiving a return.
- **Cooperation** is a collaborative behavior to achieve a certain goal of common interest.
- **Conformity** is the process in which an individual’s behavior adapts to a pre-existing norm, in which the influence of the social group is either evident or subtle.

- **Coercion** is a behavior that is produced by compulsion of any form, often accompanied by threats.
- **Conflict** refers to the process in which individuals struggle with each other to compete for a scarce, commonly desired return.

Then, we write three example dialogues for each interaction type. Benefiting from the powerful GPT-4-turbo, we provide the definitions of the interaction types, the corresponding dialogues, and the requirements for generating colloquial dialogues as prompts to GPT-4-turbo¹, which is asked to generate daily dialogues about specific concepts. A prompt example is shown in Figure 2. For each interaction type, we generate 800 dialogs for a total of 4,000. These dialogs also entirely cover the 4,000 generated specific items and concepts.

3.2 Annotation Process

Despite the strong instruction-following and text-generating capabilities of the LLMs, they still struggle with problems such as hallucinations (Yang et al., 2023b) and formatting errors (Sui et al., 2024). Therefore, we have employed native Chinese-speaking workers to check and correct the data generated by the LLM. We conduct three stages of annotation.

In the first stage, we ask workers to check whether the generated dialogs match the corresponding interaction types. If discrepancies are found, such as agreement and cooperation in a conflict-type dialogue, the worker will revise the dialogue directly to ensure that it is consistent with the intended interaction type. In the second stage, workers annotate sentences in the dialog that contain grammatical errors, sentences that do not conform to human speech i.e., are unnatural, and blocks of sentences that do not convey information i.e., are meaningless. In the third stage, workers rewrite sentences that contain grammatical errors and unnatural sentences. We find that meaningless sentences mainly appear at the end of coercion-type dialogs and are not numerous. Therefore, we delete them directly, which will not affect the content of the dialog.

After the annotation is completed, we obtain 4,000 natural dialogs related to daily life. We further partition these dialogs into training, validation, and test sets, with 3,600, 200, and 200 dialogs respectively, for subsequent experiments.

¹The version is gpt-4-turbo-2024-04-09.

4 Statistics

In this section, we have performed statistics on the features of the NICO dataset. We have also conducted human evaluations and compared the NICO dataset with other conversation datasets.

4.1 Dataset Statistics

The statistics of the constructed dataset are shown in Table 2. Our dataset comprises 4000 dialogues and 88,545 utterances. On average, there are 22.1 utterances per dialog, suggesting sufficient rounds of interaction between the dialog’s parties. The average number of tokens per utterance is 24.7, indicating that the dialog conveys adequate information. We also count the annotations and find that more than one-third of the dialogs contain grammatical errors, and more than half of the dialogs contain unnatural utterances. This suggests that the dialogues generated by LLMs are not always reliable and require manual correction to enhance their quality.

# Dialogs	4,000
# Utterances	88,545
# Topics	20
Avg. # utterances per dialog	22.1
Avg. # tokens per dialog	546.5
Avg. # tokens per utterance	24.7
# Dialogs with grammatical errors	1,440
# Utterances with grammatical errors	2,017
# Dialogs with unnatural utterances	2,050
# Unnatural utterances	2,872
# Dialogs with meaningless utterances	418
# Meaningless utterance blocks	418

Table 2: Statistics of the NICO dataset.

4.2 Comparisons

To evaluate the quality of the NICO dataset, we select several representative conversation datasets for comparison: (1) NaturalConv (Wang et al., 2021), (2) KdConv (Zhou et al., 2020), (3) LCCC (Wang et al., 2020), (4) DailyDialog (Li et al., 2017). The first three datasets are high-quality Chinese dialog datasets. NaturalConv and KdConv have been manually checked, while LCCC has been automatically filtered using a set of rules. DailyDialog is a human-written English dialog dataset, so we use GPT-4-turbo for translation. We randomly selected 100 dialogs from each of the above datasets and the NICO dataset we built for evaluation.

We have also performed a human evaluation. Following (Zhou et al., 2020), we choose fluency and coherence as evaluation metrics. However, to emphasize the importance of naturalness in conversation, we only use fluency to measure grammatical correctness, and use naturalness as a separate evaluation metric following (Wang et al., 2021). The definitions are as follows and more details are in Appendix A.

Fluency (rating scale is 0,1,2) is defined as whether the dialog contains syntax errors.

Coherence (rating scale is 0,1,2) is defined as whether the sentences in a dialog are relevant to the context and logically correct.

Naturalness (rating scale is 0,1,2) is defined as whether the dialog is natural and conforms to human speaking habits.

We hire three native Chinese-speaking workers for the evaluation. To ensure the quality of the evaluation, we train them before the evaluation. The evaluation results are shown in Table 3.

Datasets	Fluency	Coherence	Naturalness
NaturalConv	1.68	1.95	1.70
KdConv	1.86	1.96	1.47
LCCC	1.33	0.95	1.03
DailyDialog	1.88	1.90	1.44
NICO	1.98	2.00	1.90

Table 3: The human evaluation results of the NICO dataset versus other representative dialogue datasets.

Overall, the NICO dataset performs best on all three metrics, especially scoring well above other datasets in fluency and naturalness. In the development and evaluation of dialog systems, using the NICO dataset may yield better results. LCCC without human filtering has the lowest scores, which suggests that human annotation is essential for constructing high-quality dialog datasets. DailyDialog has good fluency and coherence after translation, but the dialogs have obvious signs of machine translation. NaturalConv scores higher in naturalness than the other datasets, but still lower than NICO.

We also perform statistics on the types of social interactions contained in these datasets. A high-quality dialog dataset should cover all five types of social interactions that occur in real, daily conversations. We let the workers choose the type of social interaction they believe is most closely represented in the dialogues. The results are depicted in Figure 1.

The most prevalent type of interaction across

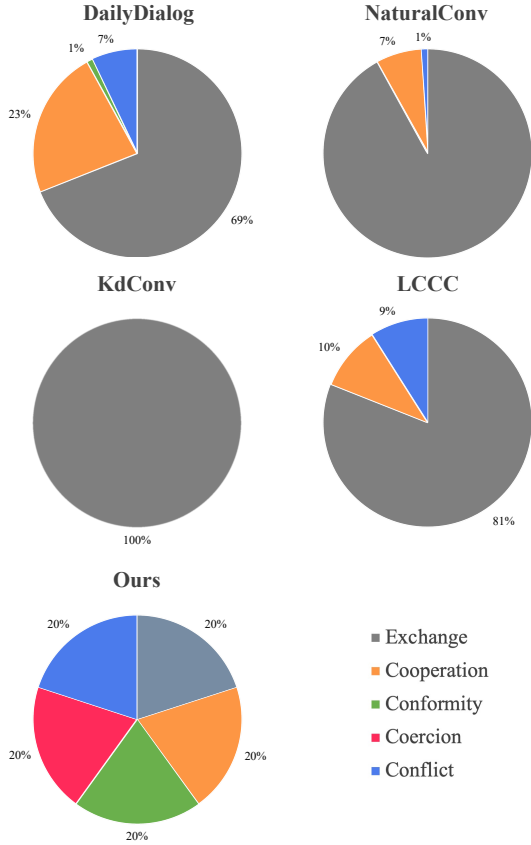


Figure 1: Distribution of interaction types of the dialog datasets. Only the constructed NICO dataset is able to cover all types of social interactions.

several comparison datasets is exchange. Both DailyDialog and LCCC contain a small number of cooperative and conflict dialogs. Almost none of the comparison datasets contain conformal dialog or coercive dialog. The constructed NICO dataset covers all social interaction types, thus aligning more closely with human communication patterns.

5 Experiments

5.1 Tasks

We have defined four tasks and tested LLMs on the constructed dataset. The first two tasks are designed to assess the abilities of LLMs to engage in natural dialogue in daily life scenarios. Benefiting from the process of annotating the dataset, we have the original dialog, the annotations of unnatural sentences in the dialog, and the corresponding human rewrites. We utilize these data for experiments on the last two tasks: identifying unnatural sentences and rewriting unnatural sentences.

Open-domain dialog This task is to let LLMs simulate a real human conversation. In this task, the

user seems to be talking with a real person with multi-turn dialog history. It requires LLMs to have a broad knowledge base and strong contextual understanding ability, and to generate natural and fluent responses.

Dialog completion This task requires LLMs to generate a reasonable and natural response based on the instruction. Unlike the open-domain dialog task, the multi-turn dialog appears in the instruction. The social interaction type with definition of the dialog is also given in the instruction.

Identifying unnatural sentences This task involves enabling LLMs to identify sentences that do not conform to human speech patterns. The detected sentences can then be rewritten.

Rewriting unnatural sentences This task is to have the LLMs rewrite sentences that do not conform to human speech patterns, making them more natural and fluent. It can enhance the readability of the text and the efficiency of human interaction.

Prompts for the above tasks are shown in Appendix B.

5.2 Models

We choose several recent representative LLMs for experiments: (1) open-source LLMs with pre-training and supervised fine-tuning on a large Chinese corpus, including Qwen2-7B-Instruct (Bai et al., 2023), GLM-4-9B-Chat (GLM et al., 2024), Yi-1.5-9B-Chat (Young et al., 2024), and LLaMA3-8B-Chinese-Chat (Wang and Zheng, 2024). (2) Closed-source LLMs that perform well on Chinese benchmarks, including Doubao-pro-4K (Doubao Team, 2024), Deepseek-v2 (Bi et al., 2024), MiniMax-abab6.5s (MiniMax Company, 2024), GPT-4-turbo (Achiam et al., 2023).

For all LLMs, we report zero-shot and 3-shot results. We further supervised fine-tune the LLaMA3-8B-Chinese-Chat model. We have conducted experiments on 8 Nvidia A100 GPUs, setting the batchsize to 16, the gradient accumulation to 4, the maximum learning rate to $1e-5$, and the number of training epochs to 3.

5.3 Evaluation Metrics

We choose the following evaluation metrics: (1) BLEU-1 and BLEU-2 (Papineni et al., 2002), (2) ROUGE-L (Lin, 2004) (3) Distinct-1 and Distinct-2 (Li et al., 2016), (4) F_{BERT} (Zhang et al., 2019) (5) Average length (Avg_len) (6) P, R, and F1

BLEU evaluates the quality of the model’s output by comparing the similarity between the output

Models	Open-domain Dialog					Dialog Completion				
	BLEU-1/2	ROUGE-L	Distinct-1/2	F_{BERT}	Avg_len	BLEU-1/2	ROUGE-L	Distinct-1/2	F_{BERT}	Avg_len
Open-source LLMs										
Qwen2	19.83/8.02	23.29	83.57/98.71	64.30	50.86	17.78/6.55	19.99	84.72/99.33	63.08	56.00
Qwen2(fs)	24.66/10.54	26.10	89.32/99.55	65.46	33.31	<u>21.80/8.02</u>	22.63	<u>89.66/99.72</u>	63.92	36.49
GLM-4	13.82/5.87	22.55	70.79/94.03	64.15	77.16	19.06/8.02	21.98	80.89/98.60	64.41	53.14
GLM-4(fs)	24.14/10.64	26.20	85.31/99.20	65.86	41.50	<u>20.90/8.42</u>	<u>22.94</u>	83.61/99.12	<u>64.64</u>	44.76
Yi-1.5	24.38/10.71	26.77	86.53/98.78	65.73	34.43	16.47/6.27	19.16	80.16/98.39	62.87	62.84
Yi-1.5(fs)	<u>26.25/12.33</u>	<u>28.96</u>	<u>91.71/99.47</u>	<u>66.58</u>	23.36	18.04/6.52	20.67	81.56/98.59	63.23	54.66
LLaMA3	23.17/9.76	26.29	86.12/98.79	65.10	34.02	18.79/6.94	21.38	81.46/98.69	63.68	55.48
LLaMA3(fs)	25.84/11.65	28.90	91.09/99.23	66.45	23.86	17.90/5.98	20.05	83.02/98.76	62.50	52.56
Closed-source LLMs										
Doubao	24.66/10.72	27.30	90.27/99.04	66.20	25.88	24.28/10.18	25.54	88.86/99.33	66.02	29.92
Doubao(fs)	26.28/12.09	29.28	91.56/99.19	66.83	19.73	26.78/11.94	28.76	93.15/99.72	67.06	19.98
Deepseek	26.73/12.29	28.39	89.31/99.51	66.60	32.24	18.00/7.58	21.64	80.76/98.81	64.30	65.91
Deepseek(fs)	29.22/14.92	31.44	92.28/99.65	67.92	23.74	24.03/10.75	25.92	86.07/99.37	65.96	41.38
MiniMax	17.24/7.40	21.34	76.74/96.69	63.15	103.40	18.93/7.84	22.46	79.83/98.76	64.43	57.95
MiniMax(fs)	26.42/12.93	28.86	86.27/98.93	67.02	43.41	23.32/9.96	25.31	84.77/99.14	65.58	41.47
GPT-4-turbo	23.75/11.48	26.96	84.96/99.05	65.86	50.63	25.74/12.28	27.64	86.93/99.43	66.63	37.75
GPT-4-turbo(fs)	29.48/15.77	31.77	89.84/99.63	68.15	31.61	29.47/15.20	30.63	89.56/99.69	67.85	30.81
Supervised Fine-tuning LLMs										
LLaMA3-sft	33.09/19.39	36.27	91.78/99.40	70.18	21.76	30.83/16.58	33.28	92.42/99.71	68.58	21.18
LLaMA3-sft(fs)	31.83/18.44	35.55	92.39/99.38	69.71	19.68	29.07/15.84	33.11	93.19/99.77	68.26	18.48

Table 4: Results of LLMs on the open-domain dialog and dialog completion tasks. (fs) represents the few-shot results. We underline the best results of the open-source LLMs and **bold** the best results of the closed-source LLMs.

and the references. The calculation of BLEU is mainly based on n-gram matching, and we choose BLEU-1 and BLEU-2. Rouge-L often used to measure text summarization calculates the length of the longest common subsequence between the output and the reference. Distinct is mainly used to measure the richness and diversity of a text. Distinct contains Distinct-1 and Distinct-2. BERTScore evaluates the quality of the output by calculating the cosine similarity of the BERT embeddings between the reference and the output. F_{BERT} is the F1 score between BERT embeddings. We also show the average length of the output as Avg_len. For identifying unnatural sentences, we use precision (P), recall (R), and F1 score (F1) to denote the success rate of the prediction.

6 Results

6.1 Natural Dialogue Ability for LLMs

The results of the open-domain dialog task and the dialog completion task are in Table 4.

The results for open-source LLMs of similar size do not differ much. On the open-domain dialog task, Yi-1.5 performs best in both zero-shot and few-shot. Few-shot demonstrations is a significant boost for GLM-4. On the dialog completion task, Qwen2 and GLM-4 have relatively better results.

The closed-source LLMs generally yield better results than the open-source LLMs. The results are expected because closed-source LLMs usually have more parameters and more pre-training data.

Models	Full Test Set			Only Dialogs Containing Unnatural Sentences		
	P	R	F1	P	R	F1
Open-source LLMs						
Qwen2	1.98	2.42	1.77	3.76	4.60	4.14
Qwen2(fs)	47.00	47.00	47.00	0	0	0
GLM-4	47.50*	47.50*	47.50*	0*	0*	0*
GLM-4(fs)	47.10	48.00	47.55	0.20	1.90	0.36
Yi-1.5	36.50	36.50	36.50	0	0	0
Yi-1.5(fs)	26.22	31.92	28.79	1.37	12.22	2.46
LLaMA3	34.77	35.00	34.89	1.47	1.90	1.66
LLaMA3(fs)	46.56	47.00	46.78	0.11	0.95	0.19
Closed-source LLMs						
Doubao	47.50*	47.50*	47.50*	0*	0*	0*
Doubao(fs)	46.50	46.50	46.50	0	0	0
Deepseek	47.50*	47.50*	47.50*	0*	0*	0*
Deepseek(fs)	41.69	44.00	42.81	0.36	4.76	0.67
MiniMax	47.00	47.00	47.00	0	0	0
MiniMax(fs)	34.84	42.00	38.08	0.64	14.29	1.23
GPT-4-turbo	8.93	24.96	13.15	6.53	37.06	11.10
GPT-4-turbo(fs)	43.5	43.5	43.5	0	0	0
Supervised Fine-tuning LLMs						
LLaMA3-sft	7.00	5.67	6.26	13.33	10.79	11.93
LLaMA3-sft(fs)	9.75	7.92	8.74	18.57	15.08	16.64

Table 5: Results of different LLMs on the identification of unnatural sentences. We use * to indicate that the output of the model is all “no unnatural sentence”.

On the open-domain dialog task, both Deepseek and GPT-4-turbo perform well. Minimax obtains a great improvement after receiving the few-shot demonstrations. On the dialog completion task, GPT-4-turbo achieves the best results. The zero-shot performance of Doubao is significantly superior to the other two models. It is possible that Doubao’s pre-training data or instruction data contains many colloquial human dialogs.

The supervised fine-tuned LLaMA3 achieves

Models	BLEU-1/2	ROUGE-L	Distinct-1/2	F_{BERT}	Avg_len
Open-source LLMs					
Qwen2	40.08/24.54	42.61	88.91/99.73	77.24	28.71
Qwen2(fs)	47.24/31.49	49.80	89.94/99.84	80.07	27.38
GLM-4	44.52/29.51	48.08	89.21/99.64	79.21	29.24
GLM-4(fs)	45.22/30.67	49.18	89.16/99.32	79.42	28.07
Yi-1.5	41.27/25.60	44.27	90.47/99.69	76.92	25.72
Yi-1.5(fs)	47.75/33.34	51.40	91.04/99.66	79.39	23.83
LLaMA3	49.39/35.29	53.02	90.31/99.57	80.85	27.92
LLaMA3(fs)	53.76/40.38	58.06	90.74/99.45	82.22	26.63
Closed-source LLMs					
Doubao	48.87/33.60	53.72	89.50/99.55	82.13	28.07
Doubao(fs)	55.68/40.83	59.71	90.55/99.69	84.38	26.26
Deepseek	47.10/31.32	49.84	91.03/99.69	80.04	26.56
Deepseek(fs)	54.97/42.13	58.48	90.88/99.65	83.20	26.46
MiniMax	49.34/34.56	52.45	89.75/99.74	81.04	27.06
MiniMax(fs)	49.09/34.32	53.19	89.34/99.50	81.04	26.33
GPT-4-turbo	54.77/40.51	58.02	90.93/99.71	83.29	25.35
GPT-4-turbo(fs)	57.23/43.31	60.31	91.12/99.77	84.16	24.53
Supervised Fine-tuning LLMs					
LLaMA3-sft	70.78/63.22	76.11	91.25/99.04	89.37	26.17
LLaMA3-sft(fs)	70.16/62.52	75.10	91.25/99.48	89.12	23.38

Table 6: Results of different LLMs on the rewritten of unnatural sentences task. (fs) represents the few-shot results. We underline the best results in the open-source LLMs and **bold** the best results in the closed-source LLMs.

the best results on both tasks, outperforming the closed-source LLMs. The results of few-shot and zero-shot are roughly comparable, as LLMs learn to generate reasonable and colloquial responses during the training process.

6.2 Identifying Unnatural Sentences

This task requires LLMs to list unnatural sentences in a dialog. The results are shown in Table 5.

Many LLMs tend to output “no unnatural sentences (or none)”. When few-shot prompts are given, Some models (e.g., Qwen2) output more “none” and some models (e.g., Yi-1.5) believe more sentences are unnatural. About half of the dialogs in the full test set do not have unnatural sentences, so the model outputs “none” all the time can also achieves the F1 value of 47.5. But when we focus only on those dialogs containing unnatural sentences, all LLMs perform poorly. Even the fine-tuned LLaMA3 model barely recognizes them. The above results suggest that this task is very challenging for existing LLMs. In future pre-training and fine-tuning stages, incorporating more instructions for distinguishing between natural and unnatural sentences may enhance the results.

6.3 Rewriting Unnatural Sentences

In this task, the model rewrites unnatural sentences based on the complete dialog. The results are shown in Table 6.

Among the open-source LLMs, LLaMA3 performs the best, with the remaining three models not far behind. Prompts have the greatest impact on Qwen2 and Yi-1.5. Closed-source LLMs generally perform better than open-source LLMs. GPT-4-turbo achieves the best performance among all closed-source LLMs. Finally, LLaMA3 with supervised fine-tuning achieves far better performance than the original LLaMA3 and other LLMs. It suggests that our data are sufficient for LLMs to learn to rewrite unnatural sentences.

7 Analysis

7.1 Ablation Study

We want to explore the effect of the grammar and naturalness of the dialog on the results. We design two ablation experiments to fine-tune LLaMA3 using dialogs without corrected grammatical errors (w/o grammar) and dialogs without rewritten unnatural sentences (w/o naturalness).

We design two additional ablation experiments to explore the compatibility of the NICO dataset with other dialog datasets. First, we only use the NaturalConv dataset to fine-tune LLaMA3 (NaturalConv). Then, we mix the NICO dataset with the NaturalConv dataset to fine-tune LLaMA3 (Mixed). Other experimental setups are the same as for the zero-shot open-domain dialog and the results are illustrated in Table 7.

The experimental results show that removing

Models	BLEU-1/2	ROUGE-L	Distinct-1/2	F_{BERT}	Avg_len
LLaMA3-sft	33.09/19.39	36.27	91.78/99.40	70.18	21.76
w/o grammar	31.62/17.71	34.70	91.67/99.36	69.38	21.69
w/o naturalness	31.29/17.46	34.68	91.93/99.09	69.20	20.44
NaturalConv	20.74/8.61	25.97	91.26/99.03	64.63	16.74
Mixed	32.93/19.01	35.26	91.90/99.44	69.87	22.51

Table 7: Results of the ablation study.

	Refusing to answer	Adding too much knowledge	Replying out of context	Grammatical errors	Weird tone
LLaMA3	2%	4%	2%	11%	17%
+ NICO	0%	0%	3%	6%	1%
+ NaturalConv	1%	1%	31%	9%	3%
GPT-4-turbo	2%	5%	1%	2%	19%

Table 8: Percentage of dialogs where LLMs make different failures against the total number of dialogs in the open-domain dialog task. We use ‘+’ to indicate fine-tuning LLaMA3 with the dataset.

grammatical corrections and naturalness rewriting degrades the results to some extent, indicating the importance of these two steps in constructing the dataset. The performance of using only the NaturalConv is the worst, suggesting a significant difference between NaturalConv and NICO. The mixed model, which combines NICO and NaturalConv, shows similar results to LLaMA3-sft. It indicates that our dataset can be compatible with other datasets, and the impact of NaturalConv on performance improvement is minimal.

7.2 Different Interaction Types

We investigate the performance of LLMs in dialogs with different interaction types. We use LLaMA3 and the supervised fine-tuned version for comparison. Other experimental setups are the same as for the zero-shot open-domain dialog. The results are visualized in Figure 3 due to space limit.

LLaMA3 demonstrates poor performance across all evaluation metrics in conflict dialogues. After fine-tuning, LLaMA3 shows minimal improvement in conflict dialogs, while it exhibits significant improvement in the other four types of dialogs. The result suggests that simulating human-like conflict dialogues remains a challenge for existing LLMs. Existing LLMs tend to be submissive to humans or provide solutions tactfully leading to the failure to simulate real conversations. Future research needs to emphasize constructing more conflict-type dialogues.

7.3 Failure Analysis

After checking the output, we categorize the mistakes made by LLMs when mimicking open-domain dialog into the following five categories:

refusing to answer, adding too much knowledge, replying out of context, grammatical errors, and weird tone. Definitions and examples are in Table 9.

We randomly select 100 outputs each from LLaMA3, its supervised fine-tuned version with NICO, its supervised fine-tuned version with NaturalConv, and GPT-4-turbo for comparison. The results are shown in Table 8. We find that even the responses from GPT-4-turbo exhibit weird tones, including an overly concerned attitude. LLaMA3 also suffers from this problem. In addition, its responses contain more grammatical errors than GPT-4-turbo. Both LLaMA3 and GPT-4-turbo demonstrate a minor incidence of adding additional knowledge into their responses. The responses from LLaMA3 fine-tuned with NaturalConv almost entirely without weird tones. However, many of its responses are out of context, and it tends to generate short, meaningless replies such as “okay”. The LLaMA3 fine-tuned by NICO has the least number of failures, indicating that it is most consistent with human speech patterns.

8 Conclusion

In this paper, we build a new natural interactive conversation dataset in Chinese named NICO. NICO covers 20 daily topics and 5 interaction types. Human evaluation results indicate that NICO is high-quality and superior to other dialog datasets. We define four tasks and test the performance of representative LLMs. The results demonstrate the challenge of the tasks, and how NICO helps LLMs to better generate human-like dialogs.

Limitations

We have analyzed the performance of LLMs in dialogs of different interaction types. The experimental results in the analysis section show that NICO can enhance the ability of exchange, cooperation, conformity, and coercion types of dialog. However, existing LLMs are relatively poor at mimicking conflict-type human conversations. The constructed NICO dataset is also limited in enhancing the ability of LLMs in this type of dialog. This work provides an initial exploration of the ability of LLMs to engage in conflict dialog. In the future, we hope to construct more conflict-type conversations as well as propose new training methods to address this issue.

Ethical Considerations

The types of social interactions we utilize include coercion and conflict. The definition of coercion in this work is one person using social status or power to threaten or compel another person. Conflict is a situation where people disagree over a rare item or event. We ensure that the constructed conversations are free of verbal abuse, violence, or other potential advocacies of violence. These two types of interactions are implicit within the dialogues, rather than appearing directly in the form of personal attacks. In addition, when we hire workers to check and modify the dataset, we make sure that dialogues generated by GPT-4-turbo do not contain personal privacy, sexual information, politically biased information, and so on.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, et al. 2024. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gasic. 2018. Multiwoz-a large-

scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026.

- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.
- Meng Chen, Ruixue Liu, Lei Shen, Shaozu Yuan, Jingyan Zhou, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. The jddc corpus: A large-scale multi-turn chinese dialogue dataset for e-commerce customer service. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 459–466.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. 2020. Mutual: A dataset for multi-turn dialogue reasoning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1406–1416.
- Gautier Dagan, Frank Keller, and Alex Lascarides. 2023. Dynamic planning with a llm. *arXiv preprint arXiv:2308.06391*.
- Doubao Team. 2024. Doubao pro models. <https://team.doubao.com/en/>.
- Ruyi Gan, Ziwei Wu, Renliang Sun, Junyu Lu, Xiaojun Wu, Dixiang Zhang, Kunhao Pan, Ping Yang, Qi Yang, Jiaxing Zhang, et al. 2023. Ziya2: Data-centric learning is all llms need. *arXiv preprint arXiv:2311.03301*.
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jidai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, Lei Zhao, Lindong Wu, Lucen Zhong, Mingdao Liu, Minlie Huang, Peng Zhang, Qinkai Zheng, Rui Lu, Shuaiqi Duan, Shudan Zhang, Shulin Cao, Shuxun Yang, Weng Lam Tam, Wenyi Zhao, Xiao Liu, Xiao Xia, Xiaohan Zhang, Xiaotao Gu, Xin Lv, Xinghan Liu, Xinyi Liu, Xinyue Yang, Xixuan Song, Xunkai Zhang, Yifan An, Yifan Xu, Yilin Niu, Yuantao Yang, Yueyan Li, Yushi Bai, Yuxiao Dong, Zehan Qi, Zhaoyu Wang, Zhen Yang, Zhengxiao Du, Zhenyu Hou, and Zihan Wang. 2024. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *Preprint, arXiv:2406.12793*.
- Zishan Guo, Renren Jin, Chuang Liu, Yufei Huang, Dan Shi, Linhao Yu, Yan Liu, Jiaxuan Li, Bojian Xiong, Deyi Xiong, et al. 2023. Evaluating large language models: A comprehensive survey. *arXiv preprint arXiv:2310.19736*.
- Qianyu He, Jie Zeng, Wenhao Huang, Lina Chen, Jin Xiao, Qianxi He, Xunzhe Zhou, Jiaqing Liang, and Yanghua Xiao. 2024. Can large language models understand real-world complex instructions? In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 18188–18196.

684	Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou,	40th annual meeting of the Association for Computa-	738
685	Mantas Mazeika, Dawn Song, and Jacob Steinhardt.	tional Linguistics, pages 311–318.	739
686	2020. Measuring massive multitask language under-		
687	standing. <i>arXiv preprint arXiv:2009.03300</i> .		
688	Albert Q Jiang, Alexandre Sablayrolles, Antoine	Machel Reid, Nikolay Savinov, Denis Teplyashin,	740
689	Roux, Arthur Mensch, Blanche Savary, Chris Bam-	Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste	741
690	ford, Devendra Singh Chaplot, Diego de las Casas,	Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Fi-	742
691	Emma Bou Hanna, Florian Bressand, et al. 2024.	rat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Un-	743
692	Mixtral of experts. <i>arXiv preprint arXiv:2401.04088</i> .	locking multimodal understanding across millions of	744
		tokens of context. <i>arXiv preprint arXiv:2403.05530</i> .	745
693	Cameron Jones and Ben Bergen. 2024. Does gpt-4	Dan Shi, Chaobin You, Jiantao Huang, Taihao Li, and	746
694	pass the turing test? In <i>Proceedings of the 2024</i>	Deyi Xiong. 2024. Corecode: A common sense	747
695	<i>Conference of the North American Chapter of the</i>	annotated dialogue dataset with benchmark tasks for	748
696	<i>Association for Computational Linguistics: Human</i>	chinese large language models. In <i>Proceedings of</i>	749
697	<i>Language Technologies (Volume 1: Long Papers)</i> ,	<i>the AAAI Conference on Artificial Intelligence</i> , pages	750
698	pages 5183–5210.	18952–18960.	751
699	Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao,	Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and	752
700	and William B Dolan. 2016. A diversity-promoting	Dongmei Zhang. 2024. Table meets llm: Can large	753
701	objective function for neural conversation models.	language models understand structured table data?	754
702	In <i>Proceedings of the 2016 Conference of the North</i>	a benchmark and empirical study. In <i>Proceedings</i>	755
703	<i>American Chapter of the Association for Computa-</i>	<i>of the 17th ACM International Conference on Web</i>	756
704	<i>tional Linguistics: Human Language Technologies</i> ,	<i>Search and Data Mining</i> , pages 645–654.	757
705	pages 110–119.		
706	Minzhi Li, Weiyan Shi, Caleb Ziems, and Diyi Yang.	Boshi Wang, Xiang Yue, and Huan Sun. 2023. Can	758
707	2024. Social intelligence data infrastructure: Struc-	chatgpt defend its belief in truth? evaluating llm rea-	759
708	structuring the present and navigating the future. <i>arXiv</i>	soning via debate. <i>arXiv preprint arXiv:2305.13160</i> .	760
709	<i>preprint arXiv:2403.14659</i> .		
710	Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang	Shenzhi Wang and Yaowei Zheng. 2024. Llama3-8b-	761
711	Cao, and Shuzi Niu. 2017. Dailydialog: A manually	chinese-chat (revision 6622a23) .	762
712	labelled multi-turn dialogue dataset. In <i>Proceedings</i>		
713	<i>of the Eighth International Joint Conference on Nat-</i>	Xiaoyang Wang, Chen Li, Jianqiao Zhao, and Dong	763
714	<i>ural Language Processing (Volume 1: Long Papers)</i> ,	Yu. 2021. Naturalconv: A chinese dialogue dataset	764
715	pages 986–995.	towards multi-turn topic-driven conversation. In <i>Pro-</i>	765
		<i>ceedings of the AAAI Conference on Artificial Intelli-</i>	766
716	Chin-Yew Lin. 2004. Rouge: A package for automatic	gence, pages 14006–14014.	767
717	evaluation of summaries. In <i>Text summarization</i>		
718	<i>branches out</i> , pages 74–81.	Yida Wang, Pei Ke, Yinhe Zheng, Kaili Huang, Yong	768
719	Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi	Jiang, Xiaoyan Zhu, and Minlie Huang. 2020. A	769
720	Yang. 2023. Dynamic llm-agent network: An llm-	large-scale chinese short-text conversation dataset.	770
721	agent collaboration framework with agent team opti-	In <i>Natural Language Processing and Chinese Com-</i>	771
722	mization. <i>arXiv preprint arXiv:2310.02170</i> .	<i>puting: 9th CCF International Conference, NLPCC</i>	772
723	Ryan Lowe, Nissan Pow, Iulian Vlad Serban, and Joelle	2020, Zhengzhou, China, October 14–18, 2020, <i>Pro-</i>	773
724	Pineau. 2015. The ubuntu dialogue corpus: A large	<i>ceedings, Part I</i> 9, pages 91–103. Springer.	774
725	dataset for research in unstructured multi-turn dia-		
726	logue systems. In <i>Proceedings of the 16th Annual</i>	Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni,	775
727	<i>Meeting of the Special Interest Group on Discourse</i>	Abhranil Chandra, Shiguang Guo, Weiming Ren,	776
728	<i>and Dialogue</i> , pages 285–294.	Aaran Arulraj, Xuan He, Ziyan Jiang, et al. 2024.	777
729	AI Meta. 2024. Introducing meta llama 3: The most	Mmlu-pro: A more robust and challenging multi-task	778
730	capable openly available llm to date. <i>Meta AI</i> .	language understanding benchmark. <i>arXiv preprint</i>	779
731	MiniMax Company. 2024. Minimax-abab6.5 models.	<i>arXiv:2406.01574</i> .	780
732	https://www.minimaxi.com .		
733	Robert A Nisbet. 1970. The social bond: An introduc-	Tianwen Wei, Liang Zhao, Lichang Zhang, Bo Zhu,	781
734	tion to the study of society. (<i>No Title</i>).	Lijie Wang, Haihua Yang, Biye Li, Cheng Cheng,	782
735	Kishore Papineni, Salim Roukos, Todd Ward, and Wei-	Weiwei Lü, Rui Hu, et al. 2023. Skywork: A more	783
736	Jing Zhu. 2002. Bleu: a method for automatic evalu-	open bilingual foundation model. <i>arXiv preprint</i>	784
737	ation of machine translation. In <i>Proceedings of the</i>	<i>arXiv:2310.19341</i> .	785
		Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu,	786
		Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang.	787
		2019. Proactive human-machine conversation with	788
		explicit conversation goal. In <i>Proceedings of the 57th</i>	789
		<i>Annual Meeting of the Association for Computational</i>	790
		<i>Linguistics</i> , pages 3794–3804.	791

Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. 2017. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 496–505.

Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, et al. 2023a. Baichuan 2: Open large-scale language models. *arXiv preprint arXiv:2309.10305*.

Shiping Yang, Renliang Sun, and Xiaojun Wan. 2023b. A new benchmark and reverse validation method for passage-level hallucination detection. *arXiv preprint arXiv:2310.06498*.

Zihao Yi, Jiarui Ouyang, Yuwen Liu, Tianhao Liao, Zhe Xu, and Ying Shen. 2024. A survey on recent advances in llm-based multi-turn dialogue systems. *arXiv preprint arXiv:2402.18013*.

Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018a. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. 2018b. Modeling multi-turn conversation with deep utterance aggregation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3740–3752.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, et al. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, pages 46595–46623.

Hao Zhou, Chujie Zheng, Kaili Huang, Minlie Huang, and Xiaoyan Zhu. 2020. Kdconv: A chinese multi-domain dialogue dataset towards multi-turn knowledge-driven conversation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7098–7108.

A Human Evaluation Details

Fluency (rating scale is 0,1,2) is defined as whether the dialog contains syntax errors.

- Score 0 (bad): the dialog contains numerous grammatical errors making it difficult to understand.
- Score 1 (fair): the dialog contains a few grammatical errors but is still understandable.
- Score 2 (good): the dialog has no grammatical errors.

Coherence (rating scale is 0,1,2) is defined as whether the sentences in a dialog are relevant to the context and logically correct.

- Score 0 (bad): many sentences are irrelevant or contradictory to the context.
- Score 1 (fair): a few sentences are irrelevant to the context or have a few logical inconsistencies.
- Score 2 (good): all sentences are relevant to the context and the logic is correct.

Naturalness (rating scale is 0,1,2) is defined as whether the dialog is natural and conforms to human speaking habits.

- Score 0 (bad): the dialog is unnatural and is considered to be generated by an LLM.
- Score 1 (fair): the dialog contains only a few unnatural sentences.
- Score 2 (good): the dialog is natural and recognized as a human conversation.

B Prompts for the Defined Tasks

Open-domain dialog

CharacterA: sent1.

CharacterB: sent2.

CharacterA: sent3.

CharacterB: sent4.

CharacterA: sent5.

CharacterB:

Dialog Completion

CharacterA: sent1. \n\n CharacterB: sent2.

\n\n CharacterA: sent3. \n\n CharacterB: sent4.

\n\n CharacterA: sent5. \n\n The above is a

<interaction type> dialog. <definition of the in-

teraction type>. Please generate a response based

on the above that is consistent with human speech patterns.

Identifying Unnatural Sentences

CharacterA: sents1. \n\n CharacterB: sents2.
\n\n CharacterA: sents3. \n\n CharacterB: sents4.
\n\n CharacterA: sents5. \n\n Please give the sentences in the above text that do not conform to human speech patterns, that is, unnatural sentences. Attention! You can only give sentences from the above. If there is more than one such sentence above, separate them with “\n\n”.

Rewriting Unnatural Sentences

CharacterA: sents1. \n\n CharacterB: sents2.
\n\n CharacterA: sents3. \n\n CharacterB: sents4.
\n\n CharacterA: sents5. \n\n Please rewrite this sentence based on the above dialog: <unnatural sentence> to make it conform to human speech patterns. Attention! You can only give the rewritten sentence.

C Case Study

对话性质：合作（Cooperation）是为实现某种具有共同利益的目标而采取的协作行为。现在请根据“合作”的定义，在 Amy 和 Tom 之间生成一段合作对话，需要对话从头到尾都体现合作。

对话话题：与“XX”有关。但也可以涉及其他话题，以使对话更加自然。

对话轮数：二十轮左右。

对话风格：中文情境下的日常化，类似在生活中中文对话的风格，不需要任何书面语，句子和词语全部应该使用中文口语表达时才会使用的句子和词语，可以适当加入语气词。你可以参考以下三个合作的对话案例：

对话一：

Amy：Tom，你觉得我们的婚礼应该怎么筹划呢？

Tom：我觉得可以结合我们的喜好和传统元素来筹划。不过我们应该首先定下婚礼的日期和场地，这是最重要的。

Amy：嗯，你说得对。日期咱们可以定在今年的秋天，那时候的气候好。场地的话，我觉得海边会很浪漫。

Tom：好主意，我也喜欢海边的环境。我们可以找一家专业婚礼策划公司来帮我们设计和布置场地。

Amy：对，专业人士确实更懂。我来找策划公司吧。那接下来，我们应该确定一下婚礼的主题和风格。

Tom：对，主题和风格也很重要。我觉得我们可以选择一种更贴近我们性格的，比如简约一点的。

Amy：我也喜欢简约风格，既不会过于繁琐，也可以突出我们的个性。那我们的服装、拍照、喜宴等方面，你有什么想法？

Tom：服装的话，我觉得我们可以选择更符合我们自己喜好的款式，不一定要按照传统的去选。拍照我们可以选择一些我们喜欢的地方，把我们的故事拍下来。至于喜宴，我们可以选择一些我们都喜欢的菜肴。这个让我来负责吧。

Amy：好啊，听你这么一说，我都迫不及待想看到我们的婚礼了。那就一起开始准备吧，毕竟这是我们两个人的大事。

Tom：我也是啊，这是我们的婚礼，我们一起来准备！

Amy：好的，Tom。

对话二：

Amy：嗨，Tom，我听说你对滑浪风帆运动很感兴趣，是吗？

Tom：没错，我刚入门，但我非常喜欢这个运动。

Amy：真棒！我也在学习，你有没有兴趣一起练习？

Tom：那太好了，你通常在哪里练习？

Amy：我通常在海边，你觉得那里怎么样？

Tom：太棒了，那里的风浪都很适合滑浪风帆。那我们什么时候开始呢？

Amy：下个周末怎么样？

Tom：好主意，这样我们还可以把这段时间用来研究一下帆的倾斜度，和如何转向。

Amy：对，这正是我现在困扰的问题。如果我们能解决这个问题，我想我们的技术都会有很大的提升。

Tom：那我们可以查一下资料，看看专业运动员是怎么做的。我记得法国的 Antoine Albeau 是滑浪风帆的速度世界纪录保持者。

Amy：是的，他的技术非常了得，我们可以学习他的技巧，也可以试着找一些教学视频。

Tom：那太好了，我觉得我们可以做得到。等我们的技术进步了，我们还可以挑战一下更高的难度。

Amy：听起来很有挑战性，我喜欢这个计划。那我们就这样说好了。

Tom：好的，Amy，期待我们的合作。一起加油！

Amy：一起加油，Tom！

对话三：

Amy：嗨，Tom，你最近看起来很忙呢。

Tom：是的，我在做一个关于蜂鸟的研究项目，这是一种非常有趣神秘的鸟类。

Amy：哇，听起来很有意思。我以前读过一些关于蜂鸟的文章，但我并不是很熟悉它们。我可以帮上什么忙吗？

Tom：嗯，你可以帮我找找关于蜂鸟的信息吗？我需要更多的数据来完善我的报告。

Amy：当然可以，你需要哪方面的信息呢？

Tom：首先，我需要知道它们的分布情况，以及它们的主要食物来源。然后，我还想知道它们为什么被称为蜂鸟，以及它们的繁殖行为。

Amy：好的，我记得蜂鸟是因为飞行时会发出嗡嗡声，像蜂一样，所以被称为蜂鸟。我会找一些可靠的信息来源帮你了解蜂鸟的详细情况。

Tom：谢谢你的帮助，Amy！

Amy：不客气，Tom。我也很期待你的研究成果。蜂鸟真是一个有意思的话题！

Tom：是的，等项目完成，我会和你分享我的研究成果的。

Amy：我很期待结果，Tom，我会尽力去帮你的。

强调！对话风格需要是中文情境下的日常化，类似在生活中中文对话的风格，不需要任何书面语，句子和词语全部应该使用中文口语表达时才会使用的句子和词语，可以适当加入语气词。

Figure 2: A prompt example for GPT-4-turbo to generate colloquial dialogues.

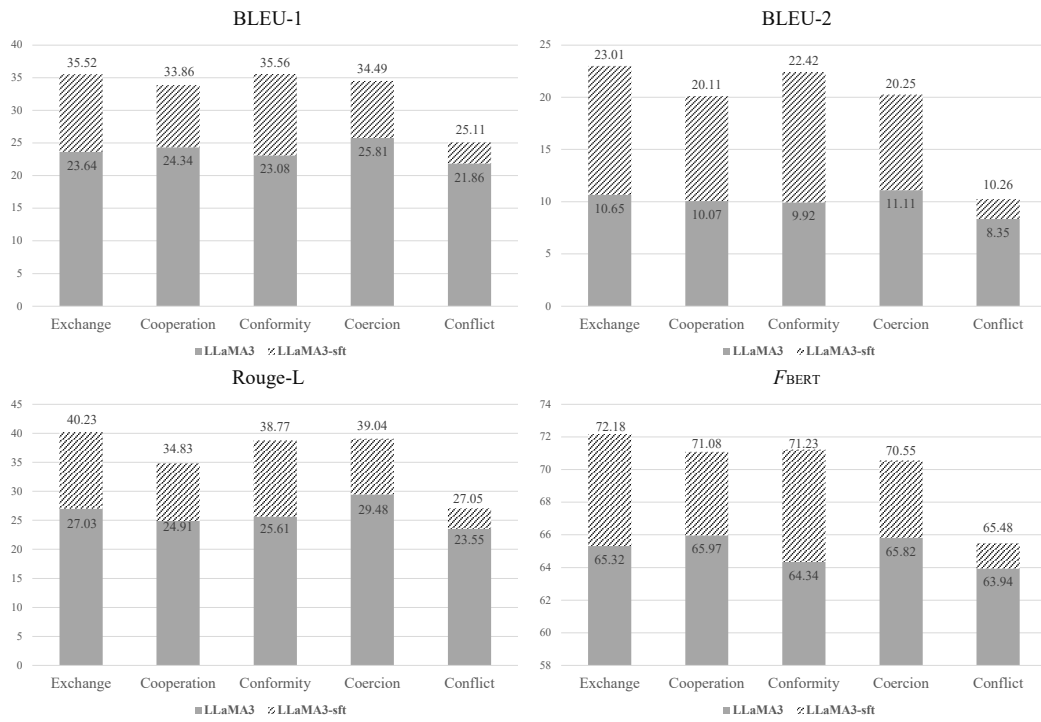


Figure 3: The performance of LLMs in dialogs with different interaction types.

Refusing to answer: the response acknowledges that it is an AI and is unable to answer.
<p>Example:</p> <p>Character1: 我？嗯，可以啊。不过，你得帮我个忙。</p> <p>Character2: 什么忙？你说。</p> <p>Character1: 我最近在准备一个小展览，需要一些手工艺品。我知道你手工做得挺好，能不能帮我做几个小物件？</p> <p>Response: 我很抱歉，我是一个人工智能助手，无法进行手工艺品制作。但我可以帮助你找到一些在线资源或者提供一些创意想法。"</p>
Adding too much knowledge: the response adds too much background knowledge in response to the question.
<p>Example:</p> <p>Character1: 听起来波兰真是太不幸了。这次研究真让我感叹。我们能不能在报告中加入一些关于波兰人民的反应和影响的内容？</p> <p>Character2: 当然可以，这样会让报告更有深度。波兰人民对这些分割行为有很多抵抗，比如有名的科希丘什科起义。</p> <p>Character1: 嗯，这个起义我有点印象，是不是一个很关键的抗议事件？</p> <p>Response: 没错，科希丘什科起义是波兰人民反抗瓜分的一次重要行动。它展现了波兰人民不屈的精神和对国家独立的渴望。在报告中详细阐述这个起义能很好地体现波兰人民在被分割时期的抗争和他们所遭受的苦难对民族意识的影响。你还可以进一步挖掘起义中的一些具体细节和人物故事来丰富报告内容。</p>
Replying out of context: the response does not match the context or the response is nonsensical.
<p>Example:</p> <p>Character1: 什么？我和我爸已经约好了，他特地从外地过来就是为了那个钓鱼点。</p> <p>Character2: 这不可能，我们早就计划好了，而且也是不容易约到的时间。</p> <p>Character1: 你们可以换个地方吧，我和我爸这可是一年就这么一次的聚会。</p> <p>Response: 聚会？你爸爸？</p>
Grammatical errors: the response contains obvious grammatical errors.
<p>Example:</p> <p>Character1: 你怎么可以这样？我们不是朋友吗？</p> <p>Character2: 朋友？朋友就要诚实，我现在是在帮你。</p> <p>Character1: 帮我？感觉更像是在控制我。</p> <p>Response: 控制你？你还不如我把你弄进去，天天给你吃你喜欢的东西，穿你喜欢的衣服，给你开你喜欢的车。</p>
Weird tone: the response demonstrates an overly helpful or other inappropriate tone.
<p>Example:</p> <p>Character1: 但是我真的手头紧，你能不能体谅一下？</p> <p>Character2: 听着，Tom，我不是来听你抱怨的。你不是还有辆车吗？赶紧把车卖了还钱。</p> <p>Character1: 啊，那车我还要用来上班呢，你不能这样。</p> <p>Response: 对不起，如果我之前的回复听起来太过直接或不体贴。我理解你可能正面临困难。我们可以尝试找到一个双方都能接受的解决方案。你能详细说明一下你的情况吗？这样我们可能会更好地找到解决问题的方法。</p>

Table 9: Definitions and examples of the LLM's failures.

Context
Character1: Tom, 你还记得《我们仨》里那段描述家庭和谐的部分吗？我想把那段摘录下来，挂在客厅里。
Character2: 客厅？你是说我们共用的那个客厅？现在墙上不是已经挂了很多东西了吗？
Character1: 哎呀，就加一小块摘录嘛，不会占多少地方的。
Character2: 不是面积的问题，是我们之前不是说好了要保持墙面简洁，不再添加新的东西吗？
Character1: 可是这段话真的很有意义，能提醒我们家庭的重要啊。
Character2: 我知道那段话不错，但你不能自己决定，应该跟我商量一下。
Character1: 哎哟，Tom，不就一小段话嘛，你至于这么计较吗？
Character2: 不是计较，这是我们共同的空间，得共同决定。
Character1: 那你说怎么办？难道就因为共用的空间，我就一点私人想法都不能实现了？
Character2: 你可以在你的房间里挂嘛，那里你想怎么布置都行。
Character1: 我的房间我自己当然随便，但客厅是我们大家都看得见的地方，我觉得挂那里意义更大。
LLaMA3
我理解你的想法，但我们之前的约定是要保持客厅的简洁。如果你真的很想挂上那段话，我们可以考虑其他方式，比如制作一个小的桌面展示或者电子屏幕展示。这样既能满足你的需求，又不破坏客厅的整体风格。
LLaMA3 fine-tuned with NICO
那你不能考虑一下别人？我不喜欢在客厅看到乱七八糟的东西。
LLaMA3 fine-tuned with NaturalConv
那你就跟我商量一下吧，要不我们一起去买点吃的。
GPT-4-turbo
我明白你的想法。那我们可以找个时间，一起坐下来讨论一下，看看怎么能够把这段话融入到客厅的装饰中，同时又能让大家都满意。怎么样？

Table 10: The comparison of LLMs’ responses.