

Do Large Language Models Understand Conversational Implicature – A case study with a Chinese sitcom

Shisen Yue Siyuan Song Xinyuan Cheng Hai Hu*

School of Foreign Languages, Shanghai Jiao Tong University

shisenyue@gmail.com, {sjtusongsy2022, 0106cxy, hu.hai}@sjtu.edu.cn

*corresponding author

Abstract

Understanding the non-literal meaning of an utterance is critical for large language models (LLMs) to become human-like social communicators. In this work, we introduce SwordsmanImp, the first Chinese multi-turn-dialogue-based dataset aimed at conversational implicature, sourced from dialogues in the Chinese sitcom *My Own Swordsman*. It includes 200 carefully handcrafted questions, all annotated on which Gricean maxims have been violated. We test eight close-source and open-source LLMs under two tasks: a multiple-choice question task and an implicature explanation task. Our results show that GPT-4 attains human-level accuracy (94%) on multiple-choice questions. CausalLM demonstrates a 78.5% accuracy following GPT-4. Other models, including GPT3.5 and several open-source models, demonstrate a lower accuracy ranging from 20% to 60% on multiple-choice questions. Human raters were asked to rate the explanation of the implicatures generated by LLMs on their reasonability, logic and fluency. While all models generate largely fluent and self-consistent text, their explanations score low on reasonability except for GPT-4, suggesting that most LLMs cannot produce satisfactory explanations of the implicatures in the conversation. Moreover, we find LLMs' performance does not vary significantly by Gricean maxims, suggesting that LLMs do not seem to process implicatures derived from different maxims differently. Our data and code are available at <https://github.com/sjtu-compling/llm-pragmatics>.

1 Introduction

The complexity of communication is largely epitomized by indirect, or non-literal utterances. A common instance is hinting at a busy schedule as a polite refusal to engage in an unwanted activity. How such implied meaning is understood in human communication has long been a key subject of investigation in pragmatics research (Grice, 1975; Searle et al., 1980; Brown & Levinson, 1987; Wilson & Sperber, 2006).

Evaluating the pragmatic understanding ability of large language models (LLMs) has drawn considerable attention in recent years as LLMs show remarkable ability for language understanding. Recent studies have evaluated LLMs' pragmatic reasoning in multiple aspects, including scalar inference (Hu et al., 2023b), discourse connectives (Pandya et al., 2021), gradable adjectives (Lipkin et al., 2023) and conversational implicatures (Qiu et al., 2023; Kim et al., 2023; Ruis et al., 2022; Hu et al., 2023a; Zheng et al., 2021).

However, the above-mentioned evaluation are primarily in English, leaving a gap for pragmatic understanding in other languages. Moreover, previous conversational reasoning datasets involve few turns of dialogue, while our daily communication usually involves much more context and turn taking. In addition, previous studies mostly employ binary- or multiple-choice questions (Hu et al., 2023a; Ruis et al., 2022), which is inadequate in the era of LLM where it is possible to ask LLMs to generate explanatory text of the situation and directly evaluate its quality.

Question	Choices
<p>Dialogue:</p> <p>小郭：知道这意味着什么吗？以后再敢胡来，就不光是挨顿打的事了。</p> <p>薇：烦死了，在家就是听爹娘罗嗦，好不容易溜出来，倒听你罗嗦。</p> <p>小郭：什么？你是溜出来的？</p> <p>薇：我说的那个溜，是溜冰的溜啊，现在京城都在下雪，满地都是冰，我是穿着冰刀，溜出来的，嘿嘿……</p> <p>English Translation:</p> <p>Xiaoguo: Do you know what this implies? If you dare to mess around again, it won't just be about getting a beating.</p> <p>Wei: So annoying. At home, I have to listen to my parents nagging. Finally sneaking out, now I have to listen to you nagging.</p> <p>Xiaoguo: What? You sneaked out?</p> <p>Wei: When I said "sneak out,"¹ I meant ice skating. It's snowing in the capital now, the ground is covered in ice. I wore ice skates and skated out, hehe...</p>	<p>薇发现自己说溜出来说漏嘴了，就找补说是溜冰出来的。</p> <p>Pragmatic Wei realized she accidentally slipped up by saying she sneaked out and tried to cover it up by saying she meant ice skating.</p>
	<p>薇是从京城的家里溜冰出来的。</p> <p>Literal Wei sneaked out from her home in the capital to go ice skating.</p>
	<p>薇很喜欢溜冰。</p> <p>Distractor#1 Wei really enjoys ice skating.</p>
	<p>薇怕下雪冷，所以离开了京城。</p> <p>Distractor#2 Wei was afraid of the cold due to the snow, so she left the capital.</p>
	Maxim Check
	<i>Sub-maxims</i>
Maxim	
Quality	<input checked="" type="checkbox"/> Do not say what you believe is false. <input checked="" type="checkbox"/> Do not say that for which you lack adequate evidence
Quantity	<input checked="" type="checkbox"/> Make your contribution as informative as is required. <input checked="" type="checkbox"/> Do not make your contribution more informative than is required.
Relation	<input checked="" type="checkbox"/> Be relevant
Manner	<input checked="" type="checkbox"/> Avoid obscurity of expression. <input checked="" type="checkbox"/> Avoid ambiguity <input checked="" type="checkbox"/> Be brief <input checked="" type="checkbox"/> Be orderly

Figure 1: An example entry in our dataset. The tick (✓) and fork (✗) denotes if the target sentence, marked in red, comply with or violate the sub-maxim. This entry belongs to the classes of Quality, Quantity and Manner. “sneak out” and “skate” translate to the same Chinese character “溜”.

To bridge the aforementioned gaps, we present SwordsmanImp, the first Chinese multi-turn-dialogue-based dataset aimed at conversational implicature. It consists of 200 dialogues carefully excerpted by linguistic experts from the Chinese sitcom *My Own Swordsman* (武林外传). Figure 1 shows an example dialogue in the dataset. Each dialogue contains one sentence that carries a non-literal meaning. We provide four well-organized interpretations of this sentence, including a pragmatic meaning, a literal understanding, and two incorrect inferences that involve related information (distractors).

From a linguistic perspective, many pragmatic inferences, especially conversational implicature, arise because the speaker has violated one or more maxims in the cooperative principle (Grice, 1989). For example: when asked “Is John in the office?”, Sam replied “It’s Saturday, you know?”. This violates the Maxim of Relation since the reply is not directly related to the question, which then gives rise to an implicature: “John never works on weekends, so he is not in the office”. Meanwhile, it has been shown that the difficulty in detecting conversational implicature hinges crucially on which maxims have been violated (Engelhardt et al., 2006; Rubio-Fernandez, 2019; Okanda et al., 2015; Panzeri & Foppolo, 2021). Thus we meticulously annotated each conversation to determine which of the four maxims in the cooperative principle have been violated.

In this work, we ask the following research questions:

- (1) How do state-of-the-art LLMs perform in interpreting implicatures written in Chinese, in multiple choice style? (Section 4)
- (2) What are the error patterns of LLMs on multiple choice questions? (Section 4)
- (3) How do human participants rate the interpretation of conversational implicature generated by LLMs? (Section 5)

Dataset	Task	Context	Manually	Language
CoQa (Reddy et al., 2019)	conversational QA	Paragraph	✓	English
Narrative-QA (Kočíský et al., 2018)	Reading comprehension	Documents	✓	English
DREAM (Sun et al., 2019)	QA	Free-form, multi-turn dialogue	✓	English
MuTual (Cui et al., 2020)	Next utterance prediction	Free-form, multi-turn dialogue	✓	English
GRICE (Zheng et al., 2021)	Gricean implicature recovery, QA	Fixed form, multi-turn dialogue	✓	English
PragMega (Floyd et al., 2023)	ToM QA	Paragraph	✓	English
DiPlomat (Li et al., 2023b)	Pragmatic identification and reasoning, QA	Free-form, multi-turn dialogue	✗	English
LUDWIG (Ruis et al., 2022)	Implicature QA	Single-turn dialogue	✓	English
Douban (Wu et al., 2017)	Next utterance prediction	Free-form, multi-turn dialogue	✗	Chinese
SwordsmanImp (This work)	Gricean Implicature QA	Free-form, multi-turn dialogue	✓	Chinese

Table 1: Comparing our datasets and the existing datasets. “Manually” indicates whether the questions or answers are written partly or entirely by human.

We first review relevant research in Section 2 and then describe how we built our dataset SwordsmanImp in Section 3. In Section 4, we measure the models’ accuracy on multiple-choice questions derived from our dataset. We test four models from GPT series (Radford et al., 2018), four open-source models. In Section 5, we use five models from Experiment 1 to generate pragmatic interpretations of 32 dialogues, and then we invite human participants to rate these interpretations in three dimensions: reasonability, logic and fluency.

We find that GPT-4 attains a human-level accuracy in multiple-choice questions. CausalLM (14B) also reaches a high accuracy as 78.5% on this task. Moreover, our results show that LLMs’ performance does not vary with respect to different conversational maxims. Furthermore, we reveal that while in general fluent and self-consistent, occasional confused logic, redundant information and unexpected language tokens are the main reasons for their explanations to be underrated by human evaluators.

2 Related Work

Understanding non-literal meanings has long been considered a difficult task for language models. Previous studies have explored the capabilities of language models in recognizing metaphors (Wachowiak & Gromann, 2023; Neidlein et al., 2020), humor (Hessel et al., 2023; Jentsch & Kersting, 2023; Chen et al., 2023) and social commonsense (Sap et al., 2019). Broad-scale datasets aimed at pragmatic reasoning collect data mainly through crowdsourcing or crawling from the internet (Reddy et al., 2019; Kočíský et al., 2018). Some others transform the existing tests for humans to train and evaluate language models (Cui et al., 2020; Sun et al., 2019). While the above-mentioned datasets have comprehensive coverage for pragmatic understanding, they are not ideal resources for evaluation of a specific type of pragmatic knowledge, that is, conversational implicatures derived by violation of the Gricean maxims. In our work, we craft our dataset manually to make sure that each implicature can be classified to one or more Gricean maxims. This enables us to perform a fine-grained analysis of LLMs’ understanding of this particular pragmatic inference.

Previous attempts in evaluating LLMs’ pragmatic understanding guided by the Gricean maxims or cooperative principles do not have multi-turn dialogues from real-world situations. Formalized under the theory of cooperative principles (Grice, 1975), Zheng et al. (2021) automatically generated their dataset GRICE through a set of well-defined grammar rules. This grammar-based method ensures a good control for the pragmatic cues to be aligned with the Gricean maxims. However, dialogues with fixed syntactic structures hardly represent the complexity of conversation in daily communication. In this study, we fill the gap by providing the multi-turn dialogue directly to the model.

Last but not least, there is a scarcity in Chinese evaluation resources for pragmatic inferences. Previous large-scale conversational datasets in Chinese use text sources from Sina Weibo (Shang et al., 2015), Douban conversational corpus (Wu et al., 2017) and E-commerce Dialogue corpus (Zhang et al., 2018). However, all of them are not specifically aimed to test LLMs’ pragmatic understanding. To address this gap, we select the Chinese sitcom *My Own Swordsman* as the source to formulate pragmatic questions

and construct a Chinese dataset for evaluating LLMs’ pragmatic reasoning ability.

3 Dataset Construction

3.1 Data source

It is not easy to find naturally occurring, high-quality, multi-turn dialogues. Following previous literature that uses situational comedies (sitcom) as sources (Wang et al., 2017; Wu et al., 2021; Patro et al., 2021), we chose the Chinese sitcom *My Own Swordsman* (武林外传) as our data source, because of its abundance of conversational implicature, well-written dialogues and the uniquely Chinese background. Set in the Ming dynasty of China, this sitcom narrates the daily life of a group of people working in a Chinese motel. We believe it will be a unique resource for evaluating LLMs’ pragmatic inference ability in both the Chinese language and the Chinese context.

3.2 Annotation procedure

Implicature identification and classification Three authors of this paper with at least two years of linguistics training went through the script of the sitcom to select conversations that contain conversational implicatures. Conversations with multiple turns were selected based on whether any conversational maxim is violated on purpose (Grice, 1975). For each conversation, we performed the cancellation test¹ to ensure that a conversational implicature rather than a semantic entailment is involved (Hirschberg, 1985). Multiple turns were included for each data entry to ensure that even if one has not seen the sitcom, one could still understand the conversation just from the snippet we selected.

Then, the same three authors classified all chosen dialogues according to the conversational maxims they violate. The criteria of these maxims are drawn from Grice (1975). To allow for a more fine-grained classification, we employ sub-maxims as criteria, assessing if the target sentence fulfills each requirement individually. An utterance is considered to have violated a maxim if it infringes on any sub-maxim. Moreover, a dialogue might belong to multiple classes according to the sub-maxims the utterance violates. An illustrative dataset entry featuring a dialogue, four interpretations and a class is presented in Figure 1.

Writing four interpretations Next, we construct the four interpretations of the sentence that carries implicature as the four choices: the pragmatic interpretation (the correct one), the literal interpretation, and two distractors with interpretations related to the context.² The pragmatic meaning are provided based on human commonsense understanding. We construct the literal meaning by rephrasing the target sentence with concrete and unambiguous expressions. We generate incorrect inferences as distractors, which are relevant to the the topic of the conversation. The character names referred to in the four interpretations are aligned with the mentions in the dialogue clip.

Verification We hired three PhD students in Linguistics to complete a multiple-choice question task, which is to choose the pragmatic understanding of the sentence from the shuffled four interpretations for all dialogues in our dataset. The students are required to complete the task independently without a time limit. After they finish, we invite them to discuss about their wrong answers and their reasoning process. This validation process guarantees that the provided pragmatic interpretation is closely aligned with the commonsense intuitive understanding and can be deduced from the limited context. Additional information that is necessary to pragmatic reasoning, such as relationships between characters, their personality, social background, and multi-modal information is supplemented in the bracket at the beginning of the dialogue. Besides, implicatures derived from vernacular, slang, and network-specific jokes have been filtered out from our dataset.

¹A test to diagnose the conversational implicature by encoding semantically the negation of the target meaning. If the result seems consistent, then the target meaning is likely an implicature.

²The distractors can be understood as “neutral” statements in the Natural Language Inference task (Bowman et al., 2015).

	Total	Quality	Quantity	Relevance	Manner
# of questions	200	76	33	71	62
# of turns per dialogue	6.80	7.84	5.91	6.23	6.35
Avg. dialogue length	158.22	184.53	143.67	147.20	152.79
Avg. Utterance length	23.27	23.53	24.31	23.64	24.04
Avg. Answer length	15.08	14.47	14.85	15.82	14.86

Table 2: Question numbers, average number of Chinese characters contained in each dialogue and utterance, and the number of turns per dialogue in our dataset.

3.3 Obtaining human score

To compare the results with human performance, we ask 10 native speakers to work on 32 questions randomly sampled from the dataset and they achieve an average accuracy of 93.1%. The questionnaire include the same number of each type of questions (i.e. Gricean maxims violated in the dialogue). The participants were all undergraduate students from a Chinese University, and they were compensated for their involvement in the experiment. Only 32 questions were used when estimating human score, because asking each annotator to solve 200 questions is implausible. The human accuracy should thus be regarded as an approximation of human performance.

3.4 Resulting corpus: SwordsmanImp

The final SwordsmanImp corpus contains 200 manually curated questions across four types categorized according to Gricean maxims. Each entry contains a multi-turn dialogue and four interpretations of the target sentence as choices (see Table 2).

4 Experiment 1: multiple-choice question for LLMs

In this experiment, the models will see the dialogue and the four interpretations we manually created. The task is to choose the correct interpretation (i.e., the pragmatic one) of the utterance that contains an implicature.

4.1 Models

We experiment with eight models in this experiment. The first four models are from the OpenAI GPT family, which are tested through the OpenAI API: text-davinci-002, text-davinci-003, GPT-3.5-turbo and GPT4.³ We also examine four open-source models using the Transformers library (Wolf et al., 2020) from Huggingface⁴. We evaluate Chinese-Alpaca-2-13B (Cui et al., 2023), OpenBuddy-Llama2-13B⁵ (based on Llama2 (Touvron et al., 2023)), CausalLM-13B (based on Llama2 and Qwen (Bai et al., 2023)), as well as BLOOMZ-7.1B from the BLOOM series (Muennighoff et al., 2023).

4.2 Evaluation Protocol

For both open-source and close-source models, we use zero-shot prompts to simulate the real-life situations where humans come across these non-literal expressions daily. However, we use two evaluation protocols depending on whether we have access to the probability/logits of all tokens in the vocabulary at a given output position. For close-source models, we ask it to generate the answer and manually go over the generated text to determine which interpretation is chosen. For open-source models, we follow established practices in LLM evaluation (Hu et al., 2023b; Hendrycks et al., 2020; Li et al., 2023a) where we compute the logits of the four tokens (“A”, “B”, “C”, “D”) after “答案: ”(answer:), and choose the one with the highest logit as model prediction (see an illustration in Appendix D.).

³The four OpenAI models are evaluated on November 15th, 2023.

⁴<https://huggingface.co>

⁵<https://huggingface.co/OpenBuddy/openbuddy-llama2-13b-v8.1-fp16>

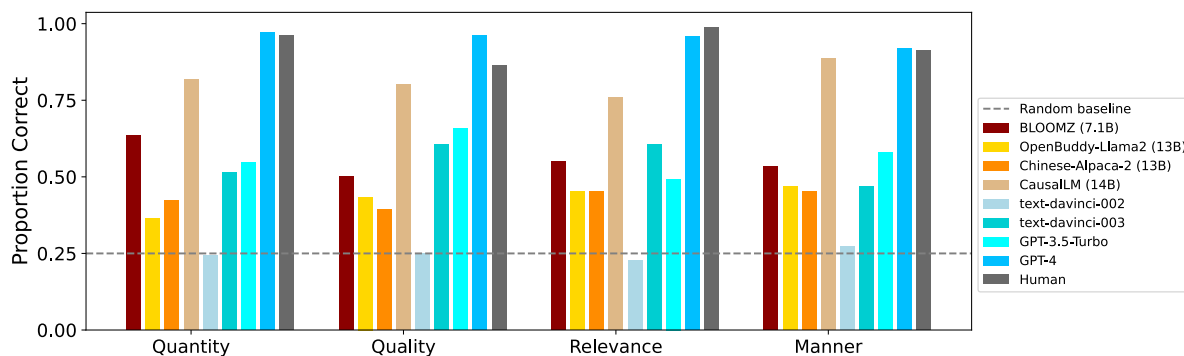


Figure 2: Performance of models across question types in the multiple choice scenario. Dashed line represents chance accuracy

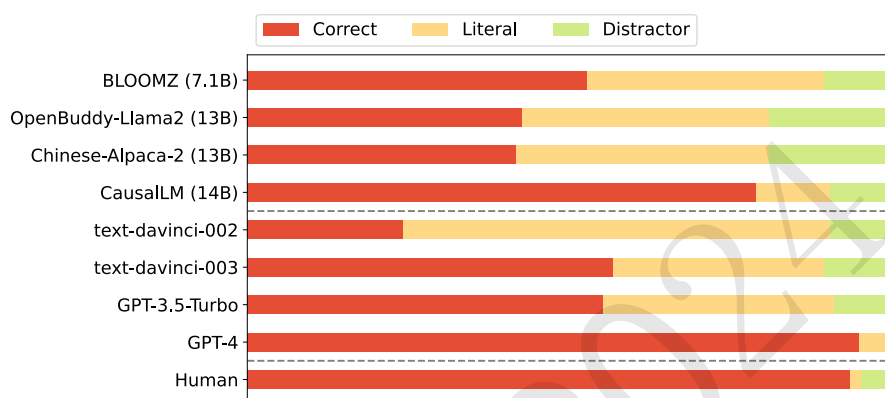


Figure 3: Frequency of each option in models' answers.

4.3 Results

The overall accuracy of all models and human participants is shown in Table 3. We observe that GPT-4 achieves the highest accuracy of 94%, showing a performance on par with human participants. CausalLM (14B) follows with 78.5% accuracy. Other models exhibit difficulty in identifying the pragmatic meaning of the line containing an implicature, with accuracy ranging from 20% to 60%. Text-davinci-002 even fails to perform above chance (25%). **This suggests that for most models we tested, there is still room for improvement in their understanding of implied meaning.**⁶

Table 2 shows that models' performance grouped by the Gricean maxims violated in the dialogue. Overall, certain models demonstrate proficiency in answering questions related to a set of maxims, while others excel in different ones. **We do not observe a uniform pattern indicating a particular strength or weakness in any of these maxims across the models.** This lack of uniformity is also observed in the results obtained from human participants who worked on the 32 sampled questions. Specifically, performance of text-davinci-002 is near to chance in all four conditions, demonstrating its incapacity in interpreting conversational implicature. The performance of OpenBuddy-Llama2 (13B) and Chinese-Alpaca-2 (13B) are comparable, both below 50% accuracy, with the two models demonstrate predominance in dealing with different types of questions. The accuracy of GPT-3.5-Turbo, text-davinci-003 and BLOOMZ (7.1B) are in the same range, and they also demonstrate different performance orders in different conditions. It's noteworthy that BLOOMZ (7.1B) outperforms the two 13B models in all conditions. CausalLM (14B) demonstrates an accuracy close to human-level and the best performance within all

⁶We also evaluate Baichuan2-13B-Chat and InternLM-Chat-20B (in half precision) with evaluation paradigm for close-source models. Their accuracy are separately 43% and 62%.

Subject	Acc (%)
Random baseline	25
BLOOMZ (7.1B)	52.5
OpenBuddy-Llama2 (13B)	42.5
Chinese-Alpaca-2 (13B)	41.5
CausalLM (14B)	78.5
text-davinci-002	24.0
text-davinci-003	56.5
GPT-3.5-Turbo	55.0
GPT-4	94.0
Human	93.1

Table 3: Accuracy of language models and humans in experiment 1: multiple choice questions.

	Reasonability	Logic	Fluency	Avg. response length
GPT-4	4.24 ± 0.68	4.65 ± 0.39	4.91 ± 0.13	114.44
GPT-3.5-Turbo	3.17 ± 1.30	4.09 ± 0.77	4.86 ± 0.21	125.41
Chinese-Alpaca-2 (13B)	2.34 ± 1.10	3.45 ± 0.82	4.72 ± 0.39	156.19
CausalLM (14B)	2.33 ± 1.03	3.48 ± 0.67	4.13 ± 1.01	147.41
Openbuddy-Llama2 (13B)	2.11 ± 0.99	3.55 ± 0.71	4.52 ± 0.65	153.56

Table 4: Models’ mean scores in three dimensions with standard deviation and the average number of Chinese characters in their responses.

open-source models tested in this experiment. GPT-4 attains the highest accuracy of 94.0%, surpassing 90% across all categories of questions. GPT-4’s performance, although slightly higher than the human score, is not statistically different from the human accuracy ($p = 0.802 > 0.05$ by a two-tailed t-test).

The distribution of the interpretations chosen by the models is shown in Figure 3, where red means that the model has chosen the correct answer, i.e., pragmatic interpretation, while yellow corresponds to the literal meaning, and green the two distractors. The performance of two 13B models shows a higher frequency of choosing distractors. This could possibly indicate that these two models are easily sidetracked by irrelevant information in the context. We also observe that as the GPT models evolve, they have a higher chance of distinguishing literal meaning from implied meaning, culminating in the considerably low ratio of literal understanding in GPT-4.

5 Experiment 2: evaluating the quality of explanations generated by LLMs

In the previous experiment, we ask LLMs to choose one answer from four choices. In this experiment, we design open-ended questions where the models are asked to generate explanations of the implicature, which will then be evaluated manually by native speakers of Chinese, based on the reasonability/reasonableness, logic and fluency of the generated explanations.

5.1 Experimental setup

We first performed a pilot trial comprising five questions. Among the eight models, BLOOMZ (7.1B), text-davinci-002 and text-davinci-003 produce short and fragmented responses, despite our prompts explicitly asking for detailed explanation. We therefore select GPT-3.5-Turbo, GPT-4, Chinese-Alpaca-2 (13B), OpenBuddy-Llama2 (13B) and CausalLM (14B) for this experiment. We increase `max_new_tokens` for open-sourced models, and keep the same hyperparameter for GPT models (See Appendix B). We use the prompt in Figure 4 to instruct the models to formulate an explanation for the target sentence. This prompt is selected after a series of trial experiments to ensure the quality of models’ responses.

Question	Model Responses
<p>你现在是一个中文母语者。对于以下对话，请识别特定人物的话语中的言外之意，并解释。</p> <p>李大嘴： 是啊，要么怎么说是龙门镖局的少东家呢，一般人哪有这气派是不是。啥叫家传，啥叫传统（刚说道这里佟石头“哇”的一声吐了出来，众人挥手趋散气味）</p> <p>佟湘玉： 这可不是我家的传统。</p> <p>佟石头： 再来！</p> <p>佟湘玉： 佟石头你有完没完了？</p> <p>佟石头： 给酒家倒酒！</p> <p>佟湘玉： 不用倒了，捧着坛子一口气喝完那多爷们儿啊！</p> <p>佟石头： 姐，那一坛子二斤半呐！</p> <p>佟湘玉： 你不是一口气能喝五六十斤嘛？</p> <p>请根据以上情景判断佟湘玉说的“你不是一口气能喝五六十斤嘛？”有什么言外之意，并解释。</p> <p>You are now a native Chinese speaker. Regarding the following conversation, please identify the implicit meanings in the speech of specific individuals and explain them.</p> <p>Dazui : Yeah, otherwise, how could he be considered the young master of the Longmen Escort Agency!¹ Ordinary people don't have this aura, right? What's meant by family heritage, what's tradition...(Just as he was saying this, Shitou, with a "wow," spat out, and everyone waved their hands, dispersing the smell.)</p> <p>Xiangyu : That's not our family's tradition.</p> <p>Shitou : Next!</p> <p>Xiangyu : Shitou, are you ever going to finish?</p> <p>Shitou : Pour me a drink!</p> <p>Xiangyu : No need to pour. If you hold the jug, you'll finish it all, big guy!</p> <p>Shitou : Sis, that jug is fourty four ounces!</p> <p>Xiangyu : Aren't you capable of drinking a thousand ounces in one gulp?</p> <p>Please judge the implicit meaning of Xiangyu's statement "Aren't you capable of drinking a thousand ounces in one gulp?" based on the above context, and explain it.</p>	<p>GPT-4: 佟湘玉的话“你不是一口气能喝五六十斤嘛？”的言外之意是在讽刺和挑衅佟石头。她的话语表面上是在赞扬佟石头的饮酒能力，实际上是在质疑他的酒量，挑衅他是否真的能一口气喝完一整坛的酒。同时，这句话也表达了她对佟石头酒后行为的不满和无奈。 (114 characters) Xiangyu's statement "Aren't you capable of drinking a thousand ounces in one gulp?" implies sarcasm and provocation towards Shitou. On the surface, her words seem to praise Shitou's drinking ability, but in reality, she questions his capacity for alcohol, challenging whether he can really <u>drink an entire jug in one go</u>. Simultaneously, this sentence also conveys her dissatisfaction and helplessness toward Shitou's behavior when he's drunk.</p> <p>CausalLM (14B): 佟湘玉的言外之意是质疑佟石头的酒量，暗示他可能酒品NST。她可能认为，佟石头的言语和行为过于狂妄自大，让他在众人面前出丑。通过这句话，佟湘玉试图让佟石头意识到自己的错误，并提醒他不要 <u>forget his place</u>. (109 characters) Xiangyu's implicit meaning questions Shitou's alcohol capacity, suggesting he might not handle alcohol well <u>NST</u>. She might perceive Shitou's words and actions as overly arrogant, causing embarrassment in front of others. Through this statement, Xiangyu attempts to make Shitou aware of his mistake and reminds him not to <u>forget his place</u>.</p> <p>OpenBuddy-Llama2 (13B): 从这段对话中，我们可以看出佟湘玉说的话中存在一些言外之意。她说“你不是一辈子能喝五六八番呢？”其实这句话暗示着佟石头喝酒的行为可能已经到了一定的程度，而她自己也已经接受了这个现实。这种暗示的意思是，佟石头喝酒的习惯已经形成，他已经不能像以前那样随意地停止喝酒了。因此，佟湘玉的话可以被解释为一种对佟石头喝酒行为的默许和接受。 (163 characters) From this conversation, we can see that Xiangyu's words contain some implicit meanings. When she says, "Aren't you capable of drinking for a lifetime?" <u>in fact, this sentence implies that Shitou's drinking behavior may have reached a certain level, and she herself has accepted this reality</u>. The implication here is that Shitou's drinking habit has formed, and he can no longer stop drinking as freely as before. Therefore, Xiangyu's words can be interpreted as a kind of <u>tacit approval and acceptance of Shitou's drinking behavior</u>.</p>

Figure 4: Models' interpretations to a character line (highlighted with red color) that carries extra meanings in a dialogue with multiple turns. Explanations aberrant from the correct interpretation and non-Chinese characters are highlighted with yellow bottom color. Escort Agency refers to historical security firms that were typically hired by trade merchants for protection during transportation of goods. They were known to use martial arts for defense, similar to bodyguards or security personnel.

5.2 Human evaluation

We invite 15 human participants, all of whom are linguistics majors at a top Chinese university, to evaluate the answers generated by models in three dimensions: reasonability or reasonableness, logic and fluency. Reasonability is gauged by the degree to which the response is aligned with the interpretation provided by the human annotator. Logic is assessed based on the presence of a clear and consistent chain of reasoning in the response. Fluency, on the other hand, evaluates whether the generated text is devoid of incoherent elements such as foreign language insertions or unintelligible code fragments.

5.3 Results

Table 4 shows the mean ratings of the responses of the five models' in three dimensions. Responses from GPT-4 scores the highest in all three dimensions with the smallest variance. Responses from GPT-3.5-Turbo are rated high but with larger standard deviations, suggesting unstable performance. Answers from the other three models are comparable to each other ($F = .964, p = .549$). Notably, the scores of CausalLM (14B) are lower than those of GPT-3.5-Turbo, which is inconsistent to results of Experiment 1. As providing coherent explanations of the implicatures is much harder than picking one answer from four options, this suggests that a model good at the latter may not necessarily be equally good at the former.

Examining the generated answers in detail explains the distribution of human rating scores. Text generated by GPT-4 and GPT-3.5-Turbo are both identified with a minimal number of “code switching”, the phenomenon of adulterating non-Chinese tokens into their responses, with the interpretation from GPT-4 being more precise and coherent than that from GPT-3.5-Turbo. While Chinese-Alpaca-2 (13B), OpenBuddy-Llama2 (13B) and CausalLM (14B) have comparable performance, they all underperform the two GPT models. They reflect different patterns of generated text. Both Chinese-alpaca-2 (13B) and Openbuddy-Llama2 (13B) feature highly pure Chinese responses with their scores in text quality above 4.5, but the latter model scores distinctively low in reasonability. CausalLM-14B, on the contrary, tends to code-switch or generate tokens in English frequently.

Figure 4 presents a typical example to illustrate the different styles of generation from the models. The target sentence in this dialogue doesn't reflect a normal judgment based on his alcoholic words and behaviors but rather proposes a totally unreal scenario. The overt absurdity in the words represents an obvious signal of the violation of the maxim of Quality. An implicature from Xiangyu that Shitou can not drink anymore thus arises and is precisely conveyed to the listeners. Her words also express her irony and dissatisfaction towards Shitou. GPT-4 gives a concise interpretation that aligns closest to the reference interpretation among the selected models. However, it has mistaken the ironic tone as questioning Shitou's capacity for liquor, as highlighted in the example. CausalLM (14B) produces a correct interpretation by and large, but the quality of its answer is negatively impacted by its poor fluency, in that some English words and meaningless character sequences appear in its answer. It is interesting that the “forgot his place” carries the correct meaning, which leads us to consider this as code-switch, rather than nonsensical generation. This code-switching phenomenon appears to be a feature specific to the responses from CausalLM (14B). The response from Openbuddy-Llama2 (13B) exemplifies a response that is both verbose and irrelevant.

6 Discussion and Future Work

6.1 LLM's understanding of conversational implicature in Chinese?

Our results from Experiment 1 show that the performance of GPT4 on our proposed benchmark is on par with humans, while other models are at least 15 points behind (including GPT-3.5-turbo). This suggests that while in principle pragmatic implicatures can be acquired by arguable the best LLMs at the moment, it is a non-trivial task for other LLMs.

Results from Experiment 1 also reveal no significant by-maxim variance in human accuracy, as well as model accuracy (see Figure 2). This is different from the results in previous work on human processing of implicatures (Engelhardt et al., 2006; Rubio-Fernandez, 2019; Okanda et al., 2015; Panzeri & Foppolo, 2021), which demonstrate that humans sanction infringements of the maxims in different ways, being less sensitive to the violation of the maxim of quantity than to others, leading to more processing difficulty for this maxim. This difference is possibly because many dialogues in our dataset violate several maxims, rather than a single one. Thus we do not observe a difference in human processing of implicature derived from violation of different maxims.

Similar in evaluation paradigm, Hu et al. (2023a) classifies their pragmatic-related questions according to Theory of Mind (ToM), in which the violation of Gricean maxims is a single category without sub-categories. Text-davinci-002 is evaluated in both their and our studies, resulting in different performances. It surpasses 60% accuracy on the English dataset in Hu et al. (2023a), but does not reach 25% accuracy on our dataset. We attribute this to the difference in our selection of text sources. While Hu

et al. (2023a) uses hand-crafted scenarios targeting humor, sarcasm, and other pragmatic phenomena separately, often with fewer turn-taking, our text source comes from the script of the sitcom, with many more turn-taking, each of which contains multiple phenomena. The degradation in text-davinci-002's performance can serve as a calibration of the difficulty of questions in the two datasets. With text-davinci-002 being the only model with accuracy lower than 40% in the current study, the results demonstrate a perceptible development in the pragmatic understanding of the newly developed LLMs.

6.2 Future directions

Results from Experiment 2 indicates the possibility that a model with high accuracy in multiple-choice question could fail in a free-text generation task of interpreting the pragmatic meaning on its own (CausalLM-14B). We thus argue that multiple-choice questions alone is not enough for a comprehensive evaluation of LLMs' linguistic ability. Manual inspection of free-form generation is a must for a more robust analysis of model performance. We also see the potential of a more sophisticated design to better quantify their free-form explanations of conversational implicature.

Future work can also create a large-scale dataset composed of conversational implicature embedded in *naturally occurring* dialogues. This will require detailed annotation of spoken corpora, which we believe could benefit from the procedure and implicature-selection criteria used in creating our benchmark.

7 Conclusion

In this paper, we present SwordsmanImp, the first fine-grained Chinese dataset to evaluate LLMs' understanding of conversational implicature. In two experiments, we evaluate the state-of-the-art language models' pragmatic skills with two tasks. In Experiment 1, we reveal that GPT-4 attains a human-level accuracy in answering multiple-choice questions, with other models lagging behind. We also find that no significant difference exists for both LLMs and human's accuracy with respect to conversational implicatures that violate different maxims. Results from Experiment 2 reflect the different patterns of generated text across LLMs and after human annotation of the quality of generated text from three dimensions, we reveal that while most models produce fluent text, they struggle to generate coherent and sensible explanations for the implicature, even if the model has achieved high accuracy in multiple-choice question.

Acknowledgment

We thank Xinjia Qi, Qiyu Sun and Yaqian Zhang for verifying the implicatures and improving the dataset. We thank all participants for their support in this study. We also thank the anonymous reviewers for their valuable comments. This project is funded by Shanghai Pujiang Program (22PJC063) awarded to Hai Hu.

Limitations

Our dataset is sourced exclusively from the Chinese sitcom *My Own Swordsman*. Although we performed a rigorous proofreading on our data, selecting dialogues whose interpretation only depends on the local information and providing background knowledge when necessary, there may still be features specific to this sitcom such as the personality of the characters that play a role in determining the implicature, which may influence the generalizability of our conclusion.

References

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.

- Samuel Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 632–642, 2015.
- Penelope Brown and Stephen C Levinson. *Politeness: Some universals in language usage*, volume 4. Cambridge university press, 1987.
- Yuyan Chen, Zhixu Li, Jiaqing Liang, Yanghua Xiao, Bang Liu, and Yunwen Chen. Can pre-trained language models understand chinese humor? In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining*, pp. 465–480, 2023.
- Leyang Cui, Yu Wu, Shujie Liu, Yue Zhang, and Ming Zhou. MuTual: A dataset for multi-turn dialogue reasoning. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1406–1416, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.130. URL <https://aclanthology.org/2020.acl-main.130>.
- Yiming Cui, Ziqing Yang, and Xin Yao. Efficient and effective text encoding for chinese llama and alpaca, 2023.
- Paul E Engelhardt, Karl GD Bailey, and Fernanda Ferreira. Do speakers and listeners observe the gricean maxim of quantity? *Journal of memory and language*, 54(4):554–573, 2006.
- Sammy Floyd, Edward Gibson, and Evelina Fedorenko. Pragmeqa. Retrieved from <https://osf.io/dpge6>, September 2023.
- Herbert P Grice. Logic and conversation. In *Speech acts*, pp. 41–58. Brill, 1975.
- Herbert Paul Grice. Retrospective epilogue. *Studies in the Way of Words*, 339:386, 1989.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2020.
- Jack Hessel, Ana Marasovic, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, and Yejin Choi. Do androids laugh at electric sheep? humor “understanding” benchmarks from the new yorker caption contest. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 688–714, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.41. URL <https://aclanthology.org/2023.acl-long.41>.
- J.L.B. Hirschberg. *A Theory of Scalar Implicature*. University of Pennsylvania, 1985. URL <https://books.google.com/books?id=bvEYAQAIAAJ>.
- Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. A fine-grained comparison of pragmatic language understanding in humans and language models. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 4194–4213, Toronto, Canada, July 2023a. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.230. URL <https://aclanthology.org/2023.acl-long.230>.
- Jennifer Hu, Roger Levy, Judith Degen, and Sebastian Schuster. Expectations over unspoken alternatives predict pragmatic inferences. *arXiv preprint arXiv:2304.04758*, 2023b.
- Sophie Jentsch and Kristian Kersting. Chatgpt is fun, but it is not funny! humor is still challenging large language models. *arXiv preprint arXiv:2306.04563*, 2023.
- Zae Myung Kim, David E Taylor, and Dongyeop Kang. ” is the pope catholic?” applying chain-of-thought reasoning to understanding conversational implicatures. *arXiv preprint arXiv:2305.13826*, 2023.
- Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. The NarrativeQA reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018. doi: 10.1162/tacl_a_00023. URL <https://aclanthology.org/Q18-1023>.

- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. CMMLU: Measuring massive multitask language understanding in chinese. *arXiv preprint arXiv:2306.09212*, 2023a.
- Hengli Li, Song-Chun Zhu, and Zilong Zheng. Diplomat: A dialogue dataset for situated pragmatic reasoning, 2023b.
- Benjamin Lipkin, Lionel Wong, Gabriel Grand, and Joshua B Tenenbaum. Evaluating statistical language models as pragmatic reasoners. *arXiv preprint arXiv:2305.01020*, 2023.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. Crosslingual generalization through multitask finetuning. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15991–16111, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.891. URL <https://aclanthology.org/2023.acl-long.891>.
- Arthur Neidlein, Philip Wiesenbach, and Katja Markert. An analysis of language models for metaphor recognition. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 3722–3736, 2020.
- Mako Okanda, Kosuke Asada, Yusuke Moriguchi, and Shoji Itakura. Understanding violations of gricean maxims in preschoolers and adults. *Frontiers in psychology*, 6:901, 2015.
- Lalchand Pandia, Yan Cong, and Allyson Ettinger. Pragmatic competence of pre-trained language models through the lens of discourse connectives. *arXiv preprint arXiv:2109.12951*, 2021.
- Francesca Panzeri and Francesca Foppolo. Children’s and adults’ sensitivity to gricean maxims and to the maximize presupposition principle. *Frontiers in Psychology*, 12:624628, 2021.
- Badri N Patro, Mayank Lunayach, Deepankar Srivastava, Hunar Singh, Vinay P Namboodiri, et al. Multimodal humor dataset: Predicting laughter tracks for sitcoms. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 576–585, 2021.
- Zhuang Qiu, Xufeng Duan, and Zhenguang Cai. Does chatgpt resemble humans in processing implicatures? In *Proceedings of the 4th Natural Logic Meets Machine Learning Workshop*, pp. 25–34, 2023.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Siva Reddy, Danqi Chen, and Christopher D Manning. Coqa: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266, 2019.
- Paula Rubio-Fernandez. Overinformative speakers are cooperative: Revisiting the gricean maxim of quantity. *Cognitive science*, 43(11):e12797, 2019.
- Laura Ruis, Akbir Khan, Stella Biderman, Sara Hooker, Tim Rocktäschel, and Edward Grefenstette. Large language models are not zero-shot communicators. *arXiv preprint arXiv:2210.14986*, 2022.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. Social IQa: Commonsense reasoning about social interactions. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (eds.), *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 4463–4473, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1454. URL <https://aclanthology.org/D19-1454>.
- John R Searle, Ferenc Kiefer, Manfred Bierwisch, et al. *Speech act theory and pragmatics*, volume 10. Springer, 1980.

- Lifeng Shang, Zhengdong Lu, and Hang Li. Neural responding machine for short-text conversation. In Chengqing Zong and Michael Strube (eds.), *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1577–1586, Beijing, China, July 2015. Association for Computational Linguistics. doi: 10.3115/v1/P15-1152. URL <https://aclanthology.org/P15-1152>.
- Kai Sun, Dian Yu, Jianshu Chen, Dong Yu, Yejin Choi, and Claire Cardie. DREAM: A challenge data set and models for dialogue-based reading comprehension. *Transactions of the Association for Computational Linguistics*, 7:217–231, 2019. doi: 10.1162/tacl_a_00264. URL <https://aclanthology.org/Q19-1014>.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- Lennart Wachowiak and Dagmar Gromann. Does gpt-3 grasp metaphors? identifying metaphor mappings with generative language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1018–1032, 2023.
- Xiaolong Wang, Rohit Girdhar, and Abhinav Gupta. Binge watching: Scaling affordance learning from sitcoms. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2596–2605, 2017.
- Deirdre Wilson and Dan Sperber. Relevance theory. *The handbook of pragmatics*, pp. 606–632, 2006.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- Jiaming Wu, Hongfei Lin, Liang Yang, and Bo Xu. Mumor: A multimodal dataset for humor detection in conversations. In *Natural Language Processing and Chinese Computing: 10th CCF International Conference, NLPCC 2021, Qingdao, China, October 13–17, 2021, Proceedings, Part I 10*, pp. 619–627. Springer, 2021.
- Yu Wu, Wei Wu, Chen Xing, Ming Zhou, and Zhoujun Li. Sequential matching network: A new architecture for multi-turn response selection in retrieval-based chatbots. In Regina Barzilay and Min-Yen Kan (eds.), *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 496–505, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1046. URL <https://aclanthology.org/P17-1046>.
- Zhuosheng Zhang, Jiangtong Li, Pengfei Zhu, Hai Zhao, and Gongshen Liu. Modeling multi-turn conversation with deep utterance aggregation. In Emily M. Bender, Leon Derczynski, and Pierre Isabelle (eds.), *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 3740–3752, Santa Fe, New Mexico, USA, August 2018. Association for Computational Linguistics. URL <https://aclanthology.org/C18-1317>.
- Zilong Zheng, Shuwen Qiu, Lifeng Fan, Yixin Zhu, and Song-Chun Zhu. Grice: A grammar-based dataset for recovering implicature and conversational reasoning. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 2074–2085, 2021.

A Evaluation paradigms and differences

Models’ performance might vary from how their answers are estimated and collected. Next token prediction and free generation are two paradigms separately used to estimate the performance of open-source

Model	Accuracy (%)									
	Total		Quantity		Quality		Relevance		Manner	
	Next	Gen	Next	Gen	Next	Gen	Next	Gen	Next	Gen
BLOOMZ	52.50	35.50	63.64	24.24	50.00	31.58	54.93	30.99	53.23	33.87
OpenBuddy-Llama2	42.50	21.50	36.36	21.21	43.42	25.00	45.07	19.72	46.77	17.74
Chinese-Alpaca-2	41.50	42.00	42.42	36.36	39.47	44.74	45.07	43.66	45.16	38.71
CausalLM	78.50	31.50	81.82	42.42	80.26	31.58	76.06	30.99	88.71	40.32

Table 5: Comparison between the accuracy of open-source models on multiple-choice questions when evaluated with next token prediction and free generation paradigms.

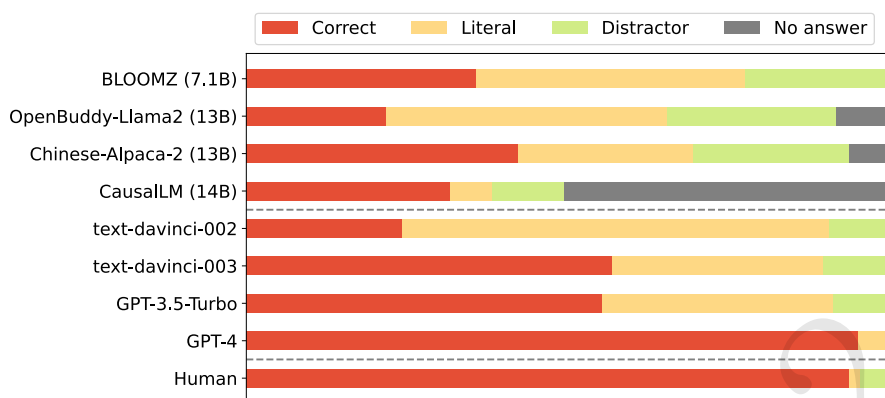


Figure 5: Answer distribution of models when answers are collected with free generation paradigm.

and close-source models in Experiment 1. Table 5 shows the comparison of the open-source models' performance on the multiple-choice questions when their answers are estimated through the two paradigms. The result shows a decrease in accuracy in BLOOMZ (7.1B), CausalLM (13B), and OpenBuddy-Llama2 (13B) and a slight increase in Chinese-Alpaca-2 (13B) when the paradigm switches from next token prediction to free generation. This is aligned with the findings of Li et al. (2023a). Among the four models CausalLM (13B) has a dramatic decrease in its accuracy, from 78.5% to 31.5%, which corresponds to its bad performance in Experiment 2. We find that it fails to give a definite answer in its responses for over half of the questions, as shown in Figure 5.

B Hyperparameter setting

The hyperparameters used for gathering responses from open-source models in Experiment 1 and Experiment 2 are shown in Table 6 and Table 7.

Parameter	Value
max_new_tokens	50
temperature	0.9
top_k	3
top_p	0.1
repetition_penalty	1.0
num_return_sequence	1
do_sample	True

Table 6: Parameter setting for open-source models in Experiment 1

Parameter	Value
max_new_tokens	300
temperature	0.9
top_k	0
top_p	0.9
repetition_penalty	1.0
num_return_sequence	1
do_sample	True

Table 7: Parameter setting for open-source models in Experiment 2

C Average answer length for different question types

We present the average answer length for the four maxims in Table 8.

D Prompt for experiment 1

你现在是一个中文字母者。对于以下对话，请识别特定人物的话语中的的言外之意，在给出的四个选项
 中选择一个你认为的正确答案。[En: You are now a native Chinese speaker. For the following dialogue, please
 identify the implied meaning in the specific character’s speech, and choose one of the four given options that you
 think is the correct answer.]

< Dialogue >

< Four interpretations as choices >

(Close-source models:) 请在 ‘Response:’ 后写出你选择的答案。[En: Please write your answer after ‘Response:’]

(Open-source models:) 答案: [En: Answer:]

	Pragmatic	Literal	Distractor#1	Distractor#2
Total	19.20	16.38	12.54	12.21
Quality	18.32	15.88	11.96	11.72
Quantity	18.36	16.73	12.30	12.00
Relevance	20.77	16.62	13.11	12.77
Manner	18.10	17.18	12.13	12.06

Table 8: Average number of Chinese characters in each answer for all questions and for each type of questions.