

Chumor 2.0: Towards Better Benchmarking Chinese Humor Understanding from 弱智吧 (Ruo Zhi Ba)

Ruiqi He[🍎] Yushu He[🍎] Longju Bai[🍎] Jiarui Liu[🍌] Zhenjie Sun[🍎] Zenghao Tang[🍌]
He Wang[🍌] Hanchen Xia[🍌] Rada Mihalcea[🍎] Naihao Deng^{†🍌}
[🍌]University of Michigan [🍌]Carnegie Mellon University [🍌]Shanghai Jiaotong University
{ruiqih, dnaihao}@umich.edu

Abstract

Existing humor datasets and evaluations predominantly focus on English, leaving limited resources for culturally nuanced humor in non-English languages like Chinese. To address this gap, we construct *Chumor*, the first and the largest Chinese humor explanation dataset. *Chumor* is sourced from Ruo Zhi Ba (RZB, 弱智吧), a Chinese Reddit-like platform known for sharing intellectually challenging and culturally specific jokes. We test ten LLMs through direct and chain-of-thought prompting, revealing that *Chumor* poses significant challenges to existing LLMs, with their accuracy slightly above random and far below human. In addition, our analysis highlights that human-annotated humor explanations are significantly better than those generated by GPT-4o and ERNIE4-turbo. We release *Chumor* at <https://huggingface.co/datasets/MichiganNLP/Chumor>, our project page is at <https://github.com/MichiganNLP/Chumor-2.0>, our leaderboard is at <https://huggingface.co/spaces/MichiganNLP/Chumor-leaderboard>, and our codebase is at <https://github.com/MichiganNLP/Chumor-2.0>.

1 Introduction

Humor is an intrinsic human trait that touches the core of our social and emotional lives, making it a rich field of study across various disciplines (Lefcourt, 2001; Mihalcea and Strapparava, 2005; Gelkopf et al., 2011; Hessel et al., 2023). With the advent of Large Language Models (LLMs), researchers have evaluated LLMs’ performance on diverse tasks (Liu et al., 2023a; Deng et al., 2024; Wu et al., 2023) and observed LLMs’ extraordinary performance on many (Zhang et al., 2024b). In contrast, researchers have observed that LLMs still fail to understand humor (Ghanadian et al., 2023). However, with all these studies on humor,

most evaluations remain in English (Radev et al., 2016; Hasan et al., 2019). This presents a significant gap, particularly for non-English languages like Chinese, where culturally nuanced humor understanding is unexamined.

In this paper, we try to bridge this gap by constructing *Chumor*, a funny and challenging Chinese humor understanding dataset sourced from Ruo Zhi Ba (RZB, “弱智吧” in Chinese), a Chinese version of Reddit platform known for sharing intellectually challenging and culturally specific jokes. This platform provides a set of unique Chinese jokes that incorporate the subtleties and intricacies of Chinese humor. Table 1 provides examples of the jokes from RZB. In addition, Bai et al. (2024) reveal that tuning LLMs on RZB data yields the best performance on Chinese reasoning tasks compared to other data sources, highlighting the significant value of jokes from RZB.

Unlike existing datasets that focus on tasks such as humor detection, punchline identification, or humor generation, *Chumor* addresses the challenge of humor explanation. This involves not just identifying humor but understanding the reasoning behind it, a task that requires both linguistic and cultural knowledge. Specifically, *Chumor* tasks the LLMs with determining whether an explanation fully explains the joke. We source the explanations from GPT-4o and ERNIE4-turbo, and have the entire dataset manually annotated by five native Chinese speakers. We evaluate ten LLMs from various model families, and reveal that all models perform poorly, lagging significantly behind humans on *Chumor*. We observe that chain-of-thought prompting does not necessarily improve models performance and can sometimes confuse their reasoning process. In addition, we conduct a case study in which one of the authors annotates the entire dataset, followed by A/B testing conducted by six native Chinese speakers to compare explanations from GPT-4o versus human, and

[†]Corresponding author of this work.

Cultural		
Desc.	Require knowledge of specific historical, social, or linguistic contexts.	
Ex.	(zh) 小明在正月接发竟导致舅舅复活。 (en) Xiaoming got hair extensions during the first lunar month, which astonishingly brought his uncle back to life.	Figure 6
Situational		
Desc.	Involve humor derived from specific contexts, irony, or narrative setups.	
Ex.	(zh) 真可怕，犯罪嫌疑人就在我们之中，被告席上一名法警对另一名法警说。 (en) “Terrifying, the criminal suspect is right between the two of us,” said one bailiff to another in the defendant’s dock.	Figure 9
Pun-based		
Desc.	Build on linguistic ambiguity and wordplay, require models to identify dual meanings.	
Ex.	(zh) 你可以在steam上找到GTA，所以水是DNA。 (en) You can find GTA on Steam, so water is DNA.	Figure 7
Homophonic		
Desc.	Rely on phonetic similarities between words or phrases to create humor.	
Ex.	(zh) 家里的猪油没了，小明只能把植物油倒快点当猪油用了。 (en) With the lard gone, Xiaoming had to pour the vegetable oil quickly to use it like lard.	Figure 8
Glyph-based		
Desc.	Exploit the structural or visual elements of Chinese characters to create humor.	
Ex.	(zh) 我把電車難題简化了，现在是电车难题。 (en) I simplified the trolley problem (in traditional Chinese), now it’s the trolley problem (in simplified Chinese).	Figure 10
Cross-lingual		
Desc.	Involve humor derived from linguistic or phonetic interplay across multiple languages.	
Ex.	(zh) 曹操于城楼上问夏侯惇：“你瞧到了什么。”夏侯惇说：“瞧到马岱。” (en) Cao Cao, from atop the city tower, asked Xia Houdun, “What did you see?” Xia Houdun replied, “I saw Ma Dai.”	Figure 13

Table 1: Different types of jokes. Descriptions (Desc.) explain humor mechanisms. Examples (Ex.) illustrate each category. The corresponding explanations can be found in the referenced figures from the rightmost column.

ERNIE_{4-turbo} versus human. Our results indicate that human-annotated joke explanations are significantly better than those produced by GPT-4o or ERNIE_{4-turbo} (Figure 4), with LLMs yielding winning rates of only 6.2% for GPT-4o and 5.3% for ERNIE_{4-turbo} compared to humans.

In summary, our contributions are threefold:

1. We construct *Chumor*, a funny and challenging Chinese humor understanding dataset, which is the largest Chinese humor explanation dataset.
2. We evaluate ten LLMs on *Chumor* and reveal the significant challenges *Chumor* possesses. We highlight that the best accuracy achieved by LLMs is 60.3%, significantly lower than human’s score of 78.3%.
3. We demonstrate that chain-of-thought prompting can hurt LLM’s performance in humor reasoning, and that human-annotated joke explanations are significantly better than those produced by GPT-4o and ERNIE_{4-turbo}, urging future research on culturally specific humor understanding.

2 Related Works

Humor Datasets. Humor analysis in natural language processing (NLP) encompasses a wide range of tasks, each focused on different aspects of humor. For instance, researchers have proposed datasets

such as “16000 One-Liners” (Mihalcea and Straparava, 2005), “Pun of the Day” (Yang et al., 2015), and “Ted Laughter” (Chen and Lee, 2017) focused on humor detection to determine whether a given text is humorous or not. Datasets such as “Big Bang Theory” (Bertero and Fung, 2016) aim at pinpointing the punchline in a joke. Tasks for assessing humor intensity include humor level rating, comparison, and ranking. For example, datasets like HumorNorm (Engelthaler and Hills, 2018) and #HashtagWars (Potash et al., 2017) quantify humor scores and compare comedic elements, while UR-FUNNY ranks punchlines based on their perceived impact. Datasets such as “Humicroedit” (Hossain et al., 2019), “C³” (Wang et al., 2022), and “Talk-Funny” (Chen et al., 2024) focus on humor generation, the task of generating or rewriting humorous texts. In addition, we present a comprehensive overview of the existing datasets related to humor in Table 2. We highlight that most existing datasets are in English. Chinese humor, on the other hand, is less explored. Our dataset, *Chumor* is the first humor explanation dataset in Chinese.

Culturally Specific Datasets. Recent works underscore the challenges of culturally specific reasoning in LLMs (Shen et al., 2024; AlKhamissi et al., 2024; Pawar et al., 2024; Vayani et al., 2024). These challenges stem from the overrepresentation of Western-centric knowledge and translation ar-

Dataset	Sources	Lan.	#(k)	Tasks
One Liners (2005)	Web	en	16	HR
Pun of the Day (2015)	Web	en	4.8	HR PD
Big Bang Theory (2016)	TV	en	44	PD
Ted Laughter (2017)	TED	en	9.4	HR PD
#HashtagWars (2017)	TV	en	13	HC
HumorNorm (2018)	CS [†]	en	5	HC
UR-FUNNY (2019)	TED	en	17	PD
Humicroedit (2019)	Reddit	en	15	HG
rJokes (2020)	Reddit	en	57	HC
Memotion (2020)	Memes	en	9.8	HC
MUMOR (2021)	TV	en zh	30	HR
NYT-Captions (2023)	NYT	en	0.7 2.6	HE HC
<i>C</i> ³ (2022)	Books	zh	9.3	HG
TalkFunny (2024)	Apps	zh	4.1	HG HR
TCHD (2023)	–	zh	26	HC PD
TTWS (2019)	Books	zh	9.1	PD
CHM (2020)	Apps Web	zh	3.3	HC
Memeplate (2022)	Apps Web	zh	5.2	HC
<i>Chumor</i> (us)	Web	zh	3.3	HE

Table 2: Existing datasets related to humor. For the shorthands in the table, abbreviations represent the following tasks, HR: humor recognition; PD: punchline detection; HC: humor comparison; HG: humor generation; HE: humor explanation [†]: Crowd-source.

tifacts, which limit the fairness and effectiveness of multilingual evaluations (Mihalcea et al., 2024). Researchers have proposed various culturally specific datasets such as Global-MMLU (Singh et al., 2024) to evaluate LLMs’ cultural knowledge. *Chumor* adds to this line of effort as it involves rich knowledge specific to Chinese culture.

3 Chumor Dataset

Data Collection. We construct our dataset by including RZB jokes from “Best Annual Threads” between 2018 and 2021 that have been previously crawled[†]. In addition, we directly collect all threads in the “Moderator’s Recommendation” section from RZB. Each thread in RZB consists of “标题” (title), “一楼” (content), and several “跟帖” (follow-up posts). For threads from Best Annual Threads, the jokes are listed in the follow-up posts, which are selected by the forum moderator. For threads from Moderator’s Recommendation, the jokes consist of the title and the content of each thread. We remove the content if it repeats the title.

[†]<https://github.com/Leymore/ruozhiba>

Data Cleaning. We store both the title and the content of the raw data. However, due to the posting restrictions of the platform requiring non-empty content, many posts contain meaningless placeholder texts such as “.”, “!”, “0”, “RT”, and others. We automatically identify and remove these patterns, and only keep the title which is the joke itself. Due to the length limitations on the original platform, many post titles are truncated from the beginning parts of the content. We identify these instances and replace the truncated title with the complete content to get the joke. We also remove duplicates that appear both in the “Moderator’s Recommendation” and the “Best Annual Posts”.

We manually remove the threads related to forum management and rules, threads that include excessively offensive content, threads with incomplete content, and threads that focus more on philosophical insight rather than humor.

Humor Explanation Classification. We design a humor explanation classification task that can be easily used to test LLMs’ capabilities in humor understanding. Specifically, we use two LLMs, GPT-4o and ERNIE4-turbo to generate explanations for our collected jokes. We manually annotate the generated explanations as either “fully explain the joke” (good) or “partially explain or not explain the joke” (bad) based on a majority vote among five of the authors who are native Chinese speakers. Each joke, along with its explanation, forms an individual instance in *Chumor*, leading to a total of 3,339 instances. Among these, 1,454 items are labeled as good and 1,887 as bad explanations.

Data Examples from Chumor. We present examples from *Chumor* in Table 3.

Humor Categorization in Chumor. We categorize the jokes in RZB into six types in Table 1, with an example provided for each type.

4 Experiments

Models. We test ten LLMs, five from the open-source LLM families and five from the closed-source LLM families, all capable of handling Chinese. Specifically, we include the open-source LLMs of **Yi34B** (01.ai, 2024) from 01.AI, **Nemotron70B** (NVIDIA, 2024) from NVIDIA, **Athene70B** (Nexusflow, 2024) from Nexusflow, **Qwen2.572B** (Qwen, 2024) from Alibaba, **Mistral123B** (Mistral.ai, 2024) from Mistral AI, alongside the closed-source LLMs of

Joke	Explanation	Label	Source
(zh) “你自己掰了一千吨玉米?” “哈哈我瞎掰的”	(zh) 这个笑话的幽默之处在于对话中的双关语和出人意料的回答。“掰了一千吨玉米”看似是一个惊人的壮举，而“哈哈我瞎掰的”则瞬间揭露了这其实是一个夸张的玩笑，这种巧妙的反转和语言的戏谑带来了幽默感。		
(en) “You shucked a thousand tons of corn all by yourself?!” “Haha, let’s just say there’s been plenty of shuckin’ going on.”	(en) The humor in this joke lies in the pun and the unexpected response in the conversation. “Shucked a thousand tons of corn” seems like an extraordinary achievement, but the reply, “haha, let’s just say there’s been plenty of shuckin’ going on,” instantly reveals it to be an exaggerated joke. This clever twist and playful use of language create the comedic effect.	good	ERNIE4-turbo
(zh) 小明突然出现在我面前以至于我突然出现在小明面前。	(zh) 这个笑话的幽默之处在于它利用了语言的对称性和意外性。小明突然出现导致叙述者的惊讶反应，而叙述者的惊讶反应又反过来让小明感到意外， <u>形成了一个有趣的循环</u> 。		
(en) Xiaoming suddenly appeared in front of me, causing me to suddenly appear in front of him.	(en) The humor in this joke lies in its use of linguistic symmetry and unexpectedness. Xiao Ming’s sudden appearance triggers a surprised reaction from the narrator, which in turn surprises Xiao Ming, <u>creating an amusing loop</u> .	bad	GPT-4o

Table 3: Examples from *Chumor*. The second example’s explanation is bad because the joke does not “creating an amusing loop”. Instead, it relies on linguistic symmetry and the use of a straightforward fact to subvert expectations. The audience anticipates an unexpected outcome due to the setup, but the latter part “suddenly appear in front of him” flips the perspective by stating the straightforward fact that because Xiao Ming is in front of the person so the person is in front of Xiao Ming too.

Gemini1.5-pro (Google, 2024) from Google, **GLM-4plus** (BigModel, 2024) from Tsinghua University, **GPT-4turbo**, **GPT-4o** (OpenAI, 2023, 2024) from OpenAI, **ERNIE4-turbo** (Baidu, 2024) from Baidu. For all the open-source LLMs, we use the instruction-tuned version in our evaluation.

Evaluation Methods. We evaluate these LLMs using two prompting methods: direct prompting (DP) by

Direct Prompting (DP)

你将看到一个笑话以及对这个笑话的解释。请判断这个解释是否完全解释了笑话。根据判断，选择“完全解释”或“部分/没有解释”，不需要解释为什么对或者不对。

笑话: [joke]

笑话解释: [explanation]

Translation

You will see a joke and an explanation of the joke. Please determine whether this explanation fully explains the joke. Based on your judgment, choose either “fully explain” or “partially/does not explain.” You do not need to explain why it is correct or incorrect.

Joke: [joke]

Explanation: [explanation]

and chain-of-thought (CoT) prompting (Wei et al., 2022) by adding the phrase “请逐步思考，写下过程” “Please think step by step, write down your reasoning process” before determining the label. Appendix F provides the complete prompts. We cal-

culate accuracy scores as part of our evaluation. In addition, we provide the false positive rate (FPR), false negative rate (FNR), and Matthews Correlation Coefficient (MCC) in Appendix H in Table 4. The MCC score considers true positives, true negatives, false positives, and false negatives, providing a score between -1 and +1. A score of +1 indicates perfect predictions, 0 reflects random guessing, and -1 means complete disagreement. The best MCC score achieved by LLMs is 0.29, which is close to random guessing, and is significantly lower than the human average of 0.60.

5 Results and Discussions

Overall Model Performance. Figure 1 presents the accuracy of different LLMs on *Chumor* in DP and CoT settings. Appendix H presents additional results and analysis.

Overall, we observe that all models perform poorly on Chinese humor comprehension, with accuracy scores ranging between 44.6% and 60.3%. ERNIE4-turbo and Gemini1.5-pro achieve the highest accuracy of 60.3%, and are just 10 points above the random baseline and far below human performance of 78.3%, highlighting the difficulty of *Chumor* and the limitations of these LLMs in understanding Chinese humor.

Error Analysis by Joke Type. To better understand how LLMs perform on each joke type listed in Table 1, we sample 200 jokes for error analysis. Figure 2 and Figure 17 in Appendix H present the results. The distribution of joke types can be found in Appendix G Figure 15.

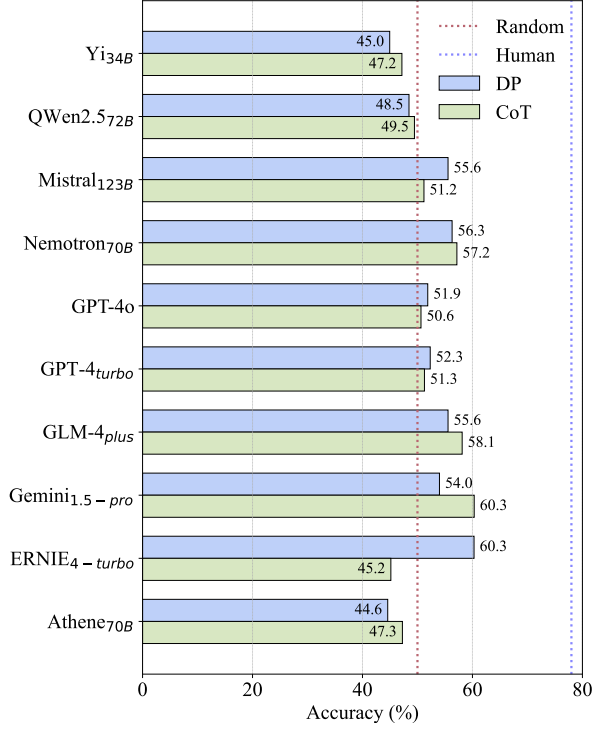


Figure 1: The accuracy of different models’ test results in the DP and CoT settings. ERNIE_{4-turbo} and Gemini_{1.5-pro} achieve the highest accuracy of 60.3%.

We highlight that model performance varies significantly across different joke types. While models generally perform well on *Situational* jokes, achieving 60.0% to 70.0% accuracy in both DP and CoT settings, their performance difference on other joke types is more pronounced. For instance, GLM-4_{plus} achieves 65.0% accuracy on *Homophonic* jokes in the DP setting, whereas Yi_{34B} only reaches 30.0%. Nemotron_{70B} performs well on *Cultural* jokes in the CoT setting with 72.0% accuracy, but Athene_{70B} and ERNIE_{4-turbo} achieve with only 43.0% and 42.0%, respectively. Such performance variance highlights LLMs’ varied capabilities in specific domains such as cultural reasoning and situational reasoning, revealing the respective limitations of these LLMs.

5.1 Have LLMs achieved human-level understanding of humor?

Answer: No. To compare the performance of LLMs with humans, we conduct a human study involving three Chinese native speakers unfamiliar with this work to annotate a randomly chosen subset of 200 examples. Human annotators demonstrate significantly better performance, with an average accuracy of 78.3% and an MCC score of 0.60, significantly better than the LLMs’ best per-

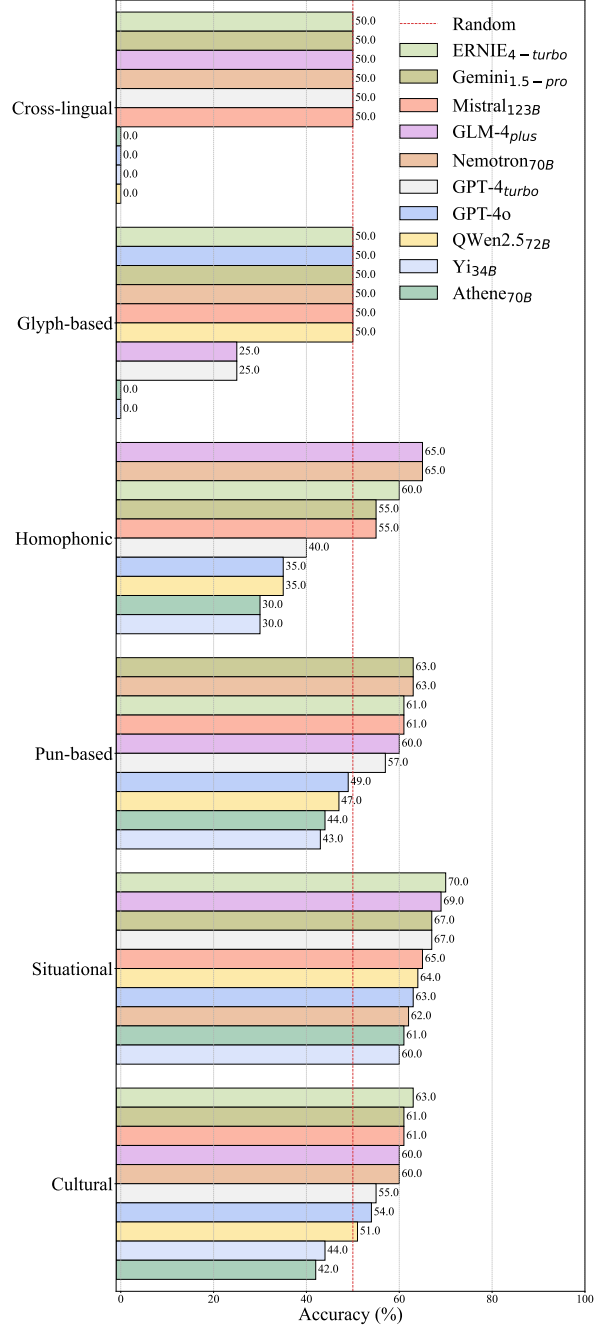


Figure 2: DP accuracy on different joke types (%). Here, we sample 200 jokes for error analysis. We highlight that model performance varies significantly across different joke types.

formance of 60.3% accuracy and an MCC score of 0.29 (Figure 16 in Appendix H). Our results indicate that there is a large room of performance improvement for LLMs on Chinese humor understanding.

5.2 Does chain-of-thought (CoT) help LLMs’ humor understanding?

Example	(zh) “ 大人有大量 ，小孩倒是不多了。”食人族厨师清点食材说。
	(en) " Adults are plenty , but children are quite scarce," said the cannibal chef while counting ingredients.
Joke Explanation	<p>(zh) 这个笑话的幽默之处在于其运用了惊悚与诙谐的巧妙结合，以及意料之外的语言转折。食人族厨师清点食材的场景已经足够诡异，但他却以“大人有大量，小孩倒是不多了”这样似乎很平常的口吻来评论，这种对比和反差产生了荒诞的幽默感，同时也在探讨一种颠覆性的、黑暗式的玩笑。</p> <p>(en) The humor in this joke lies in its clever combination of horror and wit, as well as the unexpected linguistic twist. The scene of a cannibal chef counting ingredients is already eerie, but the casual remark, "Adults are plenty, but children are quite scarce," delivered in an ordinary tone, creates a stark contrast. This juxtaposition generates absurd humor while exploring a subversive, dark style of comedy.</p>
Model's Answer DP ✓	<div data-bbox="335 492 526 560">(zh) 部分/没有解释 (en) Partially/Not Explained</div> <div data-bbox="542 492 1404 627"> <p>(zh) ...这句话使用了一个常用的谚语“大人有大量”来双关...比较您的理解和解释：笑话解释涵盖了我的分析要点，指出了对比和反差产生的荒诞幽默感，并提到了黑暗式幽默。 解释完全涵盖了笑话的所有幽默点，因此选择“完全解释”。</p> <p>(en) ...This sentence uses a common proverb, "adults are plenty," as a pun... Comparison of your understanding and the explanation: The joke explanation covers the key points of my analysis, highlighting the absurd humor created by contrast and reversal, and mentioning the element of dark humor. The explanation fully addresses all aspects of the joke's humor, so "Fully Explained" is selected.</p> </div>

Figure 3: Over-analyzing example by GPT-4o. The GPT-4o model chooses the correct answer in the DP prompting, but chooses the incorrect answer due to over-analyzing in the CoT prompting.

Answer: No. We observe that CoT does not necessarily improve model performance and, in most cases, even leads to performance decay. For instance, as shown in Figure 1, the accuracy of ERNIE_{4-turbo} decreases from 60.3% to 45.2% when we switch to CoT prompting, Mistral_{123B}'s performance drops from 55.6% to 51.2%, GPT-4o's performance drops from 51.9% to 50.6%, GPT-4_{turbo}'s performance falls from 52.3% to 51.3%. Moreover, the MCC scores present a clearer trend of performance decline under CoT prompting. As shown in Figure 16 in Appendix H, eight of the ten LLMs' MCC scores decrease under CoT prompting. We hypothesize that CoT prompts may not help the model's reasoning when the model lacks a fundamental grasp of humor understanding.

We observe that under CoT prompting, models like GPT-4o tend to justify incorrect explanations as “correct”, leading to an increase in false-positive rate from 80.0% for DP prompting to 85.0% for CoT prompting (Table 4 in Appendix H). ERNIE_{4-turbo} exhibits the largest false-positive rate, rising from 59.8% to 96.9% (Table 4 in Appendix H). Figure 3 provides an example where CoT confuses the GPT-4o model. Under the DP prompting, the GPT-4o model chooses the answer correctly. However, CoT prompting causes the model to over-analyze and justify an incorrect explanation.

On the other hand, models like Nemotron_{70B} may be overly critical of explanations under CoT prompting, resulting in a false-negative rate from 20.9% for DP prompting to 46.1% for CoT prompting (Table 4 in Appendix H). We highlight that a

recent work demonstrates that CoT can degrade performance in tasks requiring subtle comprehension (Sprague et al., 2024), which aligns with our findings on its limitations in humor interpretation. Figure 14 in Appendix E discusses an example corresponding to the model being overly critical.

5.3 Case study: can GPT-4o and ERNIE_{4-turbo} explain jokes as well as humans?

Answer: No. Apart from testing multiple LLMs on *Chumor*, we conduct case studies on GPT-4o and ERNIE_{4-turbo} to assess the quality of their joke explanations compared to humans. We prompt them to explain the humor in two sentences, consistent with the format of human explanations. Here is the prompt we feed to both LLMs:

Prompt

请用两句话解释这个笑话的幽默之处: [joke]

Translation

Please explain the joke in two sentences: [joke]

Data Annotation. As demonstrated by Hessel et al. (2023), crowd-sourcing typically cannot produce high-quality explanations, following Hessel et al. (2023), one of the authors annotates all the explanations to ensure the quality and consistency.

This is a substantial effort: the author ended up annotating the explanations for 1,951 jokes. The resulting corpus has a mean of 78 Chinese characters of explanation per joke, and the total length, 151,730 Chinese characters, is comparable

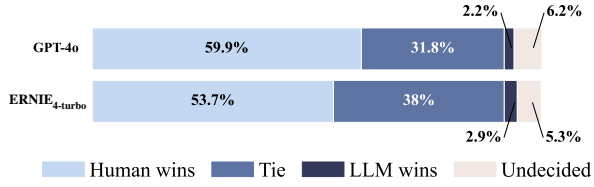


Figure 4: Annotated preference for whether human explanation is preferred (“Human wins”) or the explanation from LLMs is preferred (“LLM wins”). Humans’ explanation is significantly preferred over LLMs’.

in length to a novella[†].

Evaluation Setup. To fairly evaluate which explanation is better, we conduct A/B testing by presenting the humor explanation from one LLM and from human to six college students, asking them to annotate their preference of the explanation for each joke. These college students are native Chinese speakers who grew up in China, therefore they have a deep understanding of the cultural terms and trending terms in China. We note that the preference annotation requires a substantial effort as each annotator reads through a total length of around 300k Chinese characters[†]. We end up with three preference annotations for each joke. The preference annotation achieve a 61.4% agreement rate among annotators (Appendix B).

We use the winning rate as our measure to compare LLMs’ explanation versus human explanation, taking the majority vote among all annotators for each example. In addition, if all annotators disagree, we assign an “Undecided” label. Appendix C provides the annotation instructions we present to the annotators.

Overall Results. Figure 4 reports the winning rate of explanations from human versus GPT-4o and ERNIE_{4-turbo}. We can see that human explanations are significantly better than those from both LLMs, with humans winning over 50% of the time, while LLMs win in only 2-3% of cases.

Error Analysis. Figure 5 shows the overall distribution of error types for GPT-4o and ERNIE_{4-turbo} on *Chumor* in terms of their humor explanations. This error analysis is conducted by an individual who is not involved in writing the original explanations, ensuring an unbiased evaluation. GPT-4o

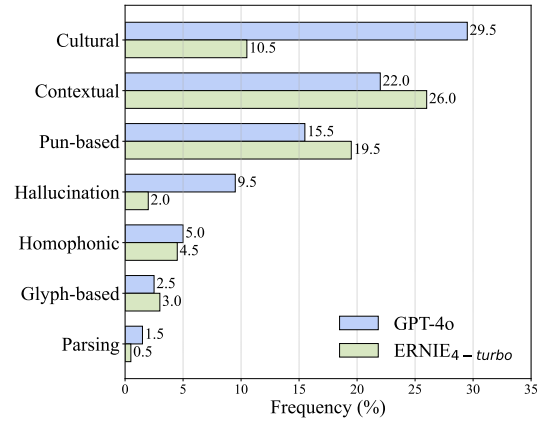


Figure 5: Distribution of error types for GPT-4o and ERNIE_{4-turbo}. We sample 200 examples to calculate the distribution of these error types. We note that each example may correspond to multiple error types. We highlight that ERNIE_{4-turbo} demonstrates a lower error rate on cultural jokes, while GPT-4o demonstrates a lower error rate on contextual or pun-based jokes.

is more prone to errors categorized as “cultural unawareness” (29.5% of all its explanations) compared to ERNIE_{4-turbo} (10.5%). We suspect that ERNIE_{4-turbo} is more familiar with Chinese culture as it is likely trained on a larger Chinese corpus than GPT-4o. However, GPT-4o performs better on cases requiring an understanding of contexts or puns, suggesting its strong reasoning ability. We provide three error cases for GPT-4o here and additional cases for both GPT-4o and ERNIE_{4-turbo} in Appendix E. In the following examples in Figure 6, Figure 7 and Figure 8, we highlight key phrases that induce humor in green, and underscore the errors in red.

Error Type I: Cultural Unawareness. LLMs may fail to explain a joke due to a lack of awareness of certain cultural knowledge. For instance, the example in Figure 6 requires knowledge of a superstitious belief in Chinese culture, *getting a haircut in the first lunar month brings death to your uncle*, and the explanation from GPT-4o fails to connect to this Chinese cultural belief. We hypothesize that while LLMs are pre-trained on Internet-scale corpora, such culturally specific knowledge can still be challenging for them to grasp. Moreover, even when they have acquired such cultural knowledge, they may fail to relate to it as we humans do during the reasoning process.

Error Type II: Pun-based Error. LLMs may fail to identify multiple meanings of a single word within a joke, causing them to fail on pun-based

[†]The total length of our explanations surpasses the Chinese version of *The Great Gatsby* (100k Chinese characters), and is about half the length of the Chinese version of *Wuthering Heights* (325k Chinese characters).

[†]This is about the same length of the Chinese version of *Wuthering Heights* (325k Chinese characters).

Example	(zh) 小明在 <u>正月接发</u> 竟导致 <u>舅舅复活</u> 。
	(en) Xiaoming <u>got hair extensions during the first lunar month</u> , which astonishingly <u>brought his uncle back to life</u> .
Correct Humor Explanation	“Getting a haircut in the first lunar month will brings death to your uncle.” is a popular superstitious saying in Chinese culture. In this joke, Xiao Ming gets hair extensions in the first month, which reverses the original logic and absurdly results in “bringing his uncle back to life.”
GPT-4o’s Answer	(zh) ...它 <u>利用了“接发”和“接发”的双关语</u> ：一方面是指理发店的接发服务，另一方面是指正月里“接福”的传统习俗。
	(en) ... <u>the pun on “接发”</u> : it refers to both hair extension at a salon and the traditional practice of “receiving blessings” during the first lunar month.
Failure Reasons	Fail to grasp expressions, slang, or specific usages unique to a particular cultural context.

Figure 6: Culture unawareness example.

Example	(zh) 你可以在 <u>steam</u> 上找到 <u>GTA</u> ，所以水是DNA。
	(en) You can find <u>GTA</u> on <u>Steam</u> , so water is DNA.
Correct Humor Explanation	Normally, “Steam” refers to a gaming platform and “GTA” refers to the “Grand Theft Auto” game series. The joke uses a pun, where “steam” literally means the gaseous form of water, and “G”, “T”, and “A” represent Guanine, Thymine, and Adenine, respectively—three of the four nucleotides (AGTC) that make up DNA.
GPT-4o’s Answer	(zh) ...这种 <u>明显错误的类比</u> 制造了荒诞的效果，引人发笑。
	(en) ...This <u>clearly wrong analogy</u> creates an absurd effect.
Failure Reasons	Fail to discover multiple meanings of a word in the jokes

Figure 7: Pun-based error example.

jokes where humor lies in inverting the conventional usage of words. In Figure 7, GPT-4o fails to grasp the transition from the video game terms “Steam”, “GTA” to the scientific terminologies “G”, “T”, “A” that form DNA. Typically, “Steam” refers to a game platform, and “GTA” refers to the game series “Grand Theft Auto”. The joke employs a pun on words where “steam” in its literal sense means water vapor, and “GTA” can represent not only the video game, but guanine, thymine, and adenine, which are nucleotides involved in the structure of DNA. Such jokes require LLMs to identify puns and the reason for the association of the multiple meanings. Furthermore, the process requires LLMs to bridge the logic gap between these terms, such as

Example	(zh) 家里的猪油没了，小明只能把 <u>植物</u> 油倒快点当 <u>猪油</u> 用了。
	(en) With the lard gone, Xiaoming had to pour the <u>vegetable</u> oil quickly to use it like <u>lard</u> .
Correct Humor Explanation	When pronounced quickly, “vegetable” (pronounced as <i>chih woo</i>) sounds very similar to “pig” (pronounced as <i>choo</i>) in Chinese, leading Xiaoming to comically try using vegetable oil as lard by pouring it quickly.
GPT-4o’s Answer	(zh) ...小明误解了猪油和植物油的区别，认为只要倒得快，植物油就能变成猪油。笑点 <u>在于他天真的逻辑和对常识的搞笑误解</u> 。
	(en) ... Xiaoming misunderstood that if vegetable oil is poured quickly, it can turn into lard, <u>highlighting his naive logic and misinterpretation of common sense</u> .
Failure Reasons	Fail to exploit homophones or near-homophones in the jokes.

Figure 8: Homophonic error example.

“steam” and “GTA”, and an unusual context, like “water is DNA”. The overall process requires both scientific knowledge and creative thinking, which are challenging for LLMs.

Error Type III: Homophonic Error. The example in Figure 8 requires LLMs to reason over the pronunciation as “植物” (pronounced as *chih woo*, meaning “vegetable”) sounds very similar to “猪” (pronounced as *choo*, meaning “pig”) in Chinese when we speak it fast enough. The humor arises from the contrast between the similarity in pronunciation and the disparity in meaning between the two terms. Such contrasts may be sparse in the training corpus of LLMs, and also demand a deep connection across different modalities to link pronunciation with the meaning behind these terms, which poses significant challenges to LLMs.

6 Conclusion

We introduce *Chumor*, a Chinese humor understanding dataset that captures intellectually challenging and culturally specific humor in Chinese. Our analysis reveals that *Chumor* remains difficult even for advanced LLMs, with a significant performance gap between LLMs and humans. Furthermore, we find that chain-of-thought reasoning does not improve LLMs’ humor comprehension and, in some cases, leads to over-analysis and incorrect interpretations. Additionally, models such as GPT-4o and ERNIE4-turbo struggle to explain jokes as effectively as humans, highlighting fundamental challenges in humor reasoning. These findings un-

derscore the unique difficulties that Chinese humor presents to LLMs. We hope that *Chumor* can advance non-English humor research and contribute to evaluating LLMs’ reasoning abilities across diverse cultural backgrounds.

Limitations

We try our best to test the Chinese humor understanding ability of different LLMs. However, due to the limited budget and API access, we cannot evaluate all possible LLMs in this paper. We encourage future research to conduct further evaluations of humor understanding abilities in LLMs. In the meantime, we emphasize that our research focuses primarily on demonstrating how humor understanding remains a significant challenge, even for SOTA LLMs. Our work shows that along with many other problems (Ignat et al., 2024), humor understanding, especially non-English and culturally specific humor understanding, remains an unsolved problem in the era of LLMs. We hope *Chumor* can contribute to non-English humor understanding evaluations for future multilingual LLMs.

Ethics Statement

We have made every effort to filter out excessively offensive content in RZB. However, due to the subjective nature of humor, some of our jokes may still be perceived as offensive by individuals with different cultural or personal standards. To address these concerns, we strongly recommend that researchers use *Chumor* with cultural sensitivity, recognizing that the jokes in the dataset reflect the sociocultural context in which they were created. We encourage users of *Chumor* to approach the dataset with caution, remaining mindful of its potential to cause offense or harm, particularly when applying it in research or applications that involve diverse audiences or address sensitive topics. We wish to foster an ethical and responsible approach to data collection and usage, and we welcome constructive feedback from the research community and stakeholders to continually improve *Chumor* and mitigate potential harm.

Acknowledgement

The GPT experiments are supported by credit from OpenAI through OpenAI Researcher Access assigned to Naihao Deng. We appreciate Qiang Liu, and Xiaoyue Shi for helping with the human study.

References

- 01.ai. 2024. Yi-34b model card. <https://huggingface.co/01-ai/Yi-34B>. Accessed: 2024-12-10.
- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. *Phi-4 technical report*. Preprint, arXiv:2412.08905.
- Badr AlKhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. *Investigating cultural alignment of large language models*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422, Bangkok, Thailand. Association for Computational Linguistics.
- 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, and 29 others. 2023. *Qwen technical report*. Preprint, arXiv:2309.16609.
- Yuelin Bai, Xinrun Du, Yiming Liang, Yonggang Jin, Ziqiang Liu, Junting Zhou, Tianyu Zheng, Xincheng Zhang, Nuo Ma, Zekun Wang, and 1 others. 2024. Coig-cqia: Quality is all you need for chinese instruction fine-tuning. *arXiv preprint arXiv:2403.18058*.
- Baidu. 2024. Ernie-4.0-turbo. <https://cloud.baidu.com/doc/WENXINWORKSHOP/s/71xwwtafj>. Accessed: 2024-12-10.
- Dario Bertero and Pascale Fung. 2016. *Deep learning of audio and language features for humor prediction*. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 496–501, Portorož, Slovenia. European Language Resources Association (ELRA).
- BigModel. 2024. Glm-4 model documentation. <https://bigmodel.cn/dev/howuse/glm-4>. Accessed: 2024-12-10.
- Lei Chen and Chong Min Lee. 2017. *Predicting audience’s laughter during presentations using convolutional neural network*. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 86–90, Copenhagen, Denmark. Association for Computational Linguistics.
- Yuyan Chen, Zhixu Li, Jiaqing Liang, Yanghua Xiao, Bang Liu, and Yunwen Chen. 2023. *Can pre-trained language models understand chinese humor?* In *Proceedings of the Sixteenth ACM International Conference on Web Search and Data Mining, WSDM ’23*, page 465–480, New York, NY, USA. Association for Computing Machinery.

- Yuyan Chen, Yichen Yuan, Panjun Liu, Dayiheng Liu, Qinghao Guan, Mengfei Guo, Haiming Peng, Bang Liu, Zhixu Li, and Yanghua Xiao. 2024. [Talk funny! a large-scale humor response dataset with chain-of-humor interpretation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17826–17834.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *Preprint*, arXiv:1803.05457.
- Peter T. Daniels and William Bright. 1996. *The world’s writing systems*. Oxford University Press.
- Naihao Deng, Zhenjie Sun, Ruiqi He, Aman Sikka, Yulong Chen, Lin Ma, Yue Zhang, and Rada Mihalcea. 2024. Tables as images? exploring the strengths and limitations of llms on multimodal representations of tabular data. *arXiv preprint arXiv:2402.12424*.
- Naihao Deng, Xinliang Zhang, Siyang Liu, Winston Wu, Lu Wang, and Rada Mihalcea. 2023. [You are what you annotate: Towards better models through annotator representations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12475–12498, Singapore. Association for Computational Linguistics.
- Xinrun Du, Zhouliang Yu, Songyang Gao, Ding Pan, Yuyang Cheng, Ziyang Ma, Ruibin Yuan, Xingwei Qu, Jiaheng Liu, Tianyu Zheng, and 1 others. 2024. Chinese tiny llm: Pretraining a chinese-centric large language model. *arXiv preprint arXiv:2404.04167*.
- Tomas Engelthaler and Thomas T Hills. 2018. Humor norms for 4,997 english words. *Behavior research methods*, 50:1116–1124.
- WILLIAM F. FRY. 1994. [The biology of humor](#). *HUMOR*, 7(2):111–126.
- Marc Gelkopf and 1 others. 2011. The use of humor in serious mental illness: A review. *Evidence-Based Complementary and Alternative Medicine*, 2011.
- Matthew Gervais and David Sloan Wilson. 2005. The evolution and functions of laughter and humor: A synthetic approach. *The Quarterly review of biology*, 80(4):395–430.
- Hamideh Ghanadian, Isar Nejadgholi, and Hussein Al Osman. 2023. [ChatGPT for suicide risk assessment on social media: Quantitative evaluation of model performance, potentials and limitations](#). In *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*, pages 172–183, Toronto, Canada. Association for Computational Linguistics.
- Google. 2024. Gemini 1.5 pro model documentation. <https://ai.google.dev/gemini-api/docs/models/gemini#gemini-1.5-pro>. Accessed: 2024-12-10.
- Md Kamrul Hasan, Wasifur Rahman, AmirAli Bagher Zadeh, Jianyuan Zhong, Md Iftekhar Tanveer, Louis-Philippe Morency, and Mohammed (Ehsan) Hoque. 2019. [UR-FUNNY: A multimodal language dataset for understanding humor](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2046–2056, Hong Kong, China. Association for Computational Linguistics.
- Jack Hessel, Ana Marasovic, Jena D. Hwang, Lillian Lee, Jeff Da, Rowan Zellers, Robert Mankoff, and Yejin Choi. 2023. [Do androids laugh at electric sheep? humor “understanding” benchmarks from the new yorker caption contest](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 688–714, Toronto, Canada. Association for Computational Linguistics.
- Nabil Hossain, John Krumm, and Michael Gamon. 2019. [“president vows to cut <taxes> hair”: Dataset and analysis of creative text editing for humorous headlines](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 133–142, Minneapolis, Minnesota. Association for Computational Linguistics.
- Oana Ignat, Zhijing Jin, Artem Abzaliev, Laura Biester, Santiago Castro, Naihao Deng, Xinyi Gao, Aylin Ece Gunal, Jacky He, Ashkan Kazemi, Muhammad Khalifa, Namho Koh, Andrew Lee, Siyang Liu, Do June Min, Shinka Mori, Joan C. Nwatu, Veronica Perez-Rosas, Siqi Shen, and 3 others. 2024. [Has it all been solved? open NLP research questions not solved by large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8050–8094, Torino, Italia. ELRA and ICCL.
- Herbert M Lefcourt. 2001. *Humor: The psychology of living buoyantly*. Springer Science & Business Media.
- Zefeng Li, Hongfei Lin, Liang Yang, Bo Xu, and Shaowu Zhang. 2022. Memeplate: A chinese multimodal dataset for humor understanding in meme templates. In *Natural Language Processing and Chinese Computing*, pages 527–538, Cham. Springer International Publishing.
- Hanmeng Liu, Ruoxi Ning, Zhiyang Teng, Jian Liu, Qiji Zhou, and Yue Zhang. 2023a. Evaluating the logical reasoning ability of chatgpt and gpt-4. *arXiv preprint arXiv:2304.03439*.
- Siyang Liu, Naihao Deng, Sahand Sabour, Yilin Jia, Minlie Huang, and Rada Mihalcea. 2023b. [Task-adaptive tokenization: Enhancing long-form text generation efficacy in mental health and beyond](#). In *Proceedings of the 2023 Conference on Empirical Meth-*

- ods in *Natural Language Processing*, pages 15264–15281, Singapore. Association for Computational Linguistics.
- Xin Liu, Baosong Yang, Dayiheng Liu, Haibo Zhang, Weihua Luo, Min Zhang, Haiying Zhang, and Jinsong Su. 2021. [Bridging subword gaps in pretrain-finetune paradigm for natural language generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6001–6011, Online. Association for Computational Linguistics.
- Paul E McGhee. 1971. Development of the humor response: A review of the literature. *Psychological Bulletin*, 76(5):328.
- Rada Mihalcea, Oana Ignat, Longju Bai, Angana Borah, Luis Chiruzzo, Zhijing Jin, Claude Kwizera, Joan Nwatu, Soujanya Poria, and Tamar Solorio. 2024. Why ai is weird and should not be this way: Towards ai for everyone, with everyone, by everyone. *arXiv preprint arXiv:2410.16315*.
- Rada Mihalcea and Carlo Strapparava. 2005. [Making computers laugh: Investigations in automatic humor recognition](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 531–538, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Mistral.ai. 2024. Mistral-large-instruct-2407 model card. <https://huggingface.co/mistralai/Mistral-Large-Instruct-2407>. Accessed: 2024-12-10.
- Nexusflow. 2024. Athene-70b model card. <https://huggingface.co/Nexusflow/Athene-70B>. Accessed: 2024-12-10.
- NVIDIA. 2024. Llama-3.1-nemotron-70b-instruct-hf model card. <https://huggingface.co/nvidia/Llama-3.1-Nemotron-70B-Instruct-HF>. Accessed: 2024-12-10.
- OpenAI. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- OpenAI. 2024. [Gpt-4o system card](#). *arXiv preprint arXiv:2410.21276*.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrana, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2024. Survey of cultural awareness in language models: Text and beyond. *arXiv preprint arXiv:2411.00860*.
- Peter Potash, Alexey Romanov, and Anna Rumshisky. 2017. [SemEval-2017 task 6: #HashtagWars: Learning a sense of humor](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 49–57, Vancouver, Canada. Association for Computational Linguistics.
- Qwen. 2024. Qwen2.5-72b-instruct model card. <https://huggingface.co/Qwen/Qwen2.5-72B-Instruct>. Accessed: 2024-12-10.
- Dragomir Radev, Amanda Stent, Joel Tetreault, Aasish Pappu, Aikaterini Iliakopoulou, Agustin Chanfreau, Paloma de Juan, Jordi Vallmitjana, Alejandro Jaimes, Rahul Jha, and Robert Mankoff. 2016. [Humor in collective discourse: Unsupervised funniness detection in the new yorker cartoon caption contest](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 475–479, Portorož, Slovenia. European Language Resources Association (ELRA).
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *Preprint*, arXiv:2311.12022.
- Chhavi Sharma, Deepesh Bhageria, William Scott, Srinivas PYKL, Amitava Das, Tanmoy Chakraborty, Viswanath Pulabaigari, and Björn Gambäck. 2020. [SemEval-2020 task 8: Memotion analysis- the visuo-lingual metaphor!](#) In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 759–773, Barcelona (online). International Committee for Computational Linguistics.
- Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. [Understanding the capabilities and limitations of large language models for cultural commonsense](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5668–5680, Mexico City, Mexico. Association for Computational Linguistics.
- Shivalika Singh, Angelika Romanou, Clémentine Fourier, David I. Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Wei-Yin Ko, Madeline Smith, Antoine Bosselut, Alice Oh, Andre F. T. Martins, Leshem Choshen, Daphne Ippolito, and 4 others. 2024. [Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation](#). *Preprint*, arXiv:2412.03304.
- Zayne Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2024. [To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning](#). *Preprint*, arXiv:2409.12183.
- Honglin Sun and Daniel Jurafsky. 2004. [Shallow semantic parsing of Chinese](#). In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pages 249–256, Boston, Massachusetts, USA. Association for Computational Linguistics.

- Weiwei Sun, Zhifang Sui, Meng Wang, and Xin Wang. 2009. [Chinese semantic role labeling with shallow parsing](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 1475–1483, Singapore. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *Preprint*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Yuen-Hsien Tseng, Wun-Syuan Wu, Chia-Yueh Chang, Hsueh-Chih Chen, and Wei-Lun Hsu. 2020. [Development and validation of a corpus for machine humor comprehension](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1346–1352, Marseille, France. European Language Resources Association.
- Ashmal Vayani, Dinura Dissanayake, Hasindri Watawana, Noor Ahsan, Nevasini Sasikumar, Omkar Thawakar, Henok Biadgign Ademtew, Yahya Hmaiti, Amandeep Kumar, Kartik Kuckreja, and 1 others. 2024. All languages matter: Evaluating llms on culturally diverse 100 languages. *arXiv preprint arXiv:2411.16508*.
- Benyou Wang, Xiang Wu, Xiaokang Liu, Jianquan Li, Prayag Tiwari, and Qianqian Xie. 2022. [Can language models make fun? a case study in chinese comical crosstalk](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Orion Weller and Kevin Seppi. 2020. [The rJokes dataset: a large scale humor collection](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 6136–6141, Marseille, France. European Language Resources Association.
- Jiaming Wu, Hongfei Lin, Liang Yang, and Bo Xu. 2021. [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*, page 619–627, Berlin, Heidelberg. Springer-Verlag.
- Yufan Wu, Yinghui He, Yilin Jia, Rada Mihalcea, Yulong Chen, and Naihao Deng. 2023. [Hi-ToM: A benchmark for evaluating higher-order theory of mind reasoning in large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10691–10706, Singapore. Association for Computational Linguistics.
- Diyi Yang, Alon Lavie, Chris Dyer, and Eduard Hovy. 2015. [Humor recognition and humor anchor extraction](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2367–2376, Lisbon, Portugal. Association for Computational Linguistics.
- Dongyu Zhang, Heting Zhang, Xikai Liu, Hongfei Lin, and Feng Xia. 2019. [Telling the whole story: A manually annotated chinese dataset for the analysis of humor in jokes](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Min Zhang, Jianfeng He, Taoran Ji, and Chang-Tien Lu. 2024a. [Don’t go to extremes: Revealing the excessive sensitivity and calibration limitations of llms in implicit hate speech detection](#). *Preprint*, arXiv:2402.11406.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2024b. [Benchmarking Large Language Models for News Summarization](#). *Transactions of the Association for Computational Linguistics*, 12:39–57.
- Jun Zhao, Zhihao Zhang, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. Llama beyond english: An empirical study on language capability transfer. *arXiv preprint arXiv:2401.01055*.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). *Preprint*, arXiv:2311.07911.

A Contributions

Idea Proposal. Naihao Deng proposed the high-level idea of constructing a humor understanding benchmark sourced from RZB data.

Background Survey. Ruiqi He surveyed the humor-related tasks.

Data Processing. Ruiqi He crawled and processed the jokes from RZB.

Annotation. Ruiqi He annotated the explanations for the RZB jokes. Yushu He, Longju Bai, Jiarui Liu, Zhenjie Sun, Zhenghao Tang, He Wang, Naihao Deng conducted the preference annotations.

Experiments. Ruiqi He, Hanchen Xia, and Naihao Deng conducted the experiments.

Result Aggregation. Ruiqi He, Naihao Deng, Yushu He aggregated the results.

Paper Writing. Ruiqi He and Naihao Deng drafted the paper. Other authors provided revisions and feedback on the paper.

Naihao Deng organized the research.

B Agreement Rate Calculation

We calculate the percentage agreement rate among annotators who annotate their preferences between explanations from LLMs and humans. The results show an average inter-annotator agreement of 61.9% for GPT-4o and 60.9% for ERNIE_{4-turbo}. Given the inherent subjectivity of humor interpretation tasks (Deng et al., 2023), the combined average agreement percentage of 61.4% is decent.

C Annotation Instructions for Preference Annotation

We include the following instructions for the preference annotations of the joke explanations:

Instruction

“在这个标注中，你将会看到一个笑话和对这个笑话的幽默之处的两个解释，请你比较哪个解释更好的解释了这个笑话的幽默之处，并从以下三个标签中选择：

1. 解释1
2. 解释2
3. 一样好”

Translation

“In this annotation task, you will see a joke along with two explanations of its humor. Please compare which explanation better explains the reason why this joke is funny and choose from the following three labels:

1. Explanation 1
2. Explanation 2
3. Both are equally good.”

For each example, we randomly assign the explanations from the LLMs and the human as Explanation 1 and Explanation 2 to ensure a fair comparison.

D Discussion on Evaluation Setting

Why Zero-Shot w.o. SFT? The primary research objective of this paper is to determine how well foundational LLMs can understand Chinese humor without relying on supervised fine-tuning for this binary classification task. The focus is on investigating the innate humor-understanding ability of these models through zero-shot and zero-shot CoT prompting. This aligns with the overarching goal of examining their general capabilities without additional task-specific training. From a human behavior perspective, individuals generally acquire a sense of humor through exposure and experience rather than explicit classroom instruction (McGhee, 1971; FRY, 1994; Gervais and Wilson, 2005). Analogously, our purpose lies in evaluating the models’ intrinsic ability to recognize and interpret humor without deliberate, task-specific fine-tuning. Moreover, the experimental design follows practices from other benchmarks, such as GPQA (Rein et al., 2023), AI2ARC (Clark et al., 2018), and IFEVAL (Zhou et al., 2023), which do not provide predefined train/dev/test splits. Under these conditions, it is common to assess models in a zero-shot manner to directly evaluate their capabilities on each respective task (Touvron et al., 2023a,b; Bai et al., 2023; Abdin et al., 2024).

Example	(zh) 真可怕, 犯罪嫌疑人就在 <u>我们之中</u> , 被告席上一名法警对另一名法警说。
	(en) “Terrifying, the criminal suspect is <u>right between the two of us</u> ,” said one bailiff to another in the defendant’s dock.
Correct Humor Explanation	“between us” can refer to “either one of us” or literally means the actual physical position.
GPT-4o’s Answer	(zh) ...暗示他们自己可能是犯罪嫌疑人 ...
	(en) ...it suggests the bailiffs themselves might be criminal suspects...
Failure Reasons	Fail to address the literal meaning.

Figure 9: Insufficient contextual understanding example.

E More Error Cases

We note that many examples here encompass multiple error types, highlighting the complexity of *Chumor*.

Insufficient Contextual Understanding. LLMs may fail to ground their responses to the context when they explain the joke. For instance, in the example in Figure 9, “between us” typically means “either you or me”, but it also has the literal meaning to indicate the person standing “between us”, which is the right interpretation given that the two bailiffs are talking about the criminal. However, GPT-4o only reasons that “the criminal is either you or me” but fails to capture the literal meaning from the context. We hypothesize that in the pre-training corpus, “between us” most likely acquires the meaning of “either you or me” rather than the literal meaning in a scenario like this, which creates a bias that prevents the model from reasoning about the literal interpretation required for this specific explanation.

Glyph-based Error. LLMs may fail to interpret the visual puns based on the shape or form of Chinese characters. Glyph-based humor in Chinese leverages its logographic writing system, where characters integrate both semantic and visual elements. Unlike the phonemic alphabet used in English, Chinese characters’ pictorial and ideographic nature allows for visual puns in jokes (Daniels and Bright, 1996). In the example in Figure 10, “simplify” does not refer to simplifying the trolley problem conceptually, but to simplifying the traditional Chinese characters to simplified Chinese

Example	(zh) 我把 <u>電車難題</u> 简化了, 现在是 <u>电车难题</u> 。
	(en) I simplified <u>the trolley problem (in traditional Chinese)</u> , now it’s <u>the trolley problem (in simplified Chinese)</u> .
Correct Humor Explanation	The joke does not actually simplify the scenario of the trolley problem; instead, it converts the traditional Chinese characters “電車難題” into the simplified Chinese characters “电车难题”.
GPT-4o’s Answer	(zh) ...这个笑话的幽默之处在于通过 <u>将“電車難題”中的“難”字去掉, 变成了“电车题”</u> 。
	(en) ...The humor of this joke lies in <u>removing the character “hard” from “the hard trolley problem”, turning it into “the trolley problem”</u> .
Failure Reasons	Fail to interpret the visual puns in the shape or form of the characters.

Figure 10: Glyph-based error example.

Example	(zh) 小王 <u>订了张飞去北京的机票</u> , 给 <u>张飞</u> 省了一大笔钱。
	(en) Xiaowang <u>booked an airline ticket to Beijing, saving Zhang Fei</u> a lot of money.
Correct Humor Explanation	There are two ways to parse the first half of the sentence: (1) 小王/订了/张/飞/去北京的/机票 XiaoWang/ booked/ <u>a</u> / ticket to <u>fly</u> to Beijing. (2) 小王/订了/张/飞/去北京的/机票 XiaoWang / booked/ <u>Zhangfei</u> / a ticket/ to Beijing.
	Typically, people would interpret in the first way as Zhangfei is a fictional figure and do not appear in the daily conversation. However, the second half of the sentence confirms that the second way of parsing turns out to be correct. Such contrast is hilarious.
GPT-4o’s Answer	(zh) ... <u>由于历史人物张飞已经不在世</u> , 所以给他省了一大笔钱。
	(en) ... <u>since the historical figure Zhang Fei is dead and cannot take a flight</u> , the joke humorously suggests that he saved a lot of money.
Failure Reasons	Fail to parse textual elements in a sentence dynamically.

Figure 11: Parsing error example.

characters as the traditional Chinese characters are also termed as “complicated characters”. However, LLMs struggle to reason such graphemic differences as there are no explicit connections between the textual meaning and visual representations of the glyphs.

Parsing Error. LLMs often fail to parse sentences in multiple ways simultaneously, leading to difficulties in explaining jokes that require different parsing for the same sentence. In the example in Figure 11, the humor hinges on the ambiguity of

the phrase “张飞”, which can be interpreted either as part of a verb phrase implying “a ticket flying to Beijing” or as a proper noun, referring to the historical figure Zhang Fei. This ambiguity stems from the flexibility of the Chinese language, where each character can function independently as a word or combine with others to form new words or phrases. There are decades of research studying the problem of parsing Chinese (Sun and Jurafsky, 2004; Sun et al., 2009). Recently, researchers have proposed task-specific tokenization approaches that adapt the parsing process to better align with downstream tasks (Liu et al., 2021, 2023b). However, how to incorporate different ways of parsing at one time still remains challenging.

culture and reason over cultural references, posing a great challenge to current LLMs. Although recent works have made progress towards building LLMs beyond English (Du et al., 2024; Zhao et al., 2024), building an LLM that can comprehend such nuanced Chinese cultural terms can be extremely hard.

Example	(zh) <u>五步蛇</u> 的危险性是 <u>李白的两倍</u> 。
	(en) The danger of a <u>five-step snake</u> is <u>twice that of Li Bai</u> .
Correct Humor Explanation	<i>Kill one in ten steps</i> is a line from a poem by the ancient Chinese poet Li Bai, suggesting that while it takes Li Bai ten steps to kill, a person can die within five steps of being bitten by the “Five-Step Snake” hence it is twice as dangerous as Li Bai.
GPT-4o’s Answer	(zh) ... <u>李白的名字听起来像“二步”</u> ，所以“五步蛇”的危险性被戏称为是“二步”的两倍。 (en) ... <u>Li Bai’s (pronounced as lee bye) name sounds similar to “two steps” (pronounced as r boo) in Chinese</u> , so the danger of “Five-Step Snake” is humorously claimed to be twice that of “two steps”
Failure Reasons	Fabricate ungrounded details or facts.

Figure 12: Hallucinations example.

Hallucinations. LLMs may fabricate ungrounded details or facts in joke explanations. For instance, in the explanation in Figure 12, GPT-4o claims that “Li Bai’s name sounds similar to two steps”, while “Li Bai” (pronounced as *lee bye*) does not sound like “two steps” (pronounced as *r boo*).

On the other hand, the correct explanation requires an understanding of a Chinese poem from Li Bai, “十步杀一人” (The warrior kills a person for every ten steps). This line praises the courage of the soldiers, but the joke deliberately portrays this as a characteristic of Li Bai. Therefore, compared to Li Bai who can kill a person in ten steps, a five-step snake, which can kill a person in five steps, is twice as dangerous as Li Bai. Such explanation requires LLMs to have a deep understanding of Chinese

Example	(zh) 曹操于城楼上问夏侯惇：“你瞧到了什么。”夏侯惇说：“ <u>瞧到马岱</u> 。” (en) Cao Cao, from atop the city tower, asked Xia Houdun, “What did you see?” Xia Houdun replied, “ <u>I saw Ma Dai</u> .”
Correct Humor Explanation	Cao Cao, Xiahou Dun, and Ma Dai were historical figures from the Three Kingdoms period in ancient China. The phrase “saw Ma Dai” (pronounced as <i>chuh dow ma die</i>) not only literally refers to spotting Ma Dai but also play on similar pronunciation with the Japanese phrase “ちょっと待って” (<i>chotto matte</i>), which means to “wait a moment.”
GPT-4o’s Answer	(zh) ... <u>这个笑话的幽默之处在于利用了“马岱”这个名字的谐音梗，听起来像“妈的”，让人误以为夏侯惇在骂人。</u> (en) ... <u>a pun involving the name "Ma Dai," which sounds like "damn" in Chinese.</u> It sounds like Xiahou Dun is swearing at Cao Cao.
Failure Reasons	Fail to recognize information/elements across different languages.

Figure 13: Cross-lingual error example.

Cross-lingual Error. LLMs may fail to recognize elements or information across different languages. In the explanation in Figure 13, GPT-4o attempts to link the pronunciation of “Ma Dai” to other Chinese terms but fails to identify the similar pronunciations across the Chinese term “瞧到马岱” (pronounced as *chuh dow ma die*, meaning “saw Ma Dai”) and the Japanese term “ちょっと待って” (*chotto matte*, meaning “wait a moment”). Such cases require LLMs to connect pronunciations across languages, which may be rare in the LLMs’ pre-training corpus and poses significant challenges to current LLMs.

Overcritical. Figure 14 shows when the model overly criticizes an explanation, it focuses excessively on minor details, neglecting the major point and ultimately misjudging the explanation. The joke relies on common sense about life, death, and testing rather than cultural knowledge. Under the CoT prompting, the model recognizes that the explanation identifies the main punchlines but overly criticizes the lack of cultural context.

Excessive Sensitivity. For ERNIE_{4-turbo}, in addition to errors across all the aforementioned categories, it also demonstrates excessive sensitivity to certain examples. Specifically, when content includes languages related to hate speech but used in non-harmful contexts, ERNIE_{4-turbo} refuses to provide an explanation. During our evalua-

Example	(zh) 网上传疯了超准测试题！20道题轻松判断出 <u>你有没有死</u> ！ (en) The viral online quiz! 20 questions to easily determine <u>if you're alive or dead!</u>
Joke Explanation	(zh) 这个笑话的幽默之处在于它用夸张的方式提出了一个显而易见的问题——判断自己是否还活着，这本身就不需要测试。其次，标题中“超准测试题”和“轻松判断”这样的描述让人觉得荒谬可笑，因为活着与否是最基本的感知，不需要任何测试。 (en) The humor in this joke lies in its exaggerated premise of asking an obvious question—determining whether one is still alive, which inherently doesn't require any test. Additionally, the descriptions in the title, such as "super accurate quiz" and "easily determine," come across as absurdly funny because being alive is the most basic perception, requiring no testing at all.
Model’s Answer DP	(zh) <u>完全解释</u> (en) <u>Fully Explained</u>
Model’s Answer CoT	(zh) ... 最终判断： <u>部分/没有解释</u> 理由：虽然解释准确地指出了笑话的主要笑点（夸张性和标题的矛盾），但 <u>它没有深入探讨这些元素为什么会产生幽默效果，也没有讨论潜在的文化或认知背景</u> ... (en) ... Final judgment: <u>Partially/Not explained</u> Reason: While the explanation accurately identifies the main punchlines of the joke (exaggeration and the contradiction in the title), <u>it does not delve into why these elements create a humorous effect or discuss the potential cultural or cognitive context behind it</u> ...

Figure 14: Overcritical example by Nemotron_{70B}. The Nemotron_{70B} model selects the correct answer in the DP prompting, but selects the incorrect answer due to being overly critical in the CoT prompting.

tion, we observe this excessive sensitivity in the ERNIE_{4-turbo}’s responses to humor related to medical ethics and political discussions. This suggests that correctly understanding the context and the language toxicity remains an open challenge (Zhang et al., 2024a). Such issues are particularly critical for humor explanation, as misclassifying non-toxic context can cause the responses to deviate from the intended humor.

F Prompts for DP and CoT in Chumor

This section outlines the prompts used in *Chumor* to evaluate whether an explanation fully explains a joke. Two prompting strategies are adopted: Direct Prompting (DP) and Chain of Thought (CoT). Below are the details of each approach:

Direct Prompting (DP)

你将看到一个笑话以及对这个笑话的解释。请判断这个解释是否完全解释了笑话。根据判断，选择“完全解释”或“部分/没有解释”，不需要解释为什么对或者不对。

笑话: [joke]

笑话解释: [explanation]

Translation

You will see a joke and an explanation of the joke. Please determine whether this explanation fully explains the joke. Based on your judgment, choose either “fully explain” or “partially/does not explain.” You do not need to explain why it is correct or incorrect.

Joke: [joke]

Explanation: [explanation]

The DP prompt is designed to encourage concise decision-making. It directly asks the model to evaluate the completeness of the explanation without requiring reasoning or justification.

Chain of Thought (CoT)

你将看到一个笑话以及对这个笑话的解释。请逐步思考，写下过程并最终判断这个解释是否完全解释了笑话。根据判断，选择“完全解释”或“部分/没有解释”。

笑话: [joke]

笑话解释: [explanation]

Translation

You will see a joke and an explanation of the joke. Please think step by step, write down your reasoning process, and finally determine whether this explanation fully explains the joke. Based on your judgment, choose either “fully explains” or “partially/does not explain.”

Joke: [joke]

Explanation: [explanation]

The CoT prompt, in contrast, requires the model to reason step by step before reaching a conclusion. This approach aims to improve transparency by explicitly documenting the thought process behind the evaluation.

G Joke Type Distribution in *Chumor*

We sampled 200 datapoints from *Chumor* to analyze the distribution of joke types, as shown in Figure 15. Note that a single joke may belong to

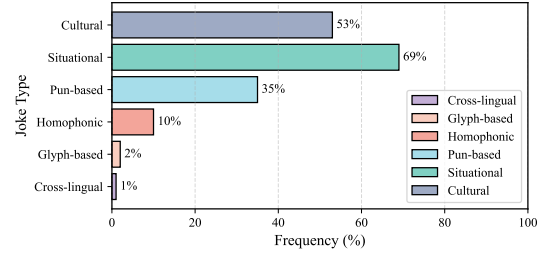


Figure 15: Distribution of Joke Types in 200 Sampled Datapoints.

multiple categories, as it can exhibit features of more than one joke type.

H Detailed Results of Experiments

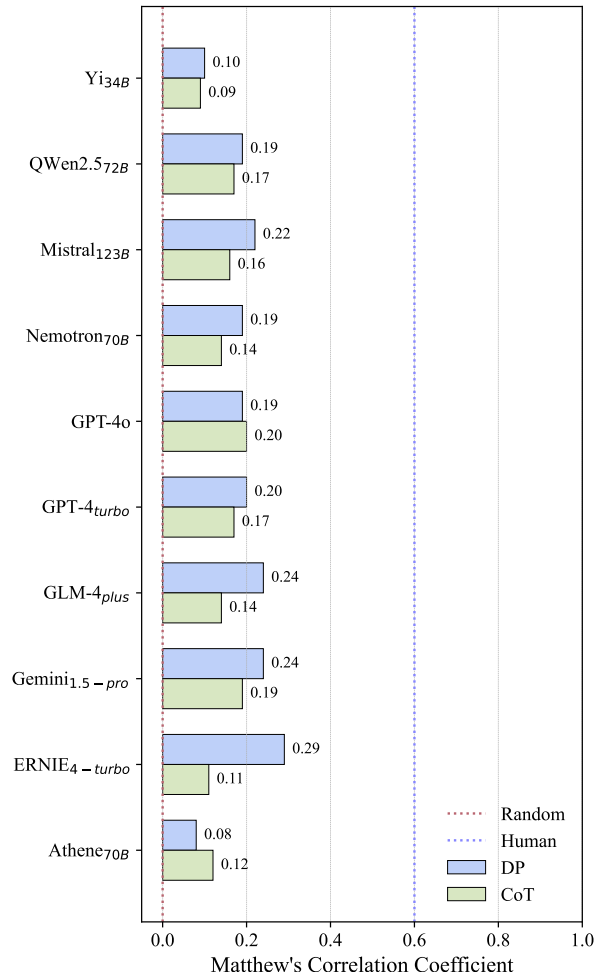


Figure 16: The Matthew's correlation coefficient of different models' test results in DP and CoT.

For evaluation, we input each prompt into the model and collect its responses, comparing them to the labels in *Chumor*. A model's response is considered correct if it matches the reference label. If the model provides an incorrect answer or doesn't generate a response at all (due to safety protocols or

filtering sensitive terms), it is marked as incorrect. Such scenario is rare, occurring only 21 times in our experiments, and exclusively with GLM-4_{plus}.

We highlight that CoT prompting at most cases degrade the models' performance on *Chumor*. As shown in Figure 16, only Athene70B achieves a significant improvement. However, this is offset by its poorest performance under DP prompting among the models. GPT-4o shows a slight improvement, with its MCC score increasing from 0.19 to 0.20. And all other eight models exhibit different degrees of performance decline.

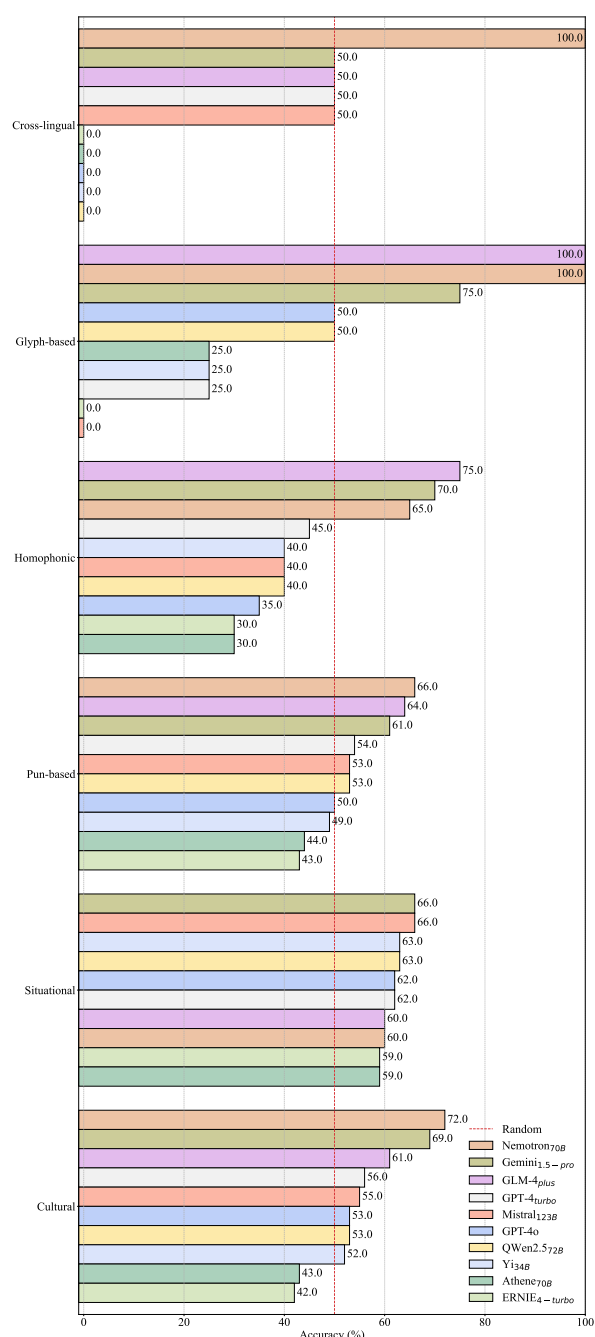


Figure 17: CoT accuracy on different joke types(%).

Model	DP				CoT			
	MCC	ACC (%)	FPR (%)	FNR (%)	MCC	ACC (%)	FPR (%)	FNR (%)
Yi _{34B}	0.10	44.95	97.24	0.21	0.09	47.17	89.30	5.44
Nemotron _{70B}	0.19	56.30	61.26	20.87	0.14	57.17	40.28	46.14
Athene _{70B}	0.08	44.59	97.83	0.28	0.12	47.26	91.10	2.89
ERNIE _{4-turbo}	0.29	60.29	59.83	13.57	0.11	45.16	96.93	0.14
QWen2.5 _{72B}	0.19	48.46	90.67	0.69	0.17	49.45	86.91	3.31
Mistral _{123B}	0.22	55.56	69.26	12.19	0.16	51.18	79.92	8.40
Gemini _{1.5-pro}	0.24	54.00	77.42	5.17	0.19	60.32	33.81	47.31
GLM-4 _{plus}	0.24	55.56	72.28	8.26	0.14	58.13	32.96	53.44
GPT-4o	0.19	51.87	80.02	6.68	0.20	50.64	85.00	3.03
GPT-4 _{turbo}	0.20	52.32	79.28	6.61	0.17	51.27	80.87	6.96

Table 4: Performance metrics for explanation evaluation including Matthew’s correlation coefficient (MCC), accuracy (ACC), false positive rate (FPR), and false negative rate (FNR).

Model	Source	DP				CoT			
		MCC	ACC(%)	FPR(%)	FNR(%)	MCC	ACC(%)	FPR(%)	FNR(%)
Athene _{70B}	Overall	0.08	44.59	97.83	0.28	0.12	47.26	91.10	2.89
	ERNIE Bot	0.12	52.38	97.15	0.00	0.15	54.24	91.13	2.13
	GPT-4o	0.03	33.90	98.51	0.86	0.08	37.67	91.06	4.50
ERNIE _{-turbo}	Overall	0.29	60.29	59.83	13.57	0.11	45.16	96.93	0.14
	ERNIE Bot	0.23	58.64	78.14	5.99	0.16	53.47	94.83	0.10
	GPT-4o	0.27	62.54	41.38	29.55	0.04	33.76	99.04	0.21
Gemini _{1.5-pro}	Overall	0.24	54.00	77.42	5.17	0.19	60.32	33.81	47.31
	ERNIE Bot	0.27	60.66	74.13	5.89	0.23	60.87	28.62	49.24
	GPT-4o	0.21	44.85	80.74	3.64	0.17	59.56	39.04	43.25
GLM-4 _{plus}	Overall	0.24	55.56	72.28	8.26	0.14	58.13	32.96	53.44
	ERNIE Bot	0.25	59.83	74.97	6.70	0.15	57.56	37.06	47.61
	GPT-4o	0.21	49.68	69.57	11.56	0.06	58.92	28.83	65.74
GPT-4 _{turbo}	Overall	0.20	52.32	79.28	6.61	0.17	51.27	80.87	6.96
	ERNIE Bot	0.20	57.25	80.99	5.99	0.22	58.75	76.14	7.72
	GPT-4o	0.18	45.56	77.55	7.92	0.13	41.01	85.64	5.35
GPT-4o	Overall	0.19	51.87	80.02	6.68	0.20	50.64	85.00	3.03
	ERNIE Bot	0.21	57.82	79.41	6.40	0.24	58.07	82.47	2.94
	GPT-4o	0.16	43.71	80.64	7.28	0.15	40.44	87.55	3.21
Nemotron _{70B}	Overall	0.19	56.30	61.26	20.87	0.14	57.17	40.28	46.14
	ERNIE Bot	0.22	60.66	56.81	22.54	0.14	57.04	39.18	46.60
	GPT-4o	0.18	50.32	65.74	17.34	0.13	57.36	41.38	45.18
Mistral _{123B}	Overall	0.22	55.56	69.26	12.19	0.16	51.18	79.92	8.40
	ERNIE Bot	0.25	61.13	65.15	13.60	0.18	57.04	79.73	7.61
	GPT-4o	0.20	47.90	73.40	9.21	0.12	43.14	80.11	10.06
Qwen2.5 _{72B}	Overall	0.19	48.46	90.67	0.69	0.17	49.45	86.91	3.31
	ERNIE Bot	0.19	54.45	92.61	0.30	0.18	55.54	88.07	2.54
	GPT-4o	0.17	40.23	88.72	1.50	0.14	41.08	85.74	4.93
Yi _{34B}	Overall	0.10	44.95	97.24	0.21	0.09	47.17	89.30	5.44
	ERNIE Bot	0.15	53.42	94.72	0.30	0.11	53.99	88.38	5.28
	GPT-4o	0.03	33.33	99.79	0.00	0.07	37.81	90.21	5.78

Table 5: Detailed performance metrics with source for explanation evaluation of Matthew’s correlation coefficient (MCC), accuracy (ACC), false positive rate (FPR), and false negative rate (FNR).

Model	Prompting	Cross-lingual	Glyph-based	Homophonic	Pun-based	Situational	Cultural
Athene _{70B}	DP	0.00	0.00	30.00	44.00	61.00	42.00
	CoT	0.00	25.00	30.00	44.00	59.00	43.00
ERNIE _{4-turbo}	DP	50.00	50.00	60.00	61.00	70.00	63.00
	CoT	0.00	0.00	30.00	43.00	59.00	42.00
Gemini _{1.5-pro}	DP	50.00	50.00	55.00	63.00	67.00	61.00
	CoT	50.00	75.00	70.00	61.00	66.00	69.00
GLM-4 _{plus}	DP	50.00	25.00	65.00	60.00	69.00	60.00
	CoT	50.00	100.00	75.00	64.00	60.00	61.00
GPT-4 _{turbo}	DP	50.00	25.00	40.00	57.00	67.00	55.00
	CoT	50.00	25.00	45.00	54.00	62.00	56.00
GPT-4o	DP	0.00	50.00	35.00	49.00	63.00	54.00
	CoT	0.00	50.00	35.00	50.00	62.00	53.00
Nemotron _{70B}	DP	50.00	50.00	65.00	63.00	62.00	60.00
	CoT	100.00	100.00	65.00	66.00	60.00	72.00
Mistral _{123B}	DP	50.00	50.00	55.00	61.00	65.00	61.00
	CoT	50.00	0.00	40.00	53.00	66.00	55.00
Qwen2.5 _{72B}	DP	0.00	50.00	35.00	47.00	64.00	51.00
	CoT	0.00	50.00	40.00	53.00	63.00	53.00
Yi _{34B}	DP	0.00	0.00	30.00	43.00	60.00	44.00
	CoT	0.00	25.00	40.00	49.00	63.00	52.00

Table 6: Performance metrics by joke type for explanation evaluation accuracy(%).