

Do Large Language Models have Problem-Solving Capability under Incomplete Information Scenarios?

Yuyan Chen¹, Tianhao Yu², Yueze Li¹, Songzhou Yan¹, Sijia Liu¹, Jiaqing Liang³, Aaron Xuxiang Tian⁴, Yanghua Xiao¹ ✉

¹Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University,

²SenseAuto, ³School of Data Science, Fudan University,

⁴Carnegie Mellon University,

{chenyuyan21, yuezeli23, szyan21, sijialiu21}@m.fudan.edu.cn,

{liangjiaqing, shawyh}@fudan.edu.cn,

avery1026@hotmail.com, aarontian00@gmail.com

Abstract

The evaluation of the problem-solving capability under incomplete information scenarios of Large Language Models (LLMs) is increasingly important, encompassing capabilities such as questioning, knowledge search, error detection, and path planning. Current research mainly focus on LLMs' problem-solving capability such as "Twenty Questions". However, these kinds of games do not require recognizing misleading cues which are necessary in the incomplete information scenario. Moreover, the existing game such as "Who is undercover" are highly subjective, making it challenging for evaluation. Therefore, in this paper, we introduce a novel game named BrainKing based on the "Who is undercover" and "Twenty Questions" for evaluating LLM capabilities under incomplete information scenarios. It requires LLMs to identify target entities with limited yes-or-no questions and potential misleading answers. By setting up easy, medium, and hard difficulty modes, we comprehensively assess the performance of LLMs across various aspects. Our results reveal the capabilities and limitations of LLMs in BrainKing, providing significant insights of LLM problem-solving levels.

1 Introduction

Incomplete information scenarios include missing information, uncertainty, and misinformation, encountered in fields such as business negotiations, military strategy, medical diagnosis, and legal judgments (Gibbons and Gibbons, 1992). The problem-solving capability under incomplete information refer to the capability to effectively handle available information, make rational inferences, and decisions in situations lacking comprehensive data. This capability is crucial in real life as we cannot

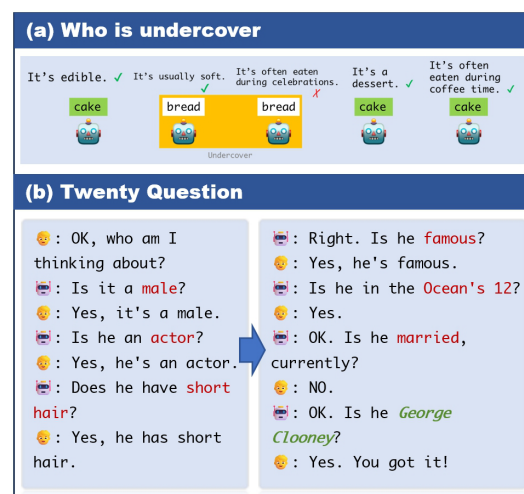


Figure 1: A sample of the "Who is undercover" game (a) and the "Twenty Questions" game (b).

possess all the necessary information for decision-making. It's also important for large language models (LLMs), which not only tests LLMs' logical reasoning capabilities but also involves adjusting strategies in constantly changing environments, significantly enhancing their robustness and quality of decision-making in various fields. Therefore, a natural question arises: *Do LLMs have problem-solving capability under incomplete information scenarios?*

Previous research on problem-solving capability under incomplete information scenarios focus on simulating complex decision-making environments through games, such as "Werewolf" (Xu et al., 2023; Ri et al., 2022; Toriumi et al., 2017), "Poker" (Brown and Sandholm, 2019) and "Avalon" (Light et al., 2023), etc. These games require players to make decisions without full information, often involving deception and strategic planning to conceal their real identities. As shown in Fig. 1(a), "Who is undercover" is another incomplete information game which requires

players to deduce whether they are the spy based on others’ descriptions. In this running example, player with “bread” is the spy in this game and he needs to hide himself against being caught through twist the facts in his description like “It typically requires more sugar, fats, and eggs.”. However, even advanced LLMs like GPT-3 and GPT-3.5, while excelling in general NLP tasks (Saad-Falcon et al., 2023; Liu et al., 2023; Kojima et al., 2022; Chung et al., 2023), revealing limitations in effective decision-making under incomplete information environments (Gigerenzer and Gaissmaier, 2011; Binz and Schulz, 2023).

Information processing is a crucial problem-solving capability under incomplete information scenarios, exemplified by the Minesweeper game (Li et al., 2023), the Twenty Questions game (Walsorth, 1882; Giordano et al., 1998), etc. As shown in Fig. 1(b), the Twenty Questions game requests the player to pose a series of yes-or-no questions to guess the given entity (i.e. “George clooney”), which can effectively evaluate LLMs’ creativity (Hu et al., 2018), knowledge retrieval (Williams and Klamen, 2015; Szymański and Duch, 2012), multi-hop reasoning capabilities (Noever and McKee, 2023; Siegler, 1977). However, the above-mentioned games, such as Twenty Questions, do not adequately assess LLMs’ capabilities in processing information and solving problems because it lacks deception and strategic complexity that require recognizing misleading cues and formulating adaptive strategies based on limited or false information. Moreover, games like “Werewolf” and “Who is undercover” are highly subjective, making it challenging to evaluate LLMs’ capabilities under incomplete information scenarios effectively.

Therefore, in this paper, we introduce a new game named BrainKing by combining the “Who is undercover” and “Twenty Questions” game to assess LLMs’ information processing and problem-solving capabilities under incomplete information scenarios. BrainKing challenges LLM participants to identify entities amidst potential misinformation through a limited set of yes-or-no questions across easy, medium, and hard difficulty modes. This game objectively assesses LLMs’ world knowledge, reverse thinking, and error detection capabilities. Our results explore five key questions regarding LLM performance in BrainKing, investigating the relationship between accuracy and rounds, the impact of starting point difficulty and the number

of wrong answers, and the correlation between accuracy and the ability to recognize confusion. This assessment, including questioning, knowledge retrieval, misinformation recognition, and planning abilities, effectively evaluates LLMs’ capabilities under incomplete information scenarios through a simple-yet-effective game with objective metrics. In summary, our research makes three key contributions:

- We propose a simple-yet-effective game named BrainKing to evaluate LLMs’ information processing and problem-solving capabilities under incomplete information scenarios.
- Based on the BrainKing game, we introduce an automated evaluation methodology that objectively measures LLMs’ performance.
- We conduct experiments to comprehensively assess existing LLMs’ capabilities under incomplete information scenarios, yielding significant conclusions for enhancing their capabilities and limitations.

2 Datasets and Task Setups

The proposed BrainKing game, inspired by the Twenty Questions game, requires LLM participants to identify an entity with a limited number of yes-or-no questions despite potentially misleading answers. It contains three difficulty modes: easy, medium, and hard, as illustrated in Fig. 2. These modes are designed to thoroughly assess an LLM’s world knowledge, reverse thinking, and error detection capabilities in identifying the target entity.

Specifically, we utilize an open-source Twenty Questions dataset¹, which comprises 78,890 entities. To ensure the task is manageable and not hindered by a lack of knowledge by all LLM participants, we adopt GPT3.5² to score each entity for commonness from 1 (less common) to 3 (more common), retaining the top 10,000 common entities. Next, we also adopt GPT3.5 to generate a hierarchical concept list for each entity. This list must include at least three concepts, where each subsequent concept is a broader or more abstract category of the preceding one. We avoid strict academic classifications to maintain the conceptual hierarchy’s logic and clarity. For instance, the hierarchical concept list for the entity “Dog”

¹<https://github.com/allenai/twentyquestions>

²<https://chat.openai.com/>

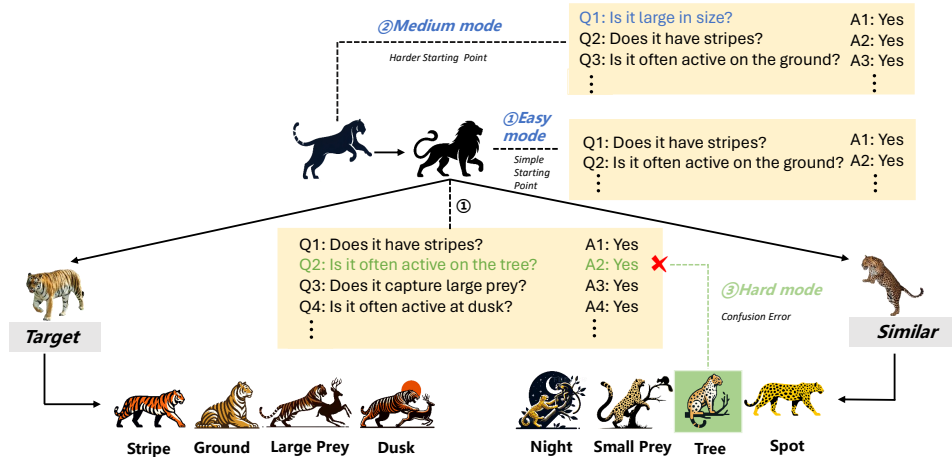


Figure 2: The overview of the proposed BrainKing benchmark, including three modes.

is “[Household pets, Terrestrial mammals, Mammals]” instead of the strict biological hierarchy “[Canis, Canidae, Carnivora]”, as starting with Canis makes it too easy to guess the dog based on prior knowledge. After that, GPT3.5 is also used to provide the most similar entity for each entity in the dataset. For example, the entity most similar to a “dog” is identified as a “wolf”.

2.1 Difficulty Modes

Easy mode is to provide a simple starting point which is the first concept from the hierarchical concept list of the target entity. It requires LLM participants to generate at most 20 yes-or-no questions to identify the target entity, which is the same as traditional Twenty Question game. For example, in Fig. 2, the simple starting point for the target entity “Tiger” is “Pantherinae”, which directly guides the LLMs to focus its inference within the realm of Pantherinae. The possible question is “Does it have stripes?”, etc.

Medium mode is to provide a harder starting point which is set as the second concept in our task from the hierarchical concept list of the target entity. It requires LLM participants to identify the target entity with a broader realm. For example, in Fig. 2, the harder starting point for the target entity tiger is Felidae, which directly guides the LLMs to focus its inference within the realm of Felidae. The possible question is “Is it large in size?” to confirm whether it belongs to Pantherinae.

Hard mode introduces a similar entity for generating wrong answers which is set as two in our task besides providing the simple starting point same as the easy mode. It requires LLM participants to generate at most 20 yes-or-no questions to identify the target entity with misleading information. For

example, in Fig. 2, a possible question is “Does it often engage in activities in trees?” and the correct answer is “No” because the target entity tiger is usually active on the ground instead of the tree which indicates the similar entity “Leopard”. LLM participants are expected not to navigate through this wrong answer and rethink the correct inference path.

2.2 Evaluation Metrics

We adopt accuracy and rounds as metrics to evaluate LLMs’ performance in the proposed BrainKing. Accuracy measures whether an LLM can infer the target entity within 20 rounds of questioning, with a successful guess scored as 1 and an unsuccessful one as 0. Rounds measures how many questions it takes to infer the entity. Once the entity is inferred, the game stops regardless of whether it is correct or incorrect, preventing the LLMs from exploring the full range of possible entities. If the entity is not inferred within Twenty Questions, we set its rounds as 30, as it is not feasible to continue BrainKing indefinitely. We multiply the accuracy and the reciprocal of the rounds by 100 to obtain the accuracy win rate and rounds win rate, respectively. The average of the accuracy win rate and rounds win rate is then calculated to determine the total win rate. The optimal score for accuracy win rate, rounds win rate, and total win rate is 100. The higher the value, the higher the win rate.

Moreover, we also introduce the ability to recognize confusion as an evaluation metric to determine whether an LLM can backtrack from misleading information in an answer and returns to the correct path of questioning. We use GPT4 to analyze the whole question-and-answer process of an LLM. If it exclusively follows the misleading direction

without returning to the correct path, it is marked as 1; otherwise, it is marked as 0.

3 Experiments

In this section, we conduct extensive experiments to evaluate different LLMs’ performance in the proposed BrainKing.

3.1 Experimental Setups

Our experiments are conducted on 8 Nvidia A100 GPUs, each with 80GB of memory, and we use PyTorch³ in Python⁴. We set the maximum sequence length for both input and output sequences to maximum 200 tokens. we use GPT4 to respond to the questions posed by each LLM.

3.2 Datasets, Baselines and Metrics

The baseline LLMs for this evaluation are BLOOM-7B (Workshop et al., 2023) BLOOM-176B (Workshop et al., 2023), Claude2 (Bai et al., 2022), Falcon-7B (Almazrouei et al., 2023), Falcon-180B (Almazrouei et al., 2023), GPT3.5 (Brown et al., 2020), GPT4 (OpenAI, 2023), LLaMA2-7B (Touvron et al., 2023), LLaMA2-70B (Touvron et al., 2023), Vicuna-7B (Chiang et al., 2023), and Vicuna-33B (Zheng et al., 2023). The prompt for playing the BrainKing is shown in Table 1.

We recruit nine volunteers to participate in BrainKing. First, we select three entities, each representing a different mode of difficulty to test each volunteer’s ability. Then, we rank their comprehensive scores from high to low. The top three scorers are assigned to hard mode, the middle three to medium mode, and the last three to easy mode. We then randomly distribute 1,000 entities, other than the initial three, among the nine volunteers and calculate their accuracy win rate, rounds win rate and total win rate. The highest total win rate from each of the three modes are averaged again to obtain the Human performance. Volunteers participate on a voluntary basis without compensation.

3.3 Main results

Question 1: Which LLM is the winner of the BrainKing? Answer 1: GPT4!

The performance of different LLMs under three difficulty modes are shown in Table 2 and Fig. 3. We observe that in the easy mode, Claude2 stands out with an accuracy of 85.4% and an average of

only 5.7 rounds. GPT4, GPT3.5, and BLOOM-176B also perform well with high accuracy above 70% and few rounds below 11. In the medium mode, GPT4 outperforms among all LLMs with an accuracy of 81.2% and an average of 10.2 rounds, suggesting its powerfulness even with a harder starting point. Claude2, GPT3.5 and BLOOM-176B, despite a decrease in accuracy, still maintain around 70%, showing good reasoning capabilities. In the hard mode, further escalates the challenge by providing the same starting point as Easy Mode and introducing a similar entity that could cause confusion. Under these conditions, GPT4 also leads with an accuracy of 78.8% and 15.8 rounds, followed by Claude2 and BLOOM-176B around 70%. This suggests that GPT4 has strong robustness when facing with misleading information. Across all modes, Falcon-7B and Vicuna-7B show poorer performance, especially in hard mode, which are below 20%. This suggests that these LLMs are less capable of handling complex reasoning and resisting misleading information.

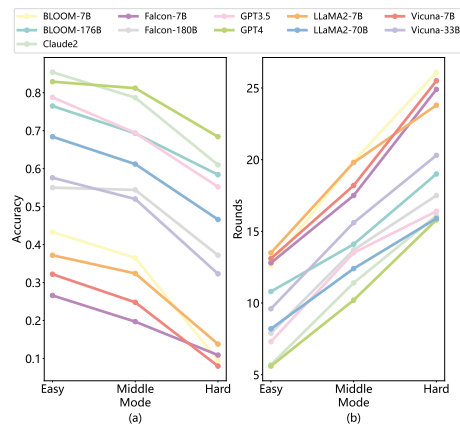


Figure 3: The comparison of the performance (accuracy and rounds) of different LLMs across three modes.

We calculate the win rates for accuracy and rounds, as well as the total win rate, and rank them accordingly. These rankings are displayed in Table 3. Figure 4 displays three tiers of performance, from strongest to weakest, across the three win rates. We observe that GPT4 consistently performs best across accuracy win rate, rounds win rate, and total win rate, demonstrating its strong capabilities. Claude2 follows closely, ranking second with a 75% accuracy win rate and a 9.0 rounds win rate, and a total win rate of 42.0%. This indicates that Claude2 is a very powerful model, effectively competing with GPT4. Human also performs well, ranking in the top three for accuracy win rate and

³<https://pytorch.org/>

⁴<https://www.python.org/>

Mode	Prompt
Easy Mode	Let’s play a game of Twenty Questions. I have a hidden thing in mind and you need to guess it. This time, the thing belongs to {a simple and straightforward starting point}. You may only ask yes/no questions for information and I will only answer with yes or no, note that I will say yes when it is likely. When you make a final guess, say "Guess:(Your guess)". You’re not allowed to make another guess after making the final guess. Now begin your first question, count the number of questions like this: "Q1:(Your question)"
Medium Mode	Let’s play a game of Twenty Questions. I have a hidden thing in mind and you need to guess it. This time, the thing belongs to {a harder starting point}. You may only ask yes/no questions for information and I will only answer with yes or no, note that I will say yes when it is likely. When you make a final guess, say "Guess:(Your guess)". You’re not allowed to make another guess after making the final guess. Now begin your first question, count the number of questions like this: "Q1:(Your question)"
Hard Mode	Let’s play a game of Twenty Questions. I have a hidden thing in mind and you need to guess it. This time, the thing belongs to {a simple and straightforward starting point}. You may only ask yes/no questions for information and I will only answer with yes or no, note that I will say yes when it is likely. During the process, I will confuse the real item with another in the field, so some questions may not have been correctly answered, especially those more detailed. You may ask more questions and propose a new guess. When you make a final guess, say "Guess:(Your guess)". You’re not allowed to make another guess after making the final guess. Now begin your first question, count the number of questions like this: "Q1:(Your question)"

Table 1: The instruction prompts of three modes in guiding LLMs to generate questions.

	Easy Mode		Medium Mode		Hard Mode	
	Accuracy	Rounds	Accuracy	Rounds	Accuracy	Rounds
BLOOM-7B	43.3	12.6	36.5	19.9	12.3	26.1
BLOOM-176B	76.5	10.8	69.3	14.1	65.5	19.0
Claude2	85.4	5.7	78.7	11.4	68.3	16.1
Falcon-7B	26.6	12.8	19.7	17.5	12.5	24.9
Falcon-180B	55.0	7.9	54.4	13.7	51.0	17.5
GPT3.5	78.8	7.3	69.4	13.5	63.7	16.4
GPT4	82.9	5.6	81.2	10.2	78.8	15.8
LLaMA2-7B	37.2	13.5	32.4	19.8	24.8	23.8
LLaMA2-70B	68.4	8.2	61.2	12.4	55.7	15.9
Vicuna-7B	32.2	13.1	24.8	18.2	17.4	25.5
Vicuna-33B	57.6	9.6	52.0	15.6	47.8	20.3

Table 2: The accuracy and rounds of different LLMs in the proposed BrainKing benchmark.

total win rate, but fall to sixth place in rounds win rate, which may suggest that while human participants can infer the entity but not much efficient compared with GPT4 and Claude2. GPT3.5, BLOOM-176B, LLaMA2-70B and Falcon-180B have moderate performances, while BLOOM-7B, Vicuna-33B, and Vicuna-7B lag behind, especially in the accuracy win rate. LLaMA2-7B and Falcon-7B are at the bottom, ranking low in all metrics, particularly in the total win rate, which have more room for optimization.

Question 2: Are accuracy and rounds in a completely inverse relationship? Answer 2: No, just inversely correlated!

We also analyze the relationship between accuracy and rounds in three modes as shown in Fig. 5. We find that the accuracy and the rounds do not show a strict inverse relationship. In the easy mode, most LLMs achieve high accuracy with fewer rounds. However, some LLMs, such as those with 7B size, still show low accuracy after many rounds, which suggests that not every round of answers is effectively used for reasoning. In the medium and hard mode, there’s a more noticeable trend of accuracy decreasing as the number of rounds increases, which suggests that LLMs need more information to make accurate reasoning and are more susceptible to being misled. Therefore, the relationship between accuracy and rounds is

	Accuracy			Rounds			Total		
	Win Rate	Rank		Win Rate	Rank		Win Rate	Rank	
GPT4	77.5	1		9.5	1		43.5	1	
Claude2	<u>75.0</u>	2		<u>9.0</u>	2		<u>42.0</u>	2	
Human	74.2	3		7.5	6		40.9	3	
GPT3.5	67.8	5		8.1	4		37.9	4	
BLOOM-176B	68.1	4		6.8	7		37.5	5	
LLaMA2-70B	58.7	6		8.2	3		33.5	6	
Falcon-180B	48.9	7		7.7	5		28.3	7	
Vicuna-33B	47.3	8		6.6	8		26.9	8	
BLOOM-7B	29.8	9		5.1	12		17.5	9	
LLaMA2-7B	27.8	10		5.3	11		16.5	10	
Vicuna-7B	21.7	11		5.3	10		13.5	11	
Falcon-7B	19.1	12		5.4	9		12.3	12	

Table 3: The win rate of different LLMs in the proposed BrainKing benchmark.

not strictly inverse but a more complex interplay where many factors can affect the outcome.

Question 3: Does the difficulty of starting points significantly affect LLMs’ performance? Answer 3: More significant for weaker LLMs!

We also analyze the change in model performance with the increase in the difficulty of starting points as shown in Table 4 and Fig. 6. As the starting point difficulty goes up, most LLMs tend to show a significant decrease in accuracy. At easy starting points, Claude2 and GPT4 demonstrate higher accuracy, and although their accuracy drops at hard starting points, they still maintain relatively high levels. Additionally, as the difficulty increases, the number of rounds needed tends to go up. For example, BLOOM-176B requires 10.8 rounds at the easy starting point, which increases to 21.9 rounds at the hard starting point. Overall, the difficulty of the starting point has a clear impact on LLMs’ performance. Strong LLMs still demonstrate their reasoning capabilities even at hard starting points, while LLMs with weaker performance are more significantly affected in both accuracy and number of rounds.

Question 4: Does the number of wrong answers significantly affect LLMs’ performance? Answer 4: Yes, most LLMs are nearing the upper limit of rounds!

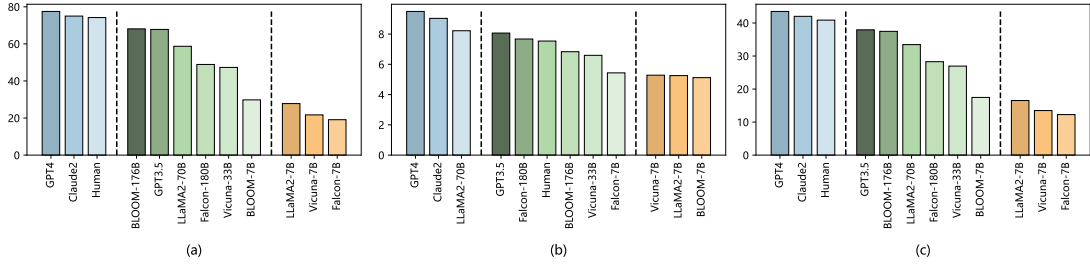


Figure 4: The win rate of different LLMs in the proposed BrainKing benchmark.

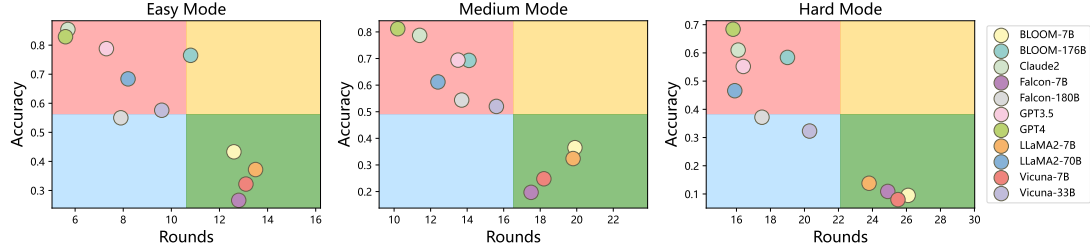


Figure 5: The relationship between accuracy and rounds across three modes of different LLMs.

	Easy Starting Point		Medium Starting Point		Hard Starting Point	
	Accuracy	Rounds	Accuracy	Rounds	Accuracy	Rounds
BLOOM-7B	43.3	12.6	36.5	19.9	11.5	29.2
BLOOM-176B	76.5	10.8	69.3	14.1	43.1	21.9
Claude2	85.4	5.7	78.7	11.4	48.8	17.3
Falcon-7B	26.6	12.8	19.7	17.5	9.8	29.1
Falcon-180B	55.0	7.9	54.4	13.7	38.8	24.2
GPT3.5	78.8	7.3	69.4	13.5	42.1	20.8
GPT4	82.9	5.6	81.2	10.2	51.2	16.9
LLaMA2-7B	37.2	13.5	32.4	19.8	14.7	29.2
LLaMA2-70B	68.4	8.2	61.2	12.4	40.7	22.5
Vicuna-7B	32.2	13.1	24.8	18.2	13.2	29.3
Vicuna-33B	57.6	9.6	52.0	15.6	32.4	26.8

Table 4: The performance (accuracy and rounds) of different LLMs with different difficulties of starting points.

	Two Wrong Answers		Three Wrong Answers		Four Wrong Answers	
	Accuracy	Rounds	Accuracy	Rounds	Accuracy	Rounds
BLOOM-7B	12.3	26.1	9.5	29.3	3.6	29.5
BLOOM-176B	65.5	19.0	58.4	23.4	53.5	27.4
Claude2	68.3	16.1	61.0	20.0	58.3	23.1
Falcon-7B	12.5	24.9	10.9	28.6	7.1	29.3
Falcon-180B	51.0	17.5	37.2	22.5	31.5	26.8
GPT3.5	63.7	16.4	55.2	22.6	45.7	26.4
GPT4	78.8	15.8	68.4	19.5	62.5	24.5
LLaMA2-7B	24.8	23.8	13.8	26.5	7.2	29.4
LLaMA2-70B	55.7	15.9	46.6	19.4	40.5	23.6
Vicuna-7B	17.4	25.5	8.0	29.0	3.3	29.5
Vicuna-33B	47.8	20.3	32.3	25.8	26.7	28.3

Table 5: The performance (accuracy and rounds) of different LLMs with different number of wrong answers.

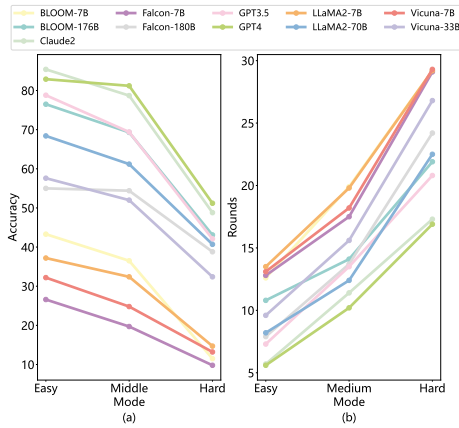


Figure 6: The performance (accuracy and rounds) of different LLMs with different difficulties of starting points.

We also analyze whether the number of wrong answers significantly affect LLMs' performance as shown in Table 5 and Fig. 7. We observe that with an increase in the number of wrong answers, the accuracy of all LLMs generally decreases, while the number of rounds needed increases or remains

relatively stable. The lack of change of rounds is mainly because the limit of rounds (that is 20) has been reached. For example, GPT4 has an accuracy rate of 78.8% with just two wrong answers, but this drops to 62.5% with four wrong answers. Particularly for poorly performing LLMs, such as Falcon-7B and Vicuna-7B, their accuracy dramatically drop almost to the point of indistinction as the number of wrong answers increases, demonstrating their vulnerability in processing misleading information. In terms of rounds, most models either need more rounds to answer or show no significant change when the number of wrong answers increases. For instance, BLOOM-176B needs 19.0 rounds with two wrong answers, but this number increases to 27.4 rounds with four wrong answers.

Question 5: In the hard mode, is there a strict proportional relationship between accuracy and the rethink ability? Answer 5: No, they are only positively correlated and the rating of rethink ability is generally higher than accuracy!

We adopt Pearson correlation coefficient that normalized to a 1-100 scale to analyze the rela-

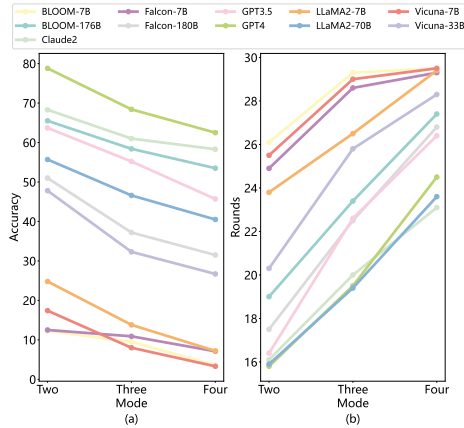


Figure 7: The performance (accuracy and rounds) of different LLMs with different number of wrong answers.

relationship between the accuracy of LLMs and their rethink abilities that backtrack from misleading information as shown in Fig. 8. The closer the colors, the higher the correlation. We find that most LLMs show similar shades, indicating a positive correlation but not a strict proportional relationship between their accuracy and their rethink abilities to detect the confusion. Specifically, GPT4 and Claude2 score high in both accuracy and rethink abilities, with a deeper blue color indicating a strong relationship between the two performances. In contrast, LLaMA2-7B, Falcon-7B, and BLOOM-7B score lower in both metrics and show lighter colors, suggesting weaker performance and lower correlation. Furthermore, the scores for the ability to rethink are generally higher than those for accuracy, which implies that even though a language model can backtrack when faced with misleading information, it does not always result in high accuracy in the final judgment. This might be because language models still struggle to effectively utilize this information, particularly in cases where the information is ambiguous.

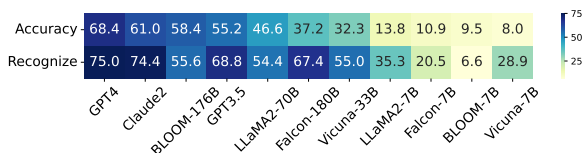


Figure 8: The Pearson correlation coefficient that normalized to a 1-100 scale between accuracy and rethink capabilities of different LLMs.

3.4 Case study

We show a good running example with questions generated by the top two LLMs (i.e. GPT4 and Claude2) and human in Fig. 9 and Fig. 10, respec-

tively. More cases are shown in Fig. 11, Fig. 12.

In Fig. 9, we find that both GPT4 and Claude2 make correct reasoning with just three questions in the easy mode. GPT4 focuses on determining the type of musical instrument and excluding percussion instruments, while Claude2 directly asks if it is a string instrument and if it belongs to classical music instruments. This shows that both are efficient in identifying with basic questions. In the medium mode with a harder starting point, GPT4 needs ten questions, while Claude2 only needs four. GPT4 starts by confirming it is a string instrument and narrows down to a guitar, including questions about whether it is played with a bow and if it is common in rock music. In contrast, Claude2's questions are more direct, quickly moving from whether it is a string instrument and played with a bow to identifying it as a guitar. This suggests that Claude2 is slightly more efficient and accurate in a more complex situation. In the hard mode, both GPT4 and Claude2 show the ability to narrow down gradually to identify the guitar. GPT4 takes ten questions, starting with confirming it is a string instrument and then asking if it is common in rock music and if it has frets to exclude the violin. Claude2 needs twelve questions, also starting with confirming it is a string instrument, but focusing more on physical characteristics like size and whether it has a hollow body, eventually asking if it is part of the guitar family. In this situation, GPT4 shows higher efficiency, although both are accurate.

Overall, the differences between them may stem from their specific ways of processing information and decision-making. However, there are also some bad cases, including wrong logic reasoning that ask same questions which have been answered before, not following instructions that asks non-yes-or-no questions or self-questioning as shown in Fig. 12.

4 Related Work

4.1 Gaming abilities of LLMs

Recent research shows the capabilities of LLMs in various gaming scenarios. For example, Zhao and Anderson (2023) evaluate LLMs in solving and creating puzzles for the NPR Sunday Puzzle game show; Jiang et al. (2023) introduce BRAIN-TEASER to assess lateral thinking in LLMs; Lorè and Heydari (2023) investigate strategic decision-making in LLMs through game theory; Egan et al. (2022) propose new summary evaluation metrics using the Shannon Game; Brookins and DeBacker

Entity: Guitar		Simple starting point: Musical instrument	Harder starting point: Music related things	Similar entity: Violin
Easy Mode		Medium Mode		Hard Mode
GPT4		Q1: Is it a stringed instrument? A1: Yes Q2: Is it typically played with a bow? A2: No Q3: Is it a guitar? Guess: Guitar. A3: Correct	Q1: Is the thing you're thinking of a musical instrument? A1: Yes Q2: Is the musical instrument you're thinking of part of the percussion family? A2: No Q3: Does the musical instrument you're thinking of a guitar? Guess: Guitar. A3: Correct	Q1: Is it a string instrument? A1: Yes Q2: Is it typically played with a bow? A2: Yes Q3: Is it larger than a violin? A3: No ... Q8: Is it commonly used in rock music? A8: Yes Q9: Does it have frets? A9: Yes Q10: Is it an electric guitar? Guess: Guitar. A10: Correct
Claude2		Q1: Is it a string instrument? A1: Yes Q2: Is it typically considered a classical music instrument? A2: No Q3: Is it a guitar? Guess: Guitar. A3: Correct	Q1: Is it an instrument? A1: Yes Q2: Is it a string instrument? A2: Yes Q3: Is it played with a bow? A3: No Q4: Is it a guitar? Guess: Guitar. A4: Correct	Q1: Is it a string instrument? A1: Yes Q2: Is it typically held against the shoulder when playing? A2: Yes Q3: Is it larger than a violin? A3: No ... Q10: Does it have a hollow body? A10: Yes Q11: Does it have a fretted neck? A11: Yes Q12: Is it one of the instruments in the guitar family (guitar, bass guitar, etc.)? Guess: Guitar. A12: Correct

Figure 9: Questions posed by the top2 LLMs, i.e. GPT4 and Claude2, among three modes for a given entity.

Human Entity: Guitar

Easy Mode (Simple starting point: Musical instrument)

Q1: Is it modern? A1: Yes
Q2: Does it have strings? A2: Yes
Q3: Does it hang on the body? A3: Yes
Q4: Does it have six strings? A4: Yes
Q5: Is it a guitar? A5: Correct

Medium Mode (Harder starting point: Music related things)

Q1: Is it a physical item?; A1: Yes
Q2: Is it an item that plays music?; A2: No
...
Q6: Is it a musical instrument?; A6: Yes
Q7: Is it a violin?; A9: No
Q8: Is it a Western musical instrument?; A10: Yes
Q9: Is it a guitar?; A9: Correct

Hard Mode (Similar entity: Violin)

Q1: Is it an abstract thing?; A1: No
Q2: Does it have a specific manifestation in the real world?; A2: Yes
Q3: Is it an audio clip?; A3: No
Q4: Is the object that produces this sound artificial?; A4: No
...
Q8: Can it be played both acoustically and electrically? A8: Yes.
Q9: Is it commonly used in rock music? A9: Yes.
Q10: Is it a guitar? Guess: Guitar. ; A10: Yes.

Figure 10: Questions posed by human among three modes for a given entity.

(2023) study LLM preferences in strategic games like the dictator game and the prisoner’s dilemma; O’Gara (2023) explore deception and lie detection in LLMs through the game Hoodwinked; Liga and Pasetto (2023) investigate the spatial reasoning of LLMs in the game of tic-tac-toe. Inspired but different from the previous work, our research evaluates the information processing and problem-solving capability of LLMs under the incomplete information scenario including but not limited to their knowledge retrieval, misleading information recognition capabilities through one simple game.

4.2 Evaluation for LLMs’ Capabilities

Recent research has extensively explored the capabilities of LLMs across various domains. For example, Ziems et al. (2023) demonstrate that LLMs can significantly contribute to Computational Social Science by classifying and explaining social phenomena; Zheng et al. (2023) explore the use of LLMs as judges to evaluate chat assistants, introducing benchmarks like MT-bench

and Chatbot Arena; Zhong et al. (2023) focus on comparing the understanding ability of ChatGPT with fine-tuned BERT-style models using the GLUE benchmark; Laskar et al. (2023) present a comprehensive evaluation of ChatGPT on diverse academic datasets, including question-answering, text summarization, code generation, commonsense reasoning; Valmeekam et al. (2023a) introduce PlanBench, a benchmark for evaluating LLMs on planning and reasoning; Del and Fishel (2023) introduce a benchmark consisting of long-form mystery narratives in assessing LLMs’ advanced reasoning abilities; Sawada et al. (2023) propose a benchmark containing advanced reasoning problems across multiple domains to evaluate the advanced reasoning capabilities of LLMs; Valmeekam et al. (2023b) investigate the planning abilities of LLMs in commonsense tasks and as heuristic guidance for other agents; While the aforementioned studies design different benchmarks in evaluating LLMs’ capabilities, there is lack of a benchmark for evaluating the information processing and problem-solving capability of LLMs under incomplete information scenarios.

5 Conclusions and Future Work

In conclusion, our study highlights the importance of a multifaceted approach to evaluating the information processing and problem-solving capability of LLMs under incomplete information scenarios. The BrainKing game, as a novel benchmark, successfully challenges LLMs in various cognitive aspects, from basic knowledge retrieval to complex reasoning and confusion recognition. Our findings offer a detailed understanding of the strengths and limitations of current LLMs, underscoring the need for LLMs to not only process information accurately but also navigate through misleading

information effectively. For future work, there is potential to extend the complexity and scope of the BrainKing, introducing more challenging scenarios to better mimic real-world conditions.

Limitations

There are a few key limitations. First, the Twenty Questions dataset used may not fully represent all possible types of entities. Moreover, selecting only the top 10,000 most common entities might limit the breadth and depth of the assessment. Second, in the hard mode of the game, LLMs are required to correctly identify the target entity based on potentially misleading information. This might be an overly challenging demand, especially for LLMs not yet fully trained to handle such complex tasks. Third, determining whether an LLM can recover from misleading information and return to the correct line of questioning can be a complex process. It involves a detailed analysis of the entire question-and-answer session, which could add to the complexity and subjectivity of the evaluation.

Acknowledgements

This work is supported by Science and Technology Commission of Shanghai Municipality Grant (No. 22511105902), Shanghai Municipal Science and Technology Major Project (No.2021SHZDZX0103), the National Natural Science Foundation of China (No.62072323), Shanghai Science and Technology Innovation Action Plan (No. 22511104700), and the Zhejiang Lab Open Research Project (NO. K2022NB0AB04).

References

- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#).
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional ai: Harmlessness from ai feedback](#).
- Marcel Binz and Eric Schulz. 2023. Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120.
- Philip Brookins and Jason Matthew DeBacker. 2023. Playing games with gpt: What can we learn about a large language model from canonical strategic games? *Available at SSRN 4493398*.
- Noam Brown and Tuomas Sandholm. 2019. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#).
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. *See <https://vicuna.lmsys.org> (accessed 14 April 2023)*.
- John Joon Young Chung, Ece Kamar, and Saleema Amershi. 2023. Increasing diversity while maintaining accuracy: Text data generation with large language models and human interventions. *arXiv preprint arXiv:2306.04140*.
- Maksym Del and Mark Fishel. 2023. True detective: a deep abductive reasoning benchmark undoable for gpt-3 and challenging for gpt-4. In *Proceedings of the The 12th Joint Conference on Lexical and Computational Semantics (* SEM 2023)*, pages 314–322.
- Nicholas Egan, Oleg Vasilyev, and John Bohannon. 2022. Play the shannon game with language models: A human-free approach to summary evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10599–10607.
- Robert Gibbons and Robert Gibbons. 1992. A primer in game theory.
- Gerd Gigerenzer and Wolfgang Gaissmaier. 2011. Heuristic decision making. *Annual review of psychology*, 62:451–482.

- H Giordano, P Panchal, B Holda, and EA Zillmer. 1998. Using the twenty questions task in neuropsychology. *Archives of Clinical Neuropsychology*, 13(1):104–104.
- Huang Hu, Xianchao Wu, Bingfeng Luo, Chongyang Tao, Can Xu, Wei Wu, and Zhan Chen. 2018. Playing 20 question game with policy-based reinforcement learning. *arXiv preprint arXiv:1808.07645*.
- Yifan Jiang, Filip Ilievski, Kaixin Ma, and Zhivar Sourati. 2023. [Brainteaser: Lateral thinking puzzles for large language models](#).
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty, and Jimmy Xiangji Huang. 2023. [A systematic study and comprehensive evaluation of chatgpt on benchmark datasets](#).
- Yinghao Li, Haorui Wang, and Chao Zhang. 2023. [Assessing logical puzzle solving in large language models: Insights from a minesweeper case study](#).
- Davide Liga and Luca Pasetto. 2023. Testing spatial reasoning of large language models: the case of tic-tac-toe. In *AIxPAC*.
- Jonathan Light, Min Cai, Sheng Shen, and Ziniu Hu. 2023. [Avalonbench: Evaluating llms playing the game of avalon](#). *arXiv e-prints*, pages arXiv–2310.
- Jiongnan Liu, Jiajie Jin, Zihan Wang, Jiehan Cheng, Zhicheng Dou, and Ji-Rong Wen. 2023. [Reta-llm: A retrieval-augmented large language model toolkit](#). *arXiv preprint arXiv:2306.05212*.
- Nunzio Lorè and Babak Heydari. 2023. [Strategic behavior of large language models: Game structure vs. contextual framing](#).
- David Noever and Forrest McKee. 2023. Chatbots as problem solvers: Playing twenty questions with role reversals. *arXiv preprint arXiv:2301.01743*.
- Aidan O’Gara. 2023. [Hoodwinked: Deception and cooperation in a text-based game for language models](#).
- OpenAI. 2023. [Gpt-4 technical report](#).
- Hong Ri, Xiaohan Kang, Mohd Nor Akmal Khalid, and Hiroyuki Iida. 2022. The dynamics of minority versus majority behaviors: A case study of the mafia game. *Information*, 13(3):134.
- Jon Saad-Falcon, Joe Barrow, Alexa Siu, Ani Nenkova, Ryan A Rossi, and Franck Dernoncourt. 2023. [Pdf-triage: Question answering over long, structured documents](#). *arXiv preprint arXiv:2309.08872*.
- Tomohiro Sawada, Daniel Paleka, Alexander Havrilla, Pranav Tadepalli, Paula Vidas, Alexander Kranias, John J. Nay, Kshitij Gupta, and Aran Komatsuzaki. 2023. [Arb: Advanced reasoning benchmark for large language models](#).
- Robert S Siegler. 1977. The twenty questions game as a form of problem solving. *Child Development*, pages 395–403.
- Julian Szymański and Włodzisław Duch. 2012. Information retrieval with semantic memory model. *Cognitive Systems Research*, 14(1):84–100.
- Fujio Toriumi, Hirotaka Osawa, Michimasa Inaba, Daisuke Katagami, Kosuke Shinoda, and Hitoshi Matsubara. 2017. Ai wolf contest—development of game ai using collective intelligence—. In *Computer Games: 5th Workshop on Computer Games, CGW 2016, and 5th Workshop on General Intelligence in Game-Playing Agents, GIGA 2016, Held in Conjunction with the 25th International Conference on Artificial Intelligence, IJCAI 2016, New York, USA, July 9-10, 2016, Revised Selected Papers 5*, pages 101–115. Springer.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Karthik Valmeekam, Matthew Marquez, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2023a. [Planbench: An extensible benchmark for evaluating large language models on planning and reasoning about change](#).
- Karthik Valmeekam, Sarath Sreedharan, Matthew Marquez, Alberto Olmo, and Subbarao Kambhampati. 2023b. [On the planning abilities of large language models \(a critical investigation with a proposed benchmark\)](#).
- Mansfield Tracy Walsorth. 1882. *Twenty Questions: A short treatise on the game to which are added a code of rules and specimen games for the use of beginners*. Holt.

Reed G Williams and Debra L Klamen. 2015. Twenty questions game performance on medical school entrance predicts clinical performance. *Medical Education*, 49(9):920–927.

BigScience Workshop, :, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klam, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Froberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Al-mubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rhea Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li,

Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanjit Gandhi, Shaden Smith, Stéphane Reuena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névél, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najaoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruo Chen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Daniel McDuff, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Ononiwu, Habib Rezanejad, Hessie Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguié, Thanh Le, Tobi Oyeade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourier, Daniel León Perrián, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Marianna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda,

- Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aaronsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yannis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2023. [Bloom: A 176b-parameter open-access multilingual language model](#).
- Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. 2023. [Exploring large language models for communication games: An empirical study on werewolf](#).
- Jingmiao Zhao and Carolyn Jane Anderson. 2023. [Solving and generating npr sunday puzzles with large language models](#).
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du, and Dacheng Tao. 2023. [Can chatgpt understand too? a comparative study on chatgpt and fine-tuned bert](#).
- Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2023. [Can large language models transform computational social science?](#)

Entity : Guitar

Simple starting point: Musical instrument

Harder starting point: Music related things

Similar entity: Violin

GPT3.5

Easy Mode:

Q1: Is the musical instrument you're thinking of commonly found in orchestras?; A1: No

Q2: Is the musical instrument you have in mind typically made of wood?; A2: No

...

Q5: Is the musical instrument you're thinking of stringed?; A5: Yes

Q6: Is the musical instrument you're thinking of commonly associated with folk music?; A6: Yes

Q7: Is the musical instrument you have in mind a type of guitar? Guess: Guitar; A7: Correct

Medium Mode:

Q1: Is the thing you're thinking of a musical instrument?; A1: Yes

Q2: Is the instrument you're thinking of part of the percussion family?; A2: No

Q3: Is the instrument you're thinking of a string instrument?; A3: Yes

...

Q11: Is the instrument you're thinking of often played with fingerpicking technique?; A11: Yes

Q12: Is the instrument you're thinking of commonly used in blues music?; A12: Yes

Q13: Is the instrument you're thinking of a resonator guitar with a wooden body? Guess: Guitar; A13: Correct

Hard Mode:

Q1: Is the musical instrument you're thinking of typically made of wood?; A1: Yes

...

Q3: Is the musical instrument you have in mind commonly used in classical music?; A3: Yes

Q4: Is this musical instrument typically played with a bow?; A4: Yes

...

Q6: Is the musical instrument you're thinking of typically larger than a violin?; A6: Yes

...

Q12: Is the musical instrument you have in mind a piano?; A12: No

Q13: Is the musical instrument you're thinking of a percussion instrument?; A13: No

Q14: Is the musical instrument you're thinking of a guitar? Guess: Guitar; A14: Correct

Falcon-180B

Easy Mode:

Q1: Is the instrument made of wood?; A1: Yes

Q2: Is it a string instrument?; A2: Yes

...

Q7: Is it smaller than a guitar?; A7: No

Q8: Is it an acoustic instrument?; A8: Yes

Q9: Is it a type of guitar? Guess: Guitar; A9: Correct

Medium Mode:

Q1: Is the music related thing a physical object?; A1: Yes

Q2: Is the music related thing a musical instrument?; A2: Yes

Q3: Is the musical instrument a string instrument?; A3: Yes

...

Q13: Is the musical instrument you have in mind commonly used in classical music?; A13: Yes

Q14: Is this musical instrument typically played with a bow?; A14: Yes

Q15: Is the string instrument a type of guitar? Guess: Guitar; A15: Correct

Hard Mode:

Q1: Is the entity a string instrument?; A1: Yes

...

Q7: Is the entity a member of the violin family?; A7: Yes

Q9: Is the entity a member of the cello family?; A9: Yes

...

Q15: Is the entity a fretted string instrument?; A15: Yes

Q16: Is the entity a plucked fretted string instrument?; A16: Yes

Q17: Is the entity an acoustic guitar? Guess: Guitar; A17: Correct

BLOOM-176B

Easy Mode:

Q1: Is the musical instrument you're thinking of a string instrument?; A1: Yes

Q2: Is this instrument typically played with a bow?; A2: No

...

Q6: Does this instrument have six strings?; A6: Yes

Q7: Is this instrument typically played while being held horizontally across the body?; A7: Yes

Q8: Based on your responses, it seems like the instrument you're thinking of might be a guitar. Is that correct? Guess: Guitar; A8: Correct

Medium Mode:

Q1: Is the thing you're thinking of a musical instrument?; A1: Yes

Q2: Is this musical instrument typically found in a classical orchestra?; A2: No

Q3: Is this musical instrument primarily used in modern or popular music genres?; A3: Yes

...

Q12: Is this musical instrument typically played with a pick?; A12: Yes

Q13: Is this musical instrument electric?; A13: Yes

Q14: It seems like you're thinking of an electric guitar. Am I correct? Guess: Guitar; A14: Correct

Hard Mode:

Q1: Is the musical instrument you're thinking of a string instrument?; A1: Yes

Q2: Is this instrument typically played with a bow?; A2: Yes

Q3: Is this instrument larger than a violin?; A3: No

...

Q12: Does this instrument have strings?; A12: Yes

Q13: Is this instrument primarily made of wood?; A13: Yes

Q14: Is the instrument you're thinking of a guitar? Guess: Guitar; A14: Correct

LLaMA2-70B

Easy Mode:

Q1: Is the musical instrument you're thinking of a stringed instrument?; A1: Yes

Q2: Does the instrument have a neck?; A2: Yes

...

Q9: Is the instrument an electric guitar? Guess: Guitar; A9: Correct

Medium Mode:

Q1: Is the thing you're thinking of a musical instrument?; A1: Yes

Q2: Does the instrument you're thinking of have strings?; A2: Yes

Q3: Does the instrument you're thinking of have a neck?; A3: Yes

...

Q14: Does the instrument you're thinking of have frets?; A14: Yes

Q15: Does the instrument you're thinking of have a pickup?; A15: Yes

Q16: Is the instrument you're thinking of an electric guitar? Guess: Guitar; A16: Correct

Hard Mode:

Q1: Is the musical instrument you're thinking of a stringed instrument?; A1: Yes

Q2: Does the instrument have a neck?; A2: Yes

Q3: Does the instrument have frets?; A3: No

Q4: Is the instrument typically played with a bow?; A4: Yes

...

Q15: This time, I'm going to guess that the musical instrument you're thinking of is a guitar. Is that correct? Guess: Guitar; A15: Correct

Figure 11: Questions posed by different LLMs among three modes for a given entity.

Entity : Guitar**Simple starting point:** Musical instrument**Harder starting point:** Music related things**Similar entity:** Violin**BLOOM-7B****Easy Mode:**

Q1: Is the musical instrument you're thinking of a stringed instrument?;

A1: Yes

Q2: Does the instrument have a neck?;

A2: Yes

...

Q8: Is this a Fender instrument?;

A8:Yes

Q9: Is it electric guitar? Guess: Guitar;

A9:Correct

Medium Mode:

Q1: Is the music thing you're thinking of a vocal-based genre?;

A1: No

Q2: Is the music thing you're thinking of a type of instrument?;

A2: Yes

...

Q11: Is the instrument a wind instrument?;

A11: No

Q12: Does the instrument have strings?;

A12: Yes

Q13: Is the instrument a guitar? Guess: Guitar;

A13:Correct

Hard Mode:

Q1: Is the musical instrument you're thinking of a wind instrument?;

A1:No

Q2: Is the musical instrument you're thinking of a stringed instrument?;

A2:Yes

Q3: Does the musical instrument you're thinking of have six strings?;

A3:No

Q4: Does the musical instrument you're thinking of have a neck?; A4:Yes

...

Q17: Does the musical instrument you're thinking of have an electronic components?; A17:No

Q18: Does the musical instrument you're thinking of have a bow?; A18:Yes

Q19:Does the musical instrument you're thinking of have a round body?; A19:No

Q20:Is that the musical instrument you're thinking of a cello? Guess: Cello;

A20: No

EXCEEDING THE LIMIT OF ROUNDS

Vicuna-33B**Easy Mode:**

Q1: Is this musical instrument played with a bow?; A1: No

Q2: Is this musical instrument a string instrument?; A2: Yes

...

Q8:Is this instrument often used in rock music?; A8:Yes

Q9:Is this instrument played by plucking the strings with fingers or a pick?; A9:Yes

Q10: Is this musical instrument an electric guitar? Guess: Guitar;

A10: Correct

Medium Mode:

Q1: Is the thing a song?; A1: No

Q2: Is the thing an album?; A2: No

Q3: Is the thing a musician or a band?; A3: No

Q4: Is the thing a music festival?; A4: No

...

Q13: Is the thing a musical instrument?; A13: Yes

Q14: Does the instrument have a neck?; A14: Yes

Q15: Is the thing an acoustic guitar? Guess: Guitar;

A15: Correct

Hard Mode:

Q1: Is this musical instrument a string instrument?; A1:Yes

Q2:Is the musical instrument you have in mind commonly used in classical music?; A2:Yes

Q3:Is this musical instrument typically played with a bow?; A3:Yes

...

Q18: Is this musical instrument a keyboard instrument?; A18:No

Q19: Is this musical instrument a vocal instrument?; A19:No

Q20: Is this musical instrument a reed instrument?; A20:No

EXCEEDING THE LIMIT OF ROUNDS

Vicuna-7B**Easy Mode:**

Q1: Is it a string instrument?; A1: Yes

SELF-QUESTIONING

Medium Mode:

Q1: Is it a singer?; A1: No

SELF-QUESTIONING

Hard Mode:

Q1: Is it a string instrument?; A1: Yes

SELF-QUESTIONING

LLaMA2-7B**Easy Mode:**

Q1: How many questions are there?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Medium Mode:

Q1: What's the thing like?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Hard Mode:

Q1: Do you like it?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Falcon-7B**Easy Mode:**

Q1: What musical instrument has 4 strings?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Medium Mode:

Q1: What is the musical related things?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Hard Mode:

Q1: What is this musical instrument?

NON-COMPLIANCE WITH YES OR NO INSTRUCTIONS.

Figure 12: Questions posed by different LLMs among three modes for a given entity.