

# Is Your Large Language Model Knowledgeable or a Choices-Only Cheater?

**Nishant Balepur**  
University of Maryland  
nbalepur@umd.edu

**Rachel Rudinger**  
University of Maryland  
rudinger@umd.edu

## Abstract

Recent work shows that large language models (LLMs) can answer multiple-choice questions using only the choices, but does this mean that MCQA leaderboard rankings of LLMs are largely influenced by abilities in choices-only settings? To answer this, we use a contrast set that probes if LLMs over-rely on choices-only shortcuts in MCQA. While previous works build contrast sets via expensive human annotations or model-generated data which can be biased, we employ graph mining to extract contrast sets from existing MCQA datasets. We use our method on UnifiedQA, a group of six commonsense reasoning datasets with high choices-only accuracy, to build an 820-question contrast set. After validating our contrast set, we test 12 LLMs, finding that these models do not exhibit reliance on choice-only shortcuts when given both the question and choices. Despite the susceptibility of MCQA to high choices-only accuracy, we argue that LLMs are not obtaining high ranks on MCQA leaderboards just due to their ability to exploit choices-only shortcuts.<sup>1</sup>

## 1 Introduction

Multiple-choice question answering (MCQA) is a popular task to test the knowledge of large language models (LLMs) (Robinson and Wingate, 2023). However, recent work shows that LLMs surpass majority class baselines in choices-only settings—when no question and just the choices are given in a prompt (Balepur et al., 2024). This raises the question: *Do models obtain high ranks in MCQA leaderboards due to their pretraining knowledge or their ability to exploit choices-only shortcuts?* Resolving this query is key to ensure that MCQA leaderboards reliably rank the knowledge of LLMs.

To answer this question, we use a variation of *contrast sets*—small datasets that test if models “pay attention” to perturbed attributes that should

<sup>1</sup>Our code is available at <https://github.com/nbalepur/mcqa-artifacts>

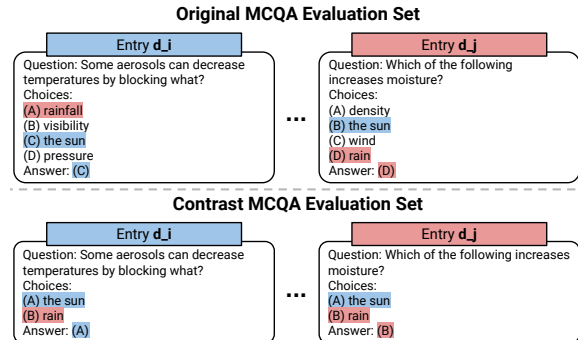


Figure 1: Example of a contrast MCQA evaluation set.

alter the model’s decision (Levesque et al., 2012; Gardner et al., 2020). For our purposes, we need a contrast set containing pairs of MC entries with identical answer choices, but varied questions that lead to distinct answers. For example, in Figure 1 (bottom), the MC entries  $d_i$  and  $d_j$  have the same choices of “the sun” and “the rain”, but  $d_i$  has a question that answered by “the sun”, and similarly for  $d_j$ . This design ensures that LLMs relying only on shortcuts or patterns in the choices, while ignoring questions, can perform no better than random chance. If a model ranks highly on an MCQA dataset but largely drops in rank on a contrast set based on this dataset, it would reveal that this model obtains a high rank on the original dataset primarily by employing choices-only shortcuts.

Contrast sets are usually built through manual annotation efforts (Gardner et al., 2020; Srikanth and Rudinger, 2022), as model-generated data can be biased. However, writing MC questions with high-quality distractors is difficult even for experts (Gierl et al., 2017). Further, rewritten questions can exhibit distributional differences from the original questions, altering the difficulty of the MCQA task.

To address this issue, we cast the creation of contrast sets for MCQA datasets to a graph mining task. We treat each MC entry  $d_i$  in the dataset as a vertex in an undirected graph, and draw edges

between entries  $d_i$  and  $d_j$  if the gold answer from  $d_i$  is semantically equivalent to a distractor in  $d_j$ , and vice versa. For instance, in Figure 1 (top), the gold answer of “rain” in  $d_i$  is semantically similar to the distractor of “rainfall” in  $d_j$  and vice versa, so we draw an edge between  $d_i$  and  $d_j$ . Thus, an edge  $(d_i, d_j)$  means that the gold answers in  $d_i$  and  $d_j$  can form a set of choices, with questions in  $d_i$  and  $d_j$  leading to distinct answers in said choices, mirroring the desired format of our contrast set. We find the maximum matching of this graph to obtain the largest contrast set of distinct MCQA questions derived from the initial dataset. This method overcomes the burden of writing contrast sets, while only minimally using models for semantic equivalence, reducing the risk of model-generated biases.

We use our approach to build an 820-question contrast set from six commonsense MCQA datasets from the UnifiedQA collection (Khashabi et al., 2020). We first ask three annotators to assess our contrast set, finding that it has questions with plausible distractors (§4.1). This finding suggests that we have built a high-quality MCQA contrast set.

After verifying the quality of our contrast set, we test 12 LLMs (Touvron et al., 2023; Penedo et al., 2023; Jiang et al., 2023; Young et al., 2024; Team et al., 2024) on the UnifiedQA evaluation set and its mined contrast set (§4.2). Our LLMs surpass random guessing using just the choices on the original evaluation set, aligning with prior work. (Balepur et al., 2024). However, when prompted with both the question *and* choices, LLM accuracy rankings between the initial evaluation set and contrast set are highly consistent, with Kendall’s  $\tau$  near 0.9.

Since no LLM rank drops markedly, we claim that our tested LLMs are not ranking highly on MCQA leaderboards solely due to their ability to exploit choices-only shortcuts. Despite the susceptibility of MCQA to high choices-only accuracy, we argue the task may still reliably rank LLM knowledge. As a result, we recommend that future works continue to explore the behavior of LLMs in choices-only settings to help explain how LLMs can adeptly perform MCQA without the question.

## 2 Automatic Contrast Set Creation

We assume we are given an MCQA dataset  $\mathcal{D}$  with data entries  $d_i = (q_i, \mathcal{C}_i, a_i)$ , where  $q_i$  is a question,  $\mathcal{C}_i$  is a list of choices, and  $a_i \in \mathcal{C}_i$  is the gold answer. Our goal is to build a contrast set  $\mathcal{D}_{contr}$  from  $\mathcal{D}$  to probe if LLMs rely on choice-only shortcuts

in MCQA. Typically, humans manually create contrast sets (Srikanth and Rudinger, 2022; Gardner et al., 2020), as model-generated data can be biased (Yu et al., 2024). However, since writing MCQA problems is challenging even for experts (Offerijns et al., 2020; Gierl et al., 2017), we seek to automatically mine a contrast set  $\mathcal{D}_{contr}$  from the original dataset  $\mathcal{D}$  without model-generated data.

To automatically build contrast sets, we need MCQA entry pairs in the style of Figure 1—pairs with the same choices  $\mathcal{C}' = \{a_i, a_j\}$ , but questions  $q_i$  and  $q_j$  leading to distinct answers  $a_i$  and  $a_j$  in  $\mathcal{C}'$ , respectively. We define this format as an **entry pair**  $p_{ij} = \langle (q_i, \{a_i, a_j\}, a_i), (q_j, \{a_i, a_j\}, a_j) \rangle$ . Creating the largest possible  $\mathcal{D}_{contr}$  with distinct questions is equivalent to finding the maximum set of unique entry pairs  $p_{ij}$  in  $\mathcal{D}$ . In the next sections, we outline our graph-based approach to mine entry pairs from the original dataset  $\mathcal{D}$  to form  $\mathcal{D}_{contr}$ .

### 2.1 Graph Representation

While a simple strategy to find an entry pair  $p_{ij}$  is to sample two entries  $(q_i, \mathcal{C}_i, a_i), (q_j, \mathcal{C}_j, a_j) \in \mathcal{D}$  and let  $\mathcal{C}' = \{a_i, a_j\}$ , this may result in low-quality questions, as there is no constraint that  $a_x$  and  $a_y$  form a plausible set of choices (§4.1). For instance, if  $a_i$  is a ratio and  $a_j$  is an integer, choices  $\{a_i, a_j\}$  are implausible and result in a low-quality question. To address this, we intuit that the original dataset  $\mathcal{D}$  reveals if two answers  $a_i$  and  $a_j$  are plausible distractors for each other. For answers  $a_i \in \mathcal{C}_x$  and  $a_j \in \mathcal{C}_y$ , if  $a_i$  is semantically equivalent to a distractor  $c \in \mathcal{C}_j \setminus \{a_j\}$  and likewise for  $a_j$  and  $\mathcal{C}_i$ , the set of choices  $\mathcal{C}' = \{a_i, a_j\}$  will be plausible.

To execute this idea, we represent the dataset  $\mathcal{D}$  as an undirected graph  $\mathcal{G}$ . Each entry  $d_i \in \mathcal{D}$  is a vertex for  $\mathcal{G}$ . We draw an edge between entries  $d_i$  and  $d_j$  if the gold answer  $a_i$  is semantically equivalent to a distractor  $c \in \mathcal{C}_j \setminus \{a_j\}$  and vice versa, meaning that the choices  $a_i$  and  $a_j$  can form a plausible set of choices based on  $\mathcal{D}$ . We create edges with semantic equivalence over exact match to consider choices with minor differences, like “rain” and “rainfall” in Figure 1, increasing the candidate size of our contrast set. We compute semantic similarity via NLI-based embeddings (Conneau et al., 2017) and set a strict cosine similarity threshold of 0.85 to determine semantic equivalence.

### 2.2 Mining Entry Pairs

We now mine entry pairs from the graph  $\mathcal{G}$  to build a contrast set  $\mathcal{D}_{contr}$ . For any edge  $(d_i, d_j)$  in  $\mathcal{G}$ ,




Figure 2: Distractor plausibility ratings across methods.

we know that  $a_i$  and  $a_j$  form a set of plausible choices. When  $\mathcal{C}' = \{a_i, a_j\}$ , entries  $(q_i, \mathcal{C}', a_i)$  and  $(q_j, \mathcal{C}', a_j)$  form an entry pair  $p_{ij}$ . Using this idea, we build  $\mathcal{D}_{contr}$  by finding the maximum matching (Boppana and Halldórsson, 1992) of  $\mathcal{G}$ , which gives the largest set of edges in  $\mathcal{G}$  where no two edges are adjacent. Each edge in the maximum matching form an entry pair for the contrast set and since no edges are adjacent, each entry pair contains two unique questions. This yields the largest possible contrast set without duplicate questions.

### 3 Experimental Setup

#### 3.1 A Contrast Set for UnifiedQA

The purpose of our contrast set  $\mathcal{D}_{contr}$  is to study whether high choices-only accuracy influences the ranking of LLMs on MCQA leaderboards. Hence,  $\mathcal{D}_{contr}$  should be based on a dataset with high accuracy in choices-only settings. The two datasets from Balepur et al. (2024) with the highest choices-only accuracy are commonsense datasets (Clark et al., 2018; Zellers et al., 2019), so we derive  $\mathcal{D}_{contr}$  from an MCQA split of UnifiedQA (Khashabi et al., 2020), which has 7611 questions from six commonsense datasets: ARC (Clark et al., 2018), OpenBookQA (Mihaylov et al., 2018), CommonsenseQA (Talmor et al., 2019), QASC (Khot et al., 2020), PIQA (Bisk et al., 2020), and SIQA (Sap et al., 2019). Using our graph mining algorithm (§2), we build an 820-question contrast set. This size aligns with contrast set sizes in prior works, ranging from 600 to 1000 (Srikanth and Rudinger, 2022; Gardner et al., 2020).

#### 3.2 Prompt Design

Few-shot prompting is currently the only method that has shown LLMs can surpass random guessing when given only the choices as input. For our experiments, we follow the few-shot format of Balepur et al. (2024) and use a full prompt (3.1) to assess when LLMs see both the questions and choices, and a choices-only (3.2) prompt for just the choices:

#### Prompt 3.1: Full Prompt

```
Question:  $q$ 
Choices: \n (A)  $c_a$  \n (B)  $c_b$  \n (C)  $c_c$  \n (D)  $c_d$ 
Answer:  $a$ 
```

#### Prompt 3.2: Choices-Only Prompt

```
Choices: \n (A)  $c_a$  \n (B)  $c_b$  \n (C)  $c_c$  \n (D)  $c_d$ 
Answer:  $a$ 
```

In the boxes above, the non-highlighted text represents the model input, while the highlighted text represents the model generation. In the few-shot prompts, exemplars follow the same format shown in the prompt box with the highlighted text replaced by the ground truth (Example in Appendix A.2).

## 4 Results

### 4.1 Qualitative Analysis

To assess the quality of the contrast set produced by our graph mining algorithm, we ask three Ph.D. students in computer science to compare 50 of our questions versus a baseline that randomly picks entry pairs (details in Appendix B.1). These methods only differ by distractors, so following Gierl et al. (2017), we ask annotators to compare the **plausibility** of the two distractors as a proxy for question quality. All three annotators find that our method has significantly more plausible distractors than the baseline (Figure 2), suggesting that our extracted contrast set from UnifiedQA is high-quality.

### 4.2 Are LLMs Knowledgeable or Choices-only Cheaters?

Following our quality checks, we use our contrast set to study if high choices-only accuracy influences the ranking of LLMs when questions *and* choices are given. We assess 6 LLM families on the UnifiedQA evaluation set and our contrast set: LLaMA-2 (Touvron et al., 2023), Falcon (Penedo et al., 2023), Mistral (Jiang et al., 2023), Mixtral (Jiang et al., 2024), Gemma (Team et al., 2024), and Yi (Young et al., 2024). We use 5-shot and 10-shot full and choices-only prompts (Prompts 3.1, 3.2). Appendix A.3 has more prompting details.

On the UnifiedQA evaluation set, our LLMs often surpass random guessing with choices-only prompts (Figure 3, left), aligning with prior work (Balepur et al., 2024). Further, LLMs with higher ranks on the UnifiedQA evaluation set using the full prompt tend to have higher accuracy when using the choices-only prompt, suggesting a correlation between an LLM’s MCQA leaderboard rank and




Figure 3: Accuracy of twelve LLMs on the UnifiedQA evaluation set (left) versus its contrast set (right), sorted by full prompt accuracy. We show 5-shot (top) and 10-shot (bottom) prompts, with 3-shot prompts in Appendix B.2.

its ability to exploit choices-only shortcuts. Simply subtracting these values cannot quantify how an LLM performs in MCQA without choices-only shortcuts, since if an LLM can answer a MC question without the question, it does not imply the model is ignoring the question when it has access to the question (Srikanth and Rudinger, 2022).

To better quantify if LLMs are obtaining high ranks on UnifiedQA due to their ability to exploit choices-only shortcuts, we compare model ranks on the original UnifiedQA evaluation set to its contrast set. We note that if a certain LLM relied on choice-only shortcuts substantially more than other models, its contrast set accuracy ranking would largely drop compared to its evaluation set accuracy ranking, as it would be penalized for ignoring the question. However, in the UnifiedQA evaluation set and its contrast set, model rankings of full prompt accuracy are consistent; the 5-shot and 10-shot rankings have Kendall’s  $\tau$  of 0.88 and 0.91, indicating high consistency. Thus, we claim that the MCQA rankings of our LLMs on UnifiedQA do not primarily stem from their ability to perform well in choices-only settings, and none of our models are considered “choices-only cheaters.”

We find that if an LLM succeeds with choices-only prompts, it does not imply that this model’s performance in MCQA solely stems from its choices-only abilities. As a result, we believe that despite high choices-only accuracy, MCQA may still be a reliable task to rank the knowledge of LLMs. Further, our results stress the need for more work in explaining how high choices-only accuracy occurs. We believe such efforts are crucial to better interpret LLM knowledge and decision-making.

## 5 Related Work

**Contrast Sets:** Contrast sets (Gardner et al., 2020) or counterfactual augmentations (Kaushik et al., 2020; Srikanth et al., 2024), are datasets that probe if models “pay attention” to desired attributes (Elazar et al., 2023). This technique has been applied to many tasks, including natural language inference (Glockner et al., 2018; Ribeiro et al., 2020), story generation (Qin et al., 2019), and ethical judgements (Hendrycks et al., 2021a). While these datasets are often created manually, many works use generation models (Wu et al., 2021; Fryer et al., 2022) to create contrast sets. Instead, we are the first to employ graph mining to build contrast sets, limiting the potential for model-generated biases.

**MCQA Evaluation:** MCQA is a popular testbed not only for benchmarking LLMs (Beeching et al., 2023; Liang et al., 2023), but also for interpreting LLM decision-making. Previous works use MCQA to study prompt sensitivity (Pezeshkpour and Hruschka, 2023; Zheng et al., 2024), logical robustness (Balepur et al., 2023), and recently, the ability to perform MCQA without using the question (Balepur et al., 2024). We give more insights into this last phenomenon by probing if LLMs ignore the question even when it is given in the prompt.

## 6 Conclusion

We find that while LLMs can perform well in MCQA without access to the question, it does not mean that model rankings on MCQA leaderboards are largely influenced by this ability. This result supports the claim that MCQA can rank the knowledge and ability of LLMs to reason over both questions and choices. Further, we are aligned with recent work that suggests that high choices-only



accuracy does not necessarily imply that models are incapable of true reasoning or comprehension, so we hope future works continue to explore what strategies LLMs may employ to perform well in choices-only settings. Our application of graph mining to MCQA sheds light on one way to do this—the automatic construction of contrast sets—and we hope similar methods can be applied to other tasks to enhance LLM interpretability.

## 7 Limitations

One limitation lies in the application of our graph mining algorithm solely to the UnifiedQA dataset collection. We choose UnifiedQA for its tendency to elicit high accuracy with choices-only prompts, as commonsense reasoning MCQA datasets have shown to be susceptible to this phenomenon. Since our results show that LLMs rankings are highly consistent on this dataset prone to high choices-only accuracy, we believe these findings will hold for other MCQA datasets like MMLU (Hendrycks et al., 2021b) with lower choices-only accuracy. However, we invite future research to apply our graph mining algorithm to other datasets, including non-MCQA datasets, to build contrast sets that can further probe LLM decision-making.

Further, we acknowledge that our contrast set contains MCQA questions limited to two choices, diverging from the original evaluation set’s range of two to eight choices. While having less options does make it more likely for a model to guess the right answer, our qualitative analysis shows that the concepts tested in our contrast set are not markedly different in plausibility (§4.1), and thus are not too easy. Further, while our contrast set is easier in theory, it still preserves LLM rankings, even on the subset used to derive the contrast set (Appendix B.3), ultimately supporting the idea that MCQA can reliably rank LLMs capabilities.

## 8 Ethical Considerations

When models heavily rely on patterns or biases present in datasets, we may overestimate model abilities and face generalizability issues during deployment. In this work, we probe the extent to which LLMs over-rely on patterns in MCQA choices when provided both the question and choices in the prompt, ultimately finding that this effect is small. However, we believe it is still critical for LLM practitioners to be aware that LLMs can outperform random guessing when using just

the choices as input, as this could have downstream effects. We encourage future research efforts in designing special datasets that can help interpret specific abilities within LLM decision-making.

Further, we note that when any model is used in a data creation pipeline, there is the possibility of models propagating their own biases. We specifically address this issue by designing a graph mining algorithm that leverages minimal model intervention, only in the form of computing semantic similarity, which greatly lowers this risk compared to synthetic data generators like LLMs. We hope future works can adopt data creation pipelines with minimal model use similar to ours to avoid the risk of generating model-specific biases or artifacts.

## 9 Acknowledgements

We would like to thank members of the CLIP lab at the University of Maryland and external collaborators for their feedback and discussions of this work, including Neha Srikanth, Shi Feng, and Jordan Boyd-Graber. We also thank Yu Hou and Dayeon Ki for their help with annotations. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 2236417. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## References

- Nishant Balepur, Shramay Palta, and Rachel Rudinger. 2023. It’s not easy being wrong: Evaluating process of elimination reasoning in large language models. *arXiv preprint arXiv:2311.07532*.
- Nishant Balepur, Abhilasha Ravichander, and Rachel Rudinger. 2024. Artifacts or abduction: How do llms answer multiple-choice questions without the question? *arXiv preprint arXiv:2402.12483*.
- Edward Beeching, Clémentine Fourier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open llm leaderboard. [https://huggingface.co/spaces/HuggingFaceH4/open\\_llm\\_leaderboard](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard).
- Yonatan Bisk, Rowan Zellers, Jianfeng Gao, Yejin Choi, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 7432–7439.

- Ravi Boppana and Magnús M Halldórsson. 1992. Approximating maximum independent sets by excluding subgraphs. *BIT Numerical Mathematics*, 32(2):180–196.
- 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. *arXiv preprint arXiv:1803.05457*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Yanai Elazar, Bhargavi Paranjape, Hao Peng, Sarah Wiegrefe, Khyathi Raghavi, Vivek Srikumar, Sameer Singh, and Noah A Smith. 2023. Measuring and improving attentiveness to partial inputs with counterfactuals. *arXiv preprint arXiv:2311.09605*.
- Zee Fryer, Vera Axelrod, Ben Packer, Alex Beutel, Jilin Chen, and Kellie Webster. 2022. [Flexible text generation for counterfactual fairness probing](#). In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 209–229, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Mark J Gierl, Okan Bulut, Qi Guo, and Xinxin Zhang. 2017. Developing, analyzing, and using distractors for multiple-choice tests in education: A comprehensive review. *Review of Educational Research*, 87(6):1082–1116.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andrew Critch, Jerry Li, Dawn Song, and Jacob Steinhardt. 2021a. [Aligning {ai} with shared human values](#). In *International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021b. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2024. [Mistral of experts](#).
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *International Conference on Learning Representations*.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Tushar Khot, Peter Clark, Michal Guerquin, Peter Jansen, and Ashish Sabharwal. 2020. Qasc: A dataset for question answering via sentence composition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8082–8090.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Alexander Cosgrove, Christopher D Manning, Christopher Re, Diana Acosta-Navas, Drew Arad Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue WANG, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekogonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri S. Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Andrew Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Ko-reeda. 2023. [Holistic evaluation of language models](#).

- Transactions on Machine Learning Research*. Featured Certification, Expert Certification.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.
- Jeroen Offerijns, Suzan Verberne, and Tessa Verhoef. 2020. Better distractions: Transformer-based distractor generation and multiple choice question filtering. *arXiv preprint arXiv:2010.09598*.
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Hamza Alobeidli, Alessandro Cappelli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for falcon LLM: Outperforming curated corpora with web data only](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Pouya Pezeshkpour and Estevam Hruschka. 2023. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*.
- Lianhui Qin, Antoine Bosselut, Ari Holtzman, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi. 2019. [Counterfactual story reasoning and generation](#). 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 5043–5053, Hong Kong, China. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Joshua Robinson and David Wingate. 2023. [Leveraging large language models for multiple choice question answering](#). In *The Eleventh International Conference on Learning Representations*.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). 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 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Neha Srikanth, Marine Carpuat, and Rachel Rudinger. 2024. How often are errors in natural language reasoning due to paraphrastic variability? *arXiv preprint arXiv:2404.11717*.
- Neha Srikanth and Rachel Rudinger. 2022. [Partial-input baselines show that NLI models can ignore context, but they don't](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4753–4763, Seattle, United States. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). 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 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, et al. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Matthew West, Geoffrey L Herman, and Craig Zilles. 2015. Prairielearn: Mastery-based online problem solving with adaptive scoring and recommendations driven by machine learning. In *2015 ASEE Annual Conference & Exposition*, pages 26–1238.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2021. [Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models](#). 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 6707–6723, Online. Association for Computational Linguistics.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Yue Yu, Yuchen Zhuang, Jieyu Zhang, Yu Meng, Alexander J Ratner, Ranjay Krishna, Jiaming Shen, and Chao Zhang. 2024. Large language model as attributed training data generator: A tale of diversity and bias. *Advances in Neural Information Processing Systems*, 36.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) In *Annual Meeting of the Association for Computational Linguistics*.

Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2024. [Large language models are not robust multiple choice selectors](#). In *The Twelfth International Conference on Learning Representations*.



## A Experimental Setup

### A.1 Dataset Details

The UnifiedQA evaluation set has questions from the evaluation sets of the following six datasets:

- **ARC:** 1172 four-choice questions drawn from grade-school science questions.
- **OpenBookQA:** 500 four-choices questions modeled after open-book exams.
- **QASC:** 926 eight-choice questions about grade school science with a focus on sentence composition.
- **CommonsenseQA:** 1221 four-choice questions meant to test commonsense knowledge from ConceptNet.
- **Physical IQa:** 1838 two-choice questions about physical commonsense reasoning.
- **Social IQa:** 1954 three-choice questions involving reasoning about everyday social interactions.

After running our algorithm, our contrast set contains 377 questions from CommonsenseQA, 285 questions from QASC, 79 questions from ARC, 53 questions from Social IQa, 22 questions from OpenBookQA, and 4 questions from Physical IQa, all of which have two choices.

### A.2 Prompt Box Example

The following subsection is adapted directly from the Appendix of [Balepur et al. \(2024\)](#) to highlight the utility of their prompt boxes.

Below, we provide a detailed example to illustrate the application of our prompt boxes. Suppose we have the full prompt (Prompt 3.1):

```
Prompt 2.1: Full Prompt
Question:  $q$ 
Choices:  $\mathcal{C}$ 
Answer:  $a$ 
```

In the above prompt, the LLM uses the question  $q$  and choices  $\mathcal{C}$  as input and is asked to generate the letter of the answer  $a$ . Suppose we have 5 few-shot examples, with questions  $q_1, \dots, q_5$ , list of choices  $\mathcal{C}_1, \dots, \mathcal{C}_5$ , and ground truth answers  $a_1, \dots, a_5$ . The expanded few-shot prompt for the prompt box is written below:

```
Prompt 2.1: Full Prompt Expanded
Question:  $q_1$ 
Choices:  $\mathcal{C}_1$ 
Answer:  $a_1$ 

Question:  $q_2$ 
Choices:  $\mathcal{C}_2$ 
Answer:  $a_2$ 

Question:  $q_3$ 
Choices:  $\mathcal{C}_3$ 
Answer:  $a_3$ 

Question:  $q_4$ 
Choices:  $\mathcal{C}_4$ 
Answer:  $a_4$ 

Question:  $q_5$ 
Choices:  $\mathcal{C}_5$ 
Answer:  $a_5$ 

Question:  $q$ 
Choices:  $\mathcal{C}$ 
Answer:
```

Using this prompt, the LLM must generate  $a$ , which is the highlighted text in the prompt box.

### A.3 Prompting Details

We design few-shot prompts following the format described by our prompt boxes. The few-shot examples were randomly selected from the training set, and we ensured that these contained a balanced distribution of output labels and that the demonstrations were shuffled. We created a few-shot prompt for each dataset. Both the UnifiedQA evaluation set and the contrast set used the exact same prompt. Even though this results in demonstrations with more than two choices, we found that this did not confuse models on the contrast set, as they never outputted an invalid letter (i.e. “(C)” when there are two choices). In the case of an invalid output, which stemmed from a non-letter choice, we marked the output as incorrect.

## B Results

### B.1 Qualitative Analysis Details

Below, we provide the exact instructions (Figure 4) and annotation interface (Figure 5) shown to our annotators. Our annotation interface is based on PrairieLearn ([West et al., 2015](#)). Our use of plausibility and relevance for this annotation task is based on existing work ([Gierl et al., 2017](#)).

The random baseline we compare against is the trivial solution described in §2.1. This baseline selects a random gold answer from the same dataset to form a set of choices. We apply this algorithm

to the same 50 sampled instances as the ones annotators evaluated with our contrast set, meaning that the questions produced by this baseline only differ by the chosen distractor; the question, choices, and gold answer are all consistent across approaches.

## **B.2 3-shot Prompting Results**

In Figure 6, we show the same results as Figure 3 but with three-shot prompting. The same trends of high choices-only accuracy and the consistency of full prompt rankings across evaluation and contrast sets both hold, with a Kendall’s  $\tau$  of 0.88. We did not test 0-shot prompting as we were working with base LLMs (i.e. unaligned and not instruction-tuned), which should not have the capability to complete tasks in a 0-shot manner. We believe that studying choices-only accuracy in 0-shot settings could be an interesting avenue for future work.

## **B.3 UnifiedQA Evaluation Subset**

Our mined contrast set only has two choices for every question, while the original evaluation set has questions ranging from 2 to 8 choices. To ensure the consistency of rankings is not confounded by the the reduction of possible choices, we also report the 10-shot accuracy on the subset of UnifiedQA that was used to derive the contrast set. This subset is essentially equivalent to the contrast set, but with additional choices on each question so that the number of choices are consistent. In Figure 7, the UnifiedQA Evaluation set and the UnifiedQA Evaluation subset have a similarly high consistency between rankings of full prompt accuracy. Thus, in our experiments, the number of choices on each question does not seem to largely influence the ranking of LLMs.

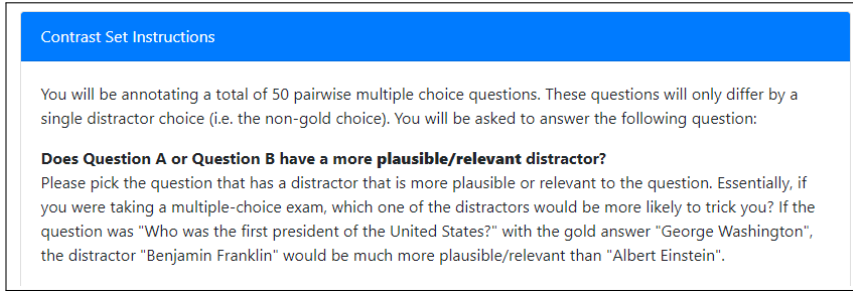


Figure 4: Instructions shown to annotators.

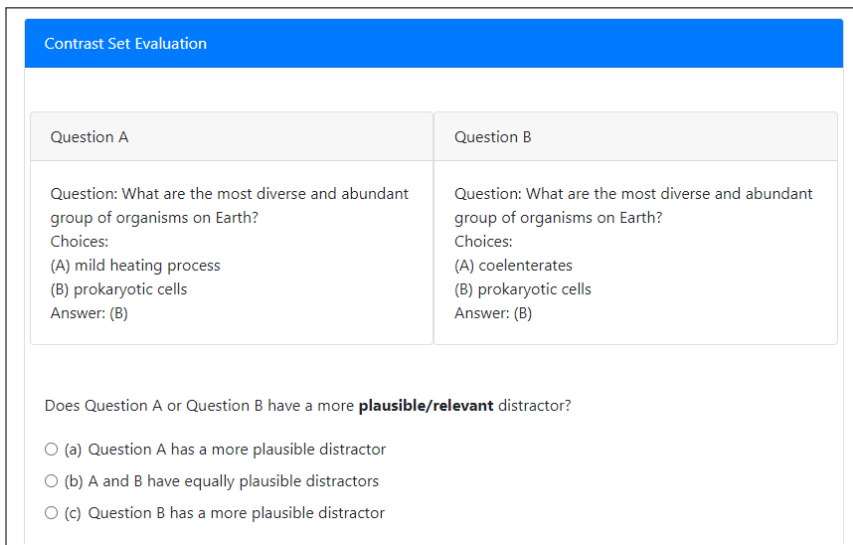


Figure 5: Interface used by annotators.




Figure 6: 3-shot benchmarking of 12 LLMs on the UnifiedQA evaluation set and the contrast set, sorted by full-prompt accuracy. The same trends found for 5-shot and 10-shot prompting hold for 3-shot prompting.

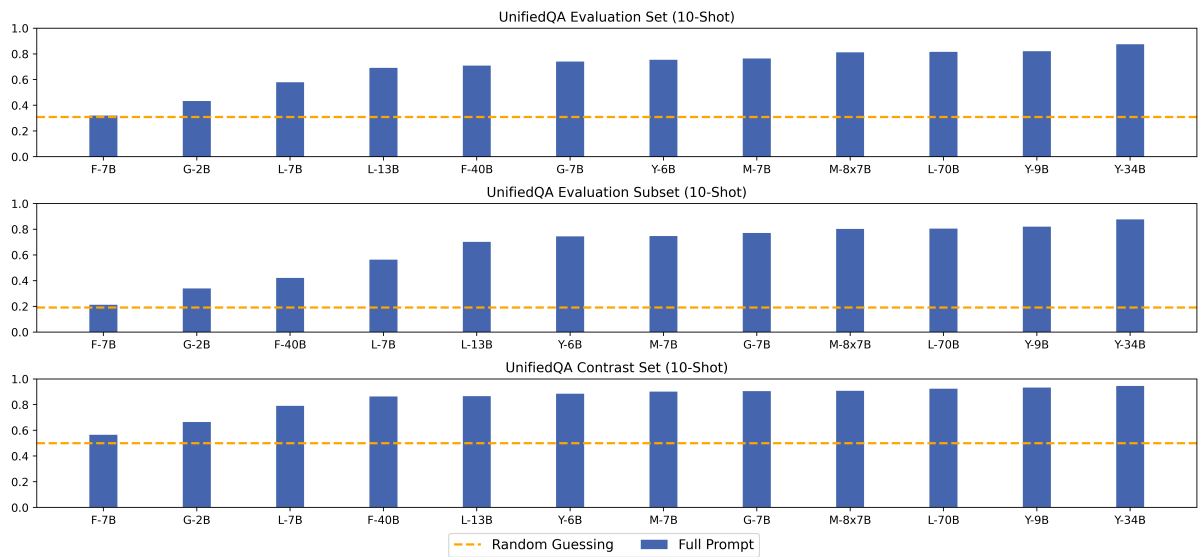


Figure 7: 10-shot benchmarking of 12 LLMs on the UnifiedQA evaluation set, the contrast set, and the subset of the full UnifiedQA evaluation split with the same questions as the contrast set.