
MMLU-Pro+: Evaluating Higher-Order Reasoning and Shortcut Learning in LLMs

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Existing benchmarks for large language models (LLMs) increasingly struggle
2 to differentiate between top-performing models, underscoring the need for more
3 challenging evaluation frameworks. We introduce MMLU-Pro+, an enhanced
4 benchmark building upon MMLU-Pro to assess shortcut learning and higher-order
5 reasoning in LLMs. By incorporating questions with multiple correct answers
6 across diverse domains, MMLU-Pro+ tests LLMs' ability to engage in complex
7 reasoning and resist simplistic problem-solving strategies. Our results show that
8 MMLU-Pro+ maintains MMLU-Pro's difficulty while providing a more rigorous
9 test of model discrimination, particularly in multi-correct answer scenarios. We
10 introduce novel metrics like shortcut selection ratio and correct pair identification
11 ratio, offering deeper insights into model behavior and anchoring bias. Evaluations
12 of six state-of-the-art LLMs reveal significant performance gaps, highlighting
13 variations in reasoning abilities and bias susceptibility. We release the dataset and
14 evaluation codes at <https://github.com/>

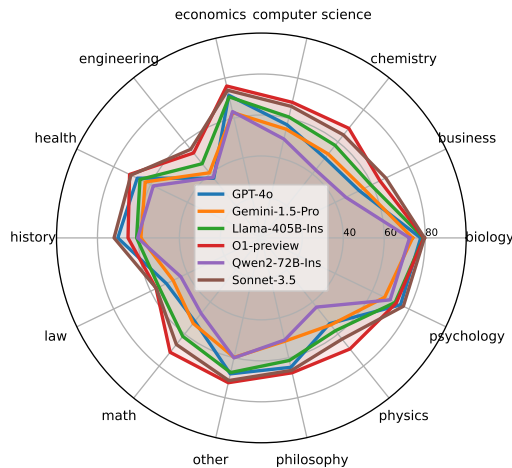


Figure 1: LLMs performance on our new MMLU-Pro+.

15 1 Introduction

16 Recent advancements in large language models (LLMs) have led to remarkable improvements in
17 various natural language processing tasks [7, 15]. State-of-the-art models such as GPT-4o, O1-

18 preview, [1], Claude-3.5 Sonnet [2], Gemini [17], and Llama [20] have demonstrated impressive
19 capabilities across a wide range of applications. However, as these models continue to evolve,
20 existing benchmarks are struggling to keep pace, often reaching performance saturation and failing to
21 effectively differentiate between model capabilities [10].

22 LLMs have been evaluated using a variety of benchmark datasets designed to test different aspects of
23 language understanding and generation. Some of the most prominent benchmarks include IFEval [26]
24 for instruction following, BBH (Big-Bench Hard) [19] for challenging reasoning tasks, MATH [25] for
25 mathematical problem-solving, GPQA [13] for general-purpose question answering, and MUSR [27]
26 for multi-task language understanding. These join other established benchmarks like the General
27 Language Understanding Evaluation (GLUE) [22] and its successor SuperGLUE [21], as well as the
28 Stanford Question Answering Dataset (SQuAD) [16] for reading comprehension.

29 Among these benchmarks, the Massive Multitask Language Understanding (MMLU) benchmark
30 [12] has been widely adopted as a standard for evaluating LLMs due to its broad coverage of
31 subjects. However, recent studies have shown that top-performing models are achieving near-identical
32 scores on MMLU, with several models scoring between 86-87% accuracy [8]. This saturation raises
33 concerns about the benchmark’s ability to measure future advancements in LLM capabilities [18].
34 In response to these limitations, researchers developed MMLU-Pro [23], an iteration of the original
35 MMLU designed to challenge LLMs with more complex, reasoning-focused questions and a greater
36 number of answer options. While MMLU-Pro made significant strides, we identified further areas for
37 enhancement that could improve the benchmark’s ability to evaluate LLMs more effectively.

38 Recent research has highlighted the challenges of shortcut learning, where models exploit superficial
39 patterns rather than developing deeper understanding[3], and the importance of evaluating higher-
40 order reasoning in language models. Geirhos et al. [11] provide a comprehensive overview of shortcut
41 learning in deep neural networks, emphasizing its prevalence and impact on model performance.
42 Wei et al. [24] demonstrate how chain-of-thought prompting can elicit more sophisticated reasoning,
43 addressing some of the limitations that lead to shortcut learning. The need for more nuanced evaluation
44 methods has been underscored by Bowman and Dahl [6], who argue for fixing benchmarking in
45 natural language understanding. Higher-order reasoning involves complex cognitive processes such
46 as analysis, synthesis, and evaluation of multiple pieces of information [4]. It requires models to
47 go beyond simple pattern matching or recall, engaging in more sophisticated thought processes that
48 mimic human-like reasoning more closely. [14].

49 In this paper, we introduce MMLU-Pro+, a new benchmark that builds upon MMLU-Pro by incor-
50 porating these insights. The novelty of MMLU-Pro+ lies in its fundamental change to the nature of
51 reasoning required from LLMs. By introducing questions with multiple correct answers, MMLU-Pro+
52 requires models to: a) Evaluate the validity of multiple statements independently; b) Recognize
53 the potential for more than one correct answer; c) Discern subtle differences between correct and
54 incorrect information; d) Resist the tendency to anchor on a single answer.

55 This approach increases the complexity of the benchmark, forcing models to engage in higher-order
56 reasoning, recognizing and evaluating nuanced or multi-faceted concepts rather than relying on
57 memorized patterns or simplistic heuristics. It tests a model’s robustness and allows for better
58 discrimination between models with varying levels of understanding and reasoning capabilities.

59 MMLU-Pro+ contributes to the field in several ways:

- 60 1. It specifically targets the reduction of shortcut learning by introducing questions with
61 multiple correct answers.
- 62 2. It provides a more realistic evaluation scenario that mirrors real-world complexity, where
63 problems often have multiple valid solutions.
- 64 3. It introduces new metrics such as the shortcut selection ratio and correct pair identification
65 ratio, offering a more nuanced understanding of model performance beyond simple accuracy.

66 By addressing these aspects, MMLU-Pro+ serves as a reliable and informative tool for tracking
67 progress in language understanding. It contributes to the ongoing efforts, highlighted by Bommasani
68 et al. [5], to better understand and evaluate the capabilities and limitations of foundation models
69 in AI, specifically targeting the reduction of shortcut learning and the promotion of higher-order
70 reasoning skills in LLMs.

71 2 Dataset Construction

72 The construction of MMLU-Pro+ involves a systematic and scalable approach to modifying the
73 original MMLU-Pro dataset, introducing multiple correct answers and various types of distractors to
74 enhance its ability to evaluate higher-order reasoning skills in LLMs.

75 We begin with the MMLU-Pro dataset [23], which encompasses questions from 14 diverse domains
76 including mathematics, physics, chemistry, law, engineering, psychology, and health. The initial
77 dataset contains over 12,000 questions, each with up to ten answer options.

78 2.1 Dataset Modification Process

79 We modify the MMLU-Pro dataset in three distinct categories:

80 **True Positive Pairs.** For these questions, we introduce a "Both X and Y are correct" option, where
81 X is the original correct answer from MMLU-Pro, and Y is a new correct option generated using
82 GPT-4o. The process for generating Y varies depending on the question type: a) For mathematical
83 questions, we prompt the LLM to rewrite numbers or equations in alternative formats. b) For other
84 types of questions, we instruct the LLM to find another correct option not already mentioned in the
85 original choices, or to present the same correct information in a different, more complex way beyond
86 simple paraphrasing. This process can be represented as:

$$Q_{TPP} = f_{LLM}(Q_{original}, LLM) \quad (1)$$

87 where Q_{TPP} is the modified question with True Positive Pairs, f_{LLM} is the LLM-based modification
88 function, and LLM refers to GPT-4o.

89 **Partial False Positive Pairs.** For these questions, we create a "Both X and Y are correct" option
90 where X is the original correct answer from MMLU-Pro, and Y is a randomly selected incorrect
91 option from the original set of options.

92 **Complete False Positive Pairs.** For these questions, we create a "Both X and Y are correct" option
93 where X and Y are two randomly selected incorrect options from the original set of options.

94 This composition allows for a comprehensive evaluation that tests more complex multi-answer
95 reasoning skills, while also challenging LLMs to identify partially or completely incorrect option
96 pairs. From the total of 12,032 questions in the dataset, 3,718 were modified using an LLM to
97 create True Positive Pairs, while 4,153 were modified without LLM intervention. Specifically, 2,029
98 questions were modified with two wrong options (Complete False Positive Pairs), and 2,124 were
99 modified with one correct and one wrong option (Partial False Positive Pairs). This ensures a robust
100 evaluation across various question types and modification strategies. Incorporating True Positive Pairs
101 tests the models ability to recognize multiple correct answers, reflecting real-world scenarios where
102 different solutions can be equally valid. Meanwhile, the Partial and Complete False Positive Pairs
103 test the models' ability to discern subtle inaccuracies and resist the tendency to assume correctness
104 when presented with familiar information. This approach not only assesses a model's knowledge
105 but also its capacity for nuanced reasoning and its robustness against potential shortcuts or biases in
106 answering multiple-choice questions.

107 2.2 Post-Processing and Quality Assurance

108 To ensure the integrity of MMLU-Pro+, we implemented a rigorous post-processing and validation
109 protocol:

110 *Human Auditing.* We conducted a comprehensive audit of 100 samples from each group (True Positive,
111 Partial False Positive, and Complete False Positive Pairs) verifying the accuracy and appropriateness
112 of new options.

113 *Consistency Checks.* We performed thorough checks across the entire dataset to ensure newly added
114 options maintain the same style as original ones, preventing unintended evaluation biases.

115 *Error Identification.* We systematically identified and flagged potential inconsistencies or errors
116 introduced during the modification process.

117 *Task Differentiation.* We ensured that the process of creating true positive pairs differs fundamentally
118 from the task of answering questions, minimizing the risk of model-specific advantages.

119 *Comprehensive Metrics.* Our evaluation metrics assess not only accuracy but also bias and shortcut
120 learning across diverse models, providing a holistic view of model performance.

121 While GPT-4o was used for creating True Positive pairs, our evaluation process is designed to be
122 model-agnostic. This approach ensures that the augmentation process genuinely increases question
123 difficulty rather than introducing biases favoring any particular model. These measures collectively
124 maintain MMLU-Pro+'s ability to differentiate between model capabilities, regardless of which LLM
125 was involved in the augmentation process. This is further validated in our experiments, where GPT-4o
126 is not the top-performing model, demonstrating the benchmark's independence from its creation
127 process.

128 **3 Experiments**

129 We evaluated several state-of-the-art LLMs, including two open-source models (LLaMA-3.1-405B-
130 Instruct [20] and Qwen-2-72B-Instruct [9]) and four closed-source models (OpenAI's GPT-4o [1] and
131 O1-preview, Claude-Sonnet-3.5 [2], and Gemini-1.5-Pro [17]), using the MMLU-Pro+ benchmark.
132 Models were assessed not only on accuracy but also on their ability to recognize and correctly select
133 the "Both X and Y are correct" option and avoid partially correct options, thereby demonstrating
134 higher-order reasoning and shortcut learning.

135 **3.1 Accuracy Analysis on MMLU-Pro+**

136 Figure 1 illustrates the overall performance of each model on the MMLU-Pro+ dataset, while Table 1
137 provides a detailed breakdown of performance across individual subject categories, including the
138 comparative drop from MMLU-Pro. These results offer insights into the models' capabilities and the
139 increased challenge presented by MMLU-Pro+.

140 O1-preview demonstrates superior performance across the majority of categories, achieving the
141 highest average accuracy. The consistent outperformance suggests a more robust capability in
142 handling the increased complexity introduced by MMLU-Pro+. The universal decrease in accuracy
143 from MMLU-Pro to MMLU-Pro+ validates the increased challenge of our new benchmark. Notably,
144 O1-preview exhibits the smallest average performance drop, indicating enhanced resilience to the
145 modified question format and the introduction of multiple correct answers.

146 The engineering category reveals an interesting divergence, with Gemini-1.5-Pro and Sonnet-3.5
147 showing markedly smaller performance drops compared to other models. Conversely, the law
148 category presents an anomaly, with minimal performance drops and even slight improvements for
149 some models, suggesting that the MMLU-Pro+ modifications may have had less impact on this
150 domain. GPT-4o exhibits the largest average performance drop, indicating a potential sensitivity to
151 the structural changes in MMLU-Pro+. This suggests that high performance on standard benchmarks
152 may not necessarily translate to robust reasoning capabilities in more complex scenarios.

153 The performance gap between different models on MMLU-Pro, particularly the superior performance
154 of models like O1-preview and Claude-3.5-Sonnet compared to GPT-4o (which was used in dataset
155 creation), further validates our dataset construction methodology. Despite GPT-4o's involvement
156 in generating True Positive pairs, it does not exhibit an advantage in the question-answering task.
157 This discrepancy in performance across models suggests that MMLU-Pro successfully challenges
158 the LLMs, requiring genuine reasoning capabilities rather than pattern matching or exploitation of
159 dataset artifacts.

160 We also measured models' performance on (1) questions with two correct options (`both_correct`),
161 (2) questions with one correct and one incorrect option (`correct_and_wrong`), and (3) questions
162 with two incorrect options (`two_wrong`). The results, as illustrated in Figure 2, reveal significant
163 variations in model performance across these question types. Interestingly, all models showed lower
164 accuracy in the `both_correct` category compared to the other two, suggesting a potential difficulty
165 in identifying multiple correct answers. GPT-4o and Gemini-1.5-Pro showed similar performance
166 patterns, with their highest accuracies in the `correct_and_wrong` category. These findings highlight
167 the varying capabilities of different models in handling nuanced multiple-choice questions

Table 1: Accuracy (%) on MMLU-Pro+ Categories with Performance Drop from MMLU-Pro

Category	Qwen2-72B-Ins	Gemini-1.5-Pro	GPT-4o	Llama-405B-Ins	Sonnet-3.5	O1-preview
biology	72.3 ^{-9.7}	73.8 ^{-9.9}	77.8 ^{-10.4}	79.2 ^{-6.0}	79.6 ^{-7.8}	79.3 ^{-9.9}
business	45.8 ^{-22.2}	55.5 ^{-16.6}	53.1 ^{-26.5}	59.4 ^{-17.6}	67.4 ^{-12.2}	64.3 ^{-23.7}
chemistry	42.5 ^{-16.3}	52.4 ^{-17.7}	49.9 ^{-25.7}	57.8 ^{-15.1}	64.3 ^{-12.3}	68.5 ^{-17.2}
computer science	49.5 ^{-18.5}	54.4 ^{-14.5}	56.6 ^{-23.6}	60.5 ^{-13.9}	65.9 ^{-14.9}	67.9 ^{-67.9}
economics	63.3 ^{-13.3}	62.8 ^{-13.9}	71.4 ^{-11.1}	70.6 ^{-9.8}	73.9 ^{-8.5}	76.1 ^{-9.6}
engineering	37.8 ^{-9.8}	40.7 ^{-3.9}	37.3 ^{-18.5}	46.3 ^{-13.5}	55.3 ^{-4.2}	53.0 ^{-15.6}
health	58.6 ^{-8.4}	63.0 ^{-4.4}	67.2 ^{-8.1}	65.8 ^{-6.5}	70.7 ^{-6.1}	71.4 ^{-7.7}
history	60.4 ^{-6.0}	58.8 ^{-7.3}	70.1 ^{-2.4}	60.9 ^{-6.6}	71.9 ^{-1.3}	65.1 ^{-9.4}
law	43.6 ^{-0.7}	47.8 ^{-0.5}	51.4 ^{-3.3}	55.4 ^{-1.2}	56.9 ^{-7.3}	57.0 ^{-11.5}
math	47.2 ^{-23.4}	53.4 ^{-8.8}	52.1 ^{-25.9}	61.5 ^{-15.7}	66.5 ^{-9.8}	71.4 ^{-18.8}
other	60.1 ^{-6.0}	59.8 ^{-10.3}	68.1 ^{-9.8}	67.4 ^{-5.7}	71.4 ^{-6.7}	72.5 ^{-8.5}
philosophy	51.1 ^{-8.2}	51.7 ^{-11.3}	64.8 ^{-6.8}	61.4 ^{-4.8}	66.3 ^{-8.2}	67.5 ^{-12.0}
physics	43.2 ^{-18.2}	54.3 ^{-14.9}	53.5 ^{-21.6}	58.0 ^{-14.2}	63.2 ^{-13.4}	69.4 ^{-17.7}
psychology	69.8 ^{-6.4}	66.8 ^{-9.7}	75.4 ^{-5.9}	72.5 ^{-4.8}	76.9 ^{-5.6}	73.1 ^{-11.8}
Average	53.2 ^{-11.9}	56.8 ^{-10.2}	60.6 ^{-14.3}	62.6 ^{-9.5}	67.9 ^{-8.5}	68.3 ^{-7.5}




Figure 2: Accuracy on the three modified groups of questions. The amount of drop w.r.t original MMLU-Pro is written on the bars.

168 3.2 Analysis of Anchoring Bias and Shortcut Learning in MMLU-Pro+

169 Figure 3 illustrates the propensity of various language models to maintain their original choices when
 170 presented with modified questions in MMLU-Pro+, specifically for True Positive Pairs. To quantify
 171 this behavior, we introduce the Shortcut Selection Ratio (SSR), defined as follows:

$$SSR_{\text{wrong}} = \frac{N_{\text{stayed_wrong}}}{N_{\text{total_TPP}}} \quad (2)$$

$$SSR_{\text{partial}} = \frac{N_{\text{stayed_partial}}}{N_{\text{total_TPP}}} \quad (3)$$

172 Where $N_{\text{stayed_wrong}}$ is the number of times the model stayed on a previously chosen wrong answer,
 173 $N_{\text{stayed_partial}}$ is the number of times the model stayed only on the previously correct answer without
 174 acknowledging the newly introduced correct option, and $N_{\text{total_TPP}}$ is the total number of True Positive
 175 Pairs.

176 This “shortcut selection ratio” provides insights into potential anchoring bias and shortcut learning
 177 behaviors. The graph reveals that all models exhibit a tendency to stick with their initial selections,
 178 both for previously wrong and partially correct options, suggesting a degree of anchoring bias. This
 179 behavior is particularly pronounced in GPT-4o and Qwen2-72B-Ins, which show higher rates of
 180 maintaining their original choices.

181 The persistence in selecting previously incorrect options (high SSR_{wrong}) is especially noteworthy,
 182 as it indicates potential limitations in these models’ ability to reassess and engage in higher-order

183 reasoning when presented with new, valid alternatives. Similarly, a high SSR_{partial} suggests a failure
 184 to recognize newly introduced correct options. Conversely, Gemini 1.5 Pro, Sonnet-3.5, and Llama
 185 1.3 405B demonstrate lower shortcut selection ratios, suggesting a greater capacity for adapting
 186 their reasoning in light of new information. Interestingly, O1-preview has improved compared to
 187 GPT-4o, but still lags behind Gemini-1.5-Pro and Sonnet-3.5. These findings highlight the challenges
 188 language models face in fully leveraging the additional correct options introduced in MMLU-Pro+,
 189 and underscore the importance of developing benchmarks that can effectively evaluate and promote
 190 higher-order reasoning.

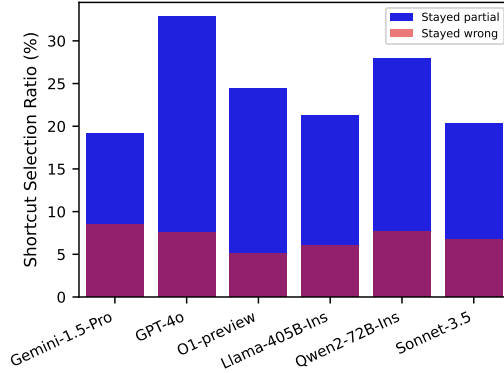


Figure 3: Shortcut Selection Ratio for True Positive Pairs in MMLU-Pro+

191 3.3 Analysis of Correct Pair Identification in MMLU-Pro+

192 In this experiment, we evaluate the models’ ability to accurately identify true positive pairs among
 193 various types of answer combinations in MMLU-Pro+, providing insights into their reasoning
 194 capabilities and resilience to distractors.

195 Figure 4 presents an error analysis for various models on MMLU-Pro+, introducing a metric called
 196 the correct pair identification ratio. This ratio is defined as:

$$\text{Correct Pair Identification (CPI) Ratio} = \frac{N_{\text{TPP}}}{N_{\text{PFPP}} + N_{\text{CFPP}}} \quad (4)$$

197 where N_{TPP} represents the number of correctly identified True Positive Pairs, N_{PFPP} is the number
 198 of times the model incorrectly predicted the Partial False Positive Pair, and N_{CFPP} is the number
 199 of times the model incorrectly predicted the Complete False Positive Pair. This ratio measures the
 200 model’s ability to identify correct pairs relative to its tendency to be misled by partially or completely
 201 incorrect pairs.

202 Sonnet-3.5 achieves the highest ratio (10.26), demonstrating superior discrimination capability in
 203 distinguishing correct answer pairs from misleading options. This suggests enhanced resistance to
 204 distractors and a more robust grasp of subject matter and question structure. The significant variation
 205 in ratios, ranging from 2.80 (Llama-405B-Ins) to 10.26 (Sonnet-3.5), reveals substantial differences
 206 in models’ higher-order reasoning capabilities. Models with higher ratios exhibit a greater capacity
 207 for nuanced understanding, correctly identifying multiple true statements while rejecting plausible
 208 but incorrect combinations. These findings highlight MMLU-Pro+’s effectiveness in differentiating
 209 models based on their ability to handle complex, multi-correct answer scenarios, underscoring the
 210 importance of sophisticated evaluation metrics in assessing advanced language models. Some samples
 211 of the True Positive Pairs from the MMLU-Pro+ dataset can be seen in Figure 5.

212 4 Interpretation and Significance of Novel Metrics

213 The Shortcut Selection Ratio (SSR) and Correct Pair Identification (CPI) Ratio provide insights into
 214 model behavior that are directly relevant to real-world applications:




Figure 4: Error Analysis: Correct Pair Identification (CPI) in MMLU-Pro+. The numbers on the bars represent the CPI ratio values. A higher CPI ratio indicates better performance in distinguishing correct answer pairs from incorrect ones.

Question (Math): What number multiplied by 4 equals 36?
Options:
[A] 11, **[B]** 10, **[C]** 12, **[D]** 6, **[E]** 15, **[F]** 9, **[G]** 7, **[H]** 8, **[I]** 13, **[J]** 14, **[K]** 36/4, **[L]** Both 9 and 36/4 are correct.
Predictions: Sonnet-3.5: **L**, O1-preview: **L**, Gemini-1.5-Pro: **F**, GPT-4o: **F**, Llama-405B-Ins: **F**, Qwen2-72B-Ins: **F**

Question (Computer Science): The disadvantage of Grid search is
Options:
[A] It cannot handle categorical variables., **[B]** It cannot be applied to non-differentiable functions., **[C]** It is not suitable for datasets with a large number of features., **[D]** It cannot handle missing data., **[E]** It cannot be used with neural networks., **[F]** It runs reasonably slow for multiple linear regression., **[G]** It can only be used for binary classification problems., **[H]** It is hard to implement., **[I]** It cannot be applied to non-continuous functions., **[J]** It is computationally expensive for large datasets., **[K]** Both It runs reasonably slow for multiple linear regression and It is computationally expensive for large datasets are correct.
Predictions: Sonnet-3.5: **K**, O1-preview: **J**, Gemini-1.5-Pro: **J**, GPT-4o: **J**, Llama-405B-Ins: **J**, Qwen2-72B-Ins: **J**

Figure 5: True Positive Pair Samples from math and computer science categories with model predictions.

215 SSR: A low SSR indicates a model’s ability to adapt its reasoning when presented with new, valid
 216 information. This is crucial in dynamic decision-making environments, such as medical diagnosis
 217 or financial analysis, where new data may necessitate re-evaluation of initial conclusions. CPI
 218 Ratio: A high CPI Ratio suggests a model’s proficiency in distinguishing between subtly different
 219 correct and incorrect information combinations. This skill is essential in fields like legal analysis,
 220 scientific research, or policy-making, where the ability to discern fine-grained differences in complex
 221 information is paramount.

222 These metrics go beyond simple accuracy, offering a more nuanced understanding of a model’s
 223 reasoning processes and its potential performance in complex, real-world tasks.

224 5 Discussion

225 In this paper, we introduced MMLU-Pro+, an enhanced benchmark designed to evaluate the higher-
 226 order reasoning capabilities of large language models. By incorporating questions with multiple
 227 correct answers and various types of distractors, MMLU-Pro+ provides a more challenging and
 228 discriminative evaluation framework than its predecessors.

229 Our experimental results demonstrate the effectiveness of MMLU-Pro+ in several key areas:

230 *Increased Difficulty.* All evaluated models showed a consistent drop in performance when moving
 231 from MMLU-Pro to MMLU-Pro+, confirming the increased challenge of our benchmark.

232 *Model Differentiation.* MMLU-Pro+ revealed substantial differences in model performance. While
 233 O1-preview outperformed other models across most categories in general assessments, its dominance
 234 diminishes when anchoring bias and higher-order reasoning are considered.

235 *Anchoring Bias and Shortcut Learning.* The shortcut selection ratio analysis exposed varying degrees
 236 of anchoring bias across models, highlighting the challenge LLMs face in adapting their reasoning
 237 when presented with new, valid alternatives. Notably, Gemini-1.5-Pro and Sonnet-3.5 demonstrated
 238 less reliance on shortcuts.

239 *Higher-Order Reasoning.* The correct pair identification ratio provided insights into models’ abilities
 240 to distinguish genuinely correct answer pairs from misleading options, with significant variations
 241 observed across different models. Sonnet-3.5 significantly outperformed others in this aspect.

242 These findings underscore the importance of new evaluation metrics in assessing advanced language
 243 models, particularly in scenarios requiring discernment between subtly different correct and incorrect
 244 information combinations.

245 MMLU-Pro+ not only serves as a more reliable and informative benchmark for tracking progress in
 246 LLM evaluation but also highlights areas for improvement in current models. The observed anchoring
 247 bias and varying abilities to identify correct pairs suggest that even top-performing models may still
 248 rely on simplistic heuristics or struggle with truly nuanced reasoning in complex scenarios.

249 Future work could explore the development of training techniques that specifically target the higher-
 250 order reasoning skills evaluated by MMLU-Pro+. Additionally, extending this benchmark approach
 251 to other domains or task types could provide a more comprehensive evaluation of LLM capabilities
 252 across diverse applications. While we used GPT-4o in our dataset construction process, our results
 253 demonstrate that this does not confer an unfair advantage to these or similar models. The significant
 254 performance drop of GPT-4o on the modified data, coupled with the superior performance of other
 255 models like Claude-3.5-Sonnet, indicates that our methodology produces a dataset that genuinely
 256 challenges LLMs. However, future work could explore alternative methods for dataset augmentation
 257 to further mitigate any potential biases introduced by LLM-assisted generation.

258 By providing a more challenging and discriminative benchmark, MMLU-Pro+ contributes to the
 259 ongoing effort to develop more capable and robust language models with human-like reasoning and
 260 understanding.

261 **References**

- 262 [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni
 263 Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4
 264 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- 265 [2] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones,
 266 Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. Constitutional ai:
 267 Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*, 2022.
- 268 [3] Emily M Bender et al. On the dangers of stochastic parrots: Can language models be too big?
 269 *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages
 270 610–623, 2021.
- 271 [4] Benjamin S Bloom et al. *Taxonomy of educational objectives. Vol. 1: Cognitive domain*. New
 272 York: McKay, 1956.
- 273 [5] Rishi et al. Bommasani. On the opportunities and risks of foundation models. *arXiv preprint*
 274 *arXiv:2108.07258*, 2021.
- 275 [6] Samuel R Bowman and George E Dahl. What will it take to fix benchmarking in natural
 276 language understanding? *arXiv preprint arXiv:2104.02145*, 2021.
- 277 [7] Tom Brown et al. Language models are few-shot learners. *Advances in neural information*
 278 *processing systems*, 33:1877–1901, 2020.

- 279 [8] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li,
280 Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E Gonzalez, et al. Chatbot arena:
281 An open platform for evaluating llms by human preference. *arXiv preprint arXiv:2403.04132*,
282 2024.
- 283 [9] Jinze Bai et al. Qwen technical report, 2023.
- 284 [10] Deep Ganguli et al. Predictability and surprise in large generative models. *arXiv preprint*
285 *arXiv:2202.07785*, 2022.
- 286 [11] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel,
287 Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature*
288 *Machine Intelligence*, 2(11):665–673, 2020.
- 289 [12] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
290 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint*
291 *arXiv:2009.03300*, 2020.
- 292 [13] Zizhun Li, Zitong Yu, Qi Yu, and Zhijie Chen. Fracgpt: Reasoning and hybrid recovery of
293 fraction arithmetic. *arXiv preprint arXiv:2307.14809*, 2023.
- 294 [14] Melanie Mitchell. Abstraction and analogy-making in artificial intelligence. *Annals of the New*
295 *York Academy of Sciences*, 1505(1):79–101, 2021.
- 296 [15] Colin Raffel et al. Exploring the limits of transfer learning with a unified text-to-text transformer.
297 *Journal of Machine Learning Research*, 21:1–67, 2020.
- 298 [16] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions
299 for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical*
300 *Methods in Natural Language Processing*, 2016.
- 301 [17] Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-
302 baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al.
303 Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv*
304 *preprint arXiv:2403.05530*, 2024.
- 305 [18] Aideen Rodriguez et al. Open-ended questions in the wild: A case study of large language
306 model evaluation. *arXiv preprint arXiv:2306.15757*, 2023.
- 307 [19] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Mohammad Shoeybi, et al. Beyond the
308 imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv*
309 *preprint arXiv:2206.04615*, 2023.
- 310 [20] Hugo et al. Touvron. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint*
311 *arXiv:2307.09288*, 2023.
- 312 [21] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill,
313 Omer Levy, and Samuel R. Bowman. Superglue: A stickier benchmark for general-purpose
314 language understanding systems. In *Advances in Neural Information Processing Systems*, 2019.
- 315 [22] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman.
316 Glue: A multi-task benchmark and analysis platform for natural language understanding. In
317 *International Conference on Learning Representations*, 2018.
- 318 [23] Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo,
319 Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex
320 Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. Mmlu-pro: A more robust and challenging
321 multi-task language understanding benchmark. *arXiv preprint arXiv:2406.01574*, 2024.
- 322 [24] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le,
323 Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models.
324 *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- 325 [25] Jonas Wiedenhoff, Jie Chen, Nafise Sadat Moosavi, and Iryna Gurevych. Solving math word
326 problems with process-and outcome-based feedback. *arXiv preprint arXiv:2308.07673*, 2023.
- 327 [26] Zihao Yue, Tianyu Lin, Yuxiao Shen, Jialu Yang, Yujie Xu, Zhihong Cheng, and Dawn
328 Song. Ifeval: Instruction following evaluation for large language models. *arXiv preprint*
329 *arXiv:2310.00567*, 2023.
- 330 [27] BOWEI ZOU and MOHIT BANSAL. Musr: Multi-task learning for ultra-fine entity typing and
331 semantic role labeling. *arXiv preprint arXiv:2305.18245*, 2023.