Relevant or Random: Can LLMs Truly Perform Analogical Reasoning?

Anonymous ACL submission

Abstract

Analogical reasoning is a unique ability of humans to address unfamiliar challenges by transferring strategies from relevant past experiences. One key finding in psychology is that 004 compared with irrelevant past experiences, recalling relevant ones can help humans better 007 handle new tasks (Ramachandran, 2012). Coincidentally, the NLP community has also recently found that self-generating relevant examples in the context can help large language models (LLMs) better solve a given problem than 012 hand-crafted prompts (Yasunaga et al., 2024). However, it is yet not clear whether relevance is the key factor eliciting such capability, i.e., can LLMs benefit more from self-generated relevant examples than irrelevant ones? In this work, we systematically explore whether LLMs 017 can truly perform analogical reasoning on a diverse set of reasoning tasks. With extensive experiments and analysis, we show that selfgenerated random examples can surprisingly achieve comparable or even better performance, e.g., 4% performance boost on GSM8K with random biological examples. We find that the accuracy of self-generated examples is the key factor and subsequently design two improved methods with significantly reduced inference 027 costs. Overall, we aim to advance a deeper understanding of LLM analogical reasoning and hope this work stimulates further research in the design of self-generated contexts.

1 Introduction

A hallmark of human intelligence is that they can solve novel problems by drawing analogy from relevant past experiences, a concept known as *analogical reasoning* in cognitive science (Vosniadou and Ortony, 1989). As indicated by the name, recalling previously acquired *relevant* experiences can facilitate humans to *better* tackle new tasks, whereas irrelevant ones are rarely beneficial and can even be distracting (Ramachandran, 2012). For instance,



Figure 1: Illustration of analogical prompting. LLMs are prompted to self-generate relevant examples as context before solving the new problem.

when faced with a novel math problem about determinants (*e.g.*, calculating the value of a given fourth-order determinant), humans can resolve this by reflecting upon the methodology employed to ascertain the value of a third-order determinant, whereas biological knowledge (*e.g.*, how the human body regulates its temperature) can generally be considered irrelevant.

With the recent advancements in scaling up model size and data, LLMs have demonstrated impressive zero-shot and few-shot performance across various reasoning tasks, especially, through advanced prompting methods like chain-of-thought or CoT (Wei et al., 2022). Compared to common approaches such as zero or few-shot CoT (Zhou et al., 2022; Kojima et al., 2022; Wang et al., 2022b), recent work claims that LLMs are able to selfgenerate examples relevant to the query as context to better solve new problems (Webb et al., 2023; Yasunaga et al., 2024); see Fig. 1 for an example. However, it remains unclear whether relevance is the key to eliciting such capability in LLMs.

In this paper, to systematically assess the capa-

064

042

bility of LLMs to perform analogical reasoning, we conduct a series of ablation experiments on a variety of reasoning tasks including problems from GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), and BIG-Bench Hard (BBH) (Suzgun et al., 2022). With extensive experiments, we aim to address the following two research questions:

065

066

067

071

081

086

087

094

096

098

100

101

103

104

105

106

107

108

109

- **Q1.** Are self-generated *relevant* examples more beneficial to LLMs than *random* ones?
- **Q2.** If not, what is the pivotal factor for LLMs' performance in analogical reasoning?

To answer these questions, we empirically analyze the analogical reasoning abilities of GPT-3.5 (turbo) and Llama 2-Chat (Touvron et al., 2023c). Surprisingly, experimental results show that prompting LLMs to self-generate random examples can achieve comparable or even better performance on certain tasks, indicating that LLMs cannot always perform analogical reasoning. As for Q2, we point out through controlled experiments that the key factor is the accuracy of selfgenerated examples. Informed by these findings, we design two approaches that can outperform existing methods with significantly reduced inference costs. Specifically, we ask LLMs to randomly generate a few problems and manually verify their correctness, then use this fixed set of problems as in-context learning demonstrations for all test samples. Consistent observations across different model types consolidate the conclusions. We summarize the major contributions of our work below:

- To the best of our knowledge, we, for the first time, extensively assess the ability of LLMs to perform analogical reasoning and explore their counterintuitive behavior on certain tasks.
- With extensive experiments and analysis, we demonstrate the effectiveness and limitations of different types of self-generated contexts. Our code base is available at <redacted>.

2 Related Work

This work mainly explores whether LLMs can truly perform analogical reasoning. In light of this, we review three lines of research that form the basis of this work: large language models, chain-of-thought prompting, and analogical prompting.

110 2.1 Large Language Models

Ever since Brown et al. (2020) demonstrated that LLMs can achieve impressive few-shot and zeroshot performance on a variety of tasks, numerous efforts have been made to advance the development of LLMs. One line of research explored the benefits of scaling up LLMs, including Gopher (Rae et al., 2021), LaMDA (Thoppilan et al., 2022), MT-NLG (Smith et al., 2022) and PaLM (Chowdhery et al., 2022). Rather than scaling up model size alone, another line of research focused on training smaller models with better strategies or objectives, such as FLAN (Wei et al., 2021), T0 (Sanh et al., 2021), Chinchilla (Hoffmann et al., 2022), UL2 (Tay et al., 2022) and Flan-T5 (Chung et al., 2022). 113

114

115

116

117

118

119

120

121

122

123

124

125

126

127

128

129

130

131

132

133

135

136

137

138

139

140

141

142

143

144

145

146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

163

The emergence of OpenAI's ChatGPT/GPT-4 has greatly changed the research focus of the NLP community (Qin et al., 2023b; OpenAI, 2023). As reinforcement learning from human feedback (RLHF) contributed significantly to the success of these models (Schulman et al., 2017; Ouyang et al., 2022), researchers have explored various methods to improve RLHF (Yuan et al., 2023; Rafailov et al., 2023; Lee et al., 2023; Azar et al., 2023; Liu et al., 2024). To promote research of open-source LLMs, Meta released Llama series models (Touvron et al., 2023a,b), which serve as the solid foundation for a large number of open-source models (Taori et al., 2023; Chiang et al., 2023; Xu et al., 2023; Zhou et al., 2023). Furthermore, numerous endeavors have been dedicated to training powerful LLMs from scratch, instead of developing LLMs based on Llama (Li et al., 2023; Jiang et al., 2023; Bi et al., 2024; Mesnard et al., 2024).

2.2 Chain-of-Thought Prompting

Chain-of-thought (CoT) prompting induces LLMs to generate intermediate reasoning steps before generating the final answer (Wei et al., 2022), greatly improving the reasoning capabilities of LLMs. Typical CoT prompting approaches include few-shot CoT (Wei et al., 2022; Zhou et al., 2022; Wang et al., 2022b; Li et al., 2022; Wang et al., 2022a), taking several labeled demonstrations of the reasoning process, and zero-shot CoT, comprising only instructions like "Let's think step by step" (Kojima et al., 2022; Zelikman et al., 2022; Zhang et al., 2023). Other ongoing research on CoT has also explored (i) optimizing the demonstration selection (Fu et al., 2022; Lu et al., 2022; Qin et al., 2023a), (ii) optimizing the quality of reasoning chains (Khot et al., 2022; Chen et al., 2022; Zhao et al., 2023; Shinn et al., 2023), and (iii) CoT in smaller language models (Magister et al., 2022; Ho et al., 2022; Fu et al., 2023).



Figure 2: Example prompts for GSM8K (mathematical reasoning). *Top*: The original prompt used in Yasunaga et al. (2024) for self-generating *relevant* math problems. *Bottom*: The prompt designed for self-generating *random* math problems. We mark the differences between these two prompts in blue and green respectively.

2.3 Analogical Prompting

164

181

182

183

185

186

While few-shot CoT can provide more detailed 165 reasoning guidance, it requires labeled examples which can be unavailable for a new task. To tackle this problem, Yasunaga et al. (2024) propose ana-168 logical prompting to guide LLMs to self-generate 169 relevant exemplars as few-shot demonstrations, which is similar to analogical reasoning, *i.e.*, humans can address new problems by drawing anal-172 ogy from relevant past experience (Vosniadou and 173 Ortony, 1989; Holyoak, 2012). In this work, we 174 step forward to explore the intrinsic principle of 175 analogical reasoning in LLMs. Specifically, we 176 aim to investigate whether LLMs can authentically 177 exhibit such reasoning capabilities and determine 178 the extent to which the relevance of self-generated examples contributes to enhancing this process. 180

3 Methodology

We mainly follow the analogical prompting approach outlined in Yasunaga et al. (2024). Specifically, for a given target problem x, analogical prompting introduces instructions like:

Problem: [x]

Relevant problems: Recall five <u>relevant</u> and diverse problems. For each problem, describe it and explain the solution. # Solve the initial problem:

The goal is to induce LLMs to self-generate *relevant* examples, aiding them to solve the target problem via in-context learning. To ensure better performance and efficiency, several key technical decisions are made in Yasunaga et al. (2024): 188

189

190

192

193

194

195

197

199

200

201

202

204

206

207

208

209

210

211

212

213

- The self-generated examples should be relevant and diverse, achieved through a specially designed instruction.
- Generate relevant problems and the solution to the initial problem in one pass.
- 3 to 5 self-generated examples perform the best.

In this work, we make minimal changes to the original prompt to guide LLMs to generate different types of *irrelevant* examples as context; see Fig. 2 and Fig. 3 for example prompts:

- *N/A*: generate problems that are N/A (not applicable) to the initial problem.
- *Random_{same}*: randomly generate examples of the same problem type (*e.g.*, math).
- *Random_{diff}*: randomly generate examples of different problem types (*e.g.*, any type except math).
- *Random_{bio}*: randomly generate biological problems.

Same as the original one, all newly designed prompts allow LLMs to generate irrelevant prob-



Figure 3: Example prompts for BBH (other reasoning). *Top*: The prompt used in Yasunaga et al. (2024) for self-generating *relevant* reasoning problems. *Bottom*: Our designed prompt for self-generating *random* reasoning problems. The differences between these two prompts are indicated in blue and green respectively.

lems and solve the initial target problem in one
pass. In addition, we show prompts for different
methods on all datasets in Appendix A.1.

4 Experiment

217

218

219

221

223

224

227

234

238

4.1 Experimental Setup

We construct the evaluation suite based on diverse reasoning-intensive tasks, including mathematical reasoning and other reasoning (*e.g.*, logical and temporal reasoning) tasks:

- Mathematical reasoning. We work with two commonly used datasets, GSM8K (Cobbe et al., 2021) and MATH (Hendrycks et al., 2021). For each dataset, we randomly sample 500 examples from the original test set and run experiments three times with different random seeds (resulting in different test samples).
- Other reasoning. Following Yasunaga et al. (2024), we evaluate five reasoning tasks in BIG-Bench Hard (BBH) (Suzgun et al., 2022): temporal sequences (temporal reasoning), logical deduction five objects and reasoning about colored objects (logical reasoning), formal fallacies (deductive reasoning) and word sorting (symbolic reasoning). For each task, we use all test samples for evaluation and run experiments three times

with different random seeds.

We mainly use GPT-3.5 (gpt-3.5-turbo) as the LLM and obtain all outputs from it with the temperature set to 0. We ask the LLM to self-generate 5 examples for GSM8K, 3 examples for MATH and BBH following Yasunaga et al. (2024). 239

240

241

242

243

246

247

248

250

251

252

253

254

255

256

257

258

260

4.2 Main Results

We now address the research questions asked in §1 with empirical results.

• Q1. Are self-generated relevant examples more beneficial to LLMs than random ones?

The results averaged over all random seeds are reported in Table 1 and Table 2; more detailed results for every seed are shown in Appendix A.2.

• Self-generated relevant examples achieve the best average performance on BBH. From the results in Table 1, we can observe that the superiority of self-generated relevant examples is empirically substantiated on BBH. Specifically, using relevant examples, denoted by 'relevant', outperforms other approaches on temporal and logical reasoning tasks. While it performs worse than 'N/A' on deductive and symbolical reasoning, it can still improve the accuracy by **1.3**% on average compared

Method	Temporal sequences	Logical deduction five objects	Reasoning about colored objects	Formal fallacies	Word sorting	Average
Relevant	60.0	51.2	76.7	51.2	76.9	63.2
N/A	57.5	45.3	75.5	53.3	77.7	61.9
Random _{same}	53.1	48.8	73.5	52.4	74.1	60.4
Random _{diff}	44.3	44.8	72.4	51.2	69.2	56.4
Random _{bio}	57.1	49.5	76.1	50.8	74.9	61.7

Table 1: Accuracy (%) of different methods on five reasoning tasks in BBH. **Bold** indicates the best results. Selfgenerated *relevant* examples achieve the best average performance. Detailed results for different seeds are reported in Appendix A.2.

Method	Task			
	GSM8K	MATH	Average	
Relevant	71.5	33.3	52.4	
N/A	75.5	36.1	55.8	
Random _{same}	75.1	36.3	55.7	
Random _{diff}	76.3	34.1	55.2	
Random _{bio}	75.3	34.6	54.9	

Table 2: Accuracy (%) of different methods on two mathematical reasoning tasks. Self-generated *irrele-vant* examples are consistently better than *relevant* ones. Table 13 in Appendix A.2 reports detailed results for different seeds.

to 'N/A'.

However, the results on mathematical reasoning tasks are quite counterintuitive as described below:

• Relevant examples do not guarantee better performance. Different from BBH, all types of self-generated irrelevant examples consistently outperform relevant ones on both mathematical reasoning datasets, showing that LLMs cannot yet perform analogical reasoning on these tasks. Interestingly, when we use randomly generated biological examples (*e.g.*, how the process of photosynthesis occurs in plants), they can yield about 2.5% better results on average compared to generating relevant math problems. Besides, 'N/A' achieves the best average result as it is the second-best on both datasets.

Problems in MATH span various subjects and difficulty levels. To investigate whether the inferior performance of relevant examples on MATH is accidentally caused by certain categories, we further report the accuracy across different subjects and difficulty levels in Table 3 and Fig. 4. The consistent performance gap between 'relevant' and other methods across different problem categories demonstrates the inherent flaws of relevant exam-



Figure 4: Comparison of all methods at different difficulty levels on the MATH dataset. Level 1 represents the easiest and level 5 is the hardest. 'relevant' clearly performs worse than other approaches at all difficulty levels.

ples, indicating that *mathematical reasoning tasks exhibit different analogical reasoning paradigms from other reasoning tasks*.

It might present challenges to prompt LLMs to accurately generate specific types of demonstrations. Therefore, given the unexpected results on mathematical reasoning tasks, one may wonder:

Q1-1. Are self-generated examples really relevant or irrelevant to the query?

To quantitatively measure the relevance between the generated examples and the query, we compute the average cosine similarity between them. Following Zhang et al. (2023), we use the sentence transformer (Reimers and Gurevych, 2019) to encode all samples. For each method, the reported result is averaged across three seeds.

As observed from Table 4, relevant examples are much more semantically similar to the query than irrelevant ones and the relevance score of 'relevant' is more biased towards 'oracle' rather than 'random' or 'N/A', demonstrating that *LLMs in*-



290

291 292

296

299

300

301

302

303

Method	Precalculus	Intermediate Algebra	Algebra	Prealgebra	Counting & Probability	Geometry	Number Theory
Relevant	10.4	9.8	51.8	56.8	22.1	24.2	37.0
N/A	9.1	15.7	55.5	61.0	28.7	25.8	34.2
Random _{same}	12.3	17.6	54.4	60.6	25.4	25.8	34.9
Random _{diff}	13.0	14.1	52.7	56.8	26.2	24.2	33.6
Random _{bio}	13.0	12.2	53.0	59.2	28.7	25.8	32.2

Table 3: Accuracy (%) across different subjects in the MATH dataset. Self-generated irrelevant examples outperform relevant ones on 6 out of 7 subjects.

Method	GSM8K	MATH	Average
Relevant	0.54	0.41	0.48
N/A	0.19	0.28	0.24
Random _{same}	0.30	0.20	0.25
Random _{diff}	0.15	0.10	0.13
Random _{bio}	0.06	0.11	0.09
Oracle	0.65	0.63	0.64

Table 4: Average relevance score (semantic similarity) between self-generated examples and the query. 'Oracle' stands for the average similarity score between the query and k most similar training samples (k is the number of self-generated examples).

	Relevant	N/A	Random _{same}
Accuracy	62.0	72.0	86.0

Table 5: Accuracy (%) of self-generated examples on the MATH dataset. The examples generated by 'relevant' are less accurate.

deed follow instructions to generate specific types of demonstrations.

305

306

307

308

310

311

312

313

315

316

317

318

319

320

We provide a case study in Table 6 to delve deeper into the demonstrations of different methods. As we can notice, the example generated by 'relevant' is more related to the query as they both involve the mathematical concept 'number bases'. In contrast, examples such as 'What is the value of x in the equation 2x + 5 = 10?' (N/A) or 'How do you bake chocolate chip cookies?' (Random_{diff}) are less relevant to the query. This comparison highlights once again that relevance may not be the key factor for analogical reasoning performance on mathematical reasoning tasks. To understand better the underlying reasons for the counterintuitive results, we then ask the following question:

Q2. If relevance is not the key factor, what is more important for the accuracy of analogical reasoning?

Looking back at Table 6, an interesting observation is that the self-generated relevant example appears to be more difficult to solve than the irrelevant ones, regardless of whether they are math problems or not. Consequently, the accuracy of relevant examples may be lower. To verify this, we conduct a pilot experiment on MATH. Specifically, we randomly select 50 samples for different types of generated math problems, *i.e.*, Relevant, N/A and Randomsame, and manually evaluate their accuracy. We exclude other methods as it is difficult to define the 'accuracy' of the examples they generate. From the results in Table 5, we can observe that while the examples generated by 'relevant' are more related to the test query, they are less accu*rate*, raising the question whether the performance of different approaches on mathematical reasoning tasks is strongly correlated with the accuracy of self-generated examples.

321

322

323

324

325

326

327

328

329

330

331

332

333

334

335

337

338

339

340

341

342

343

344

345

347

348

349

350

351

352

353

354

355

356

357

359

360

Proxy Approaches However, as the accuracy of the examples located at the output cannot be directly controlled, we meticulously design a variant called ICL, which extracts the generated examples from the model output as in-context learning (ICL) demonstrations and combines them with the query as input to LLMs, as a proxy for the original method. We also consider the following two variants: (a) GPT4-Calibration which replaces the answers of demonstrations in ICL with GPT4generated answers, and (b) Random changes the answers of demonstrations in ICL to random numbers. Our manual verification confirmed that GPT4generated answers were mostly accurate. We conduct this experiment on GSM8K and MATH with GPT-3.5 as the LLM reasoner.

From the results of different variants reported in Table 7, we can see that increasing the accuracy of generated examples can indeed improve the performance: *GPT4-Calibration* consistently outperforms *ICL* by incorporating more accurate

Query: For how many ordered pairs (A, B) where A and B are positive integers is $AAA_7 + BBB_7 = 666_7$?				
Relevant	In a certain base, the sum of two three-digit numbers is 777. If the digits of one of the numbers are reversed, the sum becomes 888. What is the base of this number system?			
N/A	What is the value of x in the equation $2x + 5 = 10$?			
Random _{same}	In a bag, there are 5 red marbles, 3 blue marbles, and 2 green marbles. If you randomly pick 2 marbles from the bag without replacement, what is the probability that both marbles are red?			
Random _{diff}	How do you bake chocolate chip cookies?			
Random _{bio}	How does the process of photosynthesis occur in plants?			
Oracle	Find the number of ordered pairs (a, b) of complex numbers such that $a^3b^5 = a^7b^2 = 1$.			

Table 6: Demonstration examples of different methods on the MATH dataset. The example generated by 'relevant' is more related to the query than other examples generated by 'N/A' or 'random'.

Variant	GSM8K			MATH		
v al faile	Relevant	N/A	Random _{same}	Relevant	N/A	Random _{same}
ICL	71.2	73.8	72.0	37.0	39.8	39.2
GPT4-Calibration	75.2	75.6	75.6	44.4	41.2	40.0
Random	70.0	72.0	68.4	36.0	38.0	37.8

Table 7: Accuracy (%) of different variants on GSM8K and MATH. When using ground truth (GPT4 generated answers), 'GPT4-Calibration' consistently outperforms 'ICL' for all methods. In contrast, 'random' always performs worse than 'ICL'.

Method	Task				
1.100mou	GSM8K	MATH	Average		
Relevant	71.5	33.3	52.4		
N/A	75.5	36.1	55.8		
Random _{same}	75.1	36.3	55.7		
Random _{diff}	76.3	34.1	55.2		
Random _{bio}	75.3	34.6	54.9		
ICL _{math}	75.7	36.8	56.3		
ICL _{bio}	77.9	34.9	56.4		

Table 8: Comparison of different methods on two mathematical reasoning tasks.

answers. In contrast, random always performs the worst among all variants. Therefore, the key factor influencing the performance on mathematical reasoning is the accuracy of self-generated examples rather than their relevance.

Given the above findings, a natural question is:

Q2-1. Can we ask the LLM to randomly generate a few math or biological problems and manually verify their correctness, then use this fixed set of problems as ICL demonstrations for all test queries?

We refer to these two methods as $\ensuremath{\text{ICL}_\text{math}}$ and ICL_{bio}, and conduct experiments with them on GSM8K and MATH (see Fig. 5 for example prompts and outputs for generating math problems). Detailed prompts and outputs for different methods are provided in Appendix A.3. Following the original setting, we ask the LLM to randomly generate 5 examples for GSM8K and 3 examples for MATH. As observed from Table 8, ensuring the accuracy of self-generated examples does lead to better performance regardless of the problem type. Moreover, both ICL variants only need to generate examples once, which significantly reduces the inference cost and further demonstrates their superiority.

371

372

373

374

376

377

378

379

381

383

384

387

389

390

391

392

393

394

395

396

397

398

4.3 Further Analysis

Repeating Problems While generating a few accurate problems as ICL demonstrations can achieve better performance, a bolder idea might be to generate one problem and repeat it multiple times as fewshot demonstrations for ICL. To investigate this, we randomly select a generated math problem and repeat it to perform ICL, denoted by ICL_{math repeat}. From the results shown in Table 11, we can see that ICL_{math_repeat} consistently performs worse than ICL_{math} on both datasets, indicating that the diversity of generated problems also matters.

Generalization to Open-source LLMs Our experiments and analysis so far used GPT-3.5 as the LLM, which is closed-source and gets updated over time. To verify whether the observations and conclusions are consistent across different models and additionally for reproducibility, we extend the ex-

370

Input
input input in the second se
Randomly recall five examples of math problems. Note that your problems should be distinct from each other (e.g., involving different numbers and names). For each problem: - After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in .
Output
Culput
O: If John has 5 apples and he huves 8 more, how many apples does he have in total?
A . To find the total number of angles lobe has we add the number of angles ha initially has to the number of angles has been added as $\pm 8 - 13$
Therefore long has been been been been been been been bee
Therefore, John has woxed (13) applies in total.
\mathbf{O} : Find the derivative of the function $f(x) = 3x/2 + 2x = 5$
a. To find the derivative of the function $f(x) = 0$ and $x = -5$.
Fig. Therefore the derivative of (A), we apply the power rule to each reference of our 2 is or, the derivative of 2 is 2, and the derivative of a constant
-5 is 0. Therefore, the derivative of $f(x) = f(x) = 0x + 2$. So, the derivative is $wolded(0x + 2)$.

Figure 5: Example prompts and outputs for randomly generating math problems. We manually verify the answers to ensure the correctness of the generated examples.

Method	Relevant	N/A	Random _{same}	Random _{diff}	Random _{bio}	ICL _{math}	ICL _{bio}
Accuracy	45.1	51.4	50.9	54.3	47.1	55.5	56.1

Table 9: Accuracy (%) of different methods on GSM8K using Llama-2-70b-Chat. Self-generated relevant examples always perform worse than irrelevant ones and both ICL variants outperform other approaches.

Variant	Method				
	Relevant	N/A	Random _{same}		
ICL	56.2	58.2	58.6		
GPT4-Calibration	60.8	61.0	60.8		
Random	53.2	54.0	59.6		

Table 10: Accuracy (%) of different variants on GSM8K using Llama-2-70b-Chat. 'GPT4-Calibration' consistently performs better than 'ICL' and 'random' for different methods.

Method	Task				
	GSM8K	MATH	Average		
ICL _{math}	75.7	36.8	56.3		
ICL_{math_repeat}	73.8	36.2	55.0		

Table 11:Comparison of two ICL variants on theGSM8K and MATH datasets.

periments to Llama-2-Chat (Touvron et al., 2023c). Specifically, we use vLLM (Kwon et al., 2023) to serve a Llama-2-70b-Chat model for experiments and report the results of different methods/variants on GSM8K in Table 9 and Table 10. We can draw similar observations: (*i*) self-generated relevant examples underperform all types of irrelevant ones, (*ii*) 'GPT4-Calibration' consistently outperforms the other two variants, and (*iii*) ICL_{math} and ICL_{bio} perform better than other approaches, demonstrating that the conclusions can be generalized to different models.

400

401

402

403

404

405

406

407

408

409

410

411 **Different Numbers of Demonstrations** While 412 we mainly follow the setting in Yasunaga et al.

Number	Relevant	N/A	Random _{same}	$Random_{diff} \\$	$Random_{bio}$
3	73.1	77.3	75.0	75.3	75.5
5	71.5	75.5	75.1	76.3	75.3

Table 12: Accuracy (%) of all methods with different numbers of demonstrations.

(2024) to ask the LLM to generate k = 5 examples for GSM8K, we further investigate the effect of the number of demonstrations. Specifically, we conduct controlled experiments with k = 3 and report the results in Table 12. We can observe that irrelevant examples consistently outperform relevant ones across different numbers of demonstrations, emphasizing their effectiveness. 413

414

415

416

417

418

419

420

421

422

423

424

425

426

427

428

429

430

431

432

433

434

435

436

5 Conclusion

In this work, we have systematically assessed the capability of LLMs to perform analogical reasoning. We have identified key research questions and empirically analyzed a representative set of LLMs on a diverse collection of reasoning tasks. Extensive experimental results and analysis show that LLMs *cannot* always perform analogical reasoning and the key influencing factor is the accuracy of self-generated examples rather than their relevance. Given these findings, we have designed two ICL-based approaches with better performance and significantly reduced inference costs. In the future, we would like to investigate additional analogical prompting methods to generate more accurate examples.

548

549

492

493

494

437 Limitations

This work has several limitations. First, due to the inference cost of ChatGPT¹, we conduct experiments on subsets of the test data for mathematical reasoning tasks. Besides, we include 3 datasets requiring different reasoning capabilities in this work. A further improvement could be to explore more diverse types of tasks.

References

445

446

447

448

449

450

451

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

480

481

482

484

485

486

487

488

489

490

491

- Mohammad Gheshlaghi Azar, Mark Rowland, Bilal Piot, Daniel Guo, Daniele Calandriello, Michal Valko, and Rémi Munos. 2023. A general theoretical paradigm to understand learning from human preferences. *arXiv preprint arXiv:2310.12036*.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiushi Du, Zhe Fu, Huazuo Gao, Kaige Gao, Wenjun Gao, Ruiqi Ge, Kang Guan, Daya Guo, Jianzhong Guo, Guangbo Hao, Zhewen Hao, Ying He, Wenjie Hu, Panpan Huang, Erhang Li, Guowei Li, Jiashi Li, Yao Li, Y. K. Li, Wenfeng Liang, Fangyun Lin, A. X. Liu, Bo Liu, Wen Liu, Xiaodong Liu, Xin Liu, Yiyuan Liu, Haoyu Lu, Shanghao Lu, Fuli Luo, Shirong Ma, Xiaotao Nie, Tian Pei, Yishi Piao, Junjie Qiu, Hui Qu, Tongzheng Ren, Zehui Ren, Chong Ruan, Zhangli Sha, Zhihong Shao, Junxiao Song, Xuecheng Su, Jingxiang Sun, Yaofeng Sun, Minghui Tang, Bingxuan Wang, Peiyi Wang, Shiyu Wang, Yaohui Wang, Yongji Wang, Tong Wu, Y. Wu, Xin Xie, Zhenda Xie, Ziwei Xie, Yiliang Xiong, Hanwei Xu, R. X. Xu, Yanhong Xu, Dejian Yang, Yuxiang You, Shuiping Yu, Xingkai Yu, B. Zhang, Haowei Zhang, Lecong Zhang, Liyue Zhang, Mingchuan Zhang, Minghua Zhang, Wentao Zhang, Yichao Zhang, Chenggang Zhao, Yao Zhao, Shangyan Zhou, Shunfeng Zhou, Qihao Zhu, and Yuheng Zou. 2024. Deepseek llm: Scaling open-source language models with longtermism.
- 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. In Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual.
 - Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. 2022. Program of thoughts

prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*.

- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. Vicuna: An opensource chatbot impressing gpt-4 with 90%* chatgpt quality.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayana Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. Specializing smaller language models towards multi-step reasoning. *arXiv preprint arXiv:2301.12726*.
- Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2022. Complexity-based prompting for multi-step reasoning. *arXiv preprint arXiv:2210.00720*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2022. Large language models are reasoning teachers. *arXiv* preprint arXiv:2212.10071.

¹experiments done between 12/2023 and 04/2024

- 550 551
- 554 559
- 571 573 574
- 575 576 577 578 579
- 581
- 585

586 588

590

- 594 595
- 596 597

599

- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. arXiv preprint arXiv:2203.15556.
- Keith J Holyoak. 2012. Analogy and relational reasoning. The Oxford handbook of thinking and reasoning, pages 234-259.
- 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.
- Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. 2022. Decomposed prompting: A modular approach for solving complex tasks. arXiv preprint arXiv:2210.02406.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS 2022).
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles.
- Harrison Lee, Samrat Phatale, Hassan Mansoor, Kellie Lu, Thomas Mesnard, Colton Bishop, Victor Carbune, and Abhinav Rastogi. 2023. Rlaif: Scaling reinforcement learning from human feedback with ai feedback. arXiv preprint arXiv:2309.00267.
- Yifei Li, Zeqi Lin, Shizhuo Zhang, Qiang Fu, Bei Chen, Jian-Guang Lou, and Weizhu Chen. 2022. On the advance of making language models better reasoners. arXiv preprint arXiv:2206.02336.
- Yuanzhi Li, Sébastien Bubeck, Ronen Eldan, Allie Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023. Textbooks are all you need II: phi-1.5 technical report. arXiv preprint arXiv:2309.05463.
- Tianqi Liu, Yao Zhao, Rishabh Joshi, Misha Khalman, Mohammad Saleh, Peter J Liu, and Jialu Liu. 2024. Statistical rejection sampling improves preference optimization. In The Twelfth International Conference on Learning Representations.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. 2022. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. arXiv preprint arXiv:2209.14610.

Lucie Charlotte Magister, Jonathan Mallinson, Jakub Adamek, Eric Malmi, and Aliaksei Severyn. 2022. Teaching small language models to reason. ArXiv preprint, abs/2212.08410.

604

605

608

609

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648

649

650

651

652

653

654

655

656

657

658

- 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.
- OpenAI. 2023. Gpt-4 technical report. arXiv preprint arXiv:2303.08774.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155.
- Chengwei Qin, Aston Zhang, Anirudh Dagar, and Wenming Ye. 2023a. In-context learning with iterative demonstration selection. arXiv preprint arXiv:2310.09881.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023b. Is ChatGPT a general-purpose natural language processing task solver? In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, pages 1339-1384, Singapore. Association for Computational Linguistics.
- Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. 2021. Scaling language models: Methods, analysis & insights from training gopher. arXiv preprint arXiv:2112.11446.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. arXiv preprint arXiv:2305.18290.
- VS Ramachandran. 2012. Encyclopedia of Human Behavior. Elsevier/Academic Press; 2012.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERTnetworks. 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 3982-3992, Hong Kong, China. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Teven Le Scao, Arun Raja, et al. 2021. Multitask prompted training enables zero-shot task generalization. arXiv preprint arXiv:2110.08207.

661

John Schulman, Filip Wolski, Prafulla Dhariwal,

Noah Shinn, Federico Cassano, Beck Labash, Ashwin

Shaden Smith, Mostofa Patwary, Brandon Norick,

Patrick LeGresley, Samyam Rajbhandari, Jared

Casper, Zhun Liu, Shrimai Prabhumoye, George

Zerveas, Vijay Korthikanti, et al. 2022. Using deep-

speed and megatron to train megatron-turing nlg

530b, a large-scale generative language model. arXiv

Mirac Suzgun, Nathan Scales, Nathanael Schärli, Se-

bastian Gehrmann, Yi Tay, Hyung Won Chung,

Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny

Zhou, et al. 2022. Challenging big-bench tasks and

whether chain-of-thought can solve them. arXiv

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann

Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,

and Tatsunori B Hashimoto. 2023. Stanford alpaca:

An instruction-following llama model. Stanford

Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Gar-

cia, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng,

Neil Houlsby, and Donald Metzler. 2022. Unify-

ing language learning paradigms. arXiv preprint

Romal Thoppilan, Daniel De Freitas, Jamie Hall,

Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze

Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du,

YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng,

Amin Ghafouri, Marcelo Menegali, Yanping Huang,

Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao

Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-

Ching Chang, Igor Krivokon, Will Rusch, Marc

Pickett, Pranesh Srinivasan, Laichee Man, Kathleen

Meier-Hellstern, Meredith Ringel Morris, Tulsee Doshi, Renelito Delos Santos, Toju Duke, Johnny So-

raker, Ben Zevenbergen, Vinodkumar Prabhakaran,

Mark Diaz, Ben Hutchinson, Kristen Olson, Ale-

jandra Molina, Erin Hoffman-John, Josh Lee, Lora

Aroyo, Ravi Rajakumar, Alena Butryna, Matthew

Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Co-

hen, Rachel Bernstein, Ray Kurzweil, Blaise Aguera-

Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc

Le. 2022. Lamda: Language models for dialog appli-

Gopinath, Karthik Narasimhan, and Shunyu Yao.

2023. Reflexion: Language agents with verbal rein-

mal policy optimization algorithms. arXiv preprint

Proxi-

Alec Radford, and Oleg Klimov. 2017.

arXiv:1707.06347.

forcement learning.

preprint arXiv:2201.11990.

preprint arXiv:2210.09261.

CRFM.

arXiv:2205.05131.

- 668
- 675

690

- 704 706
- 707 708
- 709 710

711 cations. 712 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier 713 Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal 714 Azhar, et al. 2023a. Llama: Open and effi-715 cient foundation language models. arXiv preprint 716 arXiv:2302.13971. 717

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, 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. 2023b. Llama 2: Open foundation and finetuned chat models. arXiv preprint arXiv:2307.09288.

718

719

721

722

725

726

727

728

729

730

732

733

736

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756

757

758

759

760

761

762

763

765

766

767

768

770

771

772

773

774

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023c. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.
- Stella Vosniadou and Andrew Ortony, 1989. Similarity and analogical reasoning. Cambridge University Press
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2022a. Rationaleaugmented ensembles in language models. arXiv preprint arXiv:2207.00747.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2022b. Self-consistency improves chain of thought reasoning in language models. arXiv preprint arXiv:2203.11171.
- Taylor Webb, Keith J Holyoak, and Hongjing Lu. 2023. Emergent analogical reasoning in large language models. Nature Human Behaviour, 7(9):1526-1541.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. arXiv preprint arXiv:2109.01652.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In Thirty-sixth Conference on Neural Information Processing Systems (NeurIPS 2022).
- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. arXiv preprint arXiv:2304.12244.

Michihiro Yasunaga, Xinyun Chen, Yujia Li, Panupong Pasupat, Jure Leskovec, Percy Liang, Ed H. Chi, and Denny Zhou. 2024. Large language models as analogical reasoners. In The Twelfth International Conference on Learning Representations.

776

778

781

790

793

794

795

796

797 798

801

808

809

810

811

812

814

815

816

817

818

819

820

821

822

- Zheng Yuan, Hongyi Yuan, Chuanqi Tan, Wei Wang, Songfang Huang, and Fei Huang. 2023. Rrhf: Rank responses to align language models with human feedback without tears. arXiv preprint arXiv:2304.05302.
- Eric Zelikman, Yuhuai Wu, and Noah D Goodman. 2022. Star: Bootstrapping reasoning with reasoning. arXiv preprint arXiv:2203.14465.
 - Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. Automatic chain of thought prompting in large language models. In The Eleventh International Conference on Learning Representations (ICLR 2023).
- Ruochen Zhao, Xingxuan Li, Shafiq Joty, Chengwei Qin, and Lidong Bing. 2023. Verify-and-edit: A knowledge-enhanced chain-of-thought framework. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 5823-5840, Toronto, Canada. Association for Computational Linguistics.
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srini Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2023. Lima: Less is more for alignment. arXiv preprint arXiv:2305.11206.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. 2022. Least-to-most prompting enables complex reasoning in large language models. arXiv preprint arXiv:2205.10625.

Α Appendix

A.1 Prompts for Different Methods

The prompts for different methods on all datasets are shown in Fig. $6 \sim$ Fig. 8.

A.2 Detailed Results for Different Random Seeds

We report detailed results for different random seeds in Table 13 ~ Table 14.

A.3 Prompts and Outputs for Example Generation

We show detailed prompts and outputs for randomly generating math or biological problems in Fig. 9 and Fig. 10, respectively. 823

Seed	GSM8K				MATH					
	Relevant	N/A	Random _{same}	Random _{diff}	Random _{bio}	Relevant	N/A	Random _{same}	Random _{diff}	Random _{bio}
42	71.8	76.6	73.2	74.0	74.0	37.4	42.2	41.6	39.0	39.2
100	71.2	75.2	75.2	75.8	74.8	29.0	30.6	32.6	29.4	31.2
1000	71.4	74.8	77.0	79.2	77.0	33.6	35.6	34.6	34.0	33.4
Average	71.5	75.5	75.1	76.3	75.3	33.3	36.1	36.3	34.1	34.6

Table 13: Accuracy (%) of all methods with different random seeds on two mathematical reasoning tasks.

Seed		Temporal sequences	Logical deduction five objects	Reasoning about colored objects	Formal fallacies	Word sorting	Average
42	Relevant	58.0	52.8	76.0	50.4	77.2	62.9
	N/A	56.4	44.8	77.6	54.0	76.8	61.9
	Random _{same}	52.4	48.8	74.8	51.6	72.8	60.1
	Random _{diff}	43.2	46.8	74.0	52.4	67.6	56.8
	Random _{bio}	56.8	52.0	74.0	52.0	76.4	62.2
100	Relevant	58.4	50.8	78.4	51.2	76.8	63.1
	N/A	55.2	46.0	74.8	52.8	79.2	61.6
	Random _{same}	50.8	48.4	73.6	53.2	75.2	60.2
	Random _{diff}	46.4	46.8	72.8	50.0	70.4	57.3
	Random _{bio}	58.0	48.4	78.4	51.2	73.6	61.9
1000	Relevant	63.6	50.0	75.6	52.0	76.8	63.6
	N/A	60.8	45.2	74.0	53.2	77.2	62.1
	Random _{same}	56.0	49.2	72.0	52.4	74.4	60.8
	Random _{diff}	43.2	40.8	70.4	51.2	69.6	55.0
	Random _{bio}	56.4	48.0	76.0	49.2	74.8	60.9

Table 14: Accuracy (%) of all methods with different random seeds on BBH.

Prompt: self-generate relevant examples

Your task is to tackle mathematical problems. When presented with a math problem, recall relevant problems as examples. Afterward, proceed to solve the initial problem # Initial Problem

[The target problem]

Instructions: Make sure that your response follows the instructions below.

Analogous Problems: Offer five diverse examples of math problems that are relevant or analogous to the initial problem. For each problem, elaborate on the solution and conclude with the ultimate answer (enclosed in \boxed{}). For each problem: - After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate N/A examples

Your task is to tackle mathematical problems. When presented with a math problem, recall n/a problems as examples. Afterward, proceed to solve the initial problem

Initial Problem: [The target problem]

Instructions: Make sure that your response follows the instructions below.

N/A Problems:

Offer five diverse examples of math problems that are n/a to the initial problem. For each problem, elaborate on the solution and conclude with the ultimate answer (enclosed in \boxed{}). For each problem: - After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

Copy and paste the initial problem here.
 A: Explain the solution and enclose the ultimate answer in \boxed{} here

Prompt: self-generate random math examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions Make sure that your response follows the instructions below.

Random Problem

Randomly offer five diverse examples of math problems. For each problem, elaborate on the solution and conclude with the ultimate answer (enclosed in \boxed{}). For each problem:

- After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

Copy and past the initial problem here.
 A: Explain the solution and enclose the ultimate answer in \boxed{ here.

Prompt: self-generate random no-math examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random problems (remember not to output math problems) as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions:

Make sure that your response follows the instructions below.

Random Problems

Antioun routines.
 Rendomly offer five diverse examples of any type, except math problems. For each problem, elaborate on the solution and conclude with the ultimate answer (enclosed in toxed). For each problem:
 - After 'C': "describe the problem
 - After 'A': ", explain the solution and enclose the ultimate answer in boxed}.

Solve the Initial Problem:

Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate random biological examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random biological problems (remember not to output math problems) as examples. Afterward, proceed to solve the initial problem # Initial Problem [The target problem] # Instructions: Make sure that your response follows the instructions below. ## Random Problems: Randomly offer five diverse examples of biological problems (remember not to output math problems). For each problem, elaborate on the solution and conclude with the ultimate answer (enclosed in bloxed{}). For each problem: - After 'C', 'I describe the problem - After 'C', 'I describe the problem

Solve the Initial Problem: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Figure 6: Prompts for different methods on GSM8K.

Prompt: self-generate relevant examples

Your task is to tackle mathematical problems. When presented with a math problem, recall relevant problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem [The target problem]

Instructions

Make sure to include all of the following points:

Relevant Problems

Recall three examples of math problems that are relevant to the initial problem. Note that your problems should be distinct from each other and from the initial problem (e.g., involving different numbers and names). For each problem: - After "0: ", describe the problem - After "1: explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem: Say "Let's solve the following math problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate N/A examples

Your task is to tackle mathematical problems. When presented with a math problem, recall n/a problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions: Make sure to include all of the following points:

N/A Problems

Recall three examples of math problems that are n/a to the initial problem. Note that your problems should be distinct from each other and from the initial problem (e.g., involving different numbers and names). For each problem: - After 'C': , describe the problem - After 'A': ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

As you are main booten. Say "Let's solve the following math problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate random math examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem [The target problem]

Instructions: Make sure to include all of the following points:

Random Problems:

Randomly recall three examples of math problems. Note that your problems should be distinct from each other and from the initial problem (e.g., involving different numbers and names). For each problem: - After "C: ", describe the problem: - After "A: ", explain the solution and enclose the ultimate answer in \boxed().

Solve the Initial Problem

Solve the Initial Problem: Say "Let's solve the following math problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate random no-math examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random problems (remember not to output math problems) as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions:

Make sure to include all of the following points:

Random Problems

ly recall three examples of any type, except math problems. Note that your problems should be distinct from each other and from the initial problem. For each problem

After "Q: ", describe the problem
 After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem:

Say "Let's solve the following math problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{ here.

Prompt: self-generate random biological examples

Your task is to tackle mathematical problems. When presented with a math problem, recall random biological problems (remember not to output math problems) as examples. Afterward, proceed to solve the initial problem

Initial Problem: [The target problem]

Instructions

Make sure to include all of the following points:

Random Problems:

Randomly recall three examples of biological problems (remember not to output math problems). Note that your problems should be distinct from each other and Form the initial problem. For each problem: - After "Q: ", describe the problem - After "Q: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

are solve the final problem.
 Say "Let's solve the following math problem." Then formulate your response in the following format:
 Q: Copy and paste the initial problem here.
 A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Figure 7: Prompts for different methods on MATH.

Prompt: self-generate relevant examples

Your task is to tackle reasoning problems. When presented with a problem, recall relevant problems as examples. Afterward, proceed to solve the initial problem

Initial Problem: [The target problem]

Instructions:

Make sure to include all of the following points:

Relevant Problems

Recall three examples of problems that are relevant to the initial problem. Note that your problems must be distinct from each other and from the initial problem. For each problem: - After "Q: ", describe the problem - After "A": ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

Say "Let's solve the following reasoning problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate N/A examples

Your task is to tackle reasoning problems. When presented with a problem, recall n/a problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions:

Make sure to include all of the following points:

N/A Problems

Recall three examples of problems that are n/a to the initial problem. Note that your problems must be distinct from each other and from the initial problem. For each - After "Q: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem:

Say "Let's solve the following reasoning problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here.

A: Explain the solution and enclose the ultimate answer in \boxed{} here

Prompt: self-generate random reasoning examples

Your task is to tackle reasoning problems. When presented with a problem, recall random problems as examples. Afterward, proceed to solve the initial problem.

Initial Problem: [The target problem]

Instructions Make sure to include all of the following points:

Random Problems

Randomly recall three examples of reasoning problems. Note that your problems must be distinct from each other and from the initial problem. For each proble - After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem:

Say "Let's solve the following reasoning problem." Then formulate your response in the following format:

Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Prompt: self-generate random no-reasoning examples

Your task is to tackle reasoning problems. When presented with a problem, recall random problems (remember not to ou ning problems) as examples. Afterward, proceed to solve the initial problem

Initial Problem

[The target problem]

Instructions: Make sure to include all of the following points:

Random Problems:

Randomly recall three examples of any type, except reasoning problems. Note that your problems must be distinct from each other and from the initial problem. For each problem:

- After "Q: ", describe the problem - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem: Say "Let's solve the following reasoning problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here.

A: Explain the solution and enclose the ultimate answer in \boxed{} here

Prompt: self-generate random biological examples

Your task is to tackle reasoning problems. When presented with a problem, recall random biological problems (remember not to output reasoning problems) as examples. Afterward, proceed to solve the initial problem.

Initial Problem:

[The target problem]

Instructions: Make sure to include all of the following points:

Random Problems:

Randomly recall three examples of biological problems (remember not to output reasoning problems). Note that your problems must be distinct from each other and Form the initial problem. For each problem: - After "Q: ", describe the problem: - After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Solve the Initial Problem

Solve the finitial Problem: Say "Let's solve the following reasoning problem." Then formulate your response in the following format: Q: Copy and paste the initial problem here. A: Explain the solution and enclose the ultimate answer in \boxed{} here.

Figure 8: Prompts for different methods on BBH.

Input

Randomly recall five examples of math problems. Note that your problems should be distinct from each other (e.g., involving different numbers and names) For each problem:

After "Q: ", describe the problem
 After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Output

Q: If John has 5 apples and he buys 8 more, how many apples does he have in total? A: To find the total number of apples John has, we add the number of apples he initially has to the number of apples he buys. So, we calculate 5 + 8 = 13. Therefore, John has \boxed{13} apples in total.

Q: Solve for x in the equation 2x - 4 = 10.

A. To solve for x, we first add 4 to both sides of the equation to get 2x = 14. Then, we divide both sides by 2 to isolate x, resulting in x = 7. Thus, the solution is $\log x = 7$.

Q: A rectangle has a length of 10 cm and a width of 4 cm. What is its area?

A: The area of a rectangle is found by multiplying its length by its width. Therefore, the area of this rectangle is 10 * 4 = 40 square centimeters. So, the area is \boxed{40}

Q: A fair six-sided die is rolled. What is the probability of rolling a number greater than 4?

A: A six-sided die has numbers 1 through 6. The numbers greater than 4 are 5 and 6, which are 2 out of the 6 possible outcomes. Therefore, the probability of rolling a number greater than 4 is \frac{2}{6} = \frac{1}{3}. Thus, the probability is \boxed{\frac{1}{3}}.

Q: Find the derivative of the function $f(x) = 3x^2 + 2x - 5$.

A. To find the derivative of f(x), we apply the power rule to each term. The derivative of $3x^{A_2}$ is 6x, the derivative of 2x is 2, and the derivative of a constant -5 is 0. Therefore, the derivative of f(x) is f'(x) = 6x + 2. So, the derivative is \boxed{6x + 2}.

Figure 9: Prompts and outputs for generating math problems.

Input

Randomly recall five diverse examples of biological problems (remember not to output math problems). For each problem:

- After "A: ", explain the solution and enclose the ultimate answer in \boxed{}.

Output

Q: In a certain species of plant, the allele for red flowers (R) is dominant over the allele for white flowers (r). If two heterozygous plants are crossed, what is the probability that an offspring will have white flowers?

A: To solve this problem, we use a Punnett square to determine the genotypes of the offspring. The cross of two heterozygous plants Rr x Rr produces the following genotypes: RR, Rr, Rr, and rr. The probability of an offspring having white flowers rr is 1 out of 4, since only one genotype out of the four possible genotypes results in white flowers. Therefore, the probability is \boxed{\frac{1}{4}}.

Q: A small pond is introduced with 10 fish, and the fish population doubles every month. How many fish will there be in the pond after 6 months? **A:** To find the fish population after 6 months, we use the formula for exponential growth: $P = P_0 * 2^{An}$, where P is the final population, P_0 is the initial population, and n is the number of time periods (months, in this case). Plugging in the values, we get $P = 10 * 2^{A}$. Calculating this gives P = 10 * 64 = 640. Therefore, there will be \boxed{640} fish in the pond after 6 months.

Q: What is the primary mechanism by which the human kidney concentrates urine?

A: The primary mechanism by which the human kidney concentrates urine is through the counter-current multiplier system in the loop of Henle. This system allows for the reabsorption of water and solutes from the filtrate back into the bloodstream, which concentrates the urine. The key to this process is the creation of a high osmolarity in the medullary interstitial fluid, which draws water out of the filtrate in the collecting ducts, concentrating the urine. The ultimate answer is the \boxed{counter-current multiplier system}.

Q: What is the principle behind the theory of natural selection as proposed by Charles Darwin?

A: The principle behind the theory of natural selection is that individuals within a species show variation in physical characteristics and behaviors. These variations can affect the individuals' ability to survive and reproduce in their environment. Individuals with traits that are better suited to their environment are more likely to survive and reproduce, passing on those advantageous traits to their offspring. Over time, this can lead to changes in the species as a whole. The ultimate answer is the \boxed{survival and reproduction of the fittest}.

Q: How do antibiotics kill bacteria?

A: Antibiotics can kill bacteria through several mechanisms, including inhibiting cell wall synthesis, disrupting protein synthesis, impairing nucleic acid synthesis, and interfering with metabolic pathways. The specific mechanism depends on the class of antibiotic. For example, penicillin disrupts the synthesis of the bacterial cell wall, leading to cell lysis and death. Therefore, the ultimate answer is by 'boxed{inhibiting vital processes} specific to bacteria.

Figure 10: Prompts and outputs for generating biological problems.