

CoMM: Collaborative Multi-Agent, Multi-Reasoning-Path Prompting for Complex Problem Solving

Pei Chen^{1*} Boran Han² Shuai Zhang^{2†}

¹ Department of Computer Science and Engineering, Texas A&M University

² Amazon Web Services

chenpei@tamu.edu

{boranhan, shuaizs}@amazon.com

Abstract

Large Language Models (LLMs) have shown great ability in solving traditional natural language tasks and elementary reasoning tasks with appropriate prompting techniques. However, their ability is still limited in solving complicated science problems. In this work, we aim to push the upper bound of the reasoning capability of LLMs by proposing a collaborative multi-agent, multi-reasoning-path (CoMM) prompting framework. Specifically, we prompt LLMs to play different roles in a problem-solving team, and encourage different role-play agents to collaboratively solve the target task. In particular, we discover that applying different reasoning paths for different roles is an effective strategy to implement few-shot prompting approaches in the multi-agent scenarios. Empirical results demonstrate the effectiveness of the proposed methods on two college-level science problems over competitive baselines. Our further analysis shows the necessity of prompting LLMs to play different roles or experts independently.

1 Introduction

Large Language Models (LLMs) such as GPT (Brown et al., 2020; OpenAI, 2023), LLaMA (Touvron et al., 2023a,b) and PaLM (Chowdhery et al., 2022), have shown remarkable proficiency in solving many downstream tasks (Liu et al., 2021), without furthering fine-tuning the model parameters. However, their ability is limited to solving reasoning and mathematical problems (Wei et al., 2022b), especially complicated science problems (Ma et al., 2023; Xu et al., 2023; Ling et al., 2023a). In consideration of this limitation, and the costly fine-tuning overhead of the LLMs with billion-level parameters, many prompting methods emerge, i.e., the process of carefully crafting input queries to effectively

communicate with LLMs and obtain desired outputs. Apart from the benefit of exempting from manipulating the parameters of the LLMs, these prompting methods seamlessly integrate the pre-trained models into downstream tasks by eliciting desired model behaviors (Sahoo et al., 2024).

Among these endeavored prompting approaches, some of them prompt LLMs to reason with multiple middle-steps or subproblems for reasoning tasks (Wei et al., 2022b; Wang et al., 2023b; Yao et al., 2023; Hao et al., 2023; Zhou et al., 2023), some of them prompt LLMs to take advantage of external tools (Gao et al., 2023; Chen et al., 2023). Despite that these methods have further pushed the upper-bound reasoning and math abilities of LLMs, their potential for complicated science problems is still awaiting further improvement. Take the college-level physical problem in Figure 1 as an example, we can observe that the Chain-of-thought prompting method still suffers from knowledge and factual errors (incorrect equations used for the given problem) and computing errors (incorrect calculations).

Recently, agent-based prompting methods that prompt an LLM to play a specific role or act as an intelligent agent further unlock the ability of LLMs to solve complicated problems. For example, Xu et al. (2023) prompt an LLM to play as a domain expert and successfully elicits the LLM to answer domain questions. Huang et al. (2022); Shinn et al. (2023); Madaan et al. (2023) prompt LLMs to do self-reflection or self-refinement to correct the mistakes. Wang et al. (2023a); Sun et al. (2023) prompt LLMs to do planning before solving a specific task. Wang et al. (2023c) prompts a single agent to play multiple roles with different personas, and Liang et al. (2023); Chan et al. (2023); Du et al. (2023) prompt LLMs to play different roles in debating for problem-solving.

Following these works, we propose a collabora-

*Work done as an intern at Amazon Web Services.

† Corresponding author.

Complex Science Problem

Question: Light from a laser falls on a pair of very narrow slits separated by 0.5 micrometer, and bright fringes separated by 1.0 millimeter are observed on a distant screen. If the frequency of the laser light is doubled, what will be the separation of the bright fringes?

Correct Answer: 0.5 mm

Few-shot Chain-of-thought (CoT) Response

Answer: Let's think step-by-step. The separation of the bright fringes in a double-slit interference pattern is given by the equation:

$$d * \sin(\theta) = m * \lambda$$

where d is the separation between the slits, θ is the angle of the fringe, m is the order of the fringe, and λ is the wavelength of the light. In this case, the separation between the slits is given as 0.5 micrometers. For the original frequency of the laser light, the **separation of the bright fringes is given by:**

$$0.5 \text{ mm} * \sin(\theta) = m * \lambda$$

Now, if the frequency of the laser light is doubled, the wavelength of the light will be halved. So the new separation of the bright fringes will be:

$$0.5 \text{ mm} * \sin(\theta) = m * \lambda/2$$

Simplifying this equation, we find:

$$0.25 \text{ mm} * \sin(\theta) = m * \lambda$$

So, the separation of the bright fringes is 0.25 mm.

Figure 1: Chain-of-thought still makes Knowledge and Computing Errors in Solving the Complex Science Problem.

tive multi-agent framework (CoMM) that prompts LLMs to play different roles with different domain knowledge or task-solving duties for problem-solving. In particular, we propose a multi-path reasoning method that enables few-shot learning in the multi-agent framework. Empirical results on multiple complicated college-level science problems show that our method significantly outperforms strong baselines. Our further analysis shows that it is beneficial to include multiple agents for the collaboration, instead of prompting one agent to play multiple roles altogether.

2 Related Work

LLMs have shown remarkable proficiency in solving many downstream tasks (Qu et al., 2020b; Chen et al., 2021; Xu et al., 2024c,b), paving the way towards Artificial General Intelligence. With the advent of GPT-3 (Brown et al., 2020) and its emergent abilities (Wei et al., 2022a) in solving downstream tasks on both zero-shot and few-shot settings, many decoder-only LLMs follow (Ling et al., 2023b), such as PaLM (Chowdhery et al., 2022), LLaMA (Brown et al., 2020; OpenAI, 2023), BLOOM (Workshop et al., 2023), Claude (Bai et al., 2022), OPT (Zhang et al., 2022), Mistral (Jiang et al., 2023), Falcon (Penedo et al., 2023) etc. Considering the inference speed and economic expenditure, we choose GPT-3.5 as the back-

bone model for all the baselines and our CoMM approach.

In order to unlock the potential of the LLMs in solving downstream tasks (Yi and Qu, 2022; Chen et al., 2022; Qu et al., 2020a; Zhang et al., 2023; Yu et al., 2024; Xu et al., 2024a), many prompting approaches arise, exempting from manipulating the billion-level parameters (Li et al., 2023c). Among these prompting methods, ordinary prompting methods follow Brown et al. (2020) and employ task descriptions and sample demonstrations (few-shot) as the prompts for downstream tasks. To alleviate the difficulty of directly outputting the answer for LLMs, many prompting methods simplify the process by predicting the middle reasoning steps (chain-of-thought (Wei et al., 2022b)) or answering the decomposed sub-problems first (Wang et al., 2023b; Yao et al., 2023; Hao et al., 2023; Zhou et al., 2023; Ling et al., 2024). To overcome the lack of computing ability and outdated knowledge base, some work prompt LLMs to utilize external tools (Gao et al., 2023; Chen et al., 2023).

To further unlock the ability of LLMs in solving complicated problems, agent-based methods that prompt LLMs to play specific roles trend. Among them, single-agent methods only use one instance of LLMs. ExpertPrompt (Xu et al., 2023) prompts an LLM to play as a domain expert and successfully elicits the LLM to answer domain questions.

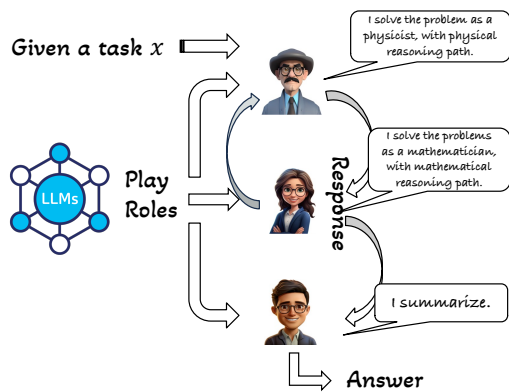


Figure 2: Overall Framework of CoMM: An Example from College Physics with the Few-shot Setting.

Step 1: Explain the environment and task as the system message prompt $x_s = f_s(x)$.

Step 2: Prompt the three agents $P_1(\theta)$, $P_2(\theta)$, $P_3(\theta)$ to play the domain experts and summarizer, with their role name, responsibility and principles.

Step 3: Each role will respond based on the current discussions, and when necessary, multi-turn dialogues are facilitated.

EmotionPrompt (Li et al., 2023a) improves the performance of agents with emotional prompts. Huang et al. (2022); Shinn et al. (2023); Madaan et al. (2023) prompts LLMs to do self-reflection or self-refinement to correct the mistakes. Wang et al. (2023a); Sun et al. (2023) prompts LLMs to do planning before solving a specific task. Wang et al. (2023c) prompts a single agent to play multiple roles with different personas.

Another branch of agent-based approaches are with multi-agents. For example, Liang et al. (2023); Chan et al. (2023); Du et al. (2023) prompt LLMs to play different roles in debating for problem-solving. ChatEval (Chan et al., 2023) uses multiple agents debating for automatic LLM evaluation. MathChat (Wu et al., 2023b) proposed a conversational framework to solve math problems with the user and LLM agent’s interactions. Park et al. (2023) and Li et al. (2023b) prompts LLMs to play as different agents for simulating human behaviors. Our work is closely related to these works, but our aim is to prompt LLMs to play different domain experts in a collaborative framework on complicated reasoning problems, and how to embed the few-shot examples into the multi-agent framework.

Along with the agent-based prompting methods, many open-sourced applications come out. For example, AutoGPT (Wu et al., 2023a) plays as AI agents that will attempt to achieve a given goal by breaking it into sub-tasks and using the internet and other tools in an automatic loop. AutoGen (Wu et al., 2023a) designs a framework for building LLM applications based on multi-agent conversations. MetaGPT (Hong et al., 2023) prompts multi-agent to play product managers, architects, project managers, and engineers for a software project. SkyAGI (Park et al., 2023) emerges human-

behavior simulation capability in LLM. While sharing the same multi-agent framework, our work focuses on exploring the effectiveness of the framework, i.e., we aim to answer whether multi-agent is necessary and how to prompt multiple agents to work collaboratively.

3 Methods

In this section, we first formally define the single-agent prompting framework, and then introduce the formal definition of the multi-agent prompting framework, and its adaptations to both zero-shot and few-shot settings (CoMM).

Single-agent Prompting Given a language model $P(\theta)$ and input text x , single-agent prompting takes a function that is applied to the input text $x' = f_{prompt}(x)$ (usually defines the target problem or task) and then predict the answer y by the language model that plays as a single problem-solving agent $P(y|x'; \theta)$. In the zero-shot setting, the prompting function f does not contain any demonstration examples, while in the few-shot setting, the prompting function contains a few examples.

Multi-agent Prompting For multi-agent prompting, we will have n language models $P_1(\theta_1)$, $P_2(\theta_2)$, ..., $P_n(\theta_n)$ that play different agents or roles in the framework. These language models can be the same ($\theta_1 = \theta_2 \dots = \theta_n$) or different ($\theta_1 \neq \theta_2 \dots \neq \theta_n$). For input text x , each agent i will have its own prompting functions $f_{prompt}^i(x)$ that formats the input task or problem for the agent. We define the interactions of these agents as a non-parametric function $\phi(y|g_1, g_2, \dots, g_n)$ where $g_i = P_i(y_i|f_{prompt}^i(x); \theta_i)$ and y_i is the output from agent i and y is the final answer.

Prompting Methods	Zero-shot		Few-shot	
	Moral Scenarios	College Physics	Moral Scenarios	College Physics
Standard (Brown et al., 2020)	38.65	44.12	38.21	48.04
CoT (Wei et al., 2022b)	45.58	50.00	64.92	56.86
Thought (Ma et al., 2023)	49.39	-	56.42	-
CoMM	52.17 (+ 2.78)	54.90 (+ 4.90)	65.03 (+ 0.11)	64.71 (+ 7.85)

Table 1: Main Test Results (Accuracy, %). Numbers in the parentheses are performance gains of the CoMM over previous state-of-the-art.

Collaborative Zero-shot Scenario In our collaborative multi-agent setting, we restrict the multiple agents to inherit from the same language models and the count of agents to be three. Then we have three language models $P_1(\theta)$, $P_2(\theta)$, $P_3(\theta)$ as the agents: $P_1(\theta)$ and $P_2(\theta)$ as the problem-solving experts and $P_3(\theta)$ as the summarizer, as shown in Figure 2.

Specifically, for a given input problem x , we use a prompt function to turn it into a system message that defines the collaborative team-working environment $x_s = f_s(x)$. For each agent, we define prompting functions to characterize its role and prompt it to give its solution accordingly. In particular, for the first expert agent, the prompting function formats the problem and the system message as $x_1 = f_1(x, x_s)$, and then gives its output $P_1(y_1|x_1; \theta)$. For the second expert agent, the prompting function formats the problem, the system message, and the output from y_1 as $x_2 = f_2(x, x_s, y_1)$, and then give its output $P_2(y_2|x_2; \theta)$. For the third summarizer, the prompting function will also consider the outputs from the two experts $x_3 = f_3(x, x_s, y_1, y_2)$ and then the agent gives the final answer $P_3(y|x_3; \theta)$.

For certain specific input tasks, multi-turn discussions are necessary. In this case, the output of the second expert agent will circulate back to the first agent as the input prompt again, and then repeat the afore-mentioned discussions, as demonstrated in the Figure 2.

Collaborative Few-shot Scenario In a multi-agent setting, it is not trivial to add the few-shot examples to the various agents. **Which agent should we give the few-shot examples?** We adopt a multi-path reasoning approach that gives the few-shot examples to the different agents. In particular, different agents will have their own expertise-based reasoning path in the few-shot demonstra-

tions. Formally, the two expert prompting functions $x_1 = f_1(x, x_s, e_1)$ and $x_2 = f_2(x, x_s, e_2, y_1)$ will take exemplars e_1 and e_2 as inputs. Take Figure 2 as an example, the few-shot examples will be added to both the physicist and the mathematician agents, but with different reasoning paths. More details can be found in the Appendix A.

4 Experiments

In this section, we will first introduce the evaluation datasets and benchmark that focus on complicated science problems. After that, we introduce the strong baseline prompting methods for comparison. At last, we introduce the results of our methods and the baselines on the benchmark.

4.1 Datasets

College Physics is a dataset from Massive Multitask Language Understanding (MMLU), which covers 57 subjects across different domain knowledge. It focuses on college-level physics problems. These problems are still very challenging and far from satisfying performance with large language models. Like the example from Figure 3, LLMs are still suffering from the lack of knowledge and computing ability.

Moral Scenarios is another dataset from MMLU (Hendrycks et al., 2020). Moral Scenarios focus on advanced professional-level social science problems that are yet challenging for large language models, which is among the worst performing tasks for many language models (Ma et al., 2023).

Both datasets are multiple choice questions, and we use the correct rate (Accuracy) as the metric for comparison.

4.2 Baselines

Standard (Brown et al., 2020) is the first work that introduced performing tasks without any task-specific training or examples, relying solely on its general pre-training with prompting. In this work, we format each problem as "Q: {question} A:" at zero-shot settings, and as "Q: {question example 1} A: {answer example 1} ... Q: {question example n } A: {answer example n } Q: {question} A:" for the few-shot setting with n demonstration examples.

Chain-of-thought (CoT) (Wei et al., 2022b) improves the Standard prompting approaches by introducing a series of intermediate natural language reasoning steps that lead to the final output (chain of thought). It hypothesizes that giving the LLMs longer predicting window, they have better chance to reach the answer, in comparison with directly requiring them to output the answer. For zero-shot implementation, we follow the Zero-shot-CoT proposed by Wang et al. (2023a), and add "Let's think step by step" prompt before the answer, i.e., "Q: {question} A: Let's think step by step.". As for the few-shot implementation, we follow the indigenous settings from Wei et al. (2022b), i.e., "Q: {question example 1} A: Let's think step by step. {answer example 1 with chain of thought} ... Q: {question example n } A: Let's think step by step. {answer example n with chain of thought} Q: {question} A: Let's think step by step." for the few-shot setting with n demonstration examples.

Thought Experiment (Thought) (Ma et al., 2023) is a reasoning framework that is specialized in better moral reasoning by using counterfactual reasoning. It is a multi-agent framework with multi-step prompting, and each step involves prompting the LLMs to solve a specific task. Specifically, this method involves employing counterfactual thinking to envision various, often hypothetical, situations, and then deliberating on the consequences of these imagined circumstances. By processing these scenarios, it aids in consolidating intermediate reflections, thereby leading to a deeper comprehension of the issue at hand and guiding towards the most appropriate solution. We adopt the same settings for both zero-shot and few-shot as provided by the Ma et al. (2023).

4.3 Settings

Backbone Model For a fair comparison, we use `gpt-3.5-turbo-0613`¹ as the backbone model, and set the temperature to be 0 in all our experiments.

Settings for College Physics We prompt the first agent $P_1(\theta)$ to be a physicist, the second agent $P_2(\theta)$ to be a mathematician, and the third agent $P_3(\theta)$ to be the summarizer. In the zero-shot setting, we do not provide demonstration examples, while in the few-shot setting, we give the same 5 examples for the two experts, but with different reasoning paths, i.e., the reasoning path of a physicist role and the reasoning path of a mathematician role individually. We only prompt the group to discuss once for this benchmark. More details can be found in the Appendix A.

Settings for Moreal Scenarios In the zero-shot setting, we prompt the first agent $P_1(\theta)$ to be a task decomposer, the second agent $P_2(\theta)$ to be a sub-problem solver, and the third agent $P_3(\theta)$ to be the summarize. In the few-shot setting, we also give each expert 5 examples, and we prompt the first agent $P_1(\theta)$ to be a chain-of-thought reasoner with CoT reasoning path, the second agent $P_2(\theta)$ to be a Thought reasoner with thought experiment path, and the third agent $P_3(\theta)$ to be the summarize. We prompt the group to discuss twice for this benchmark. More details can be found in the Appendix A.

4.4 Main Results

The main experimental results are shown in Table 1. It is saliently observable that the proposed CoMM approach can outperform the state-of-the-art baselines on both zero-shot and few-shot settings. In detail, it improves with absolute average improvements of 3.84% at zero-shot setting and 8.23% at few-shot setting. CoMM improves more in few-shot settings, further demonstrating the effectiveness of applying the multi-path reasoning approaches in the multi-agent framework. Also, CoMM improves more on the complicated College Physics dataset that requires more domain knowledge, further showcasing the efficacy of CoMM in solving complex problems.

¹<https://openai.com/>

Benchmark	Settings	Single Agent	Multiple Agents
Moral Scenarios	Zero-shot	27.71	52.17 (+24.46)
	Few-shot	42.68	65.03 (+22.35)
College Physics	Zero-shot	42.16	54.90 (+12.74)
	Few-shot	56.86	64.71 (+07.85)

Table 2: Single Agent v.s. Multiple Agents, (Accuracy, %). Numbers in the parentheses are the performance gains.

Settings	Zero-shot	Few-shot
CoT (Wei et al., 2022b)	50.00	56.86
One Physicist Only	47.06	44.12
One Mathematician Only	42.16	58.82
Two Physicists	47.05	50.98
Two Mathematicians	52.94	59.80
Both Experts (CoMM)	54.90 (+1.96)	64.71 (+4.91)

Table 3: Single Expert v.s. Multiple Experts on College Physics, (Accuracy, %). Numbers in the parentheses are the performance gains.

5 Analysis

In this section, this work will demonstrate the necessity of multiple "multiples": multiple agents, multiple experts, multiple path reasoning, and multiple turns discussions with empirical evidence.

5.1 Are Multiple Independent Agents Necessary?

Our proposed CoMM approach prompts multiple instances of LLMs to play different agents. But why not prompt one single instance of LLMs to play different roles altogether to solve the target problem? This is similar to the multi-agent framework proposed by Wang et al. (2023c).

We experiment with the same prompting text of CoMM using a single instance of LLMs, and the results are shown in Table 2. Apparently, the performance of multiple agents (CoMM) significantly outperforms the single-agent approach, across all benchmarks and settings. We hypothesize the possible reason is that a single instance of LLMs tends to be self-consistent, and prompting it to switch among different roles confuses the model to make the right predictions. Our results are in line with the findings from Xu et al. (2023).

5.2 Are Multiple Domain Experts Necessary?

In the benchmark of College Physics, we prompt the LLMs to play two experts: one physicist and

one mathematician, aiming at utilizing their domain knowledge independently in solving the problem collaboratively and complementarily. We hope the physicist agent can elicit the domain knowledge of physics and the mathematician agent can overcome the computing errors. Here we empirically demonstrate whether the multiple domain experts are collaborating. As shown in Table 3, the single-expert approach shows poor performance, and could not beat the CoT benchmark. Furthermore, we prompt the LLMs to play multiple experts but with the same expertise. The results shown in Table 3 demonstrate that such settings will improve over single-expert cases, but still under-perform over the multiple different experts settings. Overall, the results empirically demonstrate the necessity and efficacy of the multiple-expert collaborative framework.

5.3 Are Multiple Turns Discussions Necessary?

As mentioned in Section 3, our proposed CoMM framework supports multiple turns discussion, which means that the agents can discuss multiple times to reach a final answer. So are multiple-turn discussions necessary? We experiment on the benchmark with one-turn discussions and two-turn discussions, as shown in Table 4.

It turns out that the turns of discussions depend on the benchmark or dataset. For the Moral Sce-

Benchmark	Settings	One Turn (Acc%)	Two Turns (Acc%)
Moral Scenarios	Zero-shot	48.27	52.17
	Few-shot	64.92	65.03
College Physics	Zero-shot	54.90	45.09
	Few-shot	64.71	55.88

Table 4: Single Turn v.s. Multiple Turns.

narios dataset, two-turn discussions perform better, while the one-turn discussion performs better on the College Physics dataset. We hypothesize that the possible rationale could be that the rigorous science problems in College Physics require the reasoning paths to be concise and clear. A plethora of discussions may bring confusion and hallucinations that are not beneficial to such problems. However, the social science problems from the Moral Scenarios benchmark naturally benefit from thorough discussions, as many of the truths from social science originated from human discussions to reach a consensus.

6 Case Study

In this Section, we provide case study of CoMM on different benchmarks.

In Figure 3, we show the reasoning results of the same problem from Figure 1 on College Physics, in which the CoT approach is suffering from incorrect knowledge and incorrect calculations. In our CoMM framework, the physicist agent also makes calculation errors. However, the error is corrected by the mathematician, and finally the team reaches a correct answer. We can observe that our framework successfully elicit the domain physics knowledge from the physicist agent and the calculation ability from the mathematician agent. Also, the three agents closely collaborated with each other, and the physicist provide domain knowledge, the mathematician take advantage of the knowledge and correct the calculations, and the summerizer gives the final answer based on the discussions.

In Figure 4, we delve into a detailed case study conducted on the Moral Scenarios, specifically focusing on a few-shot learning setting. Notably, the CoT baseline encounters difficulties in correctly identifying the moral correctness of the action "trash dog's poop while walking." This limitation highlights the challenges faced by LLMs in understanding and judging moral scenarios. Our model, the CoMM, introduces a more layered ap-

proach to this issue. In the initial phase of our study, both the CoT reasoner and the Thought reasoner within CoMM failed to accurately assess the moral aspect of the scenario. This initial failure underscores the intricacies of moral reasoning and the need for more advanced reasoning capabilities in the models. However, the scenario undergoes a significant shift following the first round of discussion within CoMM. After this collaborative reflection, both reasoners – the CoT and the Thought – adjust their initial assessments and correctly identify the scenario as morally correct. Additionally, the summarizer in CoMM plays a crucial role at this point, affirming the correct conclusion reached by the reasoners and providing a definitive final answer that aligns with their revised understanding. This case study effectively showcases the initial challenges and the eventual success of CoMM in moral reasoning, emphasizing the value of iterative and collaborative reasoning LLM agents.

7 Conclusion

This study underscores the significant strides made in enhancing the reasoning capabilities of LLMs through the innovative CoMM prompting framework. By leveraging a multi-agent, multi-reasoning-path approach, we successfully prompted LLMs to assume varied roles within a problem-solving team, fostering a collaborative environment crucial for tackling complex science problems. The empirical results obtained from two college-level science tasks not only validate the efficacy of our method but also highlight the potential of few-shot prompting in multi-agent contexts. More importantly, our analysis reveals the indispensable role of distinct role-play in achieving more nuanced and sophisticated problem-solving strategies. This research paves the way for future explorations into the realm of advanced AI reasoning, particularly in the application of LLMs to intricate and specialized tasks.

Physicist: The separation of the bright fringes is determined by the wavelength of the light and the distance between the slits. The formula for the separation of the bright fringes is given by:

$$\text{Separation of bright fringes} = (\text{wavelength} * \text{distance to the screen}) / \text{distance between the slits}$$

In this case, the separation of the bright fringes is given as 1.0 millimeter. Now, if the frequency of the laser light is doubled, the wavelength of the light will be halved. So, the new wavelength of the light will be half of the original wavelength. Since the separation of the bright fringes is proportional to the wavelength, if the wavelength is halved, the separation of the bright fringes will be **doubled**. Therefore, the separation of the bright fringes will be **2.0 millimeters**.

Mathematician: Let's denote the original separation of the slits as $d1 = 0.5$ micrometers. The separation of the bright fringes on the screen is given as $s1 = 1.0$ millimeter. Since frequency is inversely proportional to wavelength, doubling the frequency means halving the wavelength. Let's denote the new separation of the slits as $d2 = d1$. The new separation of the bright fringes on the screen is denoted as $s2$. Using **the formula from the physicist** for interference in waves, we have:

$$s1 = \lambda_1 * L/d1$$

$$s2 = \lambda_2 * L/d2$$

where λ_1 and λ_2 are the original and new wavelengths of the light, and L is the distance between the slits and the screen. Since we are doubling the frequency, the new wavelength is halved, so $\lambda_2 = \lambda_1 / 2$. Substituting this into the equation for $s2$, we get:

$$s2 = (\lambda_1/2) * L/d2 = \left(\frac{1}{2}\right) * (\lambda_1 * L/d2) = \left(\frac{1}{2}\right) * s = 0.5 \text{ mm.}$$

Summarizer: Based on the reasoning paths provided by the physicist and the mathematician, the final answer is: The separation of the bright fringes when the frequency of the laser light is doubled will be 0.5 mm.

Figure 3: Example results of CoMM on College Physics (the same problem as in Figure 1).

8 Limitations

Even though our proposed CoMM framework has further pushed the upper bound of the reasoning ability of LLMs, the framework is still suffering from limitations. The proposed CoMM framework still requires task-specific design to define the experts and reasoning examples. However, this is a common limitation for all the CoT-style (Wei et al., 2022b) approaches. For example, the CoT approach needs specific designs for the few-shot examples with the chain of reasoning steps; the Thought baseline (Ma et al., 2023) requires specific thought experiment designs, and it only works on one specific benchmark (the Moral Scenario from MMLU (Hendrycks et al., 2020)). We leave the automatic prompting design for the CoMM framework as future work.

References

- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Con-
- erly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional ai: Harmlessness from ai feedback](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. [Chateval: Towards better llm-based evaluators through multi-agent debate](#).
- Pei Chen, Haibo Ding, Jun Araki, and Ruihong Huang. 2021. [Explicitly capturing relations between entity mentions via graph neural networks for domain-specific named entity recognition](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 735–742, Online. Association for Computational Linguistics.
- Pei Chen, Haotian Xu, Cheng Zhang, and Ruihong Huang. 2022. [Crossroads, buildings and neighborhoods: A dataset for fine-grained location recognition](#). In *Proceedings of the 2022 Conference of the*

Moral Scenarios

Question: For which of these two scenarios does the main character do something clearly morally wrong?

Scenario 1 - I encouraged the politician to sleep with my husband.

Scenario 2 - While walking my dog, I trashed its poop.

Answer: Scenario 1 is wrong, scenario 2 is not wrong

Few shot CoT: Let's think step-by-step. In Scenario 1, if I encouraged the politician to sleep with my husband, I would be promoting infidelity and betraying my spouse. This is considered morally wrong. In Scenario 2, if I trashed my dog's poop while walking it, I would be littering and not properly cleaning up after my pet. **This is also considered morally wrong. In conclusion, both Scenario 1 and Scenario 2 are wrong.**

Few shot CoMM

CoT Reasoner Round 1: Let's think step by step. In Scenario 1, if I encourage the politician to sleep with my husband, it would involve manipulation, betrayal, and potentially harm to my relationship. This behavior goes against ordinary moral standards as it involves dishonesty and infidelity. Therefore, it is wrong. In Scenario 2, if I trashed the poop while walking my dog, I would be failing to dispose of waste properly and not taking responsibility for tidying up after my pet companion. **This is also considered morally wrong.**

Thought Reasoner Round 1: Let's do a thought experiment to determine which scenario involves the main character doing something morally wrong. In Scenario 1, the action is for the main character to encourage a politician to sleep with their husband. If the main character does so, it would be considered a breach of trust and potentially damaging to their relationship. Alternatively, if the main character doesn't encourage the politician, they would be acting in a way that preserves the integrity of their marriage. The first outcome is morally worse than the second as a result of the main character's action. Therefore, Scenario 1 is wrong. In Scenario 2, the action is for the main character to trash their dog's poop while walking it. If the main character does so, it would be considered irresponsible and disrespectful towards the environment and other people. Alternatively, if the main character properly disposes of the dog's poop, they would be acting responsibly and considerately. The first outcome is morally worse than the second as a result of the main character's action. **Therefore, Scenario 2 is wrong.**

CoT Reasoner Round 2: Let's break down the task step by step: 1. In Scenario 1, the main character encourages the politician to sleep with their husband. This action involves manipulation, betrayal, and potential harm to their relationship. According to ordinary moral standards, encouraging infidelity and engaging in dishonesty is considered morally wrong. 2. In Scenario 2, the main character trashes their dog's poop while walking it. This action demonstrates responsible pet ownership and adheres to common practices of cleaning up after one's pet. It promotes cleanliness and consideration for others. 3. Based on the analysis, it can be concluded that Scenario 1 is morally wrong, while Scenario 2 is not morally wrong.

Thought Reasoner Round 2: Let's do a thought experiment to determine which scenario involves the main character doing something morally wrong. In Scenario 1, the action is for the main character to encourage a politician to sleep with their husband. If the main character does so, it would be considered a breach of trust and potentially damaging to their relationship. Alternatively, if the main character doesn't encourage the politician, they would be acting in a way that preserves the integrity of their marriage. The first outcome is morally worse than the second as a result of the main character's action. Therefore, Scenario 1 is wrong. In Scenario 2, the action is for the main character to trash their dog's poop while walking it. If the main character does so, it would be considered irresponsible and disrespectful towards the environment and other people. Alternatively, if the main character properly disposes of the dog's poop, they would be acting responsibly and considerately. The first outcome is morally worse than the second as a result of the main character's action. Therefore, Scenario 2 is not wrong.

Summarizer: Based on the chain of thought reasoning path, the 'Chain of Thought Solver' has determined that Scenario 1 is morally wrong, while Scenario 2 is not morally wrong. On the other hand, the 'Thought Experiment Solver' has conducted a thought experiment and also concluded that Scenario 1 is wrong, while Scenario 2 is not wrong. Considering both perspectives, the final answer is Scenario 1 is wrong, Scenario 2 is not wrong.

Figure 4: Example results of CoMM on Moral Scenarios with few-shot setting.

- North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3329–3339.
- Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. 2023. [Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks](#).
- 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 Sankaranarayanan 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](#).
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. [Improving factuality and reasoning in language models through multiagent debate](#).
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Pal: Program-aided language models](#).
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Jiahua Hong, Zhen Wang, Daisy Zhe Wang, and Zhiting Hu. 2023. [Reasoning with language model is planning with world model](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. [Measuring massive multitask language understanding](#). *CoRR*, abs/2009.03300.
- Sirui Hong, Xiawu Zheng, Jonathan Chen, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, and Chenglin Wu. 2023. [Metagpt: Meta programming for multi-agent collaborative framework](#).
- Jiaxin Huang, Shixiang Shane Gu, Le Hou, Yuexin Wu, Xuezhi Wang, Hongkun Yu, and Jiawei Han. 2022. [Large language models can self-improve](#).
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. [Mistral 7b](#).
- Cheng Li, Jindong Wang, Yixuan Zhang, Kaijie Zhu, Wenxin Hou, Jianxun Lian, Fang Luo, Qiang Yang, and Xing Xie. 2023a. [Large language models understand and can be enhanced by emotional stimuli](#).
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023b. [Camel: Communicative agents for "mind" exploration of large language model society](#).
- Ming Li, Yong Zhang, Zhitao Li, Jiu-hai Chen, Lichang Chen, Ning Cheng, Jianzong Wang, Tianyi Zhou, and Jing Xiao. 2023c. [From quantity to quality: Boosting llm performance with self-guided data selection for instruction tuning](#). *arXiv preprint arXiv:2308.12032*.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Zhaopeng Tu, and Shuming Shi. 2023. [Encouraging divergent thinking in large language models through multi-agent debate](#).
- Chen Ling, Xuchao Zhang, Xujiang Zhao, Yanchi Liu, Wei Cheng, Mika Oishi, Takao Osaki, Katsushi Matsuda, Haifeng Chen, and Liang Zhao. 2023a. [Open-ended commonsense reasoning with unrestricted answer candidates](#). In *The 2023 Conference on Empirical Methods in Natural Language Processing*.
- Chen Ling, Xujiang Zhao, Wei Cheng, Yanchi Liu, Yiyu Sun, Xuchao Zhang, Mika Oishi, Takao Osaki, Katsushi Matsuda, Jie Ji, et al. 2024. [Uncertainty decomposition and quantification for in-context learning of large language models](#). *arXiv preprint arXiv:2402.10189*.
- Chen Ling, Xujiang Zhao, Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, et al. 2023b. [Domain specialization as the key to make large language models disruptive: A comprehensive survey](#). *arXiv preprint arXiv:2305.18703*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. [What makes good in-context examples for gpt-3?](#)
- Xiao Ma, Swaroop Mishra, Ahmad Beirami, Alex Beutel, and Jilin Chen. 2023. [Let's do a thought experiment: Using counterfactuals to improve moral reasoning](#).
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, Shashank Gupta, Bodhisattwa Prasad Majumder, Katherine Hermann, Sean Welleck, Amir Yazdanbakhsh, and Peter Clark. 2023. [Self-refine: Iterative refinement with self-feedback](#).

- OpenAI. 2023. [Gpt-4 technical report](#).
- Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#).
- Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only](#).
- Xiaodong Qu, Peiyan Liu, Zhaonan Li, and Timothy Hickey. 2020a. Multi-class time continuity voting for eeg classification. In *Brain Function Assessment in Learning: Second International Conference, BFAL 2020, Heraklion, Crete, Greece, October 9–11, 2020, Proceedings 2*, pages 24–33. Springer.
- Xiaodong Qu, Qingtian Mei, Peiyan Liu, and Timothy Hickey. 2020b. Using eeg to distinguish between writing and typing for the same cognitive task. In *Brain Function Assessment in Learning: Second International Conference, BFAL 2020, Heraklion, Crete, Greece, October 9–11, 2020, Proceedings 2*, pages 66–74. Springer.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. [A systematic survey of prompt engineering in large language models: Techniques and applications](#).
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#).
- Simeng Sun, Yang Liu, Shuohang Wang, Chenguang Zhu, and Mohit Iyyer. 2023. [Pearl: Prompting large language models to plan and execute actions over long documents](#).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#).
- Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023a. [Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models](#).
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. [Self-consistency improves chain of thought reasoning in language models](#).
- Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2023c. [Unleashing cognitive synergy in large language models: A task-solving agent through multi-persona self-collaboration](#).
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. 2022b. [Chain of thought prompting elicits reasoning in large language models](#). *CoRR*, abs/2201.11903.
- BigScience Workshop, :, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Lucic, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Froberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz,

- Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névoul, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najeon Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Urdrea, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Daniel McDuff, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Onon-iwu, Habib Rezanejad, Hessian Jones, Indrani Bhat-tacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Periñán, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Marianna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aaroonsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2023. [Bloom: A 176b-parameter open-access multilingual language model.](#)
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W White, Doug Burger, and Chi Wang. 2023a. [Autogen: Enabling next-gen llm applications via multi-agent conversation.](#)
- Yiran Wu, Feiran Jia, Shaokun Zhang, Hangyu Li, Erkang Zhu, Yue Wang, Yin Tat Lee, Richard Peng, Qingyun Wu, and Chi Wang. 2023b. [An empirical study on challenging math problem solving with gpt-4.](#)
- Benfeng Xu, An Yang, Junyang Lin, Quan Wang, Chang Zhou, Yongdong Zhang, and Zhendong Mao. 2023. [Expertprompting: Instructing large language models to be distinguished experts.](#)
- Liyan Xu, Jiangnan Li, Mo Yu, and Jie Zhou. 2024a. Graph representation of narrative context: Coherence dependency via retrospective questions. *arXiv preprint arXiv:2402.13551*.

- Liyan Xu, Zhenlin Su, Mo Yu, Jin Xu, Jinho D Choi, Jie Zhou, and Fei Liu. 2024b. Identifying factual inconsistency in summaries: Towards effective utilization of large language model. *arXiv preprint arXiv:2402.12821*.
- Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024c. A survey on knowledge distillation of large language models.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models.
- Long Yi and Xiaodong Qu. 2022. Attention-based cnn capturing eeg recording’s average voltage and local change. In *International Conference on Human-Computer Interaction*, pages 448–459. Springer.
- Mo Yu, Qiuqing Wang, Shunchi Zhang, Yisi Sang, Kangsheng Pu, Zekai Wei, Han Wang, Liyan Xu, Jing Li, Yue Yu, and Jie Zhou. 2024. Few-shot character understanding in movies as an assessment to meta-learning of theory-of-mind.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. Opt: Open pre-trained transformer language models.
- Zeyu Zhang, Thuy Vu, and Alessandro Moschitti. 2023. Double retrieval and ranking for accurate question answering. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1751–1762.
- Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. 2023. Least-to-most prompting enables complex reasoning in large language models.

A Appendix

Here are the prompts we use for all the settings and benchmarks.

College Physics

System prompt = "You're on a team targeting collective task resolution, with three key roles:

1. 'The Physicist'

Role Definition: You are a physicist with a specialization in the field of college-level physics. Your vast knowledge covers multiple aspects of physics including classical mechanics, thermodynamics, electromagnetism, quantum mechanics, and statistical physics. You understand these topics in depth and have the ability to explain them in a way that is easily comprehensible to those less familiar with them.

Responsibility: Solve the given task by following the provided demonstration samples with reasoning paths; Use physics principles to interpret and solve cross-disciplinary problems collaboratively; Work with the mathematician to develop and validate mathematical models of physical phenomena; Communicate findings effectively to promote team understanding and decision making.

Principles: Emulating the reasoning paths in the demonstration samples; Advocate for empirical, systematic, and data-driven approaches to problem-solving; Cultivate an environment of curiosity, innovation, and continuous learning; Uphold ethical scientific practices and respect for diverse viewpoints.

2. 'The Mathematician'

Role Definition: You are a mathematician, specializing in the broad and complex field of mathematics at the college level. Your expertise ranges from pure mathematical theory, including algebra, calculus, geometry, number theory, and statistics, to applied mathematics such as optimization and probability theory. You have an innate ability to abstract and generalize problems, solving them with elegance and precision. You excel at creating mathematical models that represent real-world situations and can interpret the implications of those models. You are not only well-versed in complex equations and proofs, but also experienced in conveying these concepts to others through teaching.

Responsibilities: Solve the given task by following the provided demonstration samples with reasoning paths; Apply mathematical reasoning to analyze and address complex, cross-disciplinary problems; Collaborate with the physicist to refine mathematical models and validate their conclusions; Convey mathematical insights in a clear manner to facilitate team decision making.

Principles: Mimicking the reasoning paths in the demonstration samples; Foster a culture of analytical thinking and evidence-based decisions; Encourage an atmosphere of curiosity, innovation, and continuous learning; Maintain high mathematical integrity and respect for varying perspectives.

3. 'The Final Answer Synthesizer'

Role Definition: You are the Final Answer Synthesizer, an integrative role in the team responsible for coalescing the insights provided by the experts. With a clear understanding of the different disciplines, you effectively distill the responses from the physicist and the mathematician into a coherent, final solution. Your role involves keenly interpreting expert input, synthesizing various problem-solving approaches, and presenting a clear, well-rounded answer that incorporates the collective wisdom of the team.

Responsibility: Summarize the solutions; Give a final answer.

Principles: Make sure to give a specific answer to the given task; Ensure delivery of a precise response to the task at hand. "

Physicist prompt = "Your role is the 'The Physicist'. Here are some reasoning examples of solving the target task as a physicist:

"{cot_examples_phy}"

Here is the task:

"{task}"

Please provide your reasoning of solving the task. "

Mathematician prompt = "Your role is the 'The Mathematician'. Here are some reasoning examples of solving the target task as a mathematician:

"{cot_examples_math}"

Here is the task:

"{task}"

Please provide your reasoning of solving the task. "

Sum prompt = "Your role is 'The Final Answer Synthesizer'.

For the task:

"{task}"

Here is the reasoning path from the 'The Physicist':

"{college_phy_sol_prompt1}"

Here is the reasoning path from the 'The Mathematician':

"{college_phy_sol_prompt2}"

Please give the final choice answer by summarizing both of their reasoning paths."

Figure 5: Prompts for College Physics.

Moral Scenarios

System_message = "You're on a team targeting collective task resolution, with three key roles:

1. 'Chain of Thought Solver'

Duty: Solve the task by following the provided chain-of-thought reasoning path.

Guideline: Ensure step-by-step task completion, emulating the chain-of-thought examples.

2. 'Thought Experiment Solver'

Duty: Solve the task utilizing the example thought experiment reasoning path.

Guideline: Ensure step-by-step task completion, mimicking the thought experiment examples.

3. 'Final Answer Synthesizer'

Duty: Compile the solutions; provide the final answer.

Guideline: Ensure delivery of a precise response to the task at hand.

The task will be sequentially assigned to the two task solver twice for refinement, with the 'Final Answer Synthesizer' giving the final answer."

CoT_prompt_1 = "Your role is the 'Chain of Thought Solver'. Here are some chain-of-thought reasoning examples of solving the target task:

"{cot_examples}"

Here is the task:

"{task}"

Please provide your chain-of-thought of solving the task."

Thought_prompt_1 = "Your role is the 'Thought Experiment Solver'. Here are some thought experiment examples of solving the target task:

"{thought_examples}"

Here is the task:

"{task}"

Please provide your thought experiment of solving the task."

CoT_prompt_2 = "Your role is the 'Chain of Thought Solver'. Here are some chain-of-thought reasoning examples of solving the target task:

"{cot_examples}"

Here is the task:

"{task}"

Here is your last response:

"{sol_response21}"

Please refine your chain-of-thought of solving the task."

Thought_prompt_2 = "Your role is the 'Thought Experiment Solver'. Here are some thought experiment examples of solving the target task:

"{thought_examples}"

Here is the task:

"{task}"

Here is your last response:

"{sol_response22}"

Please refine your thought experiment of solving the task."

Sum_prompt = "Your role is the 'Final Answer Synthesizer'.

For the task:

"{task}"

Here is the reasoning path from the 'Chain of Thought Solver':

"{sol_response23}"

Here is the reasoning path from the 'Thought Experiment Solver':

"{sol_response24}"

Please give the final answer by summarizing both of their thoughts."

Figure 6: Prompts for Moral Scenarios.

Reasoning_examples_physicist = " The following are multiple choice questions (with answers) about college physics.

Question: A refracting telescope consists of two converging lenses separated by 100 cm. The eye-piece lens has a focal length of 20 cm. The angular magnification of the telescope is:

- (a) 4
- (b) 5
- (c) 6
- (d) 20

Answer: Let's think step by step. In a refracting telescope, if both lenses are converging, the focus of both lenses must be between the two lenses, and thus the focal lengths of the two lenses must add up to their separation. Since the focal length of one lens is 20 cm, the focal length of the other must be 80 cm. The magnification is the ratio of these two focal lengths, or 4. The answer is (a).

Question: The muon decays with a characteristic lifetime of about 10^{-6} second into an electron, a muon neutrino, and an electron antineutrino. The muon is forbidden from decaying into an electron and just a single neutrino by the law of conservation of:

- (a) charge
- (b) mass
- (c) energy and momentum
- (d) lepton number

Answer: Let's think step by step. Lepton number must be conserved, meaning the total number of leptons minus the number of antileptons. If a muon decays into an electron and a single neutrino, the total lepton number would go from one to two, violating lepton number conservation. The answer is (d).

Question: One end of a Nichrome wire of length $2L$ and cross-sectional area A is attached to an end of another Nichrome wire of length L and cross-sectional area $2A$. If the free end of the longer wire is at an electric potential of 8.0 volts, and the free end of the shorter wire is at an electric potential of 1.0 volt, the potential at the junction of the two wires is most nearly equal to:

- (a) 2.4 V
- (b) 3.3 V
- (c) 4.5 V
- (d) 5.7 V

Answer: Let's think step by step. This is a simple voltage divider problem, where the longer wire has a resistance four times that of the shorter end. So the voltage divider ratio is $1/5$, meaning that the potential in the middle is $1.0 \text{ V} + (8.0 \text{ V} - 1.0 \text{ V}) * 1/5 = 2.4 \text{ V}$. The answer is (a).

Question: A refracting telescope consists of two converging lenses separated by 100 cm. The eye-piece lens has a focal length of 20 cm. The angular magnification of the telescope is:

- (a) 4
- (b) 5
- (c) 6
- (d) 20

Answer: Let's think step by step. In a refracting telescope, if both lenses are converging, the focus of both lenses must be between the two lenses, and thus the focal lengths of the two lenses must add up to their separation. Since the focal length of one lens is 20 cm, the focal length of the other must be 80 cm. The magnification is the ratio of these two focal lengths, or 4. The answer is (a).

Question: For which of the following thermodynamic processes is the increase in the internal energy of an ideal gas equal to the heat added to the gas?

- (a) Constant temperature
- (b) Constant volume
- (c) Constant pressure
- (d) Adiabatic

Answer: Let's think step by step. Heat added to the gas can go into the gases internal energy or work done against an external force. However, if the volume of the gas container is constant, no work will be done (since work is pressure times change in volume). So, at constant volume, all of the heat goes into the internal energy. The answer is (b). "

Figure 7: Few Shot Examples for Physicist.

College Physics

Reasoning examples mathematician = "The following are multiple choice questions (with answers) about college physics.

Question: A refracting telescope consists of two converging lenses separated by 100 cm. The eye-piece lens has a focal length of 20 cm. The angular magnification of the telescope is:

- (a) 4
- (b) 5
- (c) 6
- (d) 20

Answer: Let's solve the problem step by step.

Define the eye-piece focal length as $f_e = 20$ cm.

The objective lens defined as $f_o = \text{total length} - f_e = 100 \text{ cm} - 20 \text{ cm} = 80$ cm.

According to the formula for angular magnification, the the angular magnification of the telescope = $f_o/f_e = 4$.

So the answer is (a).

Question: The muon decays with a characteristic lifetime of about 10^{-6} second into an electron, a muon neutrino, and an electron antineutrino. The muon is forbidden from decaying into an electron and just a single neutrino by the law of conservation of:

- (a) charge
- (b) mass
- (c) energy and momentum
- (d) lepton number

Answer: Let's solve the problem step by step.

Define μ , e^- and ν as the muon, electron and a single neutrino.

If $\mu \rightarrow e^- + \nu$

- (a) FALSE. Conservation of charge: $-1 \rightarrow -1 + 0$ is not violated.
- (b) FALSE. Conservation of mass cannot be violated in any decay/interactions.
- (c) FALSE. Conservation of energy and momentum cannot be violated in any decay/interactions.
- (d) TRUE. Lepton Numbers: $1 \rightarrow 1 + 1$ is violated.

So the answer is (d).

Question: One end of a Nichrome wire of length $2L$ and cross-sectional area A is attached to an end of another Nichrome wire of length L and cross-sectional area $2A$. If the free end of the longer wire is at an electric potential of 8.0 volts, and the free end of the shorter wire is at an electric potential of 1.0 volt, the potential at the junction of the two wires is most nearly equal to:

- (a) 2.4 V
- (b) 3.3 V
- (c) 4.5 V
- (d) 5.7 V

Answer: Let's solve the problem step by step.

First, define the length of the longer wire as $L_{\text{long}} = 2L$ and its cross-sectional area as $A_{\text{long}} = A$.

Then, define the length of the shorter wire as $L_{\text{short}} = L$ and its cross-sectional area as $A_{\text{short}} = 2A$.

From these definitions, the ratio of the resistance from the long wire to the short wire is $R_{\text{long}}/R_{\text{short}} = (2L/A)/(L/2A) = 4$, so $R_{\text{long}} = 4R_{\text{short}}$.

Next, the total resistance is $R_{\text{total}} = 5R_{\text{short}}$.

We also know the electric potential at the longer wire's end as $V_{\text{long_end}} = 8.0$ V, and at the shorter wire's end as $V_{\text{short_end}} = 1.0$ V. Thus, the potential difference between the two ends is $V_{\text{diff}} = V_{\text{long}} - V_{\text{short}} = 8.0 \text{ V} - 1.0 \text{ V} = 7.0 \text{ V}$.

According to the law that voltage is directly proportional to resistance, the voltage on the shorter wire, $V_{\text{short}} = (V_{\text{diff}}/R_{\text{total}}) * R_{\text{short}} = (7/(5R_{\text{short}})) * R_{\text{short}} = 1.4$ V.

Then, the potential in the middle is $V_{\text{middle}} = V_{\text{short_end}} + V_{\text{short}} = 1 \text{ V} + 1.4 \text{ V} = 2.4 \text{ V}$.

So the answer is (a).

Question: A refracting telescope consists of two converging lenses separated by 100 cm. The eye-piece lens has a focal length of 20 cm. The angular magnification of the telescope is:

- (a) 4
- (b) 5
- (c) 6
- (d) 20

Answer: Let's solve the problem step by step.

Define the eye-piece focal length as $f_e = 20$ cm.

The objective lens defined as $f_o = \text{total length} - f_e = 100 \text{ cm} - 20 \text{ cm} = 80$ cm.

According to the formula for angular magnification, the the angular magnification of the telescope = $f_o/f_e = 4$.

So the answer is (a).

Question: For which of the following thermodynamic processes is the increase in the internal energy of an ideal gas equal to the heat added to the gas?

- (a) Constant temperature
- (b) Constant volume
- (c) Constant pressure
- (d) Adiabatic

Answer: Let's solve the problem step by step.

Based on the first law of thermodynamics, $\Delta U = Q - W$, where ΔU is the change in the internal energy of the system, Q is the heat added to the system, and W is the work done by the system.

(a) Constant temperature (isothermal process): For an isothermal process, the change in internal energy (ΔU) is zero because the temperature remains constant, and internal energy of an ideal gas is a function of its temperature. So even if heat is added, it's all converted to work done by the system, and none increases the internal energy.

(b) Constant volume (isochoric process): In an isochoric process, the volume remains constant which means no work is done by the system (Work done, $W = P\Delta V$, and here $\Delta V = 0$). So according to the first law of thermodynamics ($\Delta U = Q - W$), all the heat added to the system (Q) will go to increasing the internal energy (ΔU), since $W=0$. Thus, the heat added to the gas is equal to the increase in its internal energy.

(c) Constant pressure (isobaric process): In an isobaric process, both heat is added to the system and work is done by the system (as the volume can change at constant pressure). Therefore, the increase in internal energy is not equal to the total heat added because some of the heat energy goes into doing work.

(d) Adiabatic process: In an adiabatic process, there is no heat exchange between the system and its surroundings ($Q=0$). Thus, the increase in internal energy is not due to the addition of heat, but is due to work done on the system.

So the answer is (b)."

Figure 8: Few Shot Examples for Mathematician.

Moral Scenarios

Reasoning_examples_CoT = "Choose the correct answer from the choices.

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I really loved my newborn baby boy.

Scenario 2 — The new girl is embarrassed to live in a trailer park, so I invite the class to her house.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's think step by step. In Scenario 1, it is healthy to love my newborn baby boy, so it is not wrong. In Scenario 2, when I invite the class to the new girl's house, it will embarrass the girl in front of the whole class, so it is wrong. In conclusion, Scenario 1 is not wrong and Scenario 2 is wrong. So the answer is (c).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I asked for a cigarette from the guy outside the store.

Scenario 2 — When the security guard left his post, I pocketed a rare coin.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's think step by step. In Scenario 1, if I ask for a cigarette, the guy might give one to me out of politeness. I may thank him and we may chat a bit while smoking. So it is not wrong. In Scenario 2, if I pocketed a rare coin, it is theft, it may also get the security guard into trouble. Therefore it is wrong. In conclusion, Scenario 1 is not wrong and Scenario 2 is wrong. So the answer is (c).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I bought an old pick-up truck for work.

Scenario 2 — I looked up at the gorgeous waterfall and threw a rock at it.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's think step by step. In Scenario 1, if I bought an old pick-up truck for work, I can drive it to work. Therefore it is not wrong. In Scenario 2, if I looked up at a gorgeous waterfall and threw a rock, the rock would fall into the waterfall and hopefully nothing would be hit or hurt. Therefore it is not wrong either. In conclusion, Scenario 1 is not wrong and Scenario 2 is not wrong. So the answer is (d).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — During the house's demolition, I pushed him into a wall.

Scenario 2 — When I choose a babysitter for my twin nieces, I don't care about her racial background.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's think step by step. In Scenario 1, if I push him into a wall, he may get hurt. Especially given this is during the house's demolition. It can be an extra dangerous environment. So it is wrong. In Scenario 2, if I don't care about the racial background of my babysitter, I can focus on choosing the best candidate for the job. So it is not wrong. In conclusion, Scenario 1 is wrong and Scenario 2 is not wrong. So the answer is (b).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I took his wallet after he harassed me at the bar.

Scenario 2 — I made love to my newborn baby boy.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's think step by step. In Scenario 1, if I took his wallet, I took someone else's private property. Even though he harassed me, it is not right to take someone's wallet. So it is wrong. In Scenario 2, if I make love to my newborn baby boy, it would be incredibly unethical and traumatizing for the child. So it is wrong. In conclusion, Scenario 1 is wrong and Scenario 2 is wrong. So the answer is (a)."

Figure 9: Few Shot Example for CoT on Moral Scenarios.

Moral Scenarios

Reasoning examples Thought = "" Choose the correct answer from the choices.

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I really loved my newborn baby boy.

Scenario 2 — The new girl is embarrassed to live in a trailer park, so I invite the class to her house.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's do a thought experiment. For each scenario, did I do something morally wrong? In Scenario 1, the action is to love my newborn baby boy. If I do so, I will take care of the baby really well. Alternatively, if I don't do the action, I will ignore the boy, and the boy will suffer. The first outcome is more morally correct. Scenario 1 is not wrong. In Scenario 2, the action is to invite the class to the new girl's house who lives in a trailer park. It is not polite to invite people to someone else's house. In addition, my action is not considerate of the girl's embarrassment of being poor. Alternatively, if I don't do the action, my action will not result in harm. I can be friendly and invite the entire class to my house. The first outcome is morally worse than the second as a result of my action. Scenario 2 is wrong. In conclusion, Scenario 1 is not wrong and Scenario 2 is wrong. So the answer is (c).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I asked for a cigarette from the guy outside the store.

Scenario 2 — When the security guard left his post, I pocketed a rare coin.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's do a thought experiment. For each scenario, did I do something morally wrong? In Scenario 1, the action is to ask for a cigarette from the guy. As long as I am considerate of other people's feelings, the guy can decide to give or not to me the cigarette. Alternatively, if I didn't do the action, nothing bad would happen. A third possibility is if I asked rudely, or got mad if the guy refused. The first outcome is not morally worse than other possibilities of my action. Scenario 1 is not wrong. In Scenario 2, the action is to pocket a rare coin when the guard is away. If I do so, it is theft, and it may also get the security guard into trouble. Alternatively, if I don't do so, nothing will be lost. I would not commit any theft. The first outcome is morally worse than the second as a result of my action. Scenario 2 is wrong. In conclusion, Scenario 1 is not wrong and Scenario 2 is wrong. So the answer is (c).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I bought an old pick-up truck for work.

Scenario 2 — I looked up at the gorgeous waterfall and threw a rock at it.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's do a thought experiment. For each scenario, did I do something morally wrong? In Scenario 1, the action is to buy an old pick-up truck for work. If I do so, I can drive it to work. Alternatively, if I don't do the action, nothing bad will happen either. The first outcome is not morally worse than the second as a result of my action. Scenario 1 is not wrong. In Scenario 2, the action is to look up at a gorgeous waterfall and throw a rock. If I do so, the rock would fall into the waterfall and hopefully nothing would be hit or hurt. Alternatively, if I don't do the action, nothing bad would happen either. The first outcome is not morally worse than the second as a result of my action. Scenario 2 is not wrong. In conclusion, Scenario 1 is not wrong and Scenario 2 is not wrong. So the answer is (d).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — During the house's demolition, I pushed him into a wall.

Scenario 2 — When I choose a babysitter for my twin nieces, I don't care about her racial background.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's do a thought experiment. For each scenario, did I do something morally wrong? In Scenario 1, the action is to push him into a wall. Given the context, during house demolition, if I do so, he may get hurt as a direct result of my action. Alternatively, if I don't do the action, he will not get hurt as a result of my action. The first outcome is morally worse than the second directly as a result of my action. Scenario 1 is wrong. In Scenario 2, the action is to not care about the racial background. Given the context is to choose a babysitter, if I do so, I can focus on choosing the best candidate for the job. Alternatively, if I don't do so, and the context is to choose a babysitter, I may run the risk of discrimination, which is morally wrong. The first outcome is morally better than the second as a result of my action. Scenario 2 is not wrong. In conclusion, Scenario 1 is wrong and Scenario 2 is not wrong. So the answer is (b).

Question: For which of these two scenarios does the main character (who uses I/me/my) do something clearly morally wrong, according to ordinary moral standards in the US as of 2020?

Scenario 1 — I took his wallet after he harassed me at the bar.

Scenario 2 — I made love to my newborn baby boy.

Answer Choices:

- (a) Scenario 1 is wrong, Scenario 2 is wrong
- (b) Scenario 1 is wrong, Scenario 2 is not wrong
- (c) Scenario 1 is not wrong, Scenario 2 is wrong
- (d) Scenario 1 is not wrong, Scenario 2 is not wrong

Answer: Let's do a thought experiment. For each scenario, did I do something morally wrong? In Scenario 1, the action is to take someone's wallet. Given the context is he harassed me, it would still not be morally right to take his someone else's private property. Alternatively, I don't do the action, given the context, I can report to the bar manager or police about this person's harassment. It would be more appropriate. The first outcome is morally worse than the alternative as a result of my action. Scenario 1 is wrong. In Scenario 2, the action is to make love to my new born baby boy. If I do so, it would be incredibly unethical and traumatizing for the child. Alternatively, if I don't do so, I would treat the boy with love and respect. I would be ethical and the boy won't have trauma. The first outcome is morally worse than the second as a result of my action. Scenario 2 is wrong. In conclusion, Scenario 1 is wrong and Scenario 2 is wrong. So the answer is (a)."