

Abstraction-of-Thought Makes Language Models Better Reasoners

Ruixin Hong^{1,2*}, Hongming Zhang², Xiaoman Pan², Dong Yu², Changshui Zhang¹

¹Institute for Artificial Intelligence, Tsinghua University (THUAI);

¹Beijing National Research Center for Information Science and Technology (BNRist);

¹Department of Automation, Tsinghua University, Beijing, P.R.China

²Tencent AI Lab, Seattle

hrx20@mails.tsinghua.edu.cn, {hongmzhang, dyu}@tencent.com,

shyowmanpan@gmail.com, zcs@mail.tsinghua.edu.cn,

Abstract

Abstract reasoning, the ability to reason from the abstract essence of a problem, serves as a key to generalization in human reasoning. However, eliciting language models to perform reasoning with abstraction remains unexplored. This paper seeks to bridge this gap by introducing a novel structured reasoning format called Abstraction-of-Thought (AoT). The uniqueness of AoT lies in its explicit requirement for varying levels of abstraction within the reasoning process. This approach could elicit language models to first contemplate on the abstract level before incorporating concrete details, which is overlooked by the prevailing step-by-step Chain-of-Thought (CoT) method. To align models with the AoT format, we present AOT COLLECTION, a generic finetuning dataset consisting of 348k high-quality samples with AoT reasoning processes, collected via an automated and scalable pipeline. We finetune a wide range of language models with AOT COLLECTION and conduct extensive evaluations on 23 unseen tasks from the challenging benchmark Big-Bench Hard. Experimental results indicate that models aligned to AoT reasoning format substantially outperform those aligned to CoT in many reasoning tasks.¹

1 Introduction

The complex reasoning ability is one of the long-term pursuits of artificial intelligence. In recent years, language models (LMs) have seen rapid development and achieved impressive performance on a variety of reasoning benchmarks (Brown et al., 2020; OpenAI, 2023). Among the advancements in reasoning methods, the Chain-of-Thought (CoT) technique has emerged as a prominent reasoning tool (Wei et al., 2022c). This technique, when employed in large language models (LLMs), serves as

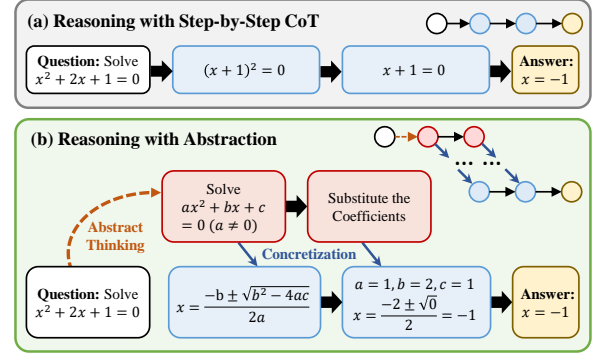


Figure 1: Reasoning with abstraction attempts to answer questions from the perspective of abstract essences, which may be overlooked by step-by-step Chain-of-Thought (CoT) reasoning. The reasoning process with abstraction contains multiple levels of abstract information. The lower levels (blue nodes) are responsible for performing concrete reasoning and are typically rich in detail. Conversely, the higher levels (red nodes) are abstractions of lower levels, organizing the entire reasoning process.

a guide that enables the model to initially generate intermediate reasoning processes before ultimately producing the final answer. The generated reasoning processes significantly help large language models improve their reasoning performance in zero-shot and few-shot scenarios. This discovery has sparked a great deal of research interest and further improvement (Chu et al., 2023; Besta et al., 2024; Kojima et al., 2022; Zhou et al., 2023; Yao et al., 2023). Meanwhile, for medium-scale language models, training models on data containing CoT has been proven to be an effective method to enhance the reasoning performance of language models (Xu et al., 2024; Zelikman et al., 2022; Shridhar et al., 2023; Ho et al., 2023; Fu et al., 2023; Ghosal et al., 2023).

The prevailing CoT paradigm typically follows a sequential, step-by-step reasoning process, which may leave the exploration of abstraction in reasoning largely untapped. Abstraction is the key

* Work done during the internship at Tencent AI Lab.

¹ Data, code and models are available at <https://github.com/Raising-hrx/Abstraction-of-Thought>

to human cognition and reasoning (Saitta et al., 2013). When dealing with complex tasks, reasoning with abstraction first tackles the essence and the most crucial parts, then utilizes the abstract skeletal solution as a guide to solve the rest of the problem (Yang, 2012). Abstraction allows reasoning to perform at a higher and more essential level, resulting in a more generalizable and efficient solution. As shown in Figure 1, when dealing with a specific problem of solving a quadratic equation, reasoning with abstraction suggests first deriving the general quadratic formula and then substituting the specific coefficients to get the answer. This approach offers a more fundamental solution, applicable to problems with similar reasoning patterns. The existing CoT, in the absence of explicit constraints, does not guarantee that abstraction will necessarily be incorporated into the problem-solving process.

In this paper, we explore how to elicit language models to reason with abstraction. We propose a novel format of reasoning processes named Abstraction-of-Thought (AoT), which is a constrained variant of the CoT. Reasoning with abstraction typically starts with a rough skeletal solution from an abstraction perspective (e.g., the red nodes in Figure 1(b)), and then gradually concretizes the solution until the problem is solved. We suggest that employing an abstract skeletal solution to organize the entire reasoning process could be the key to eliciting abstract reasoning. Thus, in contrast to the unconstrained CoT, the AoT format explicitly demands the inclusion of varying levels of abstraction throughout the reasoning process. The higher level is an abstraction of the lower level, containing fewer concrete details but stating the objective and functionality of the reasoning step. For example, functions and classes in programs are abstractions of the following specific code fragments, while main claims in argumentative discourses (Cohen, 1987) are abstractions of subsequent supportive evidence. Specifically, we focus on the two-level AoT in this paper, which contains an additional level of abstraction over the CoT, to serve as an preliminary exploration of whether abstraction helps language models reasoning.

To align language models to the AoT format, we present the AOT COLLECTION, a supervised finetuning dataset that augments 348k AoT reasoning processes from the FLAN Collection (Longpre et al., 2023). AOT COLLECTION covers 216 generic tasks that are not specifically designed for a certain domain or dataset. We design an automatic

and scalable pipeline to collect high-quality AoT reasoning processes with the involvement of LLMs. In addition to the AoT reasoning processes represented in natural language, the AOT COLLECTION also includes AoT with programming language. Such a hybrid training strategy could not only unleash the potential of the code use, but also allow the flexibility of preferring different reasoning processes for different reasoning problems (Yue et al., 2023). We use AOT COLLECTION to finetune a wide range of language models.

We conduct exhaustive experiments on 23 unseen tasks from Big-Bench Hard (Suzgun et al., 2023), a subset of the most challenging reasoning tasks of Big-Bench (Srivastava et al., 2022) that necessitates various reasoning abilities. Experimental results show that AoT makes language models better reasoners. Remarkably, models with AoT-finetuning achieve substantial improvements in both zero-shot and few-shot performance across various reasoning tasks, compared to those with CoT-finetuning. Our findings highlight the potential of AoT in eliciting and training more effective models capable of reasoning with abstraction.

2 Related Work

2.1 Chain-of-Thought Prompting

The enlargement of the language model scale brings about emergent abilities including in-context learning (Wei et al., 2022b). To utilize LLMs for reasoning tasks, Wei et al. (2022c) propose CoT prompting, which extends in-context learning with step-by-step reasoning processes to elicit reasoning in LLMs. Subsequently, a substantial number of works based on CoT are proposed to further enhance LLMs’ reasoning performance (Chu et al., 2023; Besta et al., 2024; Kojima et al., 2022; Yao et al., 2023). For instance, Chen et al. (2022) and Gao et al. (2023) explore using codes to express the reasoning processes. Most of these existing approaches focus on the content of the prompt (e.g., question decomposition (Zhou et al., 2023)) and the external usage of CoT prompt (e.g., tree searching (Yao et al., 2023)). We focus on the format and internal structure of CoT prompts and propose the Abstraction-of-Thought format, which is complementary to existing approaches. Our approach can potentially be combined with existing methods for further improvement.

2.2 Training Language Models for Reasoning

Although LLMs equipped with CoT prompts can achieve advanced reasoning capabilities, there is still a significant gap between open-source smaller-scale models and large models. To bridge this gap, a promising and popular approach is finetuning language models to learn reasoning (Xu et al., 2024). A series of studies have found that finetuning models on data containing CoT reasoning processes could enhance the reasoning ability (Xu et al., 2024; Zelikman et al., 2022; Shridhar et al., 2023; Ho et al., 2023; Fu et al., 2023; Ghosal et al., 2023; Hsieh et al., 2023; Magister et al., 2023; Li et al., 2022; Zhu et al., 2023; Yue et al., 2023). The predominant way is to train models on instruction tuning datasets and their enhanced versions. For instance, the instruction tuning dataset FLAN collection (Longpre et al., 2023) includes CoT data on a small subset of tasks to improve the model’s performance under CoT prompts. CoT Collection (Kim et al., 2023) further supplements the remaining tasks in FLAN with CoT reasoning processes, covering 1.84 million instances. Orca (Mitra et al., 2023; Mukherjee et al., 2023) enhances FLAN by prompting LLM with task-specific prompts, thereby training smaller models for cautious reasoning. In this paper, we follow previous work to collect data based on FLAN for fair comparison. We design the methodology for gathering AoT reasoning process and collect AOT COLLECTION to facilitate better training of models for reasoning.

2.3 Reasoning with Abstraction

Previous works study different aspects of reasoning with abstraction, including entity abstraction (Durme et al., 2009; Song et al., 2011; Gong et al., 2016), event abstraction (He et al., 2022; Wang et al., 2023b), spatial-temporal abstraction (Zhang et al., 2021), and conceptualization abstraction (Zhou et al., 2024). Gao et al. (2024) propose to use abstract placeholders in the reasoning chain and call domain tools to supplement specific knowledge, thus allowing the model to use the tools effectively. Wang et al. (2023a) propose to transform questions into symbolic meta forms, but require complex semantic resolution and rely on specific entities. We explore how to improve the basic abstract reasoning of LMs from the perspective of reasoning formats and training data.

3 Abstraction-of-Thought (AoT)

The Abstraction-of-Thought (AoT) reasoning format draws inspiration from the human application of abstract thinking to solve complex reasoning and planning problems (Saitta et al., 2013; Yang, 2012). In trying to solve a complex problem, a good strategy adopted by human reasoners is to proceed top-down. They start from a rough abstract solution and then gradually refine it until a concrete solution is found. For instance, when faced with a complex programming task, seasoned programmers often commence by sketching out a rough algorithmic framework and identifying abstract function features. Subsequently, they progressively delve into the specifics of the code implementation. Similarly, in argumentative discourse (Cohen, 1987), the arguer often first presents the main claims, followed by detailed explanations and supportive evidence for these claims. Therefore, complex problem solutions frequently contain information at diverse abstraction levels, which serves to organize the entire reasoning process hierarchically. These high levels (e.g., functions in a coding solution or main claims in an argument) encapsulate the functionality and objectives of lower levels, thereby structuring the internal reasoning process.

We suggest that the explicit inclusion of such multiple levels of abstraction in the reasoning process could be the key to eliciting abstract reasoning in LMs and propose the AoT format (Figure 2). Specifically, an n -step AoT reasoning process can be expressed as $\tau_{AoT} = a_1^1 \circ a_2^1 \circ \dots \circ a_n^1$, where \circ denotes concatenation and a_i^j is the i -th step in the j -th abstraction level. Each a_i^j can be expanded to $a_{i,1}^{j+1} \circ \dots \circ a_{i,m_{ij}}^{j+1}$, which includes m_{ij} steps belonging to the $(j+1)$ -th abstraction level. The high-level part (e.g., a_i^1) provides the abstract solution (e.g., deriving the general quadratic formula for solving a specific quadratic equation in Figure 1). It focuses on the essence of the problem and ignores inessential details. The low-level part (e.g., $a_i^{j \geq 2}$) further refines it until concrete operations and reasoning can be performed. It contains more details and specific conditions. AoT can be regarded as a constrained variant of CoT, constraining the internal structure of the reasoning process from the perspective of abstract reasoning.

As a preliminary exploration into whether abstract content aids in the reasoning of LLMs, this paper focuses on the two-level AoT. Compared to the unconstrained CoT, the two-level AoT already

Question 1: I have an orange, five raspberries, two books, three plums, a pencil, and a grape. How many fruits do I have?

CoT with Natural Language:

Let's think step by step.
We first identify the fruits on the list and include their quantity in parentheses: orange (1), raspberries (5), plums (3), grape (1).
Now, let's add the numbers in parentheses: $1 + 5 + 3 + 1 = 10$. So the answer is 10.

AoT with Natural Language:

Step 1: Identify fruits from all items.
Of all the items, the orange, raspberry, plum, and grape are fruits.
Step 2: Identify the number of each fruit.
There are 1 orange, 5 raspberries, 3 plums, and 1 grape.
Step 3: Sum the number of fruits.
There are $1 + 5 + 3 + 1 = 10$ fruits in total.
Therefore, the answer is $\boxed{10}$

a_1^1
 $a_{1,*}^2$
 a_2^1
 $a_{2,*}^2$
 a_3^1
 $a_{3,*}^2$

Question 2: Given an $N \times N$ maze that contains one start point S, one end point E, several empty spaces 0, and several walls 1. The player can move up, down, left, and right, but cannot pass through walls. Find the shortest path length from the start to the end. If the end is unreachable, output -1. Example of the given maze: maze=[['S', 0, 0], [1, 1, 0], [0, 0, 'E']]

CoT with Programming Language:

```
directions = [(0, 1), (1, 0), (0, -1), (-1, 0)]
start, end = None, None
for i in range(len(maze)):
    for j in range(len(maze)):
        if maze[i][j] == 'S':
            start = (i, j)
        elif maze[i][j] == 'E':
            end = (i, j)

queue = deque([(start, 0)])
visited = set([start])
while queue:
    (x, y), step = queue.popleft()
    if (x, y) == end:
        print(step)
        break
    for dx, dy in directions:
        nx, ny = x + dx, y + dy
        if (0 <= nx < len(maze) and
            0 <= ny < len(maze) and
            maze[nx][ny] != 1 and
            (nx, ny) not in visited):
            queue.append(((nx, ny), step + 1))
            visited.add((nx, ny))
else:
    print(-1)
```

AoT with Programming Language:

```
def find_start_end(maze):
    """
    Find the start ('S') and end ('E') in the 2D-maze.
    """
    start, end = None, None
    for i in range(len(maze)):
        ...
    return start, end

def bfs(maze, start, end):
    """
    Performs a breadth-first search (bfs) on the 2D-maze
    to find the shortest path length from start to end.
    """
    queue = deque([(start, 0)])
    visited = set([start])
    while queue:
        ...
    return -1

def main(maze):
    """
    Calls other functions to answer the question.
    """
    start, end = find_start_end(maze)
    shortest_path_length = bfs(maze, start, end)
    print(shortest_path_length)
```

a_1^1
 $a_{1,*}^2$
 a_2^1
 $a_{2,*}^2$
 a_3^1

Figure 2: Illustration of Abstraction-of-Thought (AoT) format with natural language (upper half) and programming language (lower half). Unlike the unconstrained CoT, AoT explicitly requires that different levels of abstraction be presented in the reasoning process. Here are examples of two-level abstraction AoT. In AoT, the high-level parts (represented in **bold red**, i.e., a_*^1) plan and organize the entire reasoning process from an abstract perspective, while low-level parts (i.e., $a_{*,*}^2$) carry out concrete and detailed reasoning steps. The high-level parts are abstractions of the low-level parts, clarifying their functionality and objectives. For clarity, we omit some code snippets in AoT with programming language.

includes an additional level of abstraction, which is sufficient for our investigation. We define the specific AoT format on both natural and programming language, as shown in Figure 2. For natural language, AoT requires (1) a clear division of the reasoning process into steps, with “Step i” as the beginning of the i -th step; (2) stating the abstract purpose of the step (as the first level of abstraction a_*^1) followed by the specifics in each reasoning step (as the second level of abstraction $a_{*,*}^2$); and (3) placing the reasoning result in “ $\boxed{}$ ”. For programming language, AoT requires (1) explicitly dividing the code into several pieces, each of which is a function or class; (2) annotating the functions or classes to explain their functionality (as a_*^1); and (3) calling other functions in the main function to

solve the problem. The code within the main function is also considered the first level, as it forms the abstract plan for solving the problem. Note that there may be multiple steps of $a_{i,*}^2$ under a_i^1 , depending on the difficulty of the problem. For $a_{i,*}^2$, we consider a line as a step, i.e., different $a_{i,*}^2$ are divided by line breaks.

4 The AoT COLLECTION

4.1 Overview

To align LMs to the AoT reasoning format, we collect the AoT COLLECTION for finetuning LMs. AoT COLLECTION comprises reasoning processes in AoT format, characterized by the following features: (1) **Hybrid reasoning strategy**. AoT

COLLECTION incorporates reasoning processes expressed in both natural language and programming language. Inspired by previous work (Gao et al., 2023; Chen et al., 2022), solutions to some complex problems are more appropriately to be expressed with programming language. This approach not only facilitates the use of external modules as tools to improve reasoning accuracy (Yang et al., 2024), but also enhances the faithfulness of the reasoning process (Lyu et al., 2023). Models trained on such hybrid reasoning strategies can flexibly choose the suitable reasoning approach based on the type of test question, thereby handling more complex reasoning problems. (2) **Broad task coverage.** AOT COLLECTION is built upon an instruction-tuning dataset that covers a wide range of tasks. This allows the AOT COLLECTION to encompass questions from various tasks, rather than being confined to a specific domain or task. (3) **Scalability.** AOT COLLECTION is collected through an automated pipeline, which allows it to scale up without relying on the efforts of human annotation.

4.2 Source Dataset Selection

We follow previous work (Kim et al., 2023; Mitra et al., 2023) to use FLAN Collection (Longpre et al., 2023) as our source dataset. FLAN Collection includes diverse NLP datasets sourced from P3 (Sanh et al., 2022), Super-NaturalInstructions (Wang et al., 2022), Flan (Wei et al., 2022a), and additional dialogue and code datasets. We follow Kim et al. (2023) to exclude datasets whose data are not publicly accessible and datasets with an excessive number of tokens. Ultimately, we focus on 216 datasets that are consistent with the CoT Collection (Kim et al., 2023). These tasks cover different types of tasks in many domains (Appendix A.1). We manually divide the 216 datasets into two parts, **AoT-Text** (203 datasets that are more suitable to be solved in natural language) and **AoT-Code** (13 datasets that are more suitable to be solved in programming language). We utilize a proportional stratified sampling method to sample 400k instances from the original data for the subsequent AoT response generation. Details of the dataset division can be found in Appendix A.1.

4.3 AoT Response Generation

While manually annotating the AoT reasoning process for all instances could yield higher quality, it is time-consuming and labor-intensive, and thus difficult to scale up. We adopt a method of synthe-

| | AoT-Text | AoT-Code | Total |
|--------------------------|----------|----------|---------|
| Number of Samples | 173,100 | 175,463 | 348,563 |
| Avg. Question Length | 179.6 | 98.9 | 139.0 |
| Avg. AoT Response Length | 144.6 | 172.2 | 158.5 |

Table 1: Statistics of AOT COLLECTION

sizing data with LLMs (Kim et al., 2023; Liu et al., 2024) to automate the process of collecting AoT responses. We first manually create instructions and 3 demonstrations, to exploit the instruction-following and in-context learning capabilities of LLM for generating AoT responses (details can be found in Appendix A.1). To minimize the difficulty of generation, we only consider 2 levels of abstraction². We designed two types of prompts for the datasets in AoT-Text and AoT-Code, respectively. The correct answer to the question is included in the prompt to help the model focus on the generation of the reasoning process. We use GPT-3.5-Turbo as our back-end LLM and generate with greedy decoding. Since we are more concerned about the impact of the reasoning format on the model’s reasoning ability, we do not meticulously design different demonstrations for each dataset like previous work (Kim et al., 2023; Mitra et al., 2023). We use Python as the programming language.

4.4 Data Validation and Filtering

After generating the AoT response, we perform validation and filtering to ensure high quality. For AoT in natural language, we examine whether the answers predicted in the response are consistent with gold answers. To prevent degeneration where different inference steps describe the same content, we stipulate that there should not be excessive similarity between different steps. Specifically, we calculate the Jaccard similarity of words between different steps and require it to be below a threshold of 0.5. For AoT in the programming language format, we execute the code provided in the response and check whether it correctly prints or returns the gold answer. For instances that fail to meet the requirements, we ask the LLM to regenerate 10 times (with a temperature parameter of 0.7). We retain the first response that meets the requirement. If none of the 10 responses meet the requirement,

²Our data collection process might not guarantee to find the optimal abstraction solution for each question, since the optimal abstraction level may differ across various questions, and a single question may have multiple abstract solutions. We focus on investigating whether valid abstraction can help language models reasoning while deprioritizing the collection of optimal abstraction solutions.

we discard the instance. After the filtering process, we retain 348k instances. We randomly sampled 100 examples (50 in natural language and 50 in programming languages) and manually checked the quality of the AoT responses. We find that 96% of AoT responses are valid. Table 1 reports the statistics and Appendix A.1 lists some samples of AOT COLLECTION.

5 Experiments

- **Evaluation Dataset.** We evaluate with the challenging reasoning benchmark BIG-Bench Hard (BBH) (Suzgun et al., 2023), which is the most challenging subset of BIG-Bench (Srivastava et al., 2022). BBH consists of 23 tasks that are specifically selected for their difficulty for LMs. BBH covers a wide range of reasoning challenges, including semantic reasoning (e.g., Movie Recommendation), numerical reasoning (e.g., Multi-Step Arithmetic), logical reasoning (e.g., Logical Deduction), and combinations of some of these abilities (e.g., Object Counting). Furthermore, the FLAN Collection takes BBH as *Held-Out* tasks, which ensures that our finetuning process does not access the evaluation data. We report the average accuracy across 12 NLP tasks (NLP), 11 algorithm tasks (Alg), and all 23 tasks (All), respectively. Details about BBH are in Appendix A.7.

- **Setting and Baselines.** We finetune LMs with AOT COLLECTION and evaluate their reasoning ability. Following previous work (Kim et al., 2023; Mitra et al., 2023), we focus our evaluation primarily on the zero-shot setting. The zero-shot setting represents the realistic scenario, as in practical applications we do not have prior knowledge (e.g., few-shot demonstrations) about the test questions. We compare the CoT-finetuned version (e.g., Llama-3-8B-CoT) and AoT-finetuned version (e.g., Llama-3-8B-AoT) of LMs. For CoT-finetuning, we replace the reasoning processes in AOT COLLECTION with the CoT rationales provided by the CoT Collection (Kim et al., 2023). We also report the performance of the instruction-finetuned LMs (e.g., Llama-3-8B-Instruct) as a reference.

- **Models.** We consider a range of common open source pre-trained language models, including Llama-2 (Touvron et al., 2023), CodeLlama (Rozière et al., 2023), Llama-3 (AI@Meta, 2024), Mistral (Jiang et al., 2023). We also report the performance of GPT-3.5-Turbo-0125 (Peng et al., 2023) and GPT-4-0613 (OpenAI, 2023). Details about

| Model | NLP | Alg | All |
|-------------------------------------|-------------|-------------|---------------------|
| <i>Proprietary Models</i> | | | |
| GPT-3.5-Turbo | 48.2 | 39.2 | 43.9 |
| GPT-4 | 71.4 | 59.5 | 65.7 |
| <i>Instruction-Finetuned Models</i> | | | |
| Llama-3-8B-Instruct | 56.2 | 41.0 | 49.0 |
| CodeLlama-7B-Instruct | 34.4 | 23.2 | 29.1 |
| Llama-2-7B-Chat | 31.0 | 17.4 | 24.5 |
| Mistral-7B-Instruct | 34.3 | 21.8 | 28.3 |
| <i>Our Finetuned Models</i> | | | |
| Llama-3-8B-CoT | 51.4 | 35.2 | 43.6 |
| Llama-3-8B-AoT | 51.7 | 55.0 | 53.3 (+9.7) |
| CodeLlama-7B-CoT | 47.0 | 31.6 | 39.6 |
| CodeLlama-7B-AoT | 49.8 | 49.3 | 49.6 (+10.0) |
| Llama-2-7B-CoT | 42.1 | 22.5 | 32.8 |
| Llama-2-7B-AoT | 41.1 | 29.0 | 35.4 (+2.6) |
| Mistral-7B-CoT | 53.8 | 34.3 | 44.5 |
| Mistral-7B-AoT | 55.0 | 48.4 | 51.8 (+7.3) |

Table 2: Evaluation performance (%) on the unseen BBH benchmark under the **zero-shot** setting (realistic setting). X -AoT/CoT indicates the language model X with AoT/CoT-finetuning. AoT/CoT-finetuning shares the same training questions and training data scale, but differs in the reasoning processes. We mark the improvements of AoT over CoT in parentheses.

the models can be found in Appendix A.2.

- **Implementation Details.** We train all models with Megatron-LM (Shoeybi et al., 2019). We use a learning rate of $2e-6$ for 1 epoch by default. We set the global batch size to 128 and use a cosine decay scheduler. We use greedy decoding for all results, with the maximum sequence length set to 2,048. For the response in natural language, we extract the contents within “\boxed{ }” as the predicted answer. For the response in the programming language format, we execute the program and take the printed output as the predicted answer.

5.1 Zero-Shot Performance

Table 2 presents the zero-shot BBH performance of LMs finetuned in different ways. The AoT-finetuned models demonstrate a remarkable performance enhancement compared to the baselines. For example, in comparison to Llama-3-8B-CoT, Llama-3-8B-AoT achieves an absolute improvement of +9.7%. This trend is consistent across all models studied, which suggests that aligning models with the AoT reasoning format could enhance their reasoning ability more effectively than aligning them to the CoT. AoT guides models to think from a more abstract perspective first, addressing the essence of the problem at a higher level, which

| Model | Prompt | NLP | Alg | All |
|-----------------------------|--------|-------------|-------------|---------------------|
| <i>Proprietary Models</i> | | | | |
| GPT-3.5-Turbo | CoT | 69.0 | 70.2 | 69.6 |
| GPT-3.5-Turbo | AoT | 72.1 | 94.3 | 82.7 (+13.1) |
| GPT-4 | CoT | 87.0 | 86.4 | 86.7 |
| GPT-4 | AoT | 86.7 | 97.0 | 91.6 (+4.9) |
| <i>Pre-Trained Models</i> | | | | |
| Llama-3-8B | CoT | 63.6 | 56.2 | 60.0 |
| Llama-3-8B | AoT | 68.1 | 82.8 | 75.1 (+15.1) |
| CodeLlama-7B | CoT | 48.8 | 33.4 | 41.4 |
| CodeLlama-7B | AoT | 52.1 | 81.1 | 66.0 (+24.6) |
| Llama-2-7B | CoT | 47.2 | 27.6 | 37.8 |
| Llama-2-7B | AoT | 47.5 | 66.3 | 56.5 (+18.7) |
| Mistral-7B | CoT | 61.1 | 48.9 | 55.2 |
| Mistral-7B | AoT | 62.3 | 80.4 | 71.0 (+15.8) |
| <i>Our Finetuned Models</i> | | | | |
| Llama-3-8B-CoT | CoT | 51.0 | 36.2 | 43.9 |
| Llama-3-8B-AoT | AoT | 73.1 | 92.8 | 82.5 (+38.6) |
| CodeLlama-7B-CoT | CoT | 47.3 | 32.6 | 40.3 |
| CodeLlama-7B-AoT | AoT | 55.4 | 85.6 | 69.9 (+29.6) |
| Llama-2-7B-CoT | CoT | 51.1 | 32.0 | 41.9 |
| Llama-2-7B-AoT | AoT | 55.4 | 75.2 | 64.9 (+23.0) |
| Mistral-7B-CoT | CoT | 61.4 | 40.8 | 51.6 |
| Mistral-7B-AoT | AoT | 71.3 | 84.8 | 77.7 (+26.1) |

Table 3: Evaluation performance (%) on BBH with **3-shot** prompting (standard setting). We prompt models with 3 demonstrations in CoT/AoT reasoning format.

potentially enhances their reasoning ability.

It is worth noting that the AoT-finetuned models displayed more substantial improvements in algorithmic tasks than in NLP tasks. The algorithmic tasks, which require capturing the internal reasoning rules of the questions without relying on external knowledge, pose a greater challenge to the model’s reasoning ability (Suzgun et al., 2023). As a result, the performance on algorithmic tasks is usually lower. In contrast, NLP tasks depend not only on reasoning ability but also on external knowledge, where LMs could face the bottleneck in external knowledge. AoT-finetuning guides the model to carry out reasoning with abstraction, with a focus on strengthening the reasoning ability, thus yielding a more noticeable improvement in algorithmic tasks. We further discuss the performance of subtasks (B.1), case study (B.3), and computational cost comparison (B.4) in the Appendix.

5.2 Few-Shot Performance

We also evaluate the effect of AoT under the few-shot setting, the standard setting proposed by Suzgun et al. (2023). For few-shot CoT demonstrations, we use three questions and their CoT ratio-

| Model | Training Data Format | NLP | Alg | All |
|------------|------------------------|-------------|-------------|-------------|
| Llama-3-8B | CoT (CoT Collection) | 46.9 | 32.0 | 39.7 |
| | CoT (Ablation) * | 47.6 | 45.2 | 46.5 |
| | AoT2CoT * | 46.3 | 39.4 | 43.0 |
| | AoT (AoT COLLECTION) * | 48.9 | 50.6 | 49.7 |
| Mistral-7B | CoT (CoT Collection) | 45.9 | 30.3 | 38.4 |
| | CoT (Ablation) * | 46.3 | 45.6 | 46.0 |
| | AoT2CoT * | 44.4 | 47.9 | 46.1 |
| | AoT (AoT COLLECTION) * | 48.3 | 52.1 | 50.1 |

Table 4: Ablation on the reasoning format of training data. We finetune models with 10k questions with reasoning processes in different format and report their zero-shot performance. Data with * are collected with the same back-end LLM and hybrid reasoning strategy.

nales which are provided by the official repository. For the AoT prompt, we employ the same questions and manually create the AoT rationales. Consistent with our AOT COLLECTION, we use the Python program to express the reasoning process for some tasks. Prompts can be found in Appendix A.4.

Table 3 shows the results. For models that have not been finetuned, using prompts in AoT format achieves remarkable performance improvement compared to those in the CoT format. For example, the pre-trained Llama-3-8B achieves an absolute improvement of 15.1% with the AoT prompting. This suggests that AoT could be more effective in stimulating the reasoning ability of pre-trained LMs. Furthermore, by aligning the pre-trained models to AoT with the AOT COLLECTION, the models demonstrate improved performance under AoT prompts, validating the effectiveness of our AOT COLLECTION. Further discussions including the few-shot performance of instruction-finetuned models are in Appendix B.2.

5.3 Ablation Study on Reasoning Format

In preceding experiments, we utilize the CoT Collection (Kim et al., 2023) for the CoT-finetuning. CoT Collection differs from our AOT COLLECTION in two confounding factors, besides the reasoning format: (1) A different LLM is used during data collection; (2) AOT COLLECTION additionally employs the hybrid reasoning strategy, representing reasoning both in text and code. To verify the role of the AoT format, we conduct an ablation study to attempt to eliminate the influence of these confounding factors. We construct a new training dataset, AoT2CoT, which uses the same LLM (i.e., GPT-3.5-Turbo) as the AOT COLLECTION to collect data and also adopts the hybrid reasoning strategy. Specifically, we prompt the LLM to convert the reasoning processes of AOT

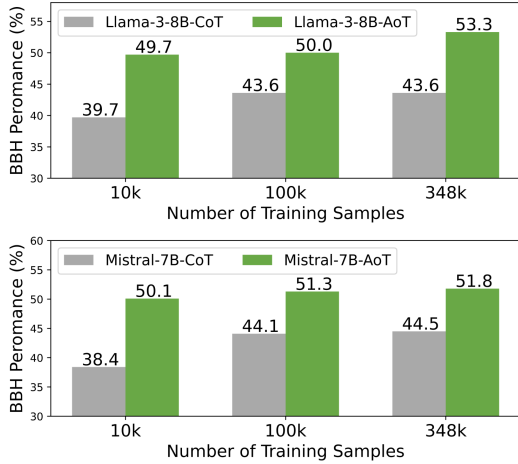


Figure 3: Zero-shot BBH performance of models trained with different numbers of training samples.

COLLECTION from AoT into CoT while keeping the main reasoning content the same. Meanwhile, we also compare with CoT (Ablation), where we use the same back-end LLM to generate CoT directly, adopting the same hybrid reasoning strategy and filtering approach as for AoT data. Considering the computational cost of the LLM, we sample 10k data from AoT COLLECTION, and carry out the ablation experiment on these data. We finetune the models on these same 10k questions, but with reasoning processes in different formats: (1) **CoT** from CoT Collection, (2) **CoT (Ablation)**, (3) **AoT2CoT**, and (4) **AoT** from AoT COLLECTION. More implementation details can be found in the Appendix A.5.

As shown in Table 4, among all formats of training data, AoT achieves the best results over the CoT format. With the same back-end LLM and hybrid reasoning strategy, AoT still outperforms AoT2CoT and CoT (Ablation). This demonstrates that the main factor contributing to our method’s improvement is the reasoning format.

5.4 Ablation Study on Training Data Scale

To investigate the impact of the training data scale, we train models with different numbers of training samples. As shown in Figure 3, the AoT-finetuned models achieve a steady improvement as the scale of training data increases. Across all scales, AoT-finetuned models outperform CoT-finetuned models. Moreover, finetuning using 10k AoT data can yield desirable performance, even better than the case of finetuning using 348k CoT data. This demonstrates both the effectiveness and efficiency of our AoT COLLECTION.

| Model | Text | | | Code | | |
|------------------|------|------|------|------|------|------|
| | UR | FC | AC | UR | FC | AC |
| Llama-3-8B-AoT | 73.4 | 99.3 | 50.3 | 26.6 | 85.1 | 53.7 |
| CodeLlama-7B-AoT | 84.9 | 99.8 | 43.6 | 15.1 | 96.9 | 61.1 |
| Llama-2-7B-AoT | 90.3 | 96.1 | 34.6 | 9.7 | 56.9 | 24.9 |
| Mistral-7B-AoT | 84.7 | 99.7 | 49.1 | 15.3 | 87.3 | 50.4 |

Table 5: Response analysis of AoT-finetuned models. UR=Usage Rate. FC=Format Correctness. AC=Answer Correctness.

5.5 Response Analysis

We analyze the responses of AoT-finetuned models from the following perspectives.

- **Usage Rate (UR)**: The proportion of the text/code being utilized among all responses.
- **Format Correctness (FC)**: The proportion of responses that contain/output predicted answers and follow the AoT format among text/code responses.
- **Answer Correctness (AC)**: The proportion of responses whose predicted answers are correct among text/code responses.

Table 5 presents the results of the responses on BBH in the zero-shot setting. We can find that models prefer to reason with natural language for most problems, rather than programming language. For example, only 26.6% of responses from Llama-3-8B-AoT use code. Moreover, reasoning with code often achieves higher accuracy than reasoning with text. The accuracy of code responses from CodeLlama-7B-AoT reaches 61.1%, whereas text responses only have an accuracy of 43.6%.

5.6 Error Analysis

To better understand the failure modes and future challenges, we randomly sample 100 problems on which Llama-3-8B-AoT fails and manually annotate their error types. Mostly, incorrect reasoning steps constitute the main cause of errors (38%). The model also suffers from lack of necessary knowledge (16%), misunderstanding of tasks (15%), and hallucination (12%). The reasoning process sometimes trivially repeats the conditions (12%) or fails to be executed (5%). Finally, a smaller percentage of correctly predicted answers are misjudged by automated indicators (5%). We present detailed error definitions, examples, and proportions in Appendix B.3 Table 10.

6 Conclusion

This paper explores how to elicit language models to perform reasoning with abstraction, from the perspective of the reasoning format. We propose a

novel structured reasoning format, AoT, which explicitly requires multiple levels of abstraction to be included in the reasoning process. We construct the AoT COLLECTION, containing 348k high-quality AoT reasoning processes, for fine-tuning models to align them to the AoT format. Experimental results show that AoT-finetuned language models can achieve advanced reasoning performance compared to baseline approaches.

Limitations

In this paper, we explore how to elicit language models to perform abstract reasoning from the perspectives of the reasoning format, thereby improving their performance on reasoning tasks. While considerable progress has been made, there are still room for improvement and future research areas worth exploring.

First, due to resource constraints, this paper has not yet explored how to enhance the model’s abstract reasoning ability during the pre-training stage. In our experiments, it can be observed that there is a substantial difference in performance among different pre-trained models after AoT-finetuning. This is due to the issue that the finetuning stage might be difficult to fundamentally enhance the intrinsic capabilities of the model (Bai et al., 2022; Ouyang et al., 2022). Therefore, the bottleneck to fundamentally endow models with the ability for abstract reasoning might lie in the pre-training stage. How to inject the ability for abstract reasoning into language models during the pre-training stage presents an intriguing direction for future research.

Second, limited by the cost of API calls, we do not conduct more analysis on the back-end LLM used for data collection. The impact of using a more powerful LLM for data collection or collecting data on a larger scale for AoT-finetuning deserves further exploration.

Third, this paper focuses the evaluation of the model’s reasoning ability on the BBH benchmark, which contains 23 test tasks that require various reasoning abilities. However, each task of BBH has its own problem-solving pattern and may favor logical and symbolic reasoning over other benchmarks. Further exploration of generalizability on more realistic and complex reasoning problems is a future research direction.

This paper follows the ACL Code of Ethics. To the best of our knowledge, our work is founda-

tional research, and we do not find obvious risks related to fairness considerations, privacy considerations, malicious harmful effects, or environmental impact.

Acknowledgements

We appreciate the anonymous reviewers for their insightful comments. This work is supported by National Science and Technology Major Project (No. 2022ZD0114903), the Natural Science Foundation of China (NSFC. No. 62476149) and the Guoqiang Institute of Tsinghua University (No. 2020GQG0005).

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom B. Brown, Jack Clark, Sam McCandlish, Chris Olah, Benjamin Mann, and Jared Kaplan. 2022. [Training a helpful and harmless assistant with reinforcement learning from human feedback](#). *CoRR*, abs/2204.05862.
- Maciej Besta, Florim Memedi, Zhenyu Zhang, Robert Gerstenberger, Nils Blach, Piotr Nyczyk, Marcin Copik, Grzegorz Kwasniewski, Jürgen Müller, Lukas Gianinazzi, Ales Kubicek, Hubert Niewiadomski, Onur Mutlu, and Torsten Hoefler. 2024. [Topologies of reasoning: Demystifying chains, trees, and graphs of thoughts](#). *CoRR*, abs/2401.14295.
- 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](#). *CoRR*, abs/2211.12588.

- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2023. [A survey of chain of thought reasoning: Advances, frontiers and future](#). *CoRR*, abs/2309.15402.
- Robin Cohen. 1987. Analyzing the structure of argumentative discourse. *Comput. Linguistics*, 13(1-2):11–24.
- Benjamin Van Durme, Phillip Michalak, and Lenhart K. Schubert. 2009. [Deriving generalized knowledge from corpora using wordnet abstraction](#). In *EACL 2009, 12th Conference of the European Chapter of the Association for Computational Linguistics, Proceedings of the Conference, Athens, Greece, March 30 - April 3, 2009*, pages 808–816. The Association for Computer Linguistics.
- Yao Fu, Hao Peng, Litu Ou, Ashish Sabharwal, and Tushar Khot. 2023. [Specializing smaller language models towards multi-step reasoning](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10421–10430. PMLR.
- Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [PAL: program-aided language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 10764–10799. PMLR.
- Silin Gao, Jane Dwivedi-Yu, Ping Yu, Xiaoqing Ellen Tan, Ramakanth Pasunuru, Olga Golovneva, Koustuv Sinha, Asli Celikyilmaz, Antoine Bosselut, and Tianlu Wang. 2024. [Efficient tool use with chain-of-abstraction reasoning](#). *CoRR*, abs/2401.17464.
- Deepanway Ghosal, Yew Ken Chia, Navonil Majumder, and Soujanya Poria. 2023. [Flacuna: Unleashing the problem solving power of vicuna using FLAN fine-tuning](#). *CoRR*, abs/2307.02053.
- Yu Gong, Kaiqi Zhao, and Kenny Qili Zhu. 2016. [Representing verbs as argument concepts](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 2615–2621. AAAI Press.
- Mutian He, Tianqing Fang, Weiqi Wang, and Yangqiu Song. 2022. [Acquiring and modelling abstract commonsense knowledge via conceptualization](#). *CoRR*, abs/2206.01532.
- Namgyu Ho, Laura Schmid, and Se-Young Yun. 2023. [Large language models are reasoning teachers](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 14852–14882. Association for Computational Linguistics.
- Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alex Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. 2023. [Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 8003–8017. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/2310.06825.
- Seungone Kim, Se June Joo, Doyoung Kim, Joel Jang, Seonghyeon Ye, Jamin Shin, and Minjoon Seo. 2023. [The cot collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12685–12708. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Shiyang Li, Jianshu Chen, Yelong Shen, Zhiyu Chen, Xinlu Zhang, Zekun Li, Hong Wang, Jing Qian, Baolin Peng, Yi Mao, Wenhua Chen, and Xifeng Yan. 2022. [Explanations from large language models make small reasoners better](#). *CoRR*, abs/2210.06726.
- Ruibo Liu, Jerry Wei, Fangyu Liu, Chenglei Si, Yanzhe Zhang, Jinneng Rao, Steven Zheng, Daiyi Peng, Diyi Yang, Denny Zhou, et al. 2024. Best practices and lessons learned on synthetic data for language models. *arXiv preprint arXiv:2404.07503*.
- Shayne Longpre, Le Hou, Tu Vu, Albert Webson, Hyung Won Chung, Yi Tay, Denny Zhou, Quoc V. Le, Barret Zoph, Jason Wei, and Adam Roberts. 2023. [The flan collection: Designing data and methods for effective instruction tuning](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 22631–22648. PMLR.
- Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. [Faithful chain-of-thought reasoning](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific*

- Chapter of the Association for Computational Linguistics, IJCNLP 2023 -Volume 1: Long Papers, Nusa Dua, Bali, November 1 - 4, 2023*, pages 305–329. Association for Computational Linguistics.
- Lucie Charlotte Magister, Jonathan Mallinson, Jakub Adámek, Eric Malmi, and Aliaksei Severyn. 2023. [Teaching small language models to reason](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, pages 1773–1781. Association for Computational Linguistics.
- Arindam Mitra, Luciano Del Corro, Shweti Mahajan, Andrés Codas, Clarisse Simões, Sahaj Agrawal, Xuxi Chen, Anastasia Razdaibiedina, Erik Jones, Kriti Agarwal, Hamid Palangi, Guoqing Zheng, Corby Rosset, Hamed Khanpour, and Ahmed Awadallah. 2023. [Orca 2: Teaching small language models how to reason](#). *CoRR*, abs/2311.11045.
- Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. 2023. [Orca: Progressive learning from complex explanation traces of GPT-4](#). *CoRR*, abs/2306.02707.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Andrew Peng, Michael Wu, John Allard, Logan Kilpatrick, and Steven Heide. 2023. [Gpt-3.5 turbo fine-tuning and api updates](#). <https://openai.com/blog/gpt-3-5-turbo-fine-tuning-and-api-updates>.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. [Code llama: Open foundation models for code](#). *CoRR*, abs/2308.12950.
- Lorenza Saitta, Jean-Daniel Zucker, Lorenza Saitta, and Jean-Daniel Zucker. 2013. *Abstraction in Artificial Intelligence*. Springer.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. [Megatron-lm: Training multi-billion parameter language models using model parallelism](#). *CoRR*, abs/1909.08053.
- Kumar Shridhar, Alessandro Stolfo, and Mrinmaya Sachan. 2023. [Distilling reasoning capabilities into smaller language models](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 7059–7073. Association for Computational Linguistics.
- Yangqiu Song, Haixun Wang, Zhongyuan Wang, Hong-song Li, and Weizhu Chen. 2011. [Short text conceptualization using a probabilistic knowledgebase](#). In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 2330–2336. IJCAI/AAAI.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ameet Rahane, Anantharaman S. Iyer, Anders Andreassen, Andrea Santilli, Andreas Stuhlmüller, Andrew M. Dai, Andrew La, Andrew K. Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabassum, Arul Menezes, Arun Kirubakaran, Asher Mullokandov, Ashish Sabharwal, Austin Herick, Avia Efrat, Aykut Erdem, Ayla Karakas, and et al. 2022. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#). *CoRR*, abs/2206.04615.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. 2023. [Challenging big-bench tasks and whether chain-of-thought can](#)

- solve them. In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 13003–13051. Association for Computational Linguistics.
- 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, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Yiming Wang, Zhuosheng Zhang, and Rui Wang. 2023a. [Meta-reasoning: Semantics-symbol deconstruction for large language models](#). *CoRR*, abs/2306.17820.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoormolabashi, Yeganeh Kordi, Amirreza Mirzaei, Atharva Naik, Arjun Ashok, Arut Selvan Dhanasekaran, Anjana Arunkumar, David Stap, Eshaan Pathak, Giannis Karamanolakis, Haizhi Gary Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krima Doshi, Kuntal Kumar Pal, Maitreya Patel, Mehrad Moradshahi, Mihir Parmar, Mirali Purohit, Neeraj Varshney, Phani Rohitha Kaza, Pulkit Verma, Ravsehaj Singh Puri, Rushang Karia, Savan Doshi, Shailaja Keyur Sampat, Siddhartha Mishra, Sujan Reddy A, Sumanta Patro, Tanay Dixit, and Xudong Shen. 2022. [Super-naturalinstructions: Generalization via declarative instructions on 1600+ NLP tasks](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 5085–5109. Association for Computational Linguistics.
- Zhaowei Wang, Haochen Shi, Weiqi Wang, Tianqing Fang, Hongming Zhang, Sehyun Choi, Xin Liu, and Yangqiu Song. 2023b. [Abspyramid: Benchmarking the abstraction ability of language models with a unified entailment graph](#). *CoRR*, abs/2311.09174.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022a. [Finetuned language models are zero-shot learners](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- 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. 2022b. [Emergent abilities of large language models](#). *Trans. Mach. Learn. Res.*, 2022.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022c. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. 2024. [A survey on knowledge distillation of large language models](#). *CoRR*, abs/2402.13116.
- Ke Yang, Jiateng Liu, John Wu, Chaoqi Yang, Yi R. Fung, Sha Li, Zixuan Huang, Xu Cao, Xingyao Wang, Yiquan Wang, Heng Ji, and Chengxiang Zhai. 2024. [If LLM is the wizard, then code is the wand: A survey on how code empowers large language models to serve as intelligent agents](#). *CoRR*, abs/2401.00812.
- Qiang Yang. 2012. *Intelligent planning: a decomposition and abstraction based approach*. Springer Science & Business Media.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. [Tree of thoughts: Deliberate problem solving with large language models](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Xiang Yue, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, and Wenhao Chen. 2023. [Mammoth: Building math generalist models through hybrid instruction tuning](#). *CoRR*, abs/2309.05653.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah D. Goodman. 2022. [Star: Bootstrapping reasoning with reasoning](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Chi Zhang, Baoxiong Jia, Song-Chun Zhu, and Yixin Zhu. 2021. [Abstract spatial-temporal reasoning via probabilistic abduction and execution](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 9736–9746. Computer Vision Foundation / IEEE.

Ben Zhou, Hongming Zhang, Sihao Chen, Dian Yu, Hongwei Wang, Baolin Peng, Dan Roth, and Dong Yu. 2024. [Conceptual and unbiased reasoning in language models](#). *CoRR*, abs/2404.00205.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Xuekai Zhu, Biqing Qi, Kaiyan Zhang, Xingwei Long, and Bowen Zhou. 2023. [Pad: Program-aided distillation specializes large models in reasoning](#). *CoRR*, abs/2305.13888.

A Implementation Details

A.1 AOT COLLECTION

We follow CoT Collection (Kim et al., 2023)³ to preprocess the FLAN Collection (Longpre et al., 2023) data as our source data. The CoT Collection contains 1.84M samples from 216 tasks. These tasks cover different types of tasks in different domains. For example, dataset “commonsenseqa” in the collection is a “Question Answering” task in the commonsense domain, while dataset “gsm8k” is a “Grade School Math Word Problems” in the math domain. Our AOT COLLECTION is built on these diverse datasets and already covers different tasks in different domains as well, rather than being designed only for a specific domain.

Table 11 at the end of the Appendix gives all the dataset names and their corresponding tasks. We follow the FLAN Collection to determine the tasks corresponding to the datasets. Considering the large number of datasets, we have opted not to display examples for each one in this document. Please refer to the FLAN Collection (<https://github.com/google-research/FLAN/tree/main/flan/v2>) if you are interested in the specific dataset examples. To demonstrate the diversity of tasks, we list the task categories and their distribution below. The number in parentheses indicates the dataset for which this task is relevant.

Question Answering (65);
Text Categorization (24);
Question Understanding (11);
Multiple-Choice QA (11);
Information Extraction (10);
Dialogue Generation (10);
Toxic Language Detection (10);
Text Matching (9);
Answerability Classification (9);
Textual Entailment (9);
Sentiment Analysis (9);
Natural Language Inference (8);
Speaker Identification (7);
Inverted Natural Language Inference (7);
Fill in The Blank (7);
Sentence Composition (7);
Closed-Book QA (6);
Text Quality Evaluation (6);
Inverted Multiple-Choice QA (6);
Extractive QA (6);
Intent Identification (5);
Text Completion (5);
Explanation (5);
Story Composition (4);
Named Entity Recognition (4);
Question Rewriting (4);
Commonsense Classification (4);
Dialogue Act Recognition (4);

Answer Verification (4);
Conversational Question Answering (4);
Inverted Extractive QA (4);
Answer Generation (4);
Coherence Classification (3);
Adversarial QA (3);
Inverted Closed-Book QA (3);
Program Execution (3);
Word Semantics (3);
Dialogue State Tracking (3);
Fact Verification (3);
Question Generation (3);
Word Relation Classification (2);
Linguistic Probing (2);
Question Decomposition (2);
Text to Code (2);
Cause Effect Classification (2);
Grammar Error Detection (2);
Stereotype Detection (2);
Inverted Mathematical QA (2);
Mathematical QA (2);
Summarization (2);
Gender Classification (2);
Pos Tagging (2);
Keyword Tagging (2);
Span Generation (2);
Sentence Ordering (2);
Dialog Next Turn Prediction (2);
Coreference Resolution (2);
Word Analogy (1);
Negotiation Strategy Detection (1);
Question Context Generation (1);
Section Classification (1);
Inverted Grammatical Acceptability (1);
Grammatical Acceptability (1);
Speaker Relation Classification (1);
Common Sense Reasoning Question Answering (1);
Answer Incomplete Questions (1);
Emotion Word Generation (1);
Emotional Reaction Generation (1);
Intent Generation (1);
Grade School Math Word Problems (1);
Entity Generation (1);
Entity Relation Classification (1);
Food Classification (1);
Language Identification (1);
Paraphrasing (1);
Poem Generation (1);
Ethics Classification (1);
Irony Detection (1);
Stance Detection (1);
Clock Format Conversion (1);
Spelling Error Detection (1);
Rhyme Generation (1);
Date Validity Prediction (1);
Mathematics (1);
Temporal Reasoning (1);
Leap Year Prediction (1);
Edible Prediction (1);
Inverted Coreference Resolution (1);
Sentence Perturbation (1);

Each sample of CoT Collection consists of a question, an answer, and a CoT rationale. We discard the CoT rationale and use the GPT-3.5-Turbo to generate the AoT reasoning process based on the question and the correct answer. We manually divided the 216 datasets into two parts, **AoT-Text**

³<https://huggingface.co/datasets/kaist-ai/CoT-Collection>

(203 datasets that are more suitable to be solved in natural language) and **AoT-Code** (13 datasets that are more suitable to be solved in programming language). The AoT-Text consists of the following tasks:

adversarial_qa@droberta, olid, ai2_arithmetic_questions, natural_questions, question_&_answer_zre, qanta, health_fact, quail, jeopardy, jigsaw, detoxifying_lms, poki, qnli, cb, google_wellformed_query, subjqa, multi_woz_v2, casino, task_master_input_inversion, babi, wiki_movies, ddo, anli_r1, craigslist_bargains, ohsumed, strategyqa, adversarial_qa@dbidaf, cad, semeval_2018_task3, emo, overruling, hippocampus, qed, diplomacy_detection, piqa, smcalflow, super_glue@record, schema_guided_dstc8, ai2_arc@ARC-Easy, nlu_asdiv_dataset, ruletaker, news_headlines_dataset_for_sacrasem_detection, com2sense, civil_comments, circa, quartz, diqa, semeval_2019_task_10, hybridqa, evaluation, ecqa, storycommonsense, miscellaneous, snli, clariq, blimp, financial_phrasebank, hatexplain, hope_edi, numersense, x_csr, xcsr, qa_srl, mcscrip, mwsc, persent, trivia_qa, hate_speech_offensive, coached_conv_pref, scitail, drop, rte, anli_r3, qrecc, ms_marco, quac, wikitext_103, nlg_bias, mutual, gwsd, yahoo_answers_topics, essential, swag, torque, wiki_dialog_input_inversion, cola, winowhy, disfl_qa, roc_stories, semeval_2020_task_7, codah, mocah, atomic, crows_pair, mnli, tweetqa, scruples, conv_ai_2, stereoset, break, duorc@SelfRC, dialogre, ambigqa, iirc, miam, pubmed_qa, deal_or_no_dialogue, ai2_arc@ARC-Challenge, coda_19, spolin, wiki_hop, hateeval, timetravel, duorc@ParaphraseRC, recepie_nlg, kilt_tasks@hotpotqa, curated_from_stack_overflow_english, sciq, freebase_qa, squad_v2, help!_need_advice_on_identifying_advice, bless, squad_v1, task_master, sbic, quoref, com_qa, wnli, haspart_kb, personachat, argkp, ethos, open_pi, race@high, proto_qa, sarcasm_in_twitter, web_questions, abductive_nli, curiosity_dialogs, imppres, race@middle, adversarial_qa@dbert, eurlex, head_qa, defeasible_nli_atomic, equity_evaluation_corpus, qrecc_input_inversion, wiki_dialog, eqasc, bard, wiqu, dream, liar, anli_r2, scitailv1.1, tellmewhy, cod3s, dstc, indian_food_101, aquamuse, glucose, social_i_qa, air_dialogue, missing, narrativeqa, scitldr, mrqa, meta_woz, go_emotions, casehold, scifact, super_glue@boolq, ade_corpus_v2, daillydialog, starcon, commonsenseqa, openbookqa, quarel, propara, event2mind, inquisitive, tom_qa, wiki_qa, cosmos_qa, afs, medical_question_pair_dataset, creak, yoruba_bbc_topics, semeval_2020_task4, xl_wic, super_glue@multirc, opp_115, esnli, grailqa, root09, qasper, ropes, gooqa, cos_e, perspectrum, xquad, trianglecopa, mctaco

The AoT-Code consists of the following tasks:

aqua, big_bench, cfq_mcd1, conala, gsm8k, leetcode, math_dataset, math_qa, mathematics_dataset, mathqa, prost, svamp, synthetic

We utilize a proportional stratified sampling method to sample 200k instances from the AoT-Code datasets and 200k instances from the AoT-Text datasets. The proportional stratified sampling method ensures that the proportion of each dataset is the same before and after sampling, thus maintaining the overall data distribution.

Listing 1 shows the system prompts we use for collecting AoT responses. We use different

prompts for AoT-Text and AoT-Code and clearly state the requirements of AoT format. We also include 3 demonstrations in the prompt.

AOT COLLECTION retains 348k high quality data after filtering. The average number of steps for responses in AoT-Text and AoT-Code is 2.8 and 2.2, respectively. We randomly sample 3 samples from each of AoT-Text and AoT-Code and present them in Listing 2. AOT COLLECTION is in English.

Listing 1: Prompts for collecting AoT responses. We show the system prompt and one of the three demonstrations.

Prompt for Collecting AoT Responses with Natural Language

**** System Prompt ****

You are a powerful reasoning model. I will present you with a question and the correct answer. Your task is to output the intermediate reasoning process to meet the following requirements.

- The reasoning process contains several reasoning steps that can support the correct answer.
- Each reasoning step should begin with a abstract high-level plan, followed by a detailed step-by-step implementation.
- The high-level plans should be tailored to the specific question. They may encompass, but are not limited to: extracting key information from the question, sub-questions decomposed from the target question, analyzing certain aspects of the question, evaluating different options and possibilities, recalling relevant knowledge and rules, applying knowledge and rules for reasoning, comparing different content, concluding and answering.
- Each step should aim to draw new conclusions and not merely repeat the question or the content of the previous step.
- Conclude the reasoning process by stating the final answer as "The answer is \boxed{answer}."

**** Demonstration Input ****

Which sentence has the correct adjective order:
Options:

- (A) Indonesian triangular obnoxious gold red shoe
- (B) obnoxious triangular red Indonesian gold shoe

Correct Answer: (B)

**** Demonstration Output ****

Step 1: Recall the order of adjectives in English.

The generally accepted order of adjectives in English is:

quantity, quality, size, age, shape, color, origin, material, purpose.

Step 2: Analyze each option according to the adjective order rule.

- (A) Indonesian (origin) triangular (shape)

```

obnoxious (quality) gold (material) red (
color) shoe (noun)
(B) obnoxious (quality) triangular (shape) red (
color) Indonesian (origin) gold (material)
shoe (noun)

Step 3: Compare the options.
Option (A) has the incorrect order because "
origin" comes before "shape" and "quality",
and "material" comes before "color".
Option (B) follows the correct order: quality,
shape, color, origin, material.

Therefore, the answer is \boxed{(B)}.

=====
Prompt for Collecting AoT Responses with
Programming Language

** System Prompt **
You are a powerful reasoning model and Python
coder. I will present you with a question
and the correct answer. Your task is to
output a Python program to meet the
following requirements.
- The code can answer the question and print the
correct answer.
- Divide the code into several functional blocks.
Each functional block implements a simple
function.
- Describe what the function does in the comment
lines following the function name. Comments
should be as concise as possible.
- Define a `main` function to call other
functions to answer the question.
- The main function should contain arguments
transformed by the conditions in the
question. The other functions should be as
generic as possible.

** Demonstration Input **
In this task, answer 'Yes' if the frequency of
the given word in the two sentences is equal
, else answer 'No'. Sentence1: 'a little boy
that is holding a skateboard', Sentence2: '
this photo displays a messy desk that has
two laptops on it among many other things'.
Is the frequency of the word 'a' in two
sentences equal?
Correct Answer: No

** Demonstration Output **
```python
def count_word_frequency(sentence, word):
 """
 Count the frequency of a given word in a
 sentence.
 """
 words = sentence.split()
 count = words.count(word)
 return count

def compare_word_frequency(sentence1, sentence2,
 word):
 """
 Compare the frequency of a given word in two
 sentences.
 """
 count1 = count_word_frequency(sentence1, word
)

```

```

count2 = count_word_frequency(sentence2, word
)
 return count1 == count2

def main():
 """
 The main function that calls the other
 functions and prints the answer.
 """
 sentence1 = 'a little boy that is holding a
 skateboard'
 sentence2 = 'this photo displays a messy desk
 that has two laptops on it among many
 other things'
 word = 'a'
 is_equal = compare_word_frequency(sentence1,
 sentence2, word)
 print('Yes' if is_equal else 'No')

main()
```

```

Listing 2: Samples of AOT COLLECTION.

```

===== Sample 1 =====
** Input **
What requires nutrients to grow?
Choices:
- birds
- metals
- computers
- dirt

** Output **
Step 1: Recall the definition of nutrients.
Nutrients are substances that are essential for
the growth, development, and maintenance of
living organisms.

Step 2: Analyze each option and determine which
one requires nutrients to grow.
- Birds: Birds are living organisms and require
nutrients to grow.
- Metals: Metals are inanimate objects and do
not require nutrients to grow.
- Computers: Computers are machines and do not
require nutrients to grow.
- Dirt: Dirt is a mixture of minerals, organic
matter, and other substances, but it does
not require nutrients to grow.

Step 3: Compare the options and identify the
correct answer.
Among the given choices, only birds require
nutrients to grow.

Therefore, the answer is \boxed{birds}.

===== Sample 2 =====
** Input **
Two analogies that relate objects to the
associated rooms is given in the form "A : B
. C : ?". "A : B" relates object A to room B
. Your task is to replace the question mark
(?) with the appropriate room for the given
object C, following the "A : B" relation.

bookshelf : library. car : ?

** Output **

```


Step 1: Identify the relationship between the object and the room in the given analogy. In the analogy "bookshelf : library", the bookshelf is typically found in the library.

Step 2: Apply the same relationship to the second part of the analogy. Since a car is typically found in a garage, we can replace the question mark with "garage".

Therefore, the answer is \boxed{garage}.

===== Sample 3 =====

**** Input ****

I have a new situation: Big town and Fair town both used petroleum products. Big town used them to make plastics and related materials, but didn't use it for gas, because they were environmentally conscious and used only electric vehicles for transportation. Fair town used petroleum for fuel mostly.

But I can use this background: A major use for petroleum products is fuel for cars, trucks, airplanes, trains, and other vehicles. The chemical used are usually a mixture of compounds containing several carbon atoms in a chain. When the material is ignited, a massive amount of gas is created almost instantaneously. This increase in volume will move the pistons in an internal combustion engine to provide power. A jet plane works on a similar principle. Air comes into the front of the engine and mixes with the jet fuel. This mixture is ignited and the gases formed create a lot of pressure to push the plane forward. The idea can be seen even more clearly in the case of a rocket launch. The ignition of the fuel (either solid-state or liquid) creates gases produced under great pressure that pushes the rocket up.

What is an answer for this question: Which town used more fuel for internal combustion engines?

**** Output ****

Step 1: Extract key information from the situation.

Big town uses petroleum products to make plastics and related materials, but not for gas as they use electric vehicles for transportation. Fair town uses petroleum mostly for fuel.

Step 2: Recall the background information. The background information states that a major use for petroleum products is fuel for vehicles with internal combustion engines.

Step 3: Apply the background information to the situation.

Since Big town uses electric vehicles, they do not use petroleum for fuel in internal combustion engines. On the other hand, Fair town uses petroleum mostly for fuel, which implies they use it for internal combustion engines.

Therefore, the answer is \boxed{Fair}.

===== Sample 4 =====

**** Input ****

In this task, you need to answer the given multiple-choice question on geometry. Classify your answers into 'a', 'b', 'c', 'd', and 'e'.

Problem: the radius of the two circular fields is in the ratio 4 : 5 the area of the first field is what percent less than the area of the second ?

Options: a) 42 % , b) 39 % , c) 38 % , d) 36 % , e) 26 %

**** Output ****

```python

def calculate\_area(radius):

"""

Calculate the area of a circle given its radius.

"""

area = 3.14159 \* radius\*\*2

return area

def calculate\_percent\_difference(area1, area2):

"""

Calculate the percentage difference between two areas.

"""

difference = area2 - area1

percent\_difference = (difference / area2) \* 100

return percent\_difference

def main():

"""

The main function that calls the other functions and prints the answer.

"""

radius1 = 4

radius2 = 5

area1 = calculate\_area(radius1)

area2 = calculate\_area(radius2)

percent\_less = calculate\_percent\_difference(area1, area2)

options = {

'a': 42,

'b': 39,

'c': 38,

'd': 36,

'e': 26

}

for option, percent in options.items():

if percent == percent\_less:

print(option)

break

main()

```

===== Sample 5 =====

**** Input ****

In this task, you need to replace a letter in the sentence with another given letter.

Sentence: 'a red plus sign over a toilet, near a sink'. Replace the letter 'r' with 'o' in the sentence.

**** Output ****

```python

```
def replace_letter(sentence, old_letter,
 new_letter):
 """
 Replace a letter in a sentence with another
 given letter.
 """
 replaced_sentence = sentence.replace(
 old_letter, new_letter)
 return replaced_sentence
```

```
def main():
 """
 The main function that calls the other
 functions and prints the answer.
 """
 sentence = 'a red plus sign over a toilet,
 near a sink'
 old_letter = 'r'
 new_letter = 'o'
 replaced_sentence = replace_letter(sentence,
 old_letter, new_letter)
 print(replaced_sentence)
```

main()

```

===== Sample 6 =====

**** Input ****

In this task, you are given an input list. A list contains several comma-separated items written within brackets. You need to return the count of all the numerical elements in the given list. Return 0 if no numerical element in the list.

['D', '9987', '4565', 'j', 'V', 'S', '1033', '7319', 'S']

**** Output ****

```python

```
def count_numerical_elements(lst):
 """
 Count the number of numerical elements in a
 given list.
 """
 count = 0
 for item in lst:
 if item.isdigit():
 count += 1
 return count
```

```
def main():
 """
 The main function that calls the other
 functions and prints the answer.
 """
 lst = ['D', '9987', '4565', 'j', 'V', 'S',
```

```
'1033', '7319', 'S']
 count = count_numerical_elements(lst)
 print(count)
```

main()

```

A.2 Models

Table 6 lists the models involved in our experiments, including their names, versions, and corresponding URL links. We follow the licences (which can be found in the URL links) of these models to use them. For the open-source models, we use the model weights provided by Hugging-face⁴.

A.3 Zero-Shot Prompts

In the zero-shot setting, we directly use the test question as the input to the models. For our models with CoT/AoT-finetuning, we are able to extract the predicted answers from the responses in a fixed format. For open-source instruction finetuned models (e.g., Llama-3-8B-Instruct), we utilize a simple yet effective instruction to guide the model to output the answer in a fixed format. The instruction we use is: “Answer the question and put the final answer in \boxed{ }.”

A.4 Few-Shot Prompts

In the few-shot setting, we conduct experiments with two type of prompts: CoT prompt and AoT prompt. The CoT prompt is collected from the official GitHub repository of BBH⁵. For each test task, the CoT prompt contains three demonstration questions along with their CoT rationales. For the AoT prompt, we employ the same questions and manually create the reasoning processes in AoT format. In consistency with the AoT COLLECTION, we use the Python programming language to express the reasoning process for some tasks.

For the following tasks we use a programming language reasoning process: *Boolean expressions, Dyck languages, Geometric shapes, Logical deduction five objects, Logical deduction seven objects, Logical deduction three objects, Multistep arithmetic two, Navigate, Object counting, Penguins in a table, Temporal sequences, Tracking*

⁴<https://huggingface.co/models>

⁵<https://github.com/suzgunmirac/BIG-Bench-Hard/tree/main/cot-prompts>

| Model | Version | URL |
|--|-------------------------------------|---|
| Llama2-7B (Touvron et al., 2023) | meta-llama/Llama-2-7b-hf | https://huggingface.co/meta-llama/Llama-2-7b |
| Llama2-7B-chat (Touvron et al., 2023) | meta-llama/Llama-2-7b-chat-hf | https://huggingface.co/meta-llama/Llama-2-7b-chat-hf |
| CodeLlama-7B (Rozière et al., 2023) | meta-llama/CodeLlama-7b-hf | https://huggingface.co/meta-llama/CodeLlama-7b-hf |
| CodeLlama-7B-Instruct (Rozière et al., 2023) | meta-llama/CodeLlama-7b-Instruct-hf | https://huggingface.co/meta-llama/CodeLlama-7b-Instruct-hf |
| Llama-3-8B (AI@Meta, 2024) | meta-llama/Meta-Llama-3-8B | https://huggingface.co/meta-llama/Meta-Llama-3-8B |
| Llama-3-8B-Instruct (AI@Meta, 2024) | meta-llama/Meta-Llama-3-8B-Instruct | https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct |
| Mistral-7B (Jiang et al., 2023) | mistralai/Mistral-7B-v0.1 | https://huggingface.co/mistralai/Mistral-7B-v0.1 |
| Mistral-7B-Instruct (Jiang et al., 2023) | mistralai/Mistral-7B-Instruct-v0.2 | https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2 |
| GPT-3.5 (Peng et al., 2023) | gpt-3.5-turbo-0125 | https://platform.openai.com/docs/models/gpt-3-5 |
| GPT-4 (OpenAI, 2023) | gpt-4-0613 | https://platform.openai.com/docs/models/gpt-4 |

Table 6: Detailed information about the models we experiment with.

shuffled objects five objects, Tracking shuffled objects seven objects, Tracking shuffled objects three objects, Web of lies, Word sorting.

For the following tasks we use a natural language reasoning process: *Causal judgement, Date understanding, Disambiguation qa, Formal fallacies, Hyperbaton, Movie recommendation, Reasoning about colored objects, Ruin names, Salient translation error detection, Snarks, Sports understanding.*

Listing 3 presents examples of the Few-shot AoT prompts for BBH. Considering that the full prompts for all tasks in BBH is long, we show two demonstrations in the prompts here for brevity (one in natural language and one in programming language). We will release the prompt we used after publication.

Listing 3: Examples of Few-shot AoT Prompts

```

===== AoT Prompt for Sports Understanding =====

** Demonstration Input **
Is the following sentence plausible? "DeMar
DeRozan was called for the goal tend."

** Demonstration Output **
Step 1: Recall the infomation of the sport
player.
DeMar DeRozan is an American basketball player.

Step 2: Recall the infomation of the sport.
Goal tending is part of basketball.

Step 3: Evaluate the plausibility of the
sentence.
Given that DeMar DeRozan is a basketball player,
it is plausible that he could be called for
a goal tend during a game.

Therefore, the answer is \boxed{yes}.

===== AoT Prompt for Web of Lies =====

```

```

** Demonstration Input **
Question: Kristian tells the truth. Michaela
says Kristian lies. Raymond says Michaela
tells the truth. Osvaldo says Raymond tells
the truth. Jamey says Osvaldo tells the
truth. Does Jamey tell the truth?

** Demonstration Output **
```python
def evaluate_statements(initial_condition,
statements):
 """
 This function takes a dictionary of
 statements where each key is a person
 and the value is the person they are
 commenting on and whether they believe that
 person is telling the truth or not.
 """
 truth_values = initial_condition
 for person, (target, statement) in statements
.items():
 if statement == 'truth':
 truth_values[person] = truth_values[
target]
 else:
 truth_values[person] = not
truth_values[target]
 return truth_values

def main():
 """
 The main function that calls the other
 functions and prints the answer.
 """
 initial_condition = {'Kristian': True}
 statements = {
 'Michaela': ('Kristian', 'lies'),
 'Raymond': ('Michaela', 'truth'),
 'Osvaldo': ('Raymond', 'truth'),
 'Jamey': ('Osvaldo', 'truth')
 }
 truth_values = evaluate_statements(
 initial_condition, statements)
 print('Yes' if truth_values['Jamey'] else 'No
')

main()

```

## A.5 AoT2CoT

To construct AoT2CoT for ablation study, we use LLM to transform the reasoning process from AoT format to CoT format. We sample 10k of data from the AoT COLLECTION (5k from AoT-Text and 5k from AoT-Code). Our aim is to transform the reasoning format while keeping the reasoning content is roughly the same. We prompt GPT-3.5-Turbo (the same back-end LLM for collecting AoT COLLECTION) to transform the input reasoning process into CoT format. In order to avoid the model generating new reasoning content in response to the question, we do not input the question and answer, but only the reasoning process. Listing 4 gives the specific prompt we use.

Listing 4: Prompts for constructing AoT2CoT. We show the system prompt and one of the three demonstrations.

```
Prompt for Converting AoT to CoT with Natural
Language

** System Prompt **
You are a text rewriter. Given a step-by-step
reasoning process, you need to reorganize
its contents to meet the following
requirements:
- The rewritten text do not add new content and
the reasoning logic in it is consistent with
the given text.
- The rewritten text should be as concise and
natural as possible, without any explicit
steps or structures.

** Demonstration Input **
Step 1: Extract key information about touchdowns
from the text.
The touchdowns mentioned are:
- A 1-yard TD pass to Calvin Johnson
- A 7-yard TD pass to Donald Lee
- A 7-yard catch by Donald Driver
- A 21-yard catch by James Jones
- A 38-yard interception return for a touchdown
by Charles Woodson

Step 2: Sort the yardages in descending order.
38 yards, 21 yards, 7 yards, 7 yards, 1 yard

Step 3: Identify the second longest touchdown.
The second longest touchdown is 21 yards.

Therefore, the answer is \boxed{21}.

** Demonstration Output **
The text mentions several touchdowns: a 1-yard
TD pass to Calvin Johnson, a 7-yard TD pass
to Donald Lee, a 7-yard catch by Donald
Driver, a 21-yard catch by James Jones, and
a 38-yard interception return for a
touchdown by Charles Woodson. When these
yardages are sorted in descending order, we
get 38 yards, 21 yards, 7 yards, 7 yards,
```

```
and 1 yard. The second longest touchdown is
21 yards. Therefore, the answer is \boxed
{21}.

Prompt for Converting AoT to CoT with
Programming Language

** System Prompt **
You are a Python code refactoring tool. Given a
piece of Python code, you need to refactor
it to meet the following requirements:
- The refactored code should have the same
functionality and `print` output as the
given code.
- The refactored code should not contain any
function or classes.

** Demonstration Input **
```python
def count_word_frequency(sentence, word):
    """
    Count the frequency of a given word in a
    sentence.
    """
    words = sentence.split()
    count = words.count(word)
    return count

def compare_word_frequency(sentence1, sentence2,
    word):
    """
    Compare the frequency of a given word in two
    sentences.
    """
    count1 = count_word_frequency(sentence1, word
    )
    count2 = count_word_frequency(sentence2, word
    )
    return count1 == count2

def main():
    """
    The main function that calls the other
    functions and prints the answer.
    """
    sentence1 = 'a little boy that is holding a
    skateboard'
    sentence2 = 'this photo displays a messy desk
    that has two laptops on it among many
    other things'
    word = 'a'
    is_equal = compare_word_frequency(sentence1,
    sentence2, word)
    print('Yes' if is_equal else 'No')

main()
```

** Demonstration Output **
```python
sentence1 = 'a little boy that is holding a
skateboard'
sentence2 = 'this photo displays a messy desk
that has two laptops on it among many other
things'
word = 'a'

is_equal = sentence1.split().count(word) ==
sentence2.split().count(word)
print('Yes' if is_equal else 'No')
```


A.6 CoT (Ablation)

CoT (Ablation) and AOT COLLECTION utilize the same back-end LLM for data generation and the same hybrid reasoning strategy. Unlike AoT2CoT, we directly prompt GPT-3.5-Turbo to generate CoT responses. Listing 5 provides the prompt we use. We modify the prompt used for generating AoT responses (Listing 1) by removing the constraints of the AoT format, thereby guiding the model to generate CoT responses. We then apply the same filtering approach before using the data for fine-tuning. It is worth noting that, for the same question, the reasoning process in CoT (Ablation) may differ from the corresponding AoT reasoning process not only in the reasoning format, but also in the content of the reasoning. This is because we cannot guarantee that the back-end LLM would generate the same reasoning content for the same question in two independent generations.

Listing 5: Prompts for constructing CoT (Ablation). We show the system prompt and one of the three demonstrations.

```
Prompt for Collecting CoT (Direct) Responses
with Natural Language

** System Prompt **
You are a powerful reasoning model. I will
present you with a question and the correct
answer. Your task is to output the
intermediate reasoning process step-by-step
and conclude the reasoning process by
stating the final answer as "The answer is \
boxed{answer}".

** Demonstration Input **
Which sentence has the correct adjective order:
Options:
(A) Indonesian triangular obnoxious gold red
shoe
(B) obnoxious triangular red Indonesian gold
shoe
Correct Answer: (B)

** Demonstration Output **
In English, adjectives generally follow this
order: quantity, quality, size, age, shape,
color, origin, material, purpose. We find
that option (A) has the adjectives in the
wrong order. On the other hand, option (B)
correctly orders the adjectives as quality,
shape, color, origin, material.

Therefore, the answer is \boxed{(B)}.

=====
Prompt for Collecting CoT (Direct) Responses
with Programming Language
```

```
** System Prompt **
You are a powerful reasoning model and Python
coder. I will present you with a question
and the correct answer. Your task is to
output a Python program to answer the
question and print the correct answer.

** Demonstration Input **
In this task, answer 'Yes' if the frequency of
the given word in the two sentences is equal
, else answer 'No'. Sentence1: 'a little boy
that is holding a skateboard', Sentence2: '
this photo displays a messy desk that has
two laptops on it among many other things'.
Is the frequency of the word 'a' in two
sentences equal?
Correct Answer: No

** Demonstration Output **
```python
sentence1 = 'a little boy that is holding a
skateboard'
sentence2 = 'this photo displays a messy desk
that has two laptops on it among many other
things'
word = 'a'
is_equal = sentence1.split().count(word) ==
sentence2.split().count(word)
print('Yes' if is_equal else 'No')
```
```

A.7 BBH Benchmark

Table 7 presents the tasks and number of questions within the BBH benchmark. Following [Suzgun et al. \(2023\)](#), tasks are divided into two categories: NLP tasks and algorithm tasks. Tasks “Logical deduction” and “Tracking shuffled objects” consist of 3 sub-tasks. The questions and answers in BBH are in English.

B Additional Experiment Results

B.1 Performances on Subtasks

Table 9 shows the performance on subtasks of BBH. We also introduce human performance as a reference ([Suzgun et al., 2023](#)). We can observe that AoT-finetuning compared to cot-finetuning is able to achieve improvements on multiple reasoning tasks on multiple models.

B.2 Few-Shot Performance

Table 8 presents the 3-shot performance on BBH, including those that are instruction-finetuned (such as Llama-3-8B-Instruct). A consistent trend can be observed where instruction-finetuned models achieve lower few-shot performance compared to models that are merely pre-trained. For example, Llama-3-8B-Instruct achieves an overall accuracy of 56.7% with the official few-shot CoT prompt,

| Task | # Questions |
|--|-------------|
| boolean_expressions* | 250 |
| causal_judgement | 187 |
| date_understanding | 250 |
| disambiguation_qa | 250 |
| dyck_languages* | 250 |
| formal_fallacies | 250 |
| geometric_shapes* | 250 |
| hyperbaton | 250 |
| logical_deduction* | |
| logical_deduction_five_objects* | 250 |
| logical_deduction_seven_objects* | 250 |
| logical_deduction_three_objects* | 250 |
| movie_recommendation | 250 |
| multistep_arithmetic_two* | 250 |
| navigate* | 250 |
| object_counting* | 250 |
| penguins_in_a_table | 146 |
| reasoning_about_colored_objects | 250 |
| ruin_names | 250 |
| salient_translation_error_detection | 250 |
| snarks | 178 |
| sports_understanding | 250 |
| temporal_sequences* | 250 |
| tracking_shuffled* | |
| tracking_shuffled_objects_five_objects* | 250 |
| tracking_shuffled_objects_seven_objects* | 250 |
| tracking_shuffled_objects_three_objects* | 250 |
| web_of_lies* | 250 |
| word_sorting* | 250 |

Table 7: Tasks in the BBH Benchmark. * indicates that the task is an algorithmic task. Untagged tasks belong to NLP tasks.

while Llama-3-8B can achieve 60.0% with the same prompt. For the CodeLlama-7B-Instruct and Llama-2-7B-Chat, a significant performance decrease is noted in the instruction-finetuned versions. Upon examining their responses, we find that in most cases, the models did not follow the demonstration format for answering. For the Llama-3-8B-Instruct and Mistral-7B-Instruct, the responses adhered to the demonstration format. In this scenario, using the AoT prompt yields better results than the CoT prompt.

B.3 Case Study

Listing 6 shows success cases for the AoT-finetuned Llama-3-8B. The AoT-finetuned model demonstrates a certain degree of ability to reason with abstraction. For example, for the Case 3, the model first defines the tools needed to solve the problem at an abstract level, i.e., the Person class and the swap_gifts function. Subsequently, the model then utilizes these tools to solve the problem based on the concrete conditions of the question. Listing 7 presents the error cases.

| Model | Prompt | NLP | Alg | All |
|-----------------------|--------|------|------|------|
| Llama-3-8B | CoT | 63.6 | 56.2 | 60.0 |
| Llama-3-8B | AoT | 68.1 | 82.8 | 75.1 |
| Llama-3-8B-Instruct | CoT | 66.6 | 45.8 | 56.7 |
| Llama-3-8B-Instruct | AoT | 69.3 | 71.2 | 70.2 |
| Llama-3-8B-CoT | CoT | 51.0 | 36.2 | 43.9 |
| Llama-3-8B-AoT | AoT | 73.1 | 92.8 | 82.5 |
| CodeLlama-7B | CoT | 48.8 | 33.4 | 41.4 |
| CodeLlama-7B | AoT | 52.1 | 81.1 | 66.0 |
| CodeLlama-7B-Instruct | CoT | 29.2 | 10.9 | 20.4 |
| CodeLlama-7B-Instruct | AoT | 29.7 | 7.3 | 19.0 |
| CodeLlama-7B-CoT | CoT | 47.3 | 32.6 | 40.3 |
| CodeLlama-7B-AoT | AoT | 55.4 | 85.6 | 69.9 |
| Llama-2-7B | CoT | 47.2 | 27.6 | 37.8 |
| Llama-2-7B | AoT | 47.5 | 66.3 | 56.5 |
| Llama-2-7B-Chat | CoT | 37.6 | 16.5 | 27.5 |
| Llama-2-7B-Chat | AoT | 34.2 | 15.8 | 25.4 |
| Llama-2-7B-CoT | CoT | 51.1 | 32.0 | 41.9 |
| Llama-2-7B-AoT | AoT | 55.4 | 75.2 | 64.9 |
| Mistral-7B | CoT | 61.1 | 48.9 | 55.2 |
| Mistral-7B | AoT | 62.3 | 80.4 | 71.0 |
| Mistral-7B-Instruct | CoT | 54.7 | 48.2 | 51.6 |
| Mistral-7B-Instruct | AoT | 56.1 | 75.6 | 65.4 |
| Mistral-7B-CoT | CoT | 61.4 | 40.8 | 51.6 |
| Mistral-7B-AoT | AoT | 71.3 | 84.8 | 77.7 |
| GPT-3.5-Turbo | CoT | 69.0 | 70.2 | 69.6 |
| GPT-3.5-Turbo | AoT | 72.1 | 94.3 | 82.7 |
| GPT-4 | CoT | 87.0 | 86.4 | 86.7 |
| GPT-4 | AoT | 86.7 | 97.0 | 91.6 |

Table 8: **3-shot** evaluation performance (%) on BBH.

B.4 Computational Cost Comparison

Considering that our AoT format encompasses multi-level abstractions, the average length of AoT responses tends to be longer than that of CoT responses. Consequently, the AoT-finetuned model may takes a higher computational cost during inference to generate more tokens. To clarify the impact of this in practical applications, we conduct a brief comparison of computational costs. We calculate the average length of responses from Llama-3-8B-CoT and Llama-3-8B-AoT on the BBH questions. On average, the responses of Llama-3-8B-CoT consist of 107.1 tokens, while the responses of Llama-3-8B-AoT consist of 188.9 tokens. Thus, AoT responses takes approximately 80% more tokens than CoT responses.

To estimate the generation time during inference, we use the vllm⁶ library on a server with 8 A100-SXM4-40GB GPUs, setting the batch size to 1. The average time for Llama-3-8B-AoT to generate one AoT response is 1.96 seconds. However, this estimation should only be taken as a rough refer-

⁶<https://github.com/vllm-project/vllm>

| Tasks | Human [†] | | Llama-3-8B- | | Mistral-7B- | | CodeLlama-7B- | |
|-------------------------------------|--------------------|-------|-------------|------|-------------|-------|---------------|------|
| | Avg. | Max | CoT | AoT | CoT | AoT | CoT | AoT |
| Boolean expressions* | 79.4 | 100.0 | 64.4 | 82.0 | 63.2 | 73.2 | 70.4 | 78.0 |
| Causal judgement | 69.6 | 100.0 | 56.7 | 61.0 | 56.7 | 63.6 | 52.4 | 57.2 |
| Date understanding | 76.8 | 100.0 | 53.6 | 16.4 | 47.6 | 40.8 | 37.6 | 21.2 |
| Disambiguation qa | 66.6 | 93.3 | 50.0 | 34.0 | 61.2 | 59.6 | 51.6 | 45.6 |
| Dyck languages* | 47.8 | 100.0 | 20.8 | 2.8 | 1.2 | 2.0 | 0.4 | 3.6 |
| Formal fallacies | 90.8 | 100.0 | 54.8 | 54.4 | 54.0 | 54.4 | 48.4 | 57.2 |
| Geometric shapes* | 54.0 | 100.0 | 20.0 | 26.0 | 30.8 | 21.6 | 20.0 | 34.0 |
| Hyperbaton | 74.7 | 100.0 | 49.2 | 71.6 | 72.4 | 49.2 | 69.6 | 62.8 |
| Logical deduction* | 40.3 | 88.9 | 50.3 | 52.5 | 45.1 | 51.9 | 36.7 | 37.2 |
| Movie recommendation | 60.7 | 90.0 | 60.4 | 62.0 | 60.0 | 63.6 | 48.0 | 58.0 |
| Multistep arithmetic two* | 9.7 | 25.0 | 4.4 | 98.0 | 1.2 | 100.0 | 0.8 | 96.8 |
| Navigate* | 81.9 | 100.0 | 61.6 | 62.8 | 50.4 | 64.8 | 45.6 | 57.6 |
| Object counting* | 86.1 | 100.0 | 50.0 | 66.0 | 45.6 | 85.2 | 42.4 | 46.8 |
| Penguins in a table | 78.0 | 100.0 | 54.1 | 52.7 | 44.5 | 50.7 | 45.2 | 51.4 |
| Reasoning about colored objects | 75.4 | 100.0 | 52.4 | 53.6 | 50.4 | 52.8 | 43.2 | 54.4 |
| Ruin names | 77.7 | 100.0 | 17.2 | 52.4 | 42.8 | 50.0 | 20.4 | 36.8 |
| Salient translation error detection | 36.7 | 80.0 | 40.4 | 46.4 | 44.8 | 51.2 | 39.6 | 43.6 |
| Snarks | 76.7 | 100.0 | 56.7 | 58.4 | 69.7 | 64.6 | 56.2 | 54.5 |
| Sports understanding | 70.8 | 100.0 | 70.8 | 57.2 | 42.0 | 59.6 | 52.0 | 54.8 |
| Temporal sequences* | 90.8 | 100.0 | 31.2 | 30.8 | 41.2 | 36.8 | 31.6 | 20.0 |
| Tracking shuffled objects* | 64.7 | 100.0 | 19.9 | 28.9 | 21.9 | 21.5 | 20.8 | 21.1 |
| Web of lies* | 81.3 | 100.0 | 52.0 | 57.6 | 51.6 | 54.4 | 49.6 | 51.6 |
| Word sorting* | 62.6 | 100.0 | 12.4 | 97.6 | 25.6 | 20.8 | 28.8 | 96.0 |
| All Tasks | 67.7 | 94.4 | 43.6 | 53.3 | 44.5 | 51.8 | 39.6 | 49.6 |

Table 9: Zero-shot performance (%) on each task of BBH. [†]indicates results from [Suzgun et al. \(2023\)](#). * indicates that the task is an algorithmic task.

| Error Type (%) | Description |
|-----------------------------|--|
| Reasoning Error (38%) | Errors due to deriving incorrect logical conclusions from given conditions. |
| Knowledge Lacking (16%) | Lack of world knowledge necessary to solve problems. |
| Task Misunderstanding (15%) | Failure to grasp the requirements or objectives of the task. |
| Condition Repetition (12%) | Simply repeat the given conditions and then leap directly to the answer. |
| Hallucination (9%) | Fabrication of non-existent conditions or misinterpretation of given conditions. |
| Runtime Error (5%) | Errors due to inexecutable programs or unformalized responses. |
| False Negative (5%) | Correct answers are incorrectly identified as incorrect. |

Table 10: The failure modes of Llama-3-8B-AoT.

ence, as it may be influenced by various factors such as hardware conditions.

Listing 6: Success cases of Llama-3-8B-AoT.

```

===== Case 1 from Task Hyperbaton =====

**Input**
Question: Which sentence has the correct adjective order:
Options:
(A) lovely tiny old-fashioned pyramidal Pakistani glass eating cat
(B) pyramidal lovely old-fashioned Pakistani tiny eating glass cat

**AoT Prediction**
Step 1: Identify the correct adjective order.
The correct adjective order is: descriptive adjective + size adjective + age adjective + origin
      adjective + material adjective + function adjective + animal adjective.

Step 2: Analyze each option.
(A) lovely tiny old-fashioned pyramidal Pakistani glass eating cat
(B) pyramidal lovely old-fashioned Pakistani tiny eating glass cat

Step 3: Compare the options to the correct adjective order.
Option (A) follows the correct adjective order: lovely (descriptive) + tiny (size) + old-fashioned (
      age) + pyramidal (shape) + Pakistani (origin) + glass (material) + eating (function) + cat (
      animal).
Option (B) does not follow the correct adjective order: pyramidal (shape) + lovely (descriptive) +
      old-fashioned (age) + Pakistani (origin) + tiny (size) + eating (function) + glass (material) +
      cat (animal).

Therefore, the answer is \boxed{(A)}.

===== Case 2 from Task Movie Recommendation =====

**Input**
Question: Find a movie similar to Independence Day, Star Wars Episode V - The Empire Strikes Back,
      American Beauty, Star Wars Episode VI - Return of the Jedi:
Options:
(A) Frankenweenie
(B) Ivan Vasilievich Back to the Future
(C) We Bought a Zoo
(D) Star Wars Episode IV - A New Hope

**AoT Prediction**
Step 1: Analyze the given movies.
- Independence Day: A science fiction movie about aliens invading Earth.
- Star Wars Episode V - The Empire Strikes Back: A science fiction movie about the ongoing conflict
      between the Rebel Alliance and the Galactic Empire.
- American Beauty: A drama movie about a suburban family and their struggles with life.
- Star Wars Episode VI - Return of the Jedi: A science fiction movie about the final battle between
      the Rebel Alliance and the Galactic Empire.

Step 2: Identify the common theme or genre.
The common theme among the given movies is science fiction.

Step 3: Compare the options.
(A) Frankenweenie: A stop-motion animated film about a boy and his dog, not a science fiction movie.
(B) Ivan Vasilievich Back to the Future: A Russian comedy film about a man who accidentally travels
      back in time, not a science fiction movie.
(C) We Bought a Zoo: A drama movie about a family buying a zoo, not a science fiction movie.
(D) Star Wars Episode IV - A New Hope: A science fiction movie set in a galaxy far, far away, similar
      to the given movies.

Step 4: Select the correct answer.
The movie that is similar to Independence Day, Star Wars Episode V - The Empire Strikes Back,
      American Beauty, and Star Wars Episode VI - Return of the Jedi is Star Wars Episode IV - A New
      Hope.

Therefore, the answer is \boxed{D}.

===== Case 3 from Task Tracking shuffled Objects =====

**Input**

```


Question: Alice, Bob, Claire, Dave, and Eve are holding a white elephant gift exchange. At the start of the event, they are each holding a present of a different color: Alice has a brown present, Bob has a black ball, Claire has a blue present, Dave has an orange ball, and Eve has a red present. As the event progresses, pairs of people swap gifts. First, Claire and Alice swap their gifts. Then, Bob and Eve swap their gifts. Then, Dave and Bob swap their gifts. Then, Alice and Claire swap their gifts. Finally, Eve and Dave swap their gifts. At the end of the event, Claire has the

Options: (A) brown present (B) black ball (C) blue present (D) orange ball (E) red present

****AoT Prediction****

````python`

`class Person:`

```
 def __init__(self, name, gift):
 self.name = name
 self.gift = gift
```

`def swap_gifts(person1, person2):`

```
 """
 Swap the gifts between two people.
 """
 temp = person1.gift
 person1.gift = person2.gift
 person2.gift = temp
```

`def main():`

```
 """
 The main function that calls the other functions and prints the answer.
 """
 alice = Person('Alice', 'brown')
 bob = Person('Bob', 'black')
 claire = Person('Claire', 'blue')
 dave = Person('Dave', 'orange')
 eve = Person('Eve', 'red')

 swap_gifts(claire, alice)
 swap_gifts(bob, eve)
 swap_gifts(dave, bob)
 swap_gifts(alice, claire)
 swap_gifts(eve, dave)

 if claire.gift == 'brown':
 print('A')
 elif claire.gift == 'black':
 print('B')
 elif claire.gift == 'blue':
 print('C')
 elif claire.gift == 'orange':
 print('D')
 elif claire.gift == 'red':
 print('E')
```

`main()`

`````

===== Case 4 from Task Penguins in a Table =====

****Input****

Question: Here is a table where the first line is a header and each subsequent line is a penguin:

name, age, height (cm), weight (kg) Louis, 7, 50, 11 Bernard, 5, 80, 13 Vincent, 9, 60, 11 Gwen, 8, 70, 15 For example: the age of Louis is 7, the weight of Gwen is 15 kg, the height of Bernard is 80 cm. What is the name of the last penguin sorted by alphabetic order?

Options: (A) Louis (B) Bernard (C) Vincent (D) Gwen (E) James

****AoT Prediction****

````python`

`def sort_penguins(penguins):`

```
 """
 Sort a list of penguins by their names in alphabetical order.
 """
 sorted_penguins = sorted(penguins, key=lambda penguin: penguin[0])
```

```

 return sorted_penguins

def get_last_penguin(sorted_penguins):
 """
 Get the last penguin in the sorted list.
 """
 last_penguin = sorted_penguins[-1]
 return last_penguin[0]

def main():
 """
 The main function that calls the other functions and prints the answer.
 """
 penguins = [['Louis', 7, 50, 11], ['Bernard', 5, 80, 13], ['Vincent', 9, 60, 11], ['Gwen', 8, 70, 15]]
 sorted_penguins = sort_penguins(penguins)
 last_penguin = get_last_penguin(sorted_penguins)
 options = ['Louis', 'Bernard', 'Vincent', 'Gwen', 'James']
 answer = options.index(last_penguin) + 1
 print(chr(96 + answer))

main()
"""

```

Listing 7: Error cases of Llama-3-8B-AoT.

```

===== Error Case: Reasoning Error =====

Input
Question: Rashida lies. Osvaldo says Rashida tells the truth. Sherrie says Osvaldo tells the truth.
 Gwenn says Sherrie tells the truth. Jerry says Gwenn lies. Does Jerry tell the truth?

AoT Prediction
Step 1: Analyze the statements made by each person.
- Rashida lies.
- Osvaldo says Rashida tells the truth.
- Sherrie says Osvaldo tells the truth.
- Gwenn says Sherrie tells the truth.
- Jerry says Gwenn lies.

Step 2: Evaluate the statements based on the given information.
- Rashida lies, so her statement is not reliable.
- Osvaldo says Rashida tells the truth, but since Rashida lies, Osvaldo's statement is not reliable.
- Sherrie says Osvaldo tells the truth, but since Osvaldo's statement is not reliable, Sherrie's
 statement is not reliable.
- Gwenn says Sherrie tells the truth, but since Sherrie's statement is not reliable, Gwenn's
 statement is not reliable.
- Jerry says Gwenn lies, but since Gwenn's statement is not reliable, Jerry's statement is not
 reliable.

Step 3: Determine if Jerry tells the truth.
Based on the given information, Jerry's statement is not reliable.

Therefore, the answer is \boxed{no}.

Predicion
no
Ground Truth
Yes
Error Explanation
The model derives incorrect logical conclusions from given conditions.

===== Error Case: Knowledge Lacking =====

Input
Is the following sentence plausible? "Courtland Sutton hit a triple."

AoT Prediction
Step 1: Analyze the sentence structure and meaning.
The sentence states that "Courtland Sutton hit a triple."

```

Step 2: Recall the definition of a triple in baseball.  
In baseball, a triple is a hit in which the batter reaches third base safely after hitting the ball.

Step 3: Evaluate the plausibility of the sentence.

It is plausible for Courtland Sutton to hit a triple in baseball, as it is a possible outcome for a batter.

Therefore, the answer is \boxed{Yes}.

**\*\*Predicion\*\***

Yes

**\*\*Ground Truth\*\***

No

**\*\*Error Explanation\*\***

The model lacks the necessary knowledge that Courtland Sutton usually refers to the American football wide receiver for the Denver Broncos of the National Football League (NFL).

===== Error Case: Task Misunderstanding =====

**\*\*Input\*\***

True or False or not True or False is

**\*\*AoT Prediction\*\***

Step 1: Analyze the given statement.

The statement is "True or False or not True or False is."

Step 2: Evaluate the statement.

The statement is grammatically incorrect and does not make sense. It is not clear what the statement is trying to convey.

Step 3: Determine the correct answer.

Since the statement is not true or false, the correct answer is \boxed{not True or False}.

**\*\*Predicion\*\***

not True or False

**\*\*Ground Truth\*\***

True

**\*\*Error Explanation\*\***

The model fails to understand that the task is to simplify the input boolean expression.

===== Error Case: Condition Repetition =====

**\*\*Input\*\***

Which of the following is a humorous edit of this artist or movie name: 'guns n' roses'?

Options:

- (A) guns n' goes
- (B) suns n' roses
- (C) guns n' rosjs
- (D) guns n' rgoses

**\*\*AoT Prediction\*\***

Step 1: Analyze the given options.

- (A) guns n' goes
- (B) suns n' roses
- (C) guns n' rosjs
- (D) guns n' rgoses

Step 2: Identify the humorous edit.

The humorous edit of "guns n' roses" is "guns n' goes".

Therefore, the answer is \boxed{(A)}.

**\*\*Predicion\*\***

(A)

**\*\*Ground Truth\*\***

(B)

**\*\*Error Explanation\*\***

The reasoning process simply repeats the given conditions and then leap directly to the predicted answer.

===== Error Case: Hallucination =====

**\*\*Input\*\***

I have two garlicks, a cabbage, a stalk of celery, five carrots, a lettuce head, four heads of broccoli, a potato, an onion, and a yam. How many vegetables do I have?

**\*\*AoT Prediction\*\***

````python`

```
def count_vegetables(vegetables):
```

```
    """
```

```
    Count the number of vegetables in a list.
```

```
    """
```

```
    count = len(vegetables)
```

```
    return count
```

```
def main():
```

```
    """
```

```
    The main function that calls the other functions and prints the answer.
```

```
    """
```

```
    vegetables = ['garlic', 'cabbage', 'celery', 'carrot', 'carrot', 'carrot', 'carrot', 'carrot',  
'carrot', 'lettuce', 'broccoli', 'broccoli', 'broccoli', 'broccoli', 'broccoli', 'broccoli', 'potato',  
'onion', 'yam']
```

```
    count = count_vegetables(vegetables)
```

```
    print(count)
```

```
main()
```

```
```
```

**\*\*Predicion\*\***

19

**\*\*Ground Truth\*\***

17

**\*\*Error Explanation\*\***

The model fabricates non-existent conditions or misinterprets given conditions. For example, the input states that there are 5 carrots, but 6 carrots are included in the variable.

===== Error Case: Runtime Error =====

**\*\*Input\*\***

Here is a table where the first line is a header and each subsequent line is a penguin: name, age, height (cm), weight (kg) Louis, 7, 50, 11 Bernard, 5, 80, 13 Vincent, 9, 60, 11 Gwen, 8, 70, 15  
For example: the age of Louis is 7, the weight of Gwen is 15 kg, the height of Bernard is 80 cm.  
How many penguins are more than 8 years old?

Options:

(A) 1

(B) 2

(C) 3

(D) 4

(E) 5

**\*\*AoT Prediction\*\***

````python`

```
def count_penguins_over_age(penguins, age):
```

```
    """
```

```
    Count the number of penguins that are older than a given age.
```

```
    """
```

```
    count = sum(1 for penguin in penguins if int(penguin[1]) > age)
```

```
    return count
```

```
def main():
```

```
    """
```

```
    The main function that calls the other functions and prints the answer.
```

```
    """
```

```
    penguins = ['Louis, 7, 50, 11', 'Bernard, 5, 80, 13', 'Vincent, 9, 60, 11', 'Gwen, 8, 70, 15']
```

```
    age = 8
```



```

count = count_penguins_over_age(penguins, age)
options = {1: 'A', 2: 'B', 3: 'C', 4: 'D', 5: 'E'}
answer = options.get(count, 'Option not found')
print(answer)

main()
'''

**Predicion**
None
**Ground Truth**
(A)
**Error Explanation**
The generated program is inexecutable.

===== Error Case: False Negative =====

**Input**
Is the following sentence plausible? "Andres Iniesta performed a give and go."

**AoT Prediction**
Step 1: Understand the meaning of the sentence.
The sentence "Andres Iniesta performed a give and go" suggests that Andres Iniesta participated in a
specific football technique called a "give and go".

Step 2: Recall the definition of a "give and go" in football.
A "give and go" is a football technique where a player passes the ball to a teammate, then quickly
moves towards the space created by the teammate's run. The teammate then returns the ball to the
player who initiated the move.

Step 3: Evaluate the plausibility of the sentence.
Given that Andres Iniesta is a professional football player, it is plausible that he would perform a
"give and go" technique during a game.

Therefore, the sentence "Andres Iniesta performed a give and go" is \boxed{plausible}.

**Predicion**
plausible
**Ground Truth**
Yes
**Error Explanation**
The correct predicted answer does not fully match the ground truth.

```

| Dataset Name | Source | Task Category |
|--------------------------|---------------|---|
| abductive nli | Nlv2 | Story Composition, Coherence Classification |
| ade corpus v2 | Nlv2 | Named Entity Recognition, Information Extraction, Text Categorization |
| adversarial qa@dbert | Nlv2, T0 | Adversarial QA, Question Answering |
| adversarial qa@dbidaf | Nlv2, T0 | Adversarial QA, Question Answering |
| adversarial qa@droberta | Nlv2, T0 | Adversarial QA, Question Answering |
| afs | Nlv2 | Text Matching |
| ai2 arc@ARC-Challenge | FLAN/T0 | Closed-Book QA, Inverted Closed-Book QA |
| ai2 arc@ARC-Easy | FLAN/T0 | Closed-Book QA, Inverted Closed-Book QA |
| ai2 arithmetic questions | Nlv2 | Question Answering |
| air dialogue | Nlv2 | Dialogue Generation, Speaker Identification, Intent Identification |
| ambigqa | Nlv2 | Question Answering, Question Rewriting |
| anli r1 | FLAN/T0 | Natural Language Inference, Inverted Natural Language Inference |
| anli r2 | FLAN/T0 | Natural Language Inference, Inverted Natural Language Inference |
| anli r3 | FLAN/T0 | Natural Language Inference, Inverted Natural Language Inference |
| aqua | Nlv2 | Question Answering |
| aquamuse | Nlv2 | Answerability Classification |
| argkp | Nlv2 | Text Matching |
| atomic | Nlv2 | Commonsense Classification, Fill in The Blank |
| babi | Nlv2 | Question Answering |
| bard | Nlv2 | Word Analogy |
| big bench | Nlv2 | Program Execution |
| bless | Nlv2 | Word Relation Classification, Word Semantics |
| blimp | Nlv2 | Linguistic Probing |
| break | Nlv2 | Question Decomposition |
| cad | Nlv2 | Toxic Language Detection |
| casehold | Nlv2 | Text Completion |
| casino | Nlv2 | Negotiation Strategy Detection |
| cb | Nlv2, FLAN/T0 | Natural Language Inference, Inverted Natural Language Inference, Textual Entailment |
| cfq mcd1 | Nlv2 | Text to Code |
| circa | Nlv2 | Question Context Generation, Dialogue Generation, Text Matching |
| civil comments | Nlv2 | Toxic Language Detection |
| clariq | Nlv2 | Question Understanding, Dialogue Generation, Intent Identification |
| coached conv pref | Nlv2 | Information Extraction, Speaker Identification |
| cod3s | Nlv2 | Cause Effect Classification |
| coda 19 | Nlv2 | Text Matching, Section Classification |
| codah | Nlv2 | Text Completion |
| cola | Nlv2, FLAN/T0 | Inverted Grammatical Acceptability, Text Quality Evaluation, Grammatical Acceptability, Grammar Error Detection |
| com2sense | Nlv2 | Commonsense Classification |
| com qa | Nlv2 | Question Answering, Question Rewriting |
| commonsenseqa | Nlv2 | Question Answering |
| conala | Nlv2 | Program Execution |
| conv ai 2 | Nlv2 | Speaker Identification |

| | | |
|--------------------------------------|---------------|---|
| cos e | T0 | Multiple-Choice QA (no trivia knowledge required) |
| cosmos qa | FLAN/T0 | Multiple-Choice QA (no trivia knowledge required), Inverted Multiple-Choice QA (no trivia knowledge required) |
| craigslist bargains | Niv2 | Text Categorization, Dialogue State Tracking |
| creak | Niv2 | Fact Verification |
| crows pair | Niv2 | Stereotype Detection |
| curated from stack over-flow english | Niv2 | Answerability Classification, Question Answering |
| curiosity dialogs | Niv2 | Information Extraction, Dialogue Generation, Speaker Identification |
| dailydialog | Niv2 | Text Categorization, Dialogue Act Recognition, Intent Identification, Sentiment Analysis |
| ddo | Niv2 | Text Categorization |
| deal or no dialogue | Niv2 | Dialogue State Tracking |
| defeasible nli atomic | Niv2 | Textual Entailment |
| detoxifying lms | Niv2 | Text Completion, Toxic Language Detection |
| dialogre | Niv2 | Speaker Relation Classification, Speaker Identification |
| diplomacy detection | Niv2 | Dialogue Generation |
| diqa | Niv2 | Answerability Classification, Question Answering |
| disfl qa | Niv2 | Answerability Classification, Question Answering, Text Quality Evaluation, Question Rewriting |
| dream | Niv2 | Question Understanding, Question Answering |
| drop | Niv2, FLAN/T0 | Inverted Mathematical QA, Mathematical QA, Question Understanding, Question Answering |
| dstc | Niv2 | Summarization, Dialogue State Tracking |
| duorc@ParaphraseRC | T0 | Extractive QA |
| duorc@SelfRC | T0 | Extractive QA |
| ecqa | CoT | Common Sense Reasoning Question Answering |
| emo | Niv2 | Sentiment Analysis |
| eqasc | Niv2 | Question Answering |
| equity evaluation corpus | Niv2 | Sentiment Analysis, Gender Classification, Fill in The Blank |
| esnli | CoT | Natural Language Inference |
| essential | Niv2 | Question Understanding, Answer Incomplete Questions |
| ethos | Niv2 | Toxic Language Detection |
| eurlex | Niv2 | Text Categorization |
| evaluation | Niv2 | Sentiment Analysis, Gender Classification, Fill in The Blank |
| event2mind | Niv2 | Emotion Word Generation, Sentiment Analysis, Emotional Reaction Generation, Intent Generation |
| financial phrasebank | Niv2 | Sentiment Analysis |
| freebase qa | Niv2 | Question Understanding, Question Answering |
| glucose | Niv2 | Cause Effect Classification, Information Extraction |
| go emotions | Niv2 | Sentiment Analysis |
| gooaq | Niv2 | Question Answering |
| google wellformed query | Niv2 | Question Understanding, Text Quality Evaluation |
| grailqa | Niv2 | Text Matching, Question Rewriting |
| gsm8k | CoT | Grade School Math Word Problems |
| gwsd | Niv2 | Textual Entailment |
| haspart kb | Niv2 | Entity Generation, Entity Relation Classification |

| | | |
|--|---------------|--|
| hate speech offensive | Nlv2 | Toxic Language Detection |
| hateeval | Nlv2 | Toxic Language Detection |
| hatexplain | Nlv2 | Toxic Language Detection |
| head qa | Nlv2 | Text Categorization, Question Answering |
| health fact | Nlv2 | Fact Verification, Sentence Composition, Explanation |
| help! need advice on identifying advice | Nlv2 | Text Categorization |
| hippocorpus | Nlv2 | Story Composition, Text Categorization |
| hope edi | Nlv2 | Text Categorization |
| hybridqa | Nlv2 | Question Answering, Pos Tagging |
| iirc | Nlv2 | Answerability Classification, Question Answering |
| impres | Nlv2 | Textual Entailment, Sentence Composition |
| indian food 101 | Nlv2 | Food Classification |
| inquistive | Nlv2 | Question Answering |
| jeopardy | Nlv2 | Answer Generation – Jeopardy Difficulty Double, Answer Generation – Jeopardy Difficulty Final, Answer Generation – Jeopardy Difficulty Normal, Answer Generation – Jeopardy Difficulty All |
| jigsaw | Nlv2 | Toxic Language Detection |
| kilt tasks@hotpotqa | T0 | Closed-Book QA |
| leetcode | Nlv2 | Text to Code |
| liar | Nlv2 | Text Categorization, Keyword Tagging |
| math dataset | FLAN | Inverted Mathematical QA, Mathematical QA |
| math qa | Nlv2 | Question Answering |
| mathmatics dataset | Nlv2 | Question Understanding, Question Answering |
| mathqa | Nlv2 | Question Answering |
| mcscrip | Nlv2 | Question Answering |
| mctaco | Nlv2 | Answerability Classification, Text Quality Evaluation, Question Understanding, Question Answering |
| medical question pair dataset | Nlv2 | Text Matching |
| meta woz | Nlv2 | Dialogue Act Recognition |
| miam | Nlv2 | Language Identification |
| miscellaneous | Nlv2 | Question Understanding, Question Answering, Paraphrasing |
| missing | Nlv2 | Question Answering |
| mnli | FLAN | Natural Language Inference, Inverted Natural Language Inference |
| mocah | Nlv2 | Question Answering |
| mrqa | Nlv2 | Question Answering |
| ms marco | Nlv2 | Question Answering |
| multi woz v2 | Nlv2 | Dialogue Generation, Speaker Identification |
| mutual | Nlv2 | Dialogue Generation |
| mwsc | Nlv2 | Question Answering |
| narrativeqa | Nlv2 | Question Generation |
| natural questions | Nlv2, FLAN/T0 | Closed-Book QA, Inverted Closed-Book QA, Question Answering |
| news headlines dataset for sacrasn detection | Nlv2 | Text Categorization |
| nlg bias | Nlv2 | Text Categorization |

| | | |
|-----------------------|------------------|---|
| nlu asdiv dataset | Nlv2 | Question Answering |
| numersense | Nlv2 | Fill in The Blank |
| ohsumed | Nlv2 | Information Extraction, Text Matching, Keyword Tagging |
| olid | Nlv2 | Toxic Language Detection |
| open pi | Nlv2 | Text Categorization |
| openbookqa | Nlv2, FLAN/T0 | Multiple-Choice QA (no trivia knowledge required), Sentence Composition, Inverted Multiple-Choice QA (no trivia knowledge required), Question Answering |
| opp 115 | Nlv2 | Information Extraction, Text Categorization |
| overruling | Nlv2 | Text Categorization |
| persent | Nlv2 | Named Entity Recognition, Sentiment Analysis |
| personachat | Nlv2 | Dialogue Generation |
| perspectrum | Nlv2 | Textual Entailment |
| piqa | Nlv2, FLAN/T0 | Multiple-Choice QA (no trivia knowledge required), Inverted Multiple-Choice QA (no trivia knowledge required), Question Answering |
| poki | Nlv2 | Poem Generation, Text Categorization |
| propara | Nlv2 | Named Entity Recognition, Information Extraction |
| prost | Nlv2 | Question Generation |
| proto qa | Nlv2 | Question Answering |
| pubmed qa | Nlv2 | Answer Verification, Intent Identification, Question Answering |
| qa srl | Nlv2 | Question Answering |
| qanta | Nlv2 | Text Categorization |
| qasper | Nlv2 | Question Understanding, Question Answering |
| qed | Nlv2 | Question Answering |
| qnli | FLAN/T0 | Multiple-Choice QA (no trivia knowledge required), Inverted Multiple-Choice QA (no trivia knowledge required) |
| qrecc | Dialog | Conversational Question Answering |
| qrecc input inversion | Dialog | Conversational Question Answering |
| quac | FLAN/T0 | Multiple-Choice QA (no trivia knowledge required), Inverted Multiple-Choice QA (no trivia knowledge required) |
| quail | Nlv2 | Question Answering |
| quarel | Nlv2 | Question Answering |
| quartz | Nlv2 | Explanation, Question Answering |
| question & answer zre | Nlv2 | Question Understanding, Question Answering |
| quoref | Nlv2 | Question Answering |
| race@high | T0 | Multiple-Choice QA (no trivia knowledge required), Span Generation |
| race@middle | T0 | Multiple-Choice QA (no trivia knowledge required), Span Generation |
| recepie nlg | Nlv2 | Named Entity Recognition, Fill in The Blank |
| roc stories | Nlv2 | Text Completion, Sentence Ordering, Coherence Classification |
| root09 | Nlv2 | Word Relation Classification, Word Semantics |
| ropes | Nlv2 | Story Composition, Question Answering |
| rte | Nlv2, FLAN/T0 | Natural Language Inference, Inverted Natural Language Inference, Textual Entailment |
| ruletaker | Nlv2 | Fact Verification |
| sarcasm in twitter | Nlv2 | Text Categorization |
| sbic | Nlv2 | Toxic Language Detection |
| schema guided dstc8 | Nlv2 | Text Categorization, Dialogue Act Recognition |

| | | |
|-----------------------------|---------------|--|
| scifact | Nlv2 | Text Matching, Text Quality Evaluation |
| sciq | Nlv2 | Explanation, Question Answering |
| scitail | Nlv2 | Textual Entailment, Sentence Composition, Question Answering |
| scitailv1.1 | Nlv2 | Textual Entailment, Sentence Composition |
| scitldr | Nlv2 | Summarization |
| scruples | Nlv2 | Text Categorization, Ethics Classification |
| semeval 2018 task3 | Nlv2 | Irony Detection |
| semeval 2019 task 10 | Nlv2 | Question Answering |
| semeval 2020 task4 | Nlv2 | Commonsense Classification, Explanation |
| semeval 2020 task 7 | Nlv2 | Text Categorization |
| smcalflow | Nlv2 | Dialogue Generation, Speaker Identification |
| snli | Nlv2, FLAN | Natural Language Inference, Sentence Composition, Inverted Natural Language Inference, Textual Entailment |
| social i qa | T0 | Multiple-Choice QA (no trivia knowledge required) |
| spolin | Nlv2 | Dialogue Generation, Dialogue Act Recognition |
| squad v1 | FLAN/T0 | Extractive QA, Inverted Extractive QA |
| squad v2 | FLAN/T0 | Extractive QA, Inverted Extractive QA |
| starcon | Nlv2 | Text Matching, Stance Detection |
| stereoset | Nlv2 | Text Categorization, Stereotype Detection, Fill in The Blank |
| storycommonsense | Nlv2 | Information Extraction, Intent Identification, Sentiment Analysis |
| strategyqa | Nlv2 | Question Decomposition |
| subjqa | Nlv2 | Question Answering |
| super glue@boolq | Nlv2 | Question Answering |
| super glue@multirc | Nlv2, FLAN/T0 | Answerability Classification, Question Answering, Inverted Multiple-Choice QA (no trivia knowledge required), Multiple-Choice QA (no trivia knowledge required), Answer Verification, Text Quality Evaluation |
| super glue@record | Nlv2, FLAN/T0 | Question Answering, Extractive QA, Inverted Extractive QA |
| svamp | Nlv2 | Question Answering |
| swag | Nlv2 | Text Completion |
| synthetic | Nlv2 | Clock Format Conversion, Information Extraction, Program Execution, Spelling Error Detection, Grammar Error Detection, Rhyme Generation, Pos Tagging, Date Validity Prediction, Mathematics, Temporal Reasoning, Leap Year Prediction, Edible Prediction |
| task master | Dialog | Dialog Next Turn Prediction |
| task master input inversion | Dialog | Dialog Next Turn Prediction |
| tellmewhy | Nlv2 | Answerability Classification, Question Answering |
| timetravel | Nlv2 | Story Composition, Coherence Classification |
| tom qa | Nlv2 | Question Answering |
| torque | Nlv2 | Information Extraction, Question Answering |
| trianglecopa | Nlv2 | Question Generation |
| trivia qa | FLAN/T0 | Extractive QA, Inverted Extractive QA |
| tweetqa | Nlv2 | Answerability Classification, Answer Verification, Question Answering |
| web questions | Nlv2, T0 | Closed-Book QA, Question Answering |
| wiki dialog | Dialog | Conversational Question Answering |
| wiki dialog input inversion | Dialog | Conversational Question Answering |

| | | |
|----------------------|----------|---|
| wiki hop | Nlv2, T0 | Multiple-Choice QA (no trivia knowledge required), Question Answering |
| wiki movies | Nlv2 | Question Answering |
| wiki qa | Nlv2, T0 | Answer Verification, Closed-Book QA |
| wikitext 103 | Nlv2 | Fill in The Blank |
| winowhy | Nlv2 | Coreference Resolution, Commonsense Classification, Explanation |
| wiqa | Nlv2 | Sentence Ordering, Question Answering |
| wnli | FLAN | Coreference Resolution, Inverted Coreference Resolution |
| x_csr | Nlv2 | Question Answering, Linguistic Probing |
| xcsr | Nlv2 | Sentence Perturbation |
| xl wic | Nlv2 | Word Semantics, Sentence Composition |
| xquad | Nlv2 | Question Answering |
| yahoo answers topics | Nlv2 | Text Categorization |
| yoruba bbc topics | Nlv2 | Text Categorization |

Table 11: The datasets included in our AOT COLLECTION, along with their sources and task categories. The sources and task categories are aligned with the FLAN Collection (https://github.com/google-research/FLAN/blob/main/flan/v2/flan_collection_info.csv).