

OPT-R: Exploring the Role of Explanations in Finetuning and Prompting for Reasoning Skills of Large Language Models

Badr AlKhamissi Siddharth Verma Ping Yu Zhijing Jin
Asli Celikyilmaz Mona Diab
Meta AI

Abstract

In this paper, we conduct a thorough investigation into the reasoning capabilities of Large Language Models (LLMs), focusing specifically on the Open Pretrained Transformers (OPT) models as a representative of such models. Our study entails finetuning three different sizes of OPT on a carefully curated reasoning corpus, resulting in two sets of finetuned models: OPT-R, finetuned without explanations, and OPT-RE, finetuned with explanations. We then evaluate all models on 57 out-of-domain tasks drawn from the SUPER-NATURALINSTRUCTIONS benchmark, covering 26 distinct reasoning skills, utilizing three prompting techniques. Through a comprehensive grid of 27 configurations and 6,156 test evaluations, we investigate the dimensions of finetuning, prompting, and scale to understand the role of explanations on different reasoning skills. Our findings reveal that having explanations in the fewshot exemplar has no significant impact on the model’s performance when the model is finetuned, while positively affecting the non-finetuned counterpart. Moreover, we observe a slight yet consistent increase in classification accuracy as we incorporate explanations during prompting and finetuning, respectively. Finally, we offer insights on which skills benefit the most from incorporating explanations during finetuning and prompting, such as *Numerical* (+20.4%) and *Analogical* (+13.9%) reasoning, as well as skills that exhibit negligible or negative effects.

1 Introduction

Recently, there has been a surge in the release of Large Language Models (LLMs) by both industrial and academic institutions. These models vary from open-source releases such as OPT (Zhang et al., 2022) and LLAMA (Touvron et al., 2023) to closed-source ones like GPT-3 (Brown et al., 2020) and PALM (Chowdhery et al., 2022). In addition, researchers have developed models that are finetuned on top of these foundational models to better

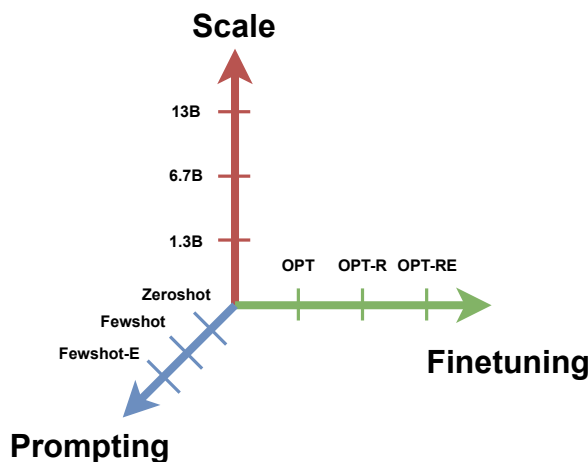


Figure 1: **Three-Dimensional Grid of Fine-Tuning, Prompting, and Scale.** Each dimension is represented as an axis, with three levels for each of finetuning, prompting, and scale plotted on each axis. The resulting grid consists of 27 different combinations evaluated on various reasoning tasks. It should be noted that there is a hidden dimension, the scoring function, comprising four components. This results in a comprehensive total of 6,156 evaluations.

follow instructions, such as OPT-IML (Iyer et al., 2022) and Alpaca (Taori et al., 2023). Despite the remarkable progress in LLMs’ performance in Natural Language Processing (NLP) tasks, reasoning remains a challenging area. For example, prior work have shown that LLMs struggle with commonsense reasoning (West et al., 2022) and arithmetic reasoning (Hendrycks et al., 2021) to name a few.

Recent efforts have attempted to improve the reasoning performance of LLMs by decomposing answers into step-by-step reasoning chains using in-context learning (Wei et al., 2022b; Kojima et al., 2022) or during finetuning (Chung et al., 2022; Wei et al., 2021a). While these approaches have shown some improvement on benchmarks such as GSM8K (Cobbe et al., 2021), it is not clear how those explanations affect finetuning, prompting, or




Figure 2: Template used during both training and inference. The model is tasked with predicting the answer followed by the explanation.

their combination. Concurrent work has investigated the generalization capability of such models to reasoning skills beyond those encountered during finetuning (Yu et al., 2022), but a comprehensive evaluation of the role of explanation during finetuning and prompting with respect to reasoning skills is still lacking.

In this paper, we aim to address this gap. We investigate OPT (Zhang et al., 2022) as a representative of such models and utilize it as our base model. Through finetuning OPT on a collection of carefully curated open-source reasoning datasets that come with explanations for each instance, we evaluate its performance on 57 tasks drawn from the SUPER-NATURALINSTRUCTIONS benchmark (Wang et al., 2022), covering 26 different reasoning skills. Our experiments are structured around three key dimensions: finetuning, prompting, and scale, each of which is comprised of three distinct components (See Figure 1). **Finetuning:** (1) a (vanilla) unfinetuned OPT model; (2) A finetuned OPT model without explanations (OPT-*R*); and, (3) A finetuned OPT model with explanations (OPT-*RE*). **Prompting:** (1) zero-shot prompting; (2) Fewshot prompting without explanations; and, (3) Fewshot prompting with explanations. Finally, **Scale:** (1) 1.3B; (2) 6.7B; and, (3) 13B. Accordingly, we create grid of 27 different components, providing a detailed analysis measuring the impact of explanations during finetuning and inference across different model scales.

Our findings reveals that finetuning on reasoning datasets leads to statistically significant improvements in seven reasoning skills, including *Numerical*, *Analogical* and *Reasoning on Objects*, with *Physical*, *Counting* and *Textual Entailment* showing a significant effect only for the OPT-*RE* model, across both fewshot prompting conditions

and model sizes, as compared to the vanilla OPT model (see Table 2). However, we also find that this approach significantly hinders the performance of three other reasoning skills (see Table 3). We also investigate the impact of incorporating explanations during fewshot prompting and find that it does not have a significant impact on the performance of the finetuned models, as measured by the variance in the difference between both prompting methods across reasoning skills for each model. However, we notice that it has a more noticeable effect on the performance of the vanilla OPT model, as shown in Table 5. Additionally, we observe a consistent increase in the average performance across all tasks from Fewshot to Fewshot-E, as well as from OPT to OPT-*R* to OPT-*RE* models, indicating that explanations do have a small effect on performance during both finetuning and prompting. Finally, Table 4 presents a summary of the results, indicating which reasoning skills demonstrate improvement due to the incorporation of explanations during either finetuning or prompting, which skills show a negative effect, and which skills have negligible effects regarding explanations.

2 OPT-*R*: Finetuning on Reasoning Skills

2.1 Reasoning Datasets with Explanations

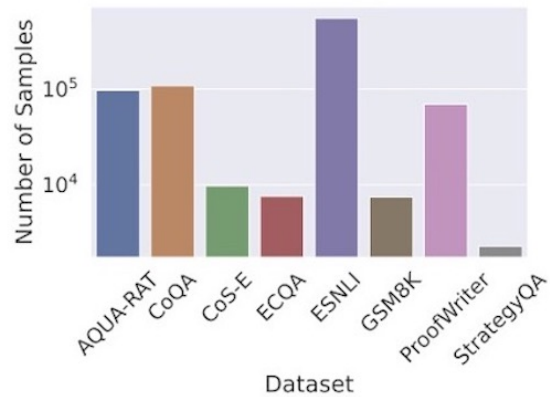


Figure 3: Number of samples in each dataset of the training corpus. Y-axis in log scale.

The finetuning corpus utilized to refine OPT is composed of various reasoning datasets, each of which includes a corresponding explanation or rationale for the answer. These rationales may consist of a sequence of smaller steps (i.e. chain-of-thought) or a free-form text that elucidates the reasoning behind the answer. As shown in Figure 2, we employ a uniform template for all tasks during

the training process. The input to the model begins with a task definition, followed by an instruction to provide an answer followed by a brief reasoning. Next, we extract two random in-context examples uniformly from the training set that remain constant throughout training for each instance. The input for the current training instance is then presented in a format specific to each task. The options for the answer are then included in the input, but not in the in-context examples (see Appendix A for further details on task-specific definitions and options). The options are pre-shuffled for each training instance. The model is finally provided with the answer prefix, "Output: The answer is", and is tasked to predict the answer, followed by an explanation if *OPT-RE* is being finetuned. Similarly, the in-context examples only comprise an explanation when training *OPT-RE*.

Below is a brief description of each dataset used during finetuning. See Figure 3 for the relative size of each dataset.

AQUA-RAT The Algebra Question Answering with Rationales dataset (Ling et al., 2017) rendering the task of solving algebraic word problems more feasible by dividing the problem into a series of smaller steps. They create a 100k-sample dataset that contains questions, answers and rationales in natural language and human-readable mathematical expressions that can be used to derive the final answer.

CoQA The Conversational Question Answering dataset Reddy et al. (2019). It consists of 127k questions and answers, compiled from 8k conversations about passages from seven different domains. Given a passage that contains a conversation, the model is tasked with answering a question by highlighting the corresponding evidence from the passage.

CoS-E The Common Sense Explanations dataset Rajani et al. (2019) to induce language models with commonsense reasoning. In this dataset, the model is given a question and a set of choices and is tasked with selecting one of the provided choices along with providing an explanation in natural language as to why that choice is correct.

ECQA The Explanations for Commonsense Question Answering dataset Aggarwal et al. (2021). It is similar to CoS-E since it requires the model to choose one of the provided options to answer the

given question, and also provide an explanation.

ESNLI The Stanford Natural Language Inference dataset with Explanations Camburu et al. (2018) to train models to provide interpretable and robust explanations for their decisions. The authors extend the SNLI dataset (Bowman et al., 2015) with human-annotated explanations. Similar to any NLI task, the model is given a premise and hypothesis and the task is to determine whether the hypothesis sentence entails, contradicts, or is neutral with respect to the given premise.

GSM8K The Grade School Math dataset Cobbe et al. (2021) to train models to better perform multi-step mathematical reasoning. It consists of 8.5k linguistically diverse grade school math word problems. Therefore, the task for the model is to answer the question by performing a series of arithmetic operations to obtain a final answer, while explaining its reasoning steps.

ProofWriter The ProofWriter dataset Tafjord et al. (2021) to generate both the implications of a theory from the RuleTaker dataset (Clark et al., 2020) and the natural language proofs that support them. Specifically, given a sequence of facts and rules, the model is tasked with answering a question using "Yes", "No", or "Unknown" and provide the reasoning path by referring to the provided facts and rules. We consider the open-world assumption subset of RuleTaker with questions that requires reasoning up to a depth of 5.

StrategyQA The Strategy Question Answering dataset Geva et al. (2021) to improve multi-hop reasoning for questions where the required reasoning steps are implicit in the question. Therefore, the task of the model is to answer the question using "Yes" or "No" then provide a strategy that explains the answer by decomposing it into a number of steps.

2.2 Finetuning Procedures

OPT The Open Pretrained Transformers (OPT) models are a suite of decoder-only pre-trained transformers ranging from 125M to 175B parameters released by Zhang et al. (2022). In this work, we use three OPT models with sizes of 1.3B, 6.7B and 13B. The details of each model architecture, pre-training corpus and training configuration (e.g. weight initialization, optimizer, tokenizer, hyperparameters, etc.) can be found in Zhang et al. (2022).

Reasoning Skill	Task IDs
Abductive Reasoning	task854
Analogical Reasoning	task1287, task1288
Argument Reasoning	task514
Causal Reasoning	task1393
Commonsense Reasoning	task279, task156, task295
Commonsense Reasoning → Numerical Commonsense ...	task1403
Commonsense Reasoning → Physical Reasoning	task084
Commonsense Reasoning → Social Situations	task580, task937, task1606
Commonsense Reasoning → Spatial Reasoning	task082, task083
Deductive Reasoning	task221, task1568, task220
Ethics	task667, task724, task723
Grammatical Reasoning	task1712, task052, task1559
Logical Reasoning	task717, task211, task268
Logical Reasoning → Reasoning with Symbols	task923, task935
Mathematics → Counting	task523, task155
Multihop Reasoning	task1297, task056
Numerical Reasoning	task621, task1333
Reasoning on Objects	task1583, task1584
Reasoning on Social Interactions	task609, task881, task875
Reasoning on Strings	task1189
Relational Reasoning	task1380, task472, task1505
Scientific Reasoning	task1431, task228, task714
Temporal Reasoning	task018, task1549, task383
Textual Entailment	task738, task890, task463
Textual Entailment → Analogical Reasoning	task1347
Textual Entailment → Deductive Reasoning	task1612, task534, task1366

Table 1: Evaluation tasks from SUP-NATINST (Wang et al., 2022) used for each reasoning skill.

Implementation Details To finetune the selected models, we utilized the metaseq¹ implementation since it enables higher training efficiency compared to other codebases (Zhang et al., 2022). Each model is finetuned twice for 10 epochs, once with explanations and once without (i.e. OPT-RE vs OPT-R, respectively). Models are evaluated at the end of each epoch on a chosen set of SUPER-NATURALINSTRUCTIONS validation tasks, and the checkpoint with the best performance is selected for evaluation on the testing tasks. The loss is calculated only on the tokens the model is tasked to predict during inference, and not the full input, what is referred to as label-loss in (Iyer et al., 2022). The samples across all datasets are shuffled during training. Further, the model is provided with two in-context examples during finetuning in addition to the task definition to match inference time following (Wang et al., 2022).

¹<https://github.com/facebookresearch/metaseq>

3 Evaluating the Models

3.1 SUPER-NATURALINSTRUCTIONS Tasks

In this study, we focus on a subset of the SUPER-NATURALINSTRUCTIONS benchmark version 2.6² (SUP-NATINST for short) proposed by Wang et al. (2022), which comprises 1,616 varied NLP tasks and includes meta-labels for each task, such as task type, domain and more importantly for this work: the underlying reasoning skills. Specifically, we select a subset of tasks that satisfy two key criteria: (i) the task focuses on a single reasoning skill, enabling us to evaluate a specific atomic skill, and (ii) the task can be tested using classification mode, as detailed in Section 3.2. Note that there is no data contamination between finetuning data and the evaluation benchmark.

²We downloaded the data from <https://github.com/allenai/natural-instructions/tree/v2.6>.

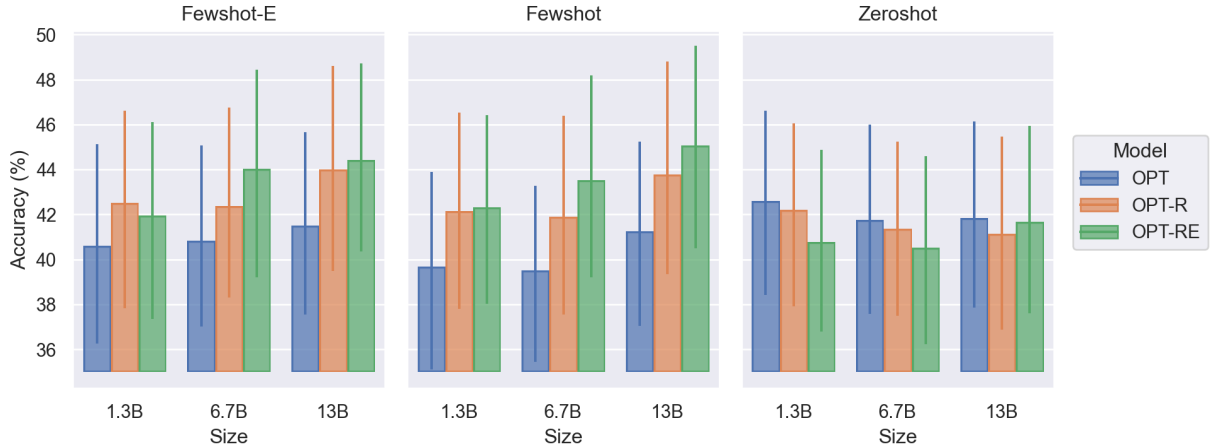


Figure 4: Results achieved across all tasks as a function of the three primary dimensions analyzed in this study: Finetuning, Prompting and Scale.

Benchmark Splits Following the task selection process, we apply a random sampling technique to ensure diversity within the testing set. Specifically, we select a maximum of three tasks from each reasoning skill, and allocate any remaining tasks to the validation set. Notably, this approach enables us to obtain a representative sample of the selected reasoning skills for testing, while also ensuring that our model’s performance is not influenced by a particular subset of tasks. Table 1 shows the complete list of tasks used for evaluating our finetuned models for each reasoning skill.

3.2 Evaluation Setup

Earlier, we mentioned that we selected 57 tasks spanning 26 reasoning skills from SUP-NATINST to evaluate our finetuned models. To meet our criteria, as detailed in Section 3.1, each task had to fulfill two conditions. The second condition required that the task can be considered a classification task. That means there is a discrete set of candidates (one of which is correct) and thereby treating it as a classification problem where the highest-scoring candidate is considered the answer. To ensure this, we utilized a straightforward heuristic: we only sampled tasks that had no more than 10 possible candidate answers.

Classification Method To determine the correct answer, we conduct a forward pass for each potential candidate answer and utilize a scoring function to measure the likelihood that the candidate tokens follows the input, similar to Brown et al. (2020). This process is repeated four times using distinct scoring functions, as detailed in the subsequent

paragraph. The highest accuracy score from the four scoring functions is considered as result of the task.

Scoring Functions This is considered the fourth dimension of this work since we evaluate each task using four different scoring functions and take the maximum accuracy as the result. The four scoring functions used are as follows: (1) **mean**, which involves computing the average of the log probabilities of candidate tokens, also referred to as token score. (2) **unconditional-norm**, which computes the difference between the sum of token scores of the candidate when unconditioned by any previous tokens and the sum of candidate token scores when conditioned by previous input. (3) **suffix**, which computes the sum of the conditioned candidate’s token scores alone. Finally, (4) **sum**, which involves calculating the sum of all the token scores passed to the model. The reason we employed different functions is that we observed significant gains in performance when using one scoring function over the other for specific tasks. Therefore, in order to ensure fairness across all tasks, we selected the highest accuracy over all scoring functions for each task.

4 Results & Findings

In this section, we present the results and findings of our experiments. First, we illustrate in Figure 4 the outcome of our evaluation on the effectiveness of finetuned models as compared to the vanilla OPT model, across three different scales when using both fewshot prompting with and without explanations. Furthermore, we observe a monotonic

increase in the performance of each model as we increase the scale under those two prompting conditions, which indicates a positive correlation between the model’s capacity and its overall performance. However, we note that this trend does not apply to the zeroshot prompting method, since we are testing out-of-distribution tasks and that the finetuned models were trained with fewshot exemplars in their context. This leads us to focus only on the fewshot prompting methods, with and without explanations, for the remaining of our evaluations. Specifically, we investigate the impact of finetuning the OPT models on reasoning datasets, as compared to the vanilla OPT model, and explore the effect of explanations during finetuning and prompting, both in terms of the reasoning skill.

4.1 Model Performance for Reasoning Skills

The results reported in this and the following section are the classification accuracy of each reasoning skill across different conditions, such as model sizes and fewshot prompting methods. Table 2 shows the reasoning skills where either OPT-RE or OPT-R are significantly better than the vanilla OPT model, as measured by Welch’s t-test, where $p < 0.05$. Conversely, Table 3 shows the reasoning skills where the vanilla OPT model performs significantly better than either of its finetuned counterparts.

Skill	OPT	OPT-R	OPT-RE
Numerical	44.8	65.2*	64.7*
Analogical	49.0	62.9*	60.8*
Counting	19.8	13.1	31.3*
Physical	38.2	37.8	49.1*
Entailment	42.6	47.2	51.6*
Social Int	34.1	43.0*	40.1
Objects	54.3	62.6*	59.9*

Table 2: Performance as a function of the reasoning skills where OPT-RE or OPT-R performs significantly better than the OPT model as measured by Welch’s t-test ($p < 0.05$) denoted by the * symbol. The performance is measured across Fewshot and Fewshot-E prompting, the three different scales and tasks under the corresponding reasoning skill. Best result indicated in **bold**.

The results reveal that the finetuned variants of the OPT model demonstrate a significant improvement on seven distinct reasoning skills, with particular emphasis on the *Numerical* and *Analogical* reasoning tasks. Specifically, for the *Mathematical*

Skill	OPT	OPT-R	OPT-RE
Argument	57.9	46.1 ⁻	48.7 ⁻
TE - Deductive	36.0	29.0 ⁻	29.4 ⁻
Commonsense	33.4	29.7	28.8 ⁻

Table 3: Performance as a function of the reasoning skill where OPT performs significantly better than either OPT-R or OPT-RE as measured by Welch’s t-test ($p < 0.05$) denoted by the ⁻ symbol. The performance is measured across Fewshot and Fewshot-E prompting, the three different scales and tasks under the corresponding reasoning skill. TE is Textual Entailment.

Counting skill, the OPT-RE variant outperforms both the OPT-R and OPT models, underscoring the criticality of incorporating explanations during the finetuning process for mathematical datasets. Likewise, the *Physical Reasoning* tasks exhibit a similar trend. On the other hand, we can see that for the *Argument*, *Deductive Textual Entailment* and *Commonsense* skills the non-finetuned version outperforms considerably.

4.2 Fine-Grained Skill Analysis

Table 4 shows the classification accuracy results obtained from the three models, in relation to the reasoning skill and few-shot prompting method used. The best accuracy value for each reasoning skill is indicated in **bold**, and the cells are shaded with colors ranging from green to white to indicate their position in the accuracy spectrum of each reasoning skill. The skills with similar performance across different models are assigned a lighter shade of green, indicating that their color spectrum ends earlier than that of other skills where the difference in performance between models is more significant. The table is divided into four blocks to distinguish effects of finetuning and prompting methods on reasoning skills: the first block showcases skills where the finetuned (OPT-RE and OPT-R) models outperform the vanilla OPT model, the second block highlights skills where OPT-RE has better accuracy than other models therefore illustrating the importance of finetuning on explanations on those skills. The third block displays skills where OPT outperforms other models showing that finetuning actually hurts performance in this case, and the fourth block identifies skills where the choice of model or prompting method has little impact on the overall performance.

Skill	OPT		OPT-R		OPT-RE	
	Fewshot	Fewshot-E	Fewshot	Fewshot-E	Fewshot	Fewshot-E
Numerical	39.9	49.7	65.1	65.3	64.7	64.8
Analogical	51.9	46.2	63.3	62.5	60.7	60.9
Objects	53.5	55.1	61.4	63.8	60.0	59.7
Social Interactions	33.6	34.7	43.8	42.3	40.2	40.0
Textual Entailment	43.3	42.0	47.1	47.3	51.9	51.2
Grammatical	54.4	55.1	61.2	60.0	62.0	63.1
Multihop	36.6	31.7	38.9	39.9	39.5	37.0
Symbols	44.2	47.2	51.7	51.8	51.9	52.4
Spatial	44.1	47.1	49.8	51.8	49.6	49.2
Social Situations	46.3	46.6	53.2	53.2	51.9	52.3
Counting	19.6	20.0	13.5	12.7	29.8	32.9
Physical	35.8	40.6	36.9	38.8	48.1	50.0
Logical	31.7	33.4	33.7	34.1	36.9	38.4
Temporal	50.7	49.7	43.4	46.5	48.5	38.5
Argument	55.8	60.1	46.3	45.9	48.6	48.8
TE - Deductive	33.7	38.3	27.9	30.1	29.0	29.9
Relational	47.4	51.1	47.6	47.9	44.8	44.6
Commonsense	35.0	31.8	29.8	29.5	28.5	29.2
TE - Analogical	16.3	18.7	18.6	20.7	18.7	18.1
Abductive	33.9	36.1	36.9	34.4	34.2	35.3
Ethics	26.8	25.8	26.5	25.9	26.2	27.6
Deductive	39.4	40.4	39.4	40.4	40.0	41.1
Causal	50.2	50.6	49.1	48.9	50.1	50.5
Scientific	23.4	23.3	24.3	24.5	25.0	24.5
Numerical Commonsense	59.5	59.2	59.0	59.0	59.2	59.4
Strings	60.7	60.7	61.1	61.2	60.7	60.7

Table 4: Classification accuracy results achieved by different models as a function of the reasoning skill and few-shot prompting method employed. The best accuracy obtained for each reasoning skill is highlighted in bold. The cells are shaded with colors ranging from green to white to indicate their position in the accuracy spectrum. Reasoning skills with smaller variance in achieved results are assigned a lighter shade of green to convey the extent of similarity between models. The first block highlights skills where the finetuned models perform notably better than the vanilla OPT. The second block emphasizes the skills where OPT-RE outperforms other models. In contrast, the third block showcases the skills where OPT outperforms the other models. Lastly, the fourth block identifies skills where the choice of model or prompting method has little impact on the overall performance.

Explanations’ Effect One of the central questions that we sought to investigate in this study is the extent to which explanations play a role in improving the reasoning capabilities of OPT models during finetuning and prompting. The results presented in Table 5 suggest that the presence or absence of explanations in the fewshot examples employed for prompting does not *significantly* impact the performance of the model when the model is finetuned on reasoning datasets. Concretely, in Table 5, we present the variance of the absolute accuracy difference for each model across reason-

ing skills by excluding the *Temporal* skill, which was identified as an outlier. Specifically, we compute the difference between the two corresponding columns for each model in Table 4. These values provide insights into the impact of including explanations during prompting on the performance of the models. Our findings reveal that the difference is negligible for OPT-R and OPT-RE models, suggesting that the choice of prompting method does not significantly affect the model’s accuracy. However, for the vanilla OPT model, the difference is more substantial, emphasizing the impor-

tance of employing explanations during fewshot prompting. However, the mean performance of each model across the distinct fewshot prompting methods demonstrates a slight yet consistent increase in classification accuracy, from Fewshot to Fewshot-E (incorporating explanations), as well as from OPT to OPT-R to OPT-RE models showing that explanations do have a small effect on performance during both finetuning and prompting.

Model	Std(IF-FE)	Avg(F)	Avg(FE)
OPT	2.31	40.68	41.82
OPT-R	0.84	43.44	43.68
OPT-RE	0.78	44.49	44.86

Table 5: The first column shows the variance of the absolute difference in accuracy for each model across different reasoning skills, when using Fewshot (F) and Fewshot-E (FE) prompting methods. The second and third columns show the average performance of each model across each prompting method. Results are obtained after dropping the outlier *Temporal* skill.

5 Related Work

Reasoning LLMs LLMs have made significant advancements in the field of NLP and related areas (Brown et al., 2020; Chowdhery et al., 2022; Chung et al., 2022), especially with the advent of the pre-train, prompt, and predict paradigm (Liu et al., 2021). This paradigm has enabled these models to solve a multitude of tasks through in-context fewshot or zeroshot learning using instructions (Wei et al., 2021b; Iyer et al., 2022). However, their reasoning abilities have been a subject of debate in recent literature (Huang and Chang, 2022; AIKhamissi et al., 2022). Several studies suggest that increasing the size of an LM trained through the same next-token prediction method can lead to the emergence of complex behaviors (Wei et al., 2022a), including reasoning. For instance, some research has demonstrated that sufficiently large LMs can use chain-of-thought prompting (Wei et al., 2022b) to simulate human-like reasoning. Other studies have shown that the addition of a simple prompt, such as "Let's think step-by-step" (Kojima et al., 2022) can elicit reasoning abilities in LLMs by generating explicit reasoning steps before decoding the final answer. However, some researchers contend that emulating the human reasoning thought process is distinct from claiming that the model can truly reason (Wei et al., 2022b).

Finetuned LLMs Concurrent studies have finetuned LLMs to follow instructions to improve their generalization ability to unseen tasks through zero and fewshot learning (Iyer et al., 2022; Chung et al., 2022). However, our approach differs in that we only finetune on a selected number of open-source datasets that provide explanations for each instance. This enables us to focus on the importance of explanations during finetuning in the context of reasoning skills. While concurrent works, such as (Iyer et al., 2022; Wang et al., 2022), have experimented with different prompting methods during finetuning and inference, our study focuses primarily on evaluating the reasoning ability of the finetuned models across a set of reasoning skills. Other concurrent studies have explored the impact of finetuning on a set of held-out reasoning tasks (Yu et al., 2022), but their evaluation approach, which involves generating answers, may be influenced by various factors such as decoding strategy, decoding parameters, and prompt templates. In contrast, we adopt a rank classification approach similar to (Brown et al., 2020), which better captures the reasoning performance of the model being evaluated, in addition to covering a larger number of reasoning skills and tasks.

6 Conclusion

In this study, we investigated the impact of incorporating explanations during finetuning and prompting on three different sizes of the OPT model. Through a systematic and comprehensive evaluation process that considered three key dimensions, we found that while explanations did provide a small improvement in performance, the effect was not significant when incorporated in the in-context demonstrations during inference for the finetuned models. Additionally, our results showed that both finetuned models exhibited significant improvements in reasoning skills such as *Numerical*, *Analogical* and *Reasoning on Objects*. Moreover, we demonstrated that skills such as *Physical*, *Counting*, and *Textual Entailment* benefited from incorporating explanations during the finetuning process. Overall, our findings provide insights into the impact of incorporating explanations on the reasoning capabilities of LLMs and offer guidance on which reasoning skills would benefit most from the inclusion and exclusion of explanations during finetuning and prompting.

Limitations

While our study provides valuable insights into the impact of finetuning on reasoning performance and the role of explanations during finetuning and prompting with respect to various reasoning skills, there are several limitations to our work. Firstly, we only consider a single LLM, OPT, as our base model. Our results may not generalize to other LLMs with different architectures or pretraining objectives. Secondly, we only use a limited set of reasoning datasets for finetuning due to the limited availability of open-source datasets with explanations. However, it is possible that our findings may not hold for models finetuned on larger closed datasets as usually seen in real-world scenarios. Thirdly, our experiments only cover a limited range of model sizes due to limitations in computational budget, therefore it is possible that our findings may not hold for much larger models. Finally, we only consider finetuning using fewshot prompting conditions in our experiments, and it is possible that our findings may not hold for models finetuned without in-context exemplars. Overall, while our study provides valuable insights into the impact of finetuning and explanations on reasoning performance, further research is needed to investigate these factors across a broader range of models, datasets, and finetuning strategies.

Ethics Statement

This work is based on analyzing and evaluating the performance of LLMs on reasoning tasks using existing public datasets. No personally identifiable information or sensitive data was collected or used in this research. We acknowledge the potential risks of developing LLMs, including their potential impact on spreading misinformation, generating unwanted content and the exacerbation of existing biases in datasets. Our work aims to contribute to improving the transparency and understanding of how LLMs can be optimized for specific reasoning skills. We hope our findings will inspire further research on developing ethical and responsible approaches for developing and deploying LLMs.

References

Shourya Aggarwal, Divyanshu Mandowara, Vishwa-jeet Agrawal, Dinesh Khandelwal, Parag Singla, and Dinesh Garg. 2021. [Explanations for CommonsenseQA: New Dataset and Models](#). In *Proceedings of the 59th Annual Meeting of the Association for*

Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 3050–3065, Online. Association for Computational Linguistics.

Badr AlKhamissi, Millicent Li, Asli Celikyilmaz, Mona T. Diab, and Marjan Ghazvininejad. 2022. A review on language models as knowledge bases. *ArXiv*, abs/2204.06031.

Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. [e-snli: Natural language inference with natural language explanations](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 9539–9549. Curran Associates, Inc.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.

Peter Clark, Oyvind Tafjord, and Kyle Richardson. 2020. Transformers as soft reasoners over language. In *IJCAI*.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *ArXiv*, abs/2110.14168.

Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *NeurIPS*.

- Jie Huang and Kevin Chen-Chuan Chang. 2022. Towards reasoning in large language models: A survey. *ArXiv*, abs/2212.10403.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Dániel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. Opt-impl: Scaling language model instruction meta learning through the lens of generalization. *arXiv preprint arXiv:2212.12017*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *arXiv preprint arXiv:2205.11916*.
- Wang Ling, Dani Yogatama, Chris Dyer, and Phil Blunsom. 2017. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 158–167, Vancouver, Canada. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55:1 – 35.
- Nazneen Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain yourself! leveraging language models for commonsense reasoning. In *ACL*.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Oyvind Tafjord, Bhavana Dalvi, and Peter Clark. 2021. ProofWriter: Generating implications, proofs, and abductive statements over natural language. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3621–3634, Online. Association for Computational Linguistics.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971.
- 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 Lai, Ishan Purohit, Ishani Mondal, Jacob Anderson, Kirby Kuznia, Krma 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*, pages 5085–5109, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021a. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021b. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed Huaihsin Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. Emergent abilities of large language models. *ArXiv*, abs/2206.07682.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022b. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena D. Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models.
- Ping Yu, Tianlu Wang, O. Yu. Golovneva, Badr AlKhamissi, Siddharth Verma, Zhijing Jin, Gargi Ghosh, Mona Diab, and Asli Celikyilmaz. 2022. Alert: Adapting language models to reasoning tasks. *ArXiv*, abs/2212.08286.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Dataset	Task Definition	Options
AQuA	You are given an algebraic word question. Questions in this task often requires executing a series of arithmetic operations to obtain a final answer. You are also given 5 answer options (associated with 'A', 'B', 'C', 'D', 'E'). Do not generate anything else apart from one of the following characters: "A", "B", "C", "D", "E" and the corresponding explanation.	-A -B -C -D -E
CoQA	You are given a passage that contains a conversation and a question. The task is to answer the question and provide an explanation that highlights the corresponding evidence in the passage.	Free-form text
CoS-E	You are given a passage that contains a sentence and a question. The task is to answer the question by selecting one of the provided choices.	Select one of the provided choices
ECQA	You are given a question that requires commonsense reasoning. The task is to answer the question by selecting one of the provided choices.	Select one of the provided choices
ESNLI	You will be presented with a premise and a hypothesis sentence. The task is to determine whether the hypothesis sentence entails (implies), contradicts (opposes), or is neutral with respect to the given premise sentence. Please answer with "Contradiction", "Neutral", or "Entailment".	-Contradiction -Neutral -Entailment
GSM8K	You will be presented with a passage that contains a grade school math word problem. The task is to answer the question by performing a series of arithmetic operations to obtain a final answer.	Number
ProofWriter	You are given a sequence of facts and rules followed by a question. The task is to answer the question using "Yes", "No" or "Unknown".	-Yes -No -Unknown
StrategyQA	You are given a sentence and a question. The required reasoning steps are implicit in the question. The task is to answer the question using "Yes" or "No" then provide a strategy that explains the answer by decomposing it into a number of steps.	-Yes -No

Table 6: Task definition and options used for each of the finetuning reasoning datasets.

A Finetuning Task Definition and Options

Table 6 shows the task definition and options provided as input to the template shown in Figure 2 during finetuning the OPT models on the reasoning datasets.