

Unveiling the Generalization Power of Fine-Tuned Large Language Models*

Haoran Yang[♣], Yumeng Zhang[♡], Jiaqi Xu[♣], Hongyuan Lu[♣],
Pheng Ann Heng[♣], Wai Lam[♣]

[♣]The Chinese University of Hong Kong [♡]Tsinghua University
{hryang, hylu, wlam}@se.cuhk.edu.hk {jqxu, pheng}@cse.cuhk.edu.hk
zhang-ym23@mails.tsinghua.edu.cn

Abstract

While Large Language Models (LLMs) have demonstrated exceptional multitasking abilities, fine-tuning these models on downstream, domain-specific datasets is often necessary to yield superior performance on test sets compared to their counterparts without fine-tuning. However, the comprehensive effects of fine-tuning on the LLMs' generalization ability are not fully understood. This paper delves into the differences between original, unmodified LLMs and their fine-tuned variants. Our primary investigation centers on whether fine-tuning affects the generalization ability intrinsic to LLMs. To elaborate on this, we conduct extensive experiments across five distinct language tasks on various datasets. Our main findings reveal that models fine-tuned on generation and classification tasks exhibit dissimilar behaviors in generalizing to different domains and tasks. Intriguingly, we observe that integrating the in-context learning strategy during fine-tuning on generation tasks can enhance the model's generalization ability. Through this systematic investigation, we aim to contribute valuable insights into the evolving landscape of fine-tuning practices for LLMs. The code and data are available at https://github.com/LHRYANG/Generalization_of_FT-LLM.

1 Introduction

The transformative impact of in-context learning (ICL) (Brown et al., 2020; Wei et al., 2022b; Rubin et al., 2022; Liu et al., 2023; Chowdhery et al., 2022; Wang et al., 2023b) in LLMs, as demonstrated by models like Llama-2 (tou, 2023) and GPT-3 (Brown et al., 2020), marks a significant advancement in the field of artificial intelligence. This learning paradigm allows LLMs to adapt to

various tasks by leveraging multiple demonstration examples presented within a prompt, without training the LLMs. However, when it comes to a specific task, fine-tuning often achieves better performance than ICL, which has been substantiated by recent studies (Shi et al., 2023; Jiao et al., 2023; Wang et al., 2023a; Zhang et al., 2023).

There are some works studying the properties of fine-tuning and ICL for LLMs. For example, Wei et al. (2022a) reveal that *multi-task* fine-tuning can enhance an LLM's zero-shot and ICL capabilities. It indicates that fine-tuning, when applied across multiple tasks, does not merely improve performance on those seen tasks but also augments the model's inherent learning abilities. The work of Mosbach et al. (2023) highlights that in terms of the out-of-domain generalization in classification tasks, *few-shot* fine-tuning and ICL exhibit similar levels of generalization. Wang et al. (2023c) find that fine-tuning may overly tailor the model to task-specific formats, potentially compromising its adaptability to other new tasks.

In this paper, we conduct a comprehensive study on how *task-specific* (not multi-task or few-shot) fine-tuning affects the generalization ability of LLMs. To provide a thorough analysis, we design a series of experiments encompassing a diverse range of datasets and tasks, covering both classification and generation tasks. For each task, we designate a specific dataset as the training set. The remaining selected datasets are subsequently divided into two groups: in-domain datasets, closely aligned with the training set in terms of content and structure, and out-of-domain datasets, which possess significant differences. With these datasets, our research investigates two critical questions: i) the ability of fine-tuned LLMs to adapt to both in-domain and out-of-domain test sets, and ii) the impact of fine-tuning on the ICL ability of LLMs across different types of tasks.

We find that models fine-tuned on text genera-

*This research/paper was partially supported by the Center for Perceptual and Interactive Intelligence (CPII) Ltd. under the Innovation and Technology Commission's InnoHK scheme.

tion and classification tasks exhibit different behaviors when evaluated on test sets. Specifically, we observe that models fine-tuned for classification tasks tend to exhibit positive transfer when applied to out-of-domain datasets of the same fine-tuning/test task type. In contrast, models fine-tuned on generation tasks frequently experience negative transfer under similar conditions. Interestingly, while fine-tuning the LLMs on generation tasks generally does not detrimentally affect their performance on classification tasks, the reverse is not true; models fine-tuned on classification tasks typically fail to work on generation tasks. Moreover, we experimentally observe that integrating the ICL strategy during fine-tuning on generation tasks can enhance an LLM’s generalization ability. We also investigate other factors, such as training data size and the number of in-context examples. We hope this study offers comprehensive insights into fine-tuning strategies for LLMs, not only in enhancing task-specific performance but also in fostering broader generalization abilities.

2 Related Work

2.1 Large Language Models

The emergence and evolution of Large Language Models (LLMs) have significantly impacted the field of natural language processing and beyond. Seminal works like BERT (Devlin et al., 2019) and GPT-2 (Radford et al., 2019) laid the foundation for understanding context and semantics in textual language. The advent of GPT-3 (Brown et al., 2020) demonstrated remarkable abilities in generating human-like text, paving the way for more advanced models like GPT-4 (OpenAI, 2023) and open-sourced Llama-2 (tou, 2023). These models, with vast number of parameters and pre-trained on gaint corpus, have shown exceptional skills in understanding and solving various tasks in a zero- or few-shot manner (Sun et al., 2023).

2.2 Fine-tuning vs. In-Context Learning

Fine-tuning (FT) has been a predominant approach for adapting Pre-trained Language Models (PLMs) to specific tasks, e.g., dialogue system (Xu et al., 2022). This process involves additional training of a PLM on a smaller, task-specific dataset. This technique has been proven effective in tailoring models like BERT and GPT-2 for specialized applications, from classification tasks (Devlin et al., 2019; Liu et al., 2019) to generation tasks (Yang et al., 2021;

Gehrmann et al., 2019). In the era of LLMs, in-context learning (ICL) (Kossen et al., 2023; Han et al., 2023) has emerged as a novel paradigm for distilling knowledge from powerful models, particularly highlighted in LLMs like GPT-3 (Brown et al., 2020) and Llama-2 (tou, 2023). ICL leverages the models’ intrinsic capabilities to understand and generate responses, most of which are learned during unsupervised pre-training (Zhou et al., 2023; Gudibande et al., 2023), based on given contexts enclosed in the prompt, without the need for explicit task-specific training. For instance, Brown et al. (2020) demonstrate the effectiveness of ICL in a wide range of tasks, by merely providing a few examples within the prompt.

The debate between FT and ICL hinges on the trade-off between specialization and generalization. While FT offers more tailored and often higher-performing models for specific tasks, it can lead to a loss of the model’s generalization abilities, as discussed by Chen et al. (2020). ICL, on the other hand, maintains the model’s broad applicability but may exhibit suboptimal performance in specific tasks, as observed by Shi et al. (2023). Recently, Mosbach et al. (2023) discovered that both few-shot FT and ICL can achieve a similar generalization on out-of-domain test sets. In contrast, our work focuses on larger-set FT instead of few-shot FT. Wang et al. (2023c) identified that format specialization is a critical factor contributing to the diminished ICL abilities in fine-tuned LLMs. In our experiment, we observed similar phenomena, particularly when models fine-tuned on classification tasks were evaluated on generation tasks. However, in other scenarios, the impact of format specialization appeared to be less pronounced. Anil et al. (2022) found that incorporating several in-context examples during FT is helpful for length generalization for text. We further expand this idea and find that this approach is also indeed valuable in preserving or even enhancing the fine-tuned models’ generalization ability.

3 Evaluation Design

This study delves into the effects of task-specific fine-tuning on the generalization ability of LLMs. We aim to uncover whether LLMs, once fine-tuned on a dataset of a particular language task, can still perform well on i) (data-level) in-domain and out-of-domain test sets of the same task type and ii) (task-level) different tasks.

Task	Train	In-domain Test	Out-of-domain Test
Summary Generation	XSum	XSum XLSum	PeerRead CNN/DailyMail
Question Generation	Socialqa	Socialqa	Tweetqa Sciqa
Sentiment Classification	Amazon	Amazon AmazonFood	SST2 Yelp
Paraphrase Detection	Paws	Paws	QQP STS-B
Natural Language Inference	MNLI	MNLI-1 MNLI-2	RTE GPTNLI

Table 1: Summary of tasks & datasets

3.1 Evaluation Taxonomy

To comprehensively assess the performance of fine-tuned LLMs across various tasks and datasets (Table 1), our study encloses three distinct settings, characterized by increasing levels of generality:

1. **Same Task, In-domain Datasets:** Given the same fine-tuning/test task, such as summary generation, we assess the fine-tuned LLMs using datasets that are aligned with their training data (in-domain), e.g., models fine-tuned on XSum and tested on XLSum (denoted as XSum \rightarrow XLSum).
2. **Same Task, Out-of-domain Datasets:** Despite being evaluated on the same task, this setting focuses on the out-of-domain generalization by testing the fine-tuned LLMs on datasets with distinct features compared to the training set, e.g., XSum \rightarrow PeerRead.
3. **Different Tasks:** In this setting, we examine the capability of LLMs, fine-tuned on one task type, to adapt across different task types, thereby evaluating the LLMs’ cross-task generalization, e.g., XSum \rightarrow Socialqa (summary to question generation) and XSum \rightarrow Amazon (generation to classification).

Through such varied settings (in-domain vs. out-of-domain data and same vs. different tasks), we can clearly understand the generalization ability of the fine-tuned LLMs and explore the boundaries of their applicability across different data distributions and task types.

3.2 Evaluation Benchmarks

To achieve a comprehensive and reliable evaluation, we carefully select the evaluation benchmarks, ensuring the validity and generality of the experimental findings. In detail, our study encompasses five widely used language tasks: summarization and

question generation, sentiment classification, natural language inference, and paraphrase detection. These tasks can be broadly categorized into *generation* and *classification*, i.e., the first two focus on text generation, and the last three on classification.

As shown in Table 1, for each task, we select three or four datasets, where one is used for fine-tuning the LLMs, and the others serve as the test sets. We endeavor to ensure that the test datasets are within the same domain or span different domains as the training set. This approach can evaluate the LLMs’ generalization in familiar contexts and to new domains (in- and out-of-domain). We briefly introduce the tasks and datasets below. The full descriptions can be found in Appx. A.

Summary Generation This task aims to generate a summary based on the given article. We select XSum (Narayan et al., 2018), XLSum (Hasan et al., 2021), PeerRead (Kang et al., 2018), and CNN/DailyMail (Hermann et al., 2015). XSum is used as the training set. The test set of XSum itself and XLSum serve as the in-domain test sets. The others are regarded as out-of-domain datasets.

Question Generation Given a paragraph and an answer, question generation infers the corresponding question. We select Socialqa (Sap et al., 2019), Tweetqa (Xiong et al., 2019), and Sciqa (Welbl et al., 2017). We choose Socialqa as the training set and its test set for in-domain testing. Tweetqa and Sciqa are used as the out-of-domain test sets.

Sentiment Classification Sentiment classification identifies the positive/negative emotions expressed in the text. This evaluation involves Amazon review (Keung et al., 2020), AmazonFood review (McAuley and Leskovec, 2013), SST2 (Wang et al., 2019), and Yelp¹. Note that the target for Yelp is a rating score ranging from 1 to 5, while the label for other datasets is positive or negative. We convert the ratings in Yelp to binary labels with negative below 3.5 and positive above 3.5. Amazon is used to fine-tune LLMs; Amazon and Amazon-food act as the in-domain test dataset, while SST2 and Yelp are the out-of-domain sets.

Paraphrase Detection This task involves classifying if two given text segments using different wordings express the same meaning. We select Paws (Zhang et al., 2019), QQP, and STS-B (Wang

¹https://huggingface.co/datasets/yelp_review_full

et al., 2018). Since STS-B is labeled using 1 to 5 similarity scores, we perform a similar processing step as YELP, i.e., if the rating is above 3.5, the two texts are paraphrased. The Paws and itself are used as the training and in-domain test set. The other datasets are the out-of-domain datasets.

Natural Language Inference Given a pair of text segments, typically referred to as premise and hypothesis, this task determines the relationship between them, i.e., if the hypothesis is entailment, contradiction, or neutral based on the premise information. We use MNLI (Williams et al., 2018) as the training set, MNLI matched (MNLI-1) and MNLI mismatched (MNLI-2) as the in-domain test sets, RTE (Wang et al., 2018) and GPTNLI² as the out-of-domain test sets.

3.3 Experimental Setup

Models & Metrics We conduct all experiments using the open-sourced Llama-2-7b (tou, 2023) due to its popularity in the NLP community. In the evaluation, we use the Rouge-L³ metric for generation tasks and the accuracy for classification tasks.

Training Details For each task-specific training set, we fine-tune the Llama-2 models using subsets of varying sizes: 2,000, 4,000, and 6,000 samples, which enables us to analyze how different training sizes impact the model’s performance.

To standardize our training process, we treat classification tasks as text generation. Specifically, we use the language modeling head to predict labels in the text form during training, as suggested in previous works (Schick and Schütze, 2021; Liu et al., 2022). During the evaluation phase, we only choose the probabilities associated with these pre-defined labels as the models’ predictions and then select the output with the highest probability as the predicted label. Details about the labels used for classification tasks and prompt formats can be found in Appx. B and Appx. C, respectively.

The models are fine-tuned with 2 epochs. We employ the AdamW (Loshchilov and Hutter, 2019) optimizer with a learning rate of 0.002. The generation length is set to 60 for generation tasks and 5 for classification tasks.

Testing Details Our primary aim is to evaluate the generalization ability of the fine-tuned LLMs.

²https://huggingface.co/datasets/pietrolesci/gpt3_nli

³<https://huggingface.co/spaces/evaluate-metric/rouge>

For this, we employ distinct approaches based on the testing datasets and task types, as in Sec. 3.1. When evaluating the models on datasets that match the task type of the fine-tuning set (Setting 1 and 2), we adopt two strategies: a 0-shot prompting approach with no in-context examples (normal testing after training), and in-context learning (ICL), where a set of in-context examples are provided. On the other hand, in scenarios where the testing task type of the test sets differs from that of the training set (Setting 3), the usage of ICL becomes necessary. This technique is essential to inform the model about the nature of the task to be performed. Notably, for classification tasks, we still show 0-shot inference performance, since only the label probabilities are considered.

Specifically, for generation tasks, the models are evaluated with 1, 2, or 4 in-context examples. For binary classification tasks, including sentiment analysis and paraphrase detection, we report results using 2, 4, and 6 in-context examples. Regarding Natural Language Inference (NLI), which contains three label categories, we showcase results with 3, 5, and 7 in-context examples. Among all these evaluation variations, we ensure that every category is presented at least once in the in-context examples.

4 Results and Findings

4.1 Same Task, In-domain Datasets

In Figure 1, we first present the results of fine-tuned Llama-2 models on in-domain testing datasets of the same fine-tune/test task type (Setting 1). The following key findings can be drawn:

Fine-tuned models without in-context learning (ICL) can generally perform better than baseline Llama-2 using ICL. The fine-tuned models, trained with varying sample sizes (2K, 4K, 6K), exhibit superior 0-shot (w/o ICL) performance compared to the original baseline Llama2 using ICL across most datasets, notably XLSum, Socialqa, MNLI-1, MNLI-2, and Paws; see Figure 1 (b) (c) (d) (g) and (h). The only exception is observed in the sentiment classification task (Amazon and AmazonFood), where fine-tuned models slightly underperform compared to baseline Llama-2 using ICL. This could be attributed to Llama-2’s inherent expertise in sentiment analysis, as indicated by its high 0-shot performance. In such cases, additional fine-tuning might not significantly enhance or could potentially impair performance, possibly due to overfitting or conflicting training data.

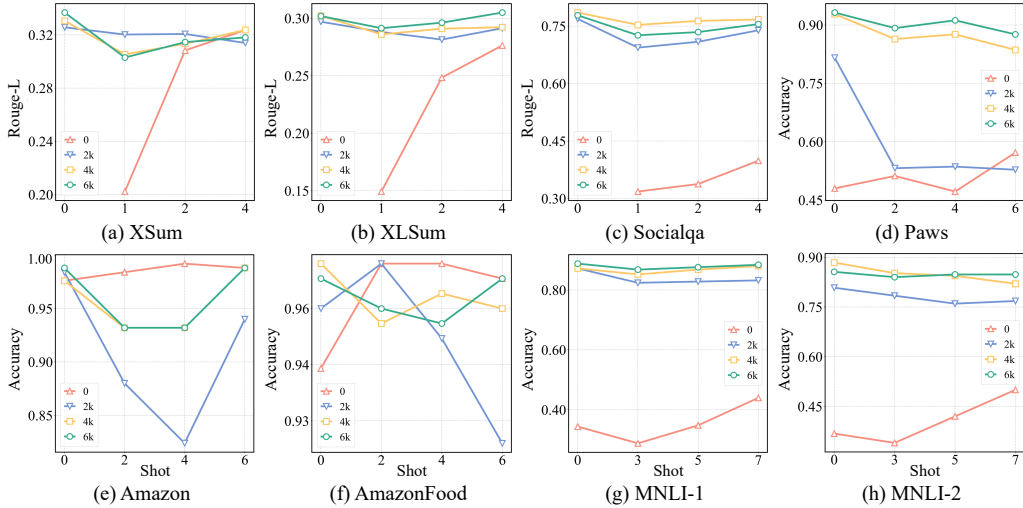


Figure 1: **In-domain** dataset testing performance comparisons of baseline Llama-2 (0 training samples, orange line) and its fine-tuned variants (2K, 4K, 6K training samples). Shot denotes the number of in-context examples. The caption for each subfigure refers to the test set. The corresponding training set can be found in Table 1. The 0-shot results of the baseline Llama model in (a), (b) and (c) are not presented since in scenarios where in-context examples are absent (0-shot), baseline models generally struggle to execute the tasks effectively, even when the prompt explicitly outlines the task requirements.

Fine-tuned LLMs often perform worse using in-context learning than the zero-shot setting.

From Figure 1, compared to baseline Llama-2, the fine-tuned models benefit little from the ICL. This trend suggests that while fine-tuning enhances a model’s 0-shot in-domain generalization, the additional in-context examples during inference are not always necessary and helpful. The drop in performance might be due to the model becoming more specialized after fine-tuning, reducing its adaptability to new contexts. The performance correlation with the number of in-context shots remains unclear for the fine-tuned models. In comparison, baseline Llama-2 shows an overall positive tendency in performance with increased in-context shots. This improvement trend suggests that LLMs without task-specific fine-tuning are more effective in leveraging the in-context examples to understand the specialized task. These observations can indicate that once LLMs have been fine-tuned on sufficient data, their ability to benefit from ICL diminishes.

Fine-tuning with more samples may not consistently improve performance on the test set.

Generally, the models fine-tuned with 4K or 6K samples outperform the 2K models on most test sets. However, the degree of improvement is not consistent. On XSum (Figure 1 (a)), for instance, fine-tuned models demonstrate only a slight performance increase with larger training sets, changing from 2K to 6K training examples. In contrast, on Paws (Figure 1 (d)), models show marked perfor-

mance gains as the number of training samples increased from 2K to 4K, with accuracy jumping from 81.6% to 93.2%. We also find that increasing the training size from 4K to 6K brings subtle or even negative impacts. These findings suggest that the relationship between the volume of fine-tuning data and in-domain test performance is task-dependent and not straightforward.

4.2 Same Task, Out-of-domain Datasets

We then show the results of the fine-tuned Llama-2 on out-of-domain testing datasets of the same task type (Setting 2). We derive the following observations based on Figure 2:

Fine-tuned models underperform compared to the baseline model on generation tasks, yet outperform on classification tasks.

For the out-of-domain testing results, a clear distinction emerges based on the task categories, i.e., between generation and classification tasks. In generation testing datasets ((a) PeerRead, (b) CNN/DailyMail, (c) Tweetqa, and (d) Sciqa in Figure 2), fine-tuned models perform worse than baseline models, and this gap persists regardless of the number of in-context examples provided. Notably, there lacks a performance growth with more training examples in datasets like CNN/DailyMail, Tweetqa and Sciqa, suggesting that the generalization ability of fine-tuned LLMs is impaired.

On the other hand, in the context of sentiment classification tasks, the best performance of the

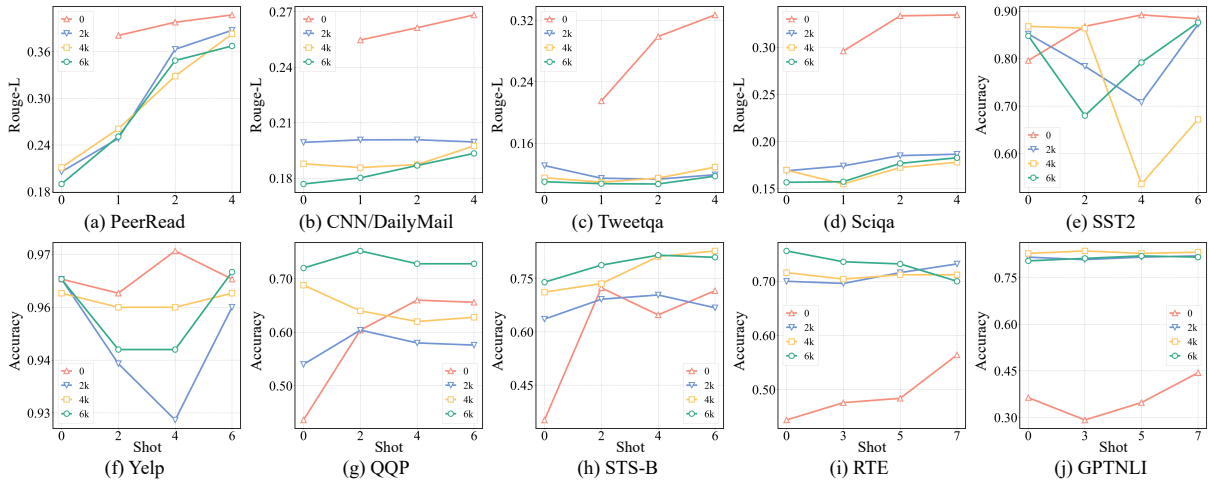


Figure 2: **Out-of-domain** dataset testing performance comparisons of baseline Llama-2 (0 training samples, orange line) and its fine-tuned variants (2K, 4K, 6K training samples).

fine-tuned models is on par with the baseline Llama-2, which means the fine-tuned LLMs’ sentiment knowledge is unlikely to be lost; see Figure 2 (e) and (f). Another notable observation is that the fine-tuned models exhibit larger variability using different ICL shots. This indicates that the performance is more sensitive to the in-context examples for the sentiment classification task. For other classification tasks (i.e., paraphrase detection and natural language inference), fine-tuned models consistently outperform the baseline Llama-2, as shown in Figure 2 (g)-(j).

The diverging effects of fine-tuning on generation and classification tasks for out-of-domain testing may originate from the difference in task output space constraints. The output space of classification tasks is inherently predefined and limited, enabling fine-tuned LLMs to apply their inherited and adapted knowledge relatively easily to new domains. In contrast, the output space of out-of-domain generation datasets largely deviates from that of the training set. Despite being given a few in-context examples, fine-tuned models may still find it challenging to reason about the expansive range of possible outputs in new domains.

4.3 Different Tasks

Last, we test whether the generalization ability of fine-tuned models is preserved on cross-task testing datasets (Setting 3). The evaluation for each testing set is performed using models fine-tuned on 2K training samples from other tasks. Due to space limitations, our analysis is confined to five test sets: XSum, Socialqa, Amazon, Paws, and MNLI-1. The following findings are mainly based on the first row of Figure 3:

The generalization through fine-tuning exhibits significant variability and highly depends on the training data.

When assessing performance on classification tasks, i.e., Amazon, MNLI, and Paws, it becomes evident that the model’s fine-tuning source greatly affects its efficacy. Fine-tuning on a dataset like Amazon negatively impacts performance on the MNLI-1 test set, whereas fine-tuning on XSum significantly boosts it; see Figure 3 (b). In parallel, for generation tasks, training on Socialqa hurt the performance on XSum while training on XSum has little impact on Socialqa; see Figure 3 (d) and (e). These intricate patterns suggest that the effectiveness of fine-tuning is not easily predictable and likely intertwines with dataset characteristics, fine-tuning procedures, etc.

Models fine-tuned on classification tasks fail to generalize to generation tasks.

From Figure 3 (d) and (e), fine-tuning the models on classification data leads to almost zero Rouge-L scores for generation tasks. Upon examining the outputs, it becomes evident that these models predominantly generate classification labels rather than coherent text, a manifestation of output space specialization, which aligns with the findings of Wang et al. (2023c). Two potential reasons exist. Firstly, the model’s output space may be constrained to the category labels seen during fine-tuning, inhibiting its ability to generate other tokens. The second may be induced by the prompt format, as listed as Prompt-1 in Table 5. The prompts for different tasks have the same start "####", which could confuse the model and cause it to misinterpret inputs from other tasks as if they belong to its training task. To further investigate the influence of prompt

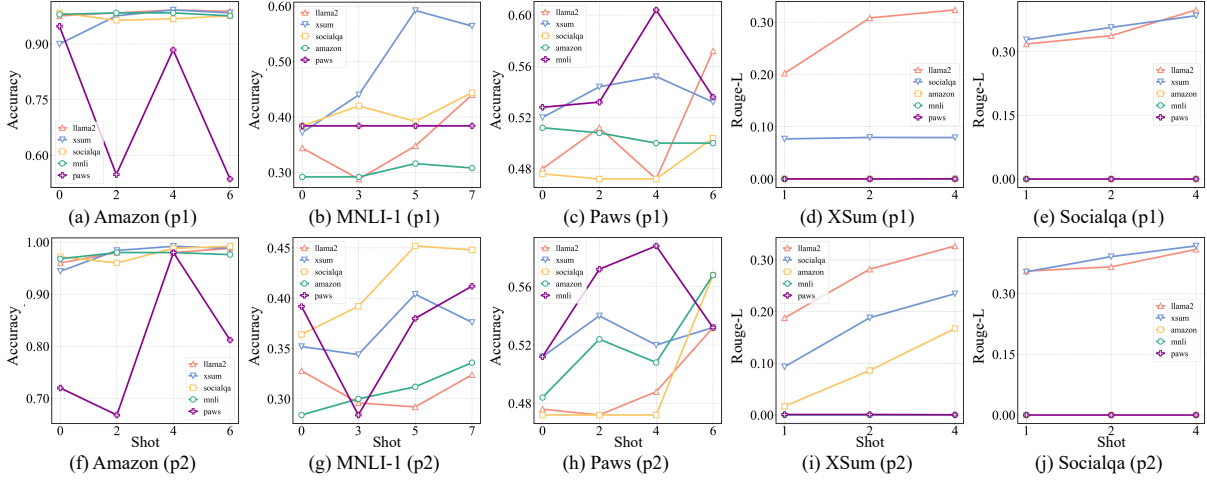


Figure 3: **Cross-task** performance comparisons of baseline Llama-2 (0 training samples, orange line) and models fine-tuned on other tasks. The caption for each subfigure refers to the test set. The legends denote the training data. The first row is the results using the Prompt-1 (p1) and the second row is the results using the Prompt-2 (p2) format. The detailed prompt formats can be found in Appx. C.

format, we introduced a distinct set of prompts that avoid uniform starting sequences (Prompt-2 in Table 5). The corresponding results are shown in the second row of Figure 3. Comparing the first and second rows, we find that the cross-task evaluation on classification tasks is more sensitive to the prompt format than the generation task evaluation. Moreover, from Figure 3 (i), the model fine-tuned on Amazon starts to work on XSum with such new prompts. Yet, for Socialqa (subfigure (j)), the fine-tuning on Amazon still fails to succeed. Hence, the prompt format is crucial for cross-task generalization, but identifying a universally effective format remains to be explored.

5 Fine-tuning with In-context Learning on Generation Tasks Helps Improve the Generalization Ability of LLMs

In Sec. 4.2, our analysis reveals that LLMs fine-tuned on classification tasks exhibit robust generalization capabilities under out-of-domain scenarios. However, for generation tasks, the fine-tuned models consistently fall short compared to baseline.

In this section, we aim to show that Fine-Tuning with In-Context Learning (FTICL) can help improve LLMs’ out-of-domain generalization for generation tasks. FTICL leverages the strengths of both fine-tuning and ICL. Specifically, instead of directly forwarding the original input into LLMs during fine-tuning, FTICL prepends in-context examples to the input and forwards the LLMs with this newly constructed input. Note that this strategy has also been adopted to mitigate other fine-tuning

issues. For example, Anil et al. (2022) find that FTICL leads to a substantial improvement in terms of input length generalization.

In our experiment, we fine-tune the FTICL models using 2,000 samples. For each generation task, we train two models with 1 or 2 in-context examples prepended in input, respectively. The results are presented in Figure 4 and Figure 5. The following findings regarding generalization are organized in terms of the same or different fine-tuning/test tasks.

5.1 Same Task

We investigate the effects of FTICL across datasets sharing the same fine-tuning/test task type. We fine-tune models on the training sets of the XSum and Socialqa datasets, respectively. These models are then evaluated both on their respective test sets and on distinct out-of-domain test sets. The main results of these evaluations are presented in Figure 4. Due to space limitations, we only report the typical results on XSum, PeerRead, Socialqa, and Sciqa. From subfigures (a) and (c), we can see that FTICL models preserve or even slightly enhance the performance on the corresponding test sets of the training set compared with vanilla fine-tuning.

When we extend our analysis to out-of-domain test sets of the same fine-tuning/test task type, we find **models fine-tuned using FTICL achieve a better out-of-domain generalization performance than the vanilla fine-tuned models**. The FTICL models can sometimes even surpass the baseline model without fine-tuning. For instance, as shown in Figure 4 (b), the FTICL model with

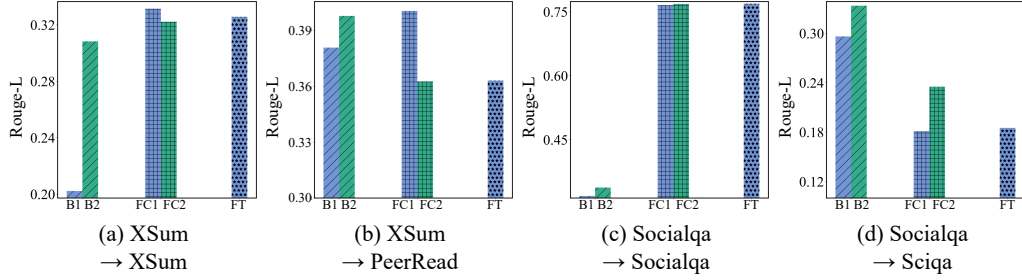


Figure 4: **Same fine-tuning/test task type evaluation of FTICL with generation tasks.** B_n represents the baseline Llama-2 model with n in-context examples during inference. FC_n denotes the FTICL models fine-tuned with n in-context examples. FT is the vanilla fine-tuned model without in-context learning. For FC_n and FT, we fine-tune with 2,000 samples, perform both 0-shot and few-shot evaluations, and report the results with the best performance.




Figure 5: **Cross-task performance of FTICL with generation tasks.** For the classification task evaluation, we also report the 0-shot performance (B0) for the baseline Llama-2.

one in-context example during fine-tuning (FC1) showcases a remarkable gain over the vanilla fine-tuned model (FT) on PeerRead, surpassing even the baseline Llama-2 model using in-context learning (B1 and B2). For Sciqa, although FC2 lags behind B1 and B2, it still performs better than FT; see Figure 4 (d). These observations reveal that FTICL may mitigate catastrophic forgetting for generation tasks by allowing the model to retain its learned capabilities more effectively than the vanilla fine-tuning method.

5.2 Different Tasks

Besides enhanced out-of-domain generalization, we also demonstrate that models utilizing FTICL exhibit superior cross-task generalization capabilities. The results are shown in Figure 5.

FTICL trained on generation tasks achieves a better cross-task generalization than vanilla fine-tuning. We can observe FTICL models fine-tuned on Socialqa achieve at least comparable performance on Amazon, shown in Figure 5 (a), and exhibit superior results over both the baseline Llama-2 and the vanilla fine-tuned models when evaluated on MNL1 and Paws; see Figure 5 (b) and (c). We can also see FC1 in subfigure (d) and FC1, FC2 in subfigure (e) outperform the vanilla fine-tuned (FT) models and can be even on par with

or surpass the baseline Llama-2 for the cross summary and question generation task evaluations. The results show that fine-tuning with ICL is better than directly fine-tuning on generation tasks.

5.3 Potential Reason

We provide one hypothesis that could drive the success of FTICL in enhancing LLMs’ generalization: **FTICL tends to deviate less from the original LLM than vanilla fine-tuning.** In other words, FTICL models preserve more general knowledge inherent in LLMs. To support this, we calculate the average parameter weight difference between the fine-tuned models (FTICL and FT) and the original Llama-2. Experimental results show consistency with our hypothesis: on Socialqa, FTICL ($7.95e-05$) vs. FT ($8.54e-05$); on XSum, FTICL ($8.03e-05$) vs. FT ($1.0e-04$). The potential reason is that the provided in-context examples encourage the LLM to leverage its existing knowledge to solve new tasks.

5.4 FTICL on Classification Tasks

We also analyse the performance of FTICL on classification tasks and find that this strategy **does not** help improve the generalization ability of LLMs as shown in Figure 6 and Figure 7.

For the classification tasks, as shown in Figure 6 (a) (c), and (e), when evaluated on the in-




Figure 6: Same fine-tuning/test task type evaluation of FTICL models fine-tuned on classification tasks.




Figure 7: Cross-task performance for FTICL models fine-tuned on classification tasks.

domain test set of the corresponding training set, fine-tuning with in-context learning (FTICL) models generally perform worse than vanilla fine-tuned models. For the out-of-domain test sets, we find FTICL models can generally outperform the baseline Llama but lag behind vanilla FT models; see Figure 6 (b) (d) (f). The reason for this phenomenon may be that classification tasks are more sensitive to the in-context examples. We also hypothesize that the optimization process plays a pivotal role. Specifically, for classification tasks, we observed that the final loss of models fine-tuned with in-context learning (FTICL) tends to be higher compared to vanilla fine-tuning (FT), with challenges in further reducing the loss. This phenomenon might stem from the model’s tendency to be lazy. Specifically, for generation tasks (e.g., summary), the model must learn to leverage contextually relevant information (e.g., the article to be summarized) to generate appropriate outputs (e.g., target summary). However, for classification tasks, in-context examples could inadvertently act as distractors since all the labels are provided in the in-context examples. It may lead the model to copy labels directly from in-context examples rather than leveraging the corresponding relevant information. This could also explain the difficulty in loss reduction. A better optimizer could possibly solve this problem. The above observations suggest that for classification tasks, it is better to adopt the vanilla fine-tuning approach instead of FTICL.

Lastly, Figure 7 shows the cross-task performance of FTICL, where the models are fine-tuned with classification tasks. In Figure 7 (a) (c), and (e), for cross-task classification tasks, we find the cross-

task transfer effects are not clear. For example, training on Amazon has a positive effect on Paws, while training on MNLI has a negative effect on Amazon. Meanwhile, as shown in Figure 7 (b) (d), and (f), we can see that FTICL models for generation task evaluation outperform vanilla fine-tuned models, which means FTICL can help alleviate the output space specialization issue mentioned in Sec. 4.3.

6 Conclusion

This study comprehensively investigates the effects of fine-tuning on the LLMs’ generalization ability. We conduct systematic experiments by evaluating the fine-tuned LLMs across various training data and language tasks. Experimental results indicate that dissimilar generalization ability after fine-tuning arises from the nature of generation or classification tasks. Further, we show that fine-tuning with in-context learning can enhance the generalization capability for generation tasks. We hope our findings can inspire future advancements in understanding and effectively utilizing LLMs to solve new tasks.

Limitations

In this work, we study the LLMs’ generalization ability between fine-tuned variants and their original counterparts and explore the effective fine-tuning strategies. This work has two limitations: i) The underlying reasons for the difference between models fine-tuned on classification tasks and generation tasks remain under-examined. ii) The intricate operational mechanisms behind the fine-tuning with in-context learning method have yet to

be exhaustively understood. iii) We do not study the recently more advanced tuning strategies such as RLHF (Stiennon et al., 2020; Kirk et al., 2024). We leave these more in-depth analyses as our future work.

References

2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Cem Anil, Yuhuai Wu, Anders Johan Andreassen, Aitor Lewkowycz, Vedant Misra, Vinay Venkatesh Ramasesh, Ambrose Slone, Guy Gur-Ari, Ethan Dyer, and Behnam Neyshabur. 2022. [Exploring length generalization in large language models](#). In *Advances in Neural Information Processing Systems*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. 2020. [Recall and learn: Fine-tuning deep pretrained language models with less forgetting](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7870–7881, Online. Association for Computational Linguistics.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sebastian Gehrmann, Zachary Ziegler, and Alexander Rush. 2019. [Generating abstractive summaries with finetuned language models](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 516–522, Tokyo, Japan. Association for Computational Linguistics.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [The false promise of imitating proprietary llms](#).
- Xiaochuang Han, Daniel Simig, Todor Mihaylov, Yulia Tsvetkov, Asli Celikyilmaz, and Tianlu Wang. 2023. [Understanding in-context learning via supportive pretraining data](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12660–12673, Toronto, Canada. Association for Computational Linguistics.
- Tahmid Hasan, Abhik Bhattacharjee, Md. Saiful Islam, Kazi Mubasshir, Yuan-Fang Li, Yong-Bin Kang, M. Sohel Rahman, and Rifat Shahriyar. 2021. [XLsum: Large-scale multilingual abstractive summarization for 44 languages](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4693–4703, Online. Association for Computational Linguistics.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 1693–1701.
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine](#).
- Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. [A dataset of peer reviews \(peerread\): Collection, insights and nlp applications](#). In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, New Orleans, USA.
- Phillip Keung, Yichao Lu, György Szarvas, and Noah A. Smith. 2020. [The multilingual Amazon reviews corpus](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*

- (EMNLP), pages 4563–4568, Online. Association for Computational Linguistics.
- Robert Kirk, Ishita Mediratta, Christoforos Nalmpantis, Jelena Luketina, Eric Hambro, Edward Grefenstette, and Roberta Raileanu. 2024. [Understanding the effects of RLHF on LLM generalisation and diversity](#). In *The Twelfth International Conference on Learning Representations*.
- Jannik Kossen, Yarin Gal, and Tom Rainforth. 2023. [In-context learning learns label relationships but is not conventional learning](#).
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#). *ACM Comput. Surv.*, 55(9).
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. [P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 61–68, Dublin, Ireland. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Julian John McAuley and Jure Leskovec. 2013. [From amateurs to connoisseurs: Modeling the evolution of user expertise through online reviews](#). In *Proceedings of the 22nd International Conference on World Wide Web, WWW '13*, page 897–908, New York, NY, USA. Association for Computing Machinery.
- Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow, and Yanai Elazar. 2023. [Few-shot fine-tuning vs. in-context learning: A fair comparison and evaluation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 12284–12314, Toronto, Canada. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. Don't give me the details, just the summary! Topic-aware convolutional neural networks for extreme summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, Seattle, United States. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan LeBras, and Yejin Choi. 2019. SocialQA: Commonsense reasoning about social interactions. In *EMNLP*.
- Timo Schick and Hinrich Schütze. 2021. [It's not just size that matters: Small language models are also few-shot learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Chufan Shi, Yixuan Su, Cheng Yang, Yujiu Yang, and Deng Cai. 2023. [Specialist or generalist? instruction tuning for specific nlp tasks](#).
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020. Learning to summarize from human feedback. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA. Curran Associates Inc.
- Jiankai Sun, Chuanyang Zheng, Enze Xie, Zhengying Liu, Ruihang Chu, Jianing Qiu, Jiaqi Xu, Mingyu Ding, Hongyang Li, Mengzhe Geng, et al. 2023. A survey of reasoning with foundation models. *arXiv preprint arXiv:2312.11562*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *International Conference on Learning Representations*.
- Xiao Wang, Weikang Zhou, Can Zu, Han Xia, Tianze Chen, Yuansen Zhang, Rui Zheng, Junjie Ye, Qi Zhang, Tao Gui, Jihua Kang, Jingsheng Yang, Siyuan Li, and Chunsai Du. 2023a. [Instructuie: Multi-task instruction tuning for unified information extraction](#).
- Yifan Wang, Qingyan Guo, Xinzhe Ni, Chufan Shi, Lema Liu, Haiyun Jiang, and Yujiu Yang. 2023b. Hint-enhanced in-context learning wakes large language models up for knowledge-intensive tasks. *arXiv preprint arXiv:2311.01949*.

- Yihan Wang, Si Si, Daliang Li, Michal Lukasik, Felix Yu, Cho-Jui Hsieh, Inderjit S Dhillon, and Sanjiv Kumar. 2023c. [Two-stage llm fine-tuning with less specialization and more generalization](#).
- Jason Wei, Maarten Bosma, Vincent 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 *International Conference on Learning Representations*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022b. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Johannes Welbl, Nelson F. Liu, and Matt Gardner. 2017. [Crowdsourcing multiple choice science questions](#). *ArXiv*, abs/1707.06209.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Wenhan Xiong, Jiawei Wu, Hong Wang, Vivek Kulka-rni, Mo Yu, Xiaoxiao Guo, Shiyu Chang, and William Yang Wang. 2019. [Tweetqa: A social media focused question answering dataset](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Chen Xu, Piji Li, Wei Wang, Haoran Yang, Siyun Wang, and Chuangbai Xiao. 2022. [Cosplay: Concept set guided personalized dialogue generation across both party personas](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 201–211.
- Haoran Yang, Wai Lam, and Piji Li. 2021. [Contrastive representation learning for exemplar-guided paraphrase generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4754–4761, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yue Zhang, Leyang Cui, Deng Cai, Xinting Huang, Tao Fang, and Wei Bi. 2023. [Multi-task instruction tuning of llama for specific scenarios: A preliminary study on writing assistance](#).
- Chunting Zhou, Pengfei Liu, Puxin Xu, Srinu Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, Susan Zhang, Gargi Ghosh, Mike Lewis, Luke Zettlemoyer, and Omer Levy. 2023. [Lima: Less is more for alignment](#).

A Dataset Information

We provide more dataset information (including the data sources and examples) in Table 2 and Table 3. It should be noted that for the summary generation task, although both XSum (Narayan et al., 2018) and CNN/DailyMail (Hermann et al., 2015) are from news articles, they are not in the same domain. There are two reasons. First, XSum is from BBC News, and CNN/DailyMail is from CNN and DailyMail news. There exist style differences for different news. Another more important reason is that XSum is a one-sentence summary dataset, while CNN/DailyMail is a multiple-sentence summary dataset; see Table 2.

B Classification Labels

In Table 4, we illustrate the tokens that need to be generated during training for the language classification tasks. We ensure the labels between different tasks have no intersection with each other. This can avoid the target specialization issue when evaluating a dataset in one task type using models trained on a dataset of a different task type.

C Prompt Format

In Table 5, we detail the prompt formats employed for each task in our experiments. Throughout the training phase, the model is trained using Prompt-1. Note that Prompt-1 is also the default choice during the testing phase. To ensure a comprehensive evaluation, particularly in cross-task scenarios where the influence of the prompt format is a crucial consideration, we additionally report results obtained using Prompt-2. This approach allows us to assess the impact of different prompt structures on model performance.

Task	Dataset	Source	Example
Summary Generation	XSum (Narayan et al., 2018)	BBC news	Input: Electrician Carl Holdsworth has set up holographic video footage of Mr and Mrs Claus behind the windows of his Chad-den house... Output: Festive revellers have travelled for miles to see Father Christmas and his wife apparently living in a Derby home.
	XLSum (Hasan et al., 2021)	BBC news	Input: Jack McLinden, who has multiple health conditions, experienced joining his heroes on the pitch before their game against Newcastle United on Monday... Output: A 14-year-old Everton fan has made history by becoming football’s first “remote” match-day mascot - with the aid of a robot.
	PeerRead (Kang et al., 2018)	Scientific peer reviews	Input: We explore techniques to maximize the effectiveness of discourse information in the task of authorship attribution... Output: Leveraging Discourse Information Effectively for Authorship Attribution.
	CNN/DailyMail (Hermann et al., 2015)	CNN news and DailyMail	Input: Chris Brown sat alone in court for 35 minutes on Friday while his lawyer talked with the judge and prosecutor behind closed doors in his probation violation case... Output: Judge orders Brown to come back to court on June 10. Prosecutors accuse Brown of not finishing 180 days of community labor.
Question Generation	Socialqa (Sap et al., 2019)	Social Commonsense	Input: Sydney is always respected by Kai. They were able to make up Kai’s mind. Answer: thank Sydney sincerely Output: What will happen to Kai?
	Tweetqa (Xiong et al., 2019)	Twitter	Input: Getting taxis is a nightmare - local drivers confused with new street layout, translations on phone app! Answer: getting taxis Output: what does ben have nightmares of?
	Sciq (Welbl et al., 2017)	School Science Textbooks	Input: Archaea live everywhere on Earth, including extreme environments. Answer: everywhere Output: Where do archea live?
Sentiment Classification	Amazon (Keung et al., 2020)	Product review	Input: I will not use it again. Made my dogs feel bad. One of my dogs lost hair because of this product. I had to end up washing my dogs to remove as much as possible. Output: negative
	AmazonFood (McAuley and Leskovec, 2013)	Food review	Input: Great taste, texture, flavor, and works well with any recipe! Wonderful pricing, too! Gluten free products are so expensive normally! Definitely recommend this! Output: positive
	SST2 (Wang et al., 2019)	Movie review	Input: In its dry and forceful way, it delivers the same message as Jiri Menzel’s Closely Watched Trains and Danis Tanovic’s No Man’s Land. Output: positive
	Yelp	Yelp review	Input: Service sucks! They’re usually understaffed and we were told to wait an hour for a table to be cleaned. There was only 1 chef and after we sat we waited for our food for 20 minutes!!!! Even the appetizers took 20 minutes. Will never recommend this place to anyone again. Big big mistake. Output: negative

Table 2: Dataset Description #1

Task	Dataset	Source	Example
Paraphrase Detection	Paws (Wang et al., 2018)	Wikipedia	Input 1: Windows XP Mode runs Windows XP in a virtual machine , and displays applications within separate windows on the Windows 7 desktop. Input 2: Windows XP - Mode executes Windows XP on a separate machine and displays applications in virtual windows on the Windows 7 desktop. Output: no
	QQP (Wang et al., 2018)	Quora Question	Input 1: How do I develop the patience to read books? Input 2: How can I develop patience and love towards reading? Output: yes
	STS-B (Wang et al., 2018)	Misc.	Input 1: US postpones missile test over N Korea tensions Input 2: Embassies staying put in N. Korea despite tension Output: no
Natural Language Inference	MNLI (Williams et al., 2018)	Misc.	Input 1: i think that's great there's a few places in Houston where they're trying that out i don't know if it's the if they've done it citywide yet or not where they have the color coded uh bags and uh bins Input 2: There are a couple places in Houston where it's being tried. Output: entailment
	RTE (Wang et al., 2018)	d Wikipedia	Input 1: The third, and most remote possibility, was considered to be a so-called Jamaican coalition, again based on the party colours, between the Christian Democrats, the FDP and the Greens. Input 2: It seems unlikely that there will be a coalition between Gerhard Schroeder's Social Democrats and Angela Merkel's Christian Democratic Union. Output: neutral
	GPTNLI	Misc.	Input 1: The Great Horned Owl has a tuft of feathers around its face that makes it look like an old man's white mustache. Input 2: The Great Horned Owl is bald. Output: contradiction

Table 3: Dataset Description #2

Task	Label
Sentiment Classification	positive, negative
Paraphrase Detection	yes, no
Natural Language Inference	entailment, contradiction, neutral

Table 4: The generation labels for classification tasks.

Task	Prompt-1	Prompt-2
Summary Generation	### Input: <i>{input}</i> ### Summary:	Please read the following text: <i>{input}</i> Provide a summary:
Question Generation	### Input: <i>{input}</i> ### Answer: <i>{answer}</i> ### Question:	Given the context: <i>{input}</i> And the answer: <i>{answer}</i> Generate a suitable question:
Sentiment Classification	### Input: <i>{input}</i> ### Sentiment:	Analyze the sentiment of the following text: <i>{input}</i> Sentiment:
Paraphrase Detection	### Input_1: <i>{input_1}</i> ### Input_2: <i>{input_2}</i> ### Paraphrase Classification:	Let's compare the two sentences: Sentence_1: <i>{input_1}</i> Sentence_2: <i>{input_2}</i> Are they paraphrasing?:
Natural Language Inference	### Input_1: <i>{input_1}</i> ### Input_2: <i>{input_2}</i> ### Inference:	Consider the following texts: Text 1: <i>{input_1}</i> Text 2: <i>{input_2}</i> The relation is

Table 5: Prompt formats for each task