

# Language Models can Exploit Cross-Task In-context Learning for Data-Scarce Novel Tasks

Anwoy Chatterjee\*, Eshaan Tanwar\*, Subhabrata Dutta, Tanmoy Chakraborty  
Indian Institute of Technology Delhi

anwoy.chatterjee@ee.iitd.ac.in, eshaantanwar2000@gmail.com  
subha0009@gmail.com, tanchak@iitd.ac.in

## Abstract

Large Language Models (LLMs) have transformed NLP with their remarkable In-context Learning (ICL) capabilities. Automated assistants based on LLMs are gaining popularity; however, adapting them to novel tasks is still challenging. While colossal models excel in zero-shot performance, their computational demands limit widespread use, and smaller language models struggle without context. This paper investigates whether LLMs can generalize from labeled examples of predefined tasks to novel tasks. Drawing inspiration from biological neurons and the mechanistic interpretation of the Transformer architecture, we explore the potential for information sharing across tasks. We design a cross-task prompting setup with three LLMs and show that LLMs achieve significant performance improvements despite no examples from the target task in the context. Cross-task prompting leads to a remarkable performance boost of 107% for LLaMA-2 7B, 18.6% for LLaMA-2 13B, and 3.2% for GPT 3.5 on average over zero-shot prompting, and performs comparable to standard in-context learning. The effectiveness of generating pseudo-labels for in-task examples is demonstrated, and our analyses reveal a strong correlation between the effect of cross-task examples and model activation similarities in source and target input tokens. This paper offers a first-of-its-kind exploration of LLMs' ability to solve novel tasks based on contextual signals from different task examples.

## 1 Introduction

Large Language Models (LLMs) have revolutionized the state of Natural Language Processing for the past few years. With the ability of In-Context Learning (ICL), one can adopt an LLM to almost any task without costly gradient updates (Brown et al., 2020). At the same time, automated assistants, built on top of foundational LLMs, are

being popularized (Pahune and Chandrasekharan, 2023). A crucial challenge with this escalating usage popularity is handling novel tasks. Humongous models like GPT-4 are able to deliver up-to-par performance even in a zero-shot regime (OpenAI, 2023). However, the computational requirements of deploying such large-scale models counteract the practicality of their usage en masse. Relatively smaller LMs, on the other hand, suffer drastically in the absence of in-context examples. The availability of labeled examples usually varies across the use cases of the language model. For example, in an NLP research setup, expert users can quickly come up with a few handwritten examples. However, when we consider mass-scale usage of average users who are not experienced prompt engineers or need quick answers, the zero-shot performance of a model becomes extremely crucial. For example, assuming the popular usage of ChatGPT, very few non-expert users would opt to write down examples while asking ChatGPT to perform some tasks.

This naturally raises the question of whether one can make an LLM generalize from labeled examples of a predefined set of tasks to an input defining a novel task. In the world of biological neurons, such abilities are commonplace: inculcating specific limb usage into an untrained limb while training the opposite limb (Ruddy and Carson, 2013), or relatively easier adoption of newer skills from the culminated experience of older skills (Qin et al., 2013; Márton et al., 2021). Drawing a blunt parallel between biological neurons and LLMs would be naive. However, one can find a supporting intuition in the mechanistic interpretation of the Transformer architecture (Elhage et al., 2021; Conmy et al., 2023; Wang et al., 2022a). One can argue that if information pathways necessary to solve a novel task are similar to those corresponding to some different task from a task library, an LLM may gather useful information across tasks. Earlier

\*Equal contribution named in alphabetical order

evidence found by Tanwar et al. (2023) also elicits intuitive motivation as they showed that LLMs can learn to infer from cross-lingual examples if proper alignment is provided.

In this work<sup>1</sup>, we design a cross-task prompting setup (Section 2) using three different LLMs: LLaMA-2 7B and 13B, and GPT 3.5; we select 50 different pairs of tasks where one serves as a *source* (i.e., context example task) and the other as *target*. Despite no examples from the target task presented in the context, LLMs can produce a staggering improvement over the zero-shot regime; on average, cross-task prompting improves performance by 107% for LLaMA-2 7B, 18.6% for LLaMA-2 13B and 3.2% for GPT 3.5 (Section 3). With multiple source tasks, cross-task performance is even better than, if not comparable, usual in-task prompting (**Contribution #1**). However, learning from examples of different tasks is heavily sensitive to the choice of source task for a given target and the LLM is prone to copy the label space of the source task into the target. To circumvent this, we propose a pseudo-labeling based approach: in a data-scarce setup, cross-task prompting with majority voting is first employed to generate noisy, in-task examples; these are subsequently used for standard few-shot prompting (**Contribution #2**). Finally, we provide introductory analysis towards interpreting cross-task signal transfer by dissecting the model activations. We find that the cross-task signal transfer is abrupt and happens at later layers, with the effective layers widely varying for different target tasks (**Contribution #3**). In a nutshell, this is the first exploration of LLMs’ ability to learn to solve novel tasks based on the contextual signals of different task examples.

## 2 Prompting techniques

**Task definition** for a given task is a natural language instruction for the LM describing what is asked of it (see Figure 1). Since a cross-task setup would result in in-context examples from different tasks (with different label spaces), such a definition is necessary to discriminate.

Next, given two datasets  $D_s$  and  $D_t$  corresponding to two different tasks with task definitions  $d_s$  and  $d_t$ , respectively, we formalize the cross-task prompting as inferring the output  $\hat{y}_t$  for an input  $x_t \in D_t$  conditioned upon a demonstration from

dataset  $D_s$ :

$$\hat{y}_t = \operatorname{argmax}_y p(y | (d_s \oplus x_s \oplus y_s \oplus d_t \oplus x_t))$$

In this setup, we denote  $D_s$  and  $D_t$  as source and target task datasets, respectively. Note that the formalization is equivalent to 1-shot prompting where the only provided example input-output pair comes from a different dataset. Table 14 of Appendix F presents illustrative examples of prompts in this setup. Our experiments suggest that >1-shot setup in cross-task prompting does not improve (and often deteriorates) performance.

Liu et al. (2022) showed that for a target input  $x_t$ , generating context  $C$  from semantically similar examples leads to not only better results but also a more robust method of prompting. Similarly, Tanwar et al. (2023) showed that ICL could be done in a cross-lingual setup by aligning the semantic and task signals of the source and target languages.

Drawing inspiration from these works, we set up the cross-task prompting regime with sampling examples  $x_s$  from the source task dataset  $D_s$  that are semantically similar to the target input  $x_t$ . To extract semantically similar examples, we first utilize Sentence-BERT (Reimers and Gurevych, 2019) to extract the sentence embedding of target and source inputs<sup>2</sup>. Following this, based on the cosine similarity between the embeddings, we select top source examples.

**In-task examples combined with cross-task prompting.** So far, we have assumed the unavailability of any labeled target dataset. But, if we do have a labeled target example, could prepending a source task boost its prompting performance?

To emulate such a scenario, a labeled example from the target dataset  $T_t$  is sampled. This labeled example  $(x_{lt}, y_{lt})$ , where  $x_{lt} \in T_t$ , is then used to construct prompt. This mixed setup can be formalized as,

$$\hat{y}_t = \operatorname{argmax}_y p(y | (d_s \oplus x_s \oplus y_s \oplus d_t \oplus x_{lt} \oplus y_{lt} \oplus x_t))$$

Examples of prompts are provided in Table 15 of Appendix F.

## 3 Results and Analysis

**Datasets and experimental setup.** Our corpus of tasks consists of ten source and five target tasks.

<sup>1</sup>Code available at <https://github.com/C-anwoy/Cross-Task-ICL>

<sup>2</sup>We also experimented with E5 and LLaMA-2 7B last layer outputs, see Appendix B.

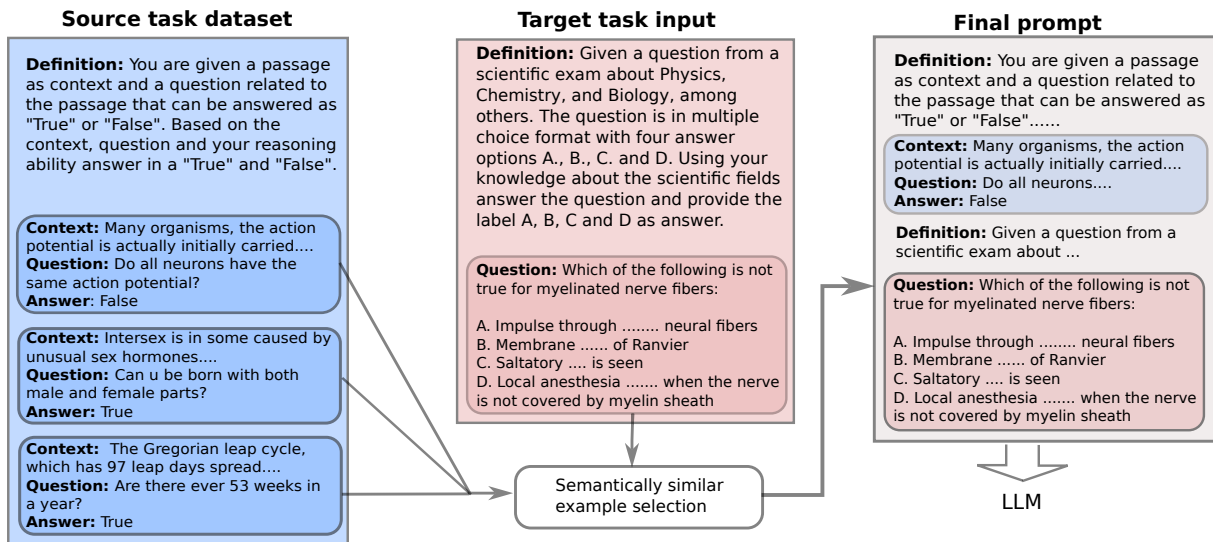


Figure 1: In this working example, we aim to solve a question from MedMCQA using demonstrations from BoolQ. To do so, we sample a semantically similar demonstration from the source task and then use this demonstration along with task descriptors of the source and target tasks to generate a cross-task prompt that is fed to an LLM.

TAR \ SRC	SRC											
	Zero-shot	ARC-Easy	AG-news	BoolQ	Com-QA	C-POS	C-NER	MNLI	QQP	RACE	SST2	
LLaMA-2 7B												
ARC-Challenge	4.6	<b>43.6</b>	33.6	35.0	<b>43.6</b>	33.8	34.2	34.2	36.4	42.8	33.2	
Financial-Phrasebank	34.1	43.5	62.1	40.9	14.6	1.4	0.4	62.7	44.7	53.4	<b>65.0</b>	
MedMCQA	4.2	31.4	26.8	28.0	<b>33.0</b>	26.2	24.0	23.0	26.0	31.6	23.2	
SciQ	8.0	59.0	45.4	49.0	<b>65.6</b>	34.4	29.7	44.8	25.4	64.4	39.0	
Social-i-QA	41.1	44.3	40.1	40.3	48.5	38.9	39.1	38.9	39.7	<b>49.1</b>	39.3	
Average	18.4	44.3	41.6	38.6	40.71	26.9	25.5	40.7	34.4	<b>48.3</b>	39.9	
LLaMA-2 13B												
ARC-Challenge	52.0	<b>59.2</b>	50.6	54.6	57.4	50.6	49.8	49.2	52.8	56.8	50.8	
Financial-Phrasebank	65.4	66.6	78.6	65.8	61.7	2.2	13.17	79.0	<b>82.4</b>	66.5	72.8	
MedMCQA	9.2	37.2	34.2	36.6	38.8	31.6	32.6	30.2	30.6	<b>39.0</b>	31.4	
SciQ	55.8	<b>83.4</b>	76.2	80.2	<b>83.4</b>	76.4	78.6	77.2	71.2	82.8	72.6	
Social-i-QA	55.3	60.8	55.8	56.5	63.5	56.8	55.9	53.5	55.7	<b>63.7</b>	53.1	
Average	47.5	61.4	59.1	58.7	60.9	43.5	46.0	57.8	58.5	<b>61.7</b>	56.1	
GPT3.5												
ARC-Challenge	74.6	77.2	74.4	77.8	76.2	70.2	70.8	75.2	74.0	<b>78.2</b>	72.6	
Financial-Phrasebank	57.5	79.8	<b>83.6</b>	78.6	82.0	50.3	64.3	72.2	74.2	73.4	75.6	
MedMCQA	49.6	49.8	48.5	<b>50.0</b>	48.0	47.4	45.0	46.0	49.0	47.6	48.0	
SciQ	91.2	91.4	89.9	90.6	91	87.8	82.0	88.8	89.8	<b>92.2</b>	89.0	
Social-i-QA	76.0	76.2	75.8	74.0	75.2	73.2	74.0	75.8	73.0	<b>77.2</b>	74.8	
Average	69.8	<b>74.9</b>	74.5	74.2	74.5	65.8	67.2	71.6	72.0	73.7	72.0	

Table 1: Accuracy for cross-task setup using one source example. Source tasks are mentioned in columns, and target tasks are in rows. Cross-task prompting brings improvement over zero-shot prompting for certain source-target pairs. (Abbreviations: Com-QA → Commonsense-QA, C-POS → Conll2003-POS, C-NER → Conll2003-NER).

We consider ARC-Easy (Clark et al., 2018), AG-news (Zhang et al., 2015), BoolQ (Clark et al., 2019), Commonsense-QA (Talmor et al., 2019), Conll2003-POS (Tjong Kim Sang and De Meulder, 2003), Conll2003-NER (Tjong Kim Sang and De Meulder, 2003), MNLI (Williams et al., 2018), QQP (Sharma et al., 2019), RACE (Lai et al., 2017) and SST2 (Socher et al., 2013) to be our source tasks. Our motivation remains to incorporate domain diversity and difficulty of the problems while choosing target tasks to emulate the “novel task” phenomenon as closely as possible

since learning from cross-task examples makes sense only when the target task is truly data-scarce. Our target tasks are Social-i-QA (Sap et al., 2019), SciQ (Auer et al., 2023), MedMCQA (Pal et al., 2022), Financial-Phrasebank (Malo et al., 2014), and ARC-Challenge (Clark et al., 2018); first four of these require domain expertise, while the last one is a more challenging version of the one present in the source tasks. In total, this gives us 50 unique cross-task setups (see Appendix A for dataset details. Table 16 and Table 17 of Appendix F show task definitions corresponding to source and target

tasks, respectively). The dataset size is standardised by sampling 10,000 and 500 examples for each source and target task, i.e.,  $|D_s| = 10,000$  and  $|D_t| = 500$ . For our experiments, we used the 7-billion and 13-billion variants of LLaMA-2 (Touvron et al., 2023), and text-davinci-003 (Brown et al., 2020) referred to as GPT3.5. We experiment with greedy and force decoding setup and selected greedy decoding as our standard for all experiments (see Appendix D for more details). We set the number of examples used to create cross-task context as one for all our experiments, unless mentioned explicitly.

**Does cross-task prompting work?** As evident from Table 1, cross-task prompting significantly improves performance compared to zero-shot prompting (see Table 20 in Appendix for results of significance testing). The best overall source-target pair improves performance by 162% for LLaMA-2 7B, 30% for LLaMA-2 13B and 7% for GPT3.5. We note that different models give the best performance in different source-target pairs and not all coupling of source-target tasks seem to work; e.g., Commonsense-QA as a source task decreases performance on Financial-Phrasebank, but is the best source task for ARC-Challenge and MedMCQA (for LLaMA-2 7B). Token classification tasks (POS and NER) seem to depreciate performance for LLaMA-2 13B and GPT3.5; their performance for LLaMA-2 7B is also sub-par compared to other source-target pairings. RACE and ARC-Easy are robust source tasks that improve performance for all target tasks in all models. ARC-Easy, MNLI and BoolQ can also be considered to have this robustness to some degree, as they only hurt the performance in one or two cases.

On average, cross-task prompting improves performance by 107% for LLaMA-2 7B, 18.6% for LLaMA-2 13B and 3.2% for GPT3.5. This is a strong argument for prompting in a cross-task manner when we lack labeled target data, especially using small models, which have poor zero-shot abilities (Wei et al., 2022). We also experiment with LLaMA-2 7B Chat model to analyse the performance of cross-task prompting for instruction-tuned models (see Appendix C).

**Performance with dissimilar source tasks.** Few of the chosen source tasks have slight overlap in their broad topic or format with some of the target tasks, for example, source datasets like ARC-Easy, Commonsense-QA, RACE consists of ques-

tions in multiple-choice format, similar to target datasets like ARC-Challenge, Social-i-QA, SciQ and MedMCQA. To analyse the efficacy of cross-task prompting with completely dissimilar source tasks, let us consider only the source tasks whose topic and format are absolutely disjoint from all target tasks. The following five source tasks can be considered dissimilar to the target tasks – BoolQ, Conll2003-POS, Conll2003-NER, MNLI and QQP. As reported in Table 1, even with these dissimilar source tasks, we observe a substantial performance improvement of 80.5% for LLaMA-2 7B, 11.4% for LLaMA-2 13B, and 0.5% for GPT3.5 on average over zero-shot prompting. Moreover, if we consider only BoolQ, MNLI and QQP, the average performance gains become 105.9% for LLaMA-2 7B, 22.8% for LLaMA-2 13B, and 4% for GPT3.5. This shows that cross-task prompting can especially enable the relatively smaller LMs to substantially boost their zero-shot performance even using tasks dissimilar to the target.

**Importance of source task definitions.** To further investigate the role of task definitions in cross-task prompting, we check the effect of removing source task definitions on performance. We note an average drop of 11% for LLaMA-2 7B and 8% for LLaMA-2 13B when we prompt without using source task definitions (see Table 13 in Appendix). Hence, definitions play a crucial role in cross-task prompting.

**Increasing number of examples for cross-task prompting.** Unlike in-task prompting (Brown et al., 2020), where increasing the number of examples increases prompting performance, in cross-task prompting, performance does not improve with an increase in source examples (c.f. Fig 3). For most target tasks, increasing source examples does not affect performance, while for some, the performance decreases.

**Semantically similar vs random example selection.** As shown in Table 2, in a cross-task setup, choosing the examples randomly leads to substantially poorer performance than when we generate the context  $C$  with semantically similar examples; concerning is the fact that, in many cases, it causes 0% accuracy. This may be caused by the model getting confused without any semantic alignment in the prompt. Furthermore, random labeling of in-context examples (Wei et al., 2023) result in near-random performance (see Table 12 in Appendix E).

**Mixed cross-task prompting.** So far, we have seen

SRC \ TAR	AG-news	BoolQ	Com-QA	MNLI	QQP
Random					
ARC-Challenge	26.6 <sub>3.0</sub>	3.7 <sub>0.7</sub>	0.0 <sub>0.0</sub>	29.5 <sub>3.7</sub>	26.8 <sub>2.9</sub>
Financial-Phrasebank	0 <sub>0</sub>	1.4 <sub>1.3</sub>	13.9 <sub>14.75</sub>	0.0 <sub>0.0</sub>	0.0 <sub>0.0</sub>
MedMCQA	26.1 <sub>1.5</sub>	20.8 <sub>1.1</sub>	0.0 <sub>0.0</sub>	25.1 <sub>0.9</sub>	25.0 <sub>0.2</sub>
SciQ	34.5 <sub>3.2</sub>	10.0 <sub>4.25</sub>	0.0 <sub>0.0</sub>	34.0 <sub>4.1</sub>	28.7 <sub>0.8</sub>
Social-i-QA	36.8 <sub>1.3</sub>	0.6 <sub>0.1</sub>	0.0 <sub>0.0</sub>	24.9 <sub>2.7</sub>	29.2 <sub>7.9</sub>
Semantically Similar					
ARC-Challenge	36.6	35.0	43.6	34.2	36.4
Financial-Phrasebank	62.1	40.9	14.6	62.7	44.7
MedMCQA	26.8	28.0	33.0	23.0	26.0
SciQ	45.4	49.0	65.6	44.8	25.4
Social-i-QA	40.1	40.3	48.5	38.9	39.7

Table 2: Accuracy when cross-task examples are picked randomly vs when they are pickled based on semantic similarity. Experiments were done on the LLaMA-2 7B model. For random prompting, we generated the results over three seeds and reported the average performance (Com-QA: Commonsense-QA).

	Best source cross-task	Random mixed cross-task	Best mixed cross-task
LLaMA-2 7B			
ARC-Challenge	51.6	42.8	44.0
Financial-Phrasebank	64.7	48.3	56.0
MedMCQA	34	28.6	29.0
SciQ	46.2	60.5	67.2
Social-i-QA	42.7	42.9	42.9
Average	47.84	44.62	47.82
LLaMA-2 13B			
ARC-Challenge	66.4	57.2	60.4
Financial-Phrasebank	76.6	71.4	61.5
MedMCQA	38.4	38.3	41.2
SciQ	84.6	80.5	84.2
Social-i-QA	49.3	61.9	61.27
Average	63.06	61.86	61.71

Table 3: Accuracy in mixed cross-task prompting. Four examples are used to create context  $C$  in all setups.

that cross-task prompting works for single source-target task pairs. Next, we experiment on using multiple source tasks to construct the prompt context. To explore such a setup, we prompt LLaMA models using three methods:

- Best source cross-task:** We select the best source task for every target task using Table 1 and sample four semantically similar examples from that source task.
- Random mixed cross-task:** To see if a diverse set of tasks is beneficial, we randomly sample four source tasks and construct the prompt using most semantically similar examples.
- Best mixed cross-task:** This method is a combination of the first two; we use the top four best source tasks from Table 1 and sample a semantically similar source example from each task.

Table 3 shows that a mixed prompting mechanism does not perform better than the “best source cross-task” prompting method. On the contrary, it seems to hurt the performance of the model; in fact, single-source task prompting with only one example (Table 1) seems to do better than mixed prompting. Diversity seems to lead the model to

get more confused, thus hurting its performance.

### Combining in-task with cross-task prompting.

Combining cross-task prompting with labeled target examples improves performance, as seen in Table 4; apart from Financial-Phrasebank in LLaMA-2 13B and MedMCQA in GPT3.5, for all other instances there exist a source task that improves the performance when coupled with the in-task example. However, we see that this improvement is immensely dependent on the source-target task pair chosen and unlike Table 1, we are unable to find robust source datasets that improve the performance throughout the setup. Nevertheless, in multiple source-target instances, there is a noteworthy improvement.

To study the interaction between heterogeneous tasks in the context, we experiment by varying the number of source tasks and target task demonstrations in the context. Figure 2 shows the variation of accuracy for 8-shots, as we gradually move from an entirely in-task context to a complete cross-task context. We observe that for all target tasks, apart from Financial-Phrasebank, the accuracy with an 8-shot in-task prompt and an 8-shot cross-task prompt (with *best-mixed cross-task* strategy) is almost the same in all three LLMs.

## 4 Pseudo-label generation using cross-task prompting

Thus far we observe that cross-task prompting, though sometimes capable of even outperforming standard in-task prompting, is particularly sensitive to the choice of source task. Furthermore, the performance does not scale with the number of source task examples provided. Drawing inspiration from earlier works on pseudo-labeled examples to construct prompts (He et al., 2022; Vu et al., 2021), we propose a more practical method for potential usage of cross-task prompting.

Given a small unlabeled dataset  $D_{pl} \subset D_t$ , we assign a pseudo-label to the example using cross-task prompting. This is done using all source tasks available to us, and then a final  $y_{pl}$  is assigned to  $x_{pl} \in D_{pl}$  based on a majority vote from all the generated answers. Finally, this pseudo-labeled dataset  $D_{pl}$  is used to construct the prompt context in an in-context prompting setup.

**Cross task generated pseudo-examples vs gold label.** To see the efficacy of our proposed method, we utilise a  $D_{pl}$  of size 8 and have three setups to create examples for context: 1. *Gold-label:* We

SRC		TAR										
		In-task	+ARC-Easy	+AG-news	+BoolQ	+Com-QA	+C-POS	+C-NER	+MNLI	+QQP	+RACE	+SST2
LLaMA-2 7B												
ARC-Challenge		49.8	<b>51.2</b>	41.6	45.8	49.6	43.6	40.2	41.6	45.2	50.0	42.8
Financial-Phrasebank		33.3	33.3	33.3	33.3	33.5	33.3	34.5	35.7	33.3	34.5	<b>50.3</b>
MedMCQA		32.4	33.6	29.2	30.6	32.8	31.0	30.8	28.2	30.2	<b>34.2</b>	31.2
SciQ		67.4	<b>73.4</b>	60.4	68.6	71.6	59.0	57.4	58.6	57.9	71.2	62.4
Social-i-QA		43.3	49.9	46.7	44.5	<b>58.3</b>	44.7	44.3	41.5	43.9	51.7	43.1
Average		45.2	48.3	42.2	44.5	48.1	42.3	41.4	41.1	42.1	<b>48.3</b>	45.9
LLaMA-2 13B												
ARC-Challenge		61.8	<b>63.0</b>	60.8	61.6	61.0	61.2	60.0	60.0	59.6	61.2	58.8
Financial-Phrasebank		<b>86.0</b>	77.0	71.2	77.6	82.4	65.1	69.4	75.2	64.8	81.8	85.6
MedMCQA		40.2	40.8	37.6	40.4	41.0	38.2	39.0	38.6	38.6	<b>41.8</b>	37.2
SciQ		84.8	<b>85.8</b>	83.0	83.8	85.6	81.4	81.4	80.2	79.2	84.0	79.6
Social-i-QA		54.8	<b>62.1</b>	52.09	52.3	60.9	53.3	53.7	54.1	53.5	61.6	53.9
Average		65.5	65.7	60.9	63.1	66.2	59.8	60.7	61.6	59.1	<b>66.0</b>	63.0
GPT3.5												
ARC-Challenge		78.6	78.2	78.2	78.8	<b>79.8</b>	77.4	78.4	78.4	78.4	77.6	79.0
Financial-Phrasebank		75.4	80.2	79.4	79.6	83.8	70.4	80.2	80.2	81.0	<b>81.6</b>	77.6
MedMCQA		<b>52.0</b>	51.2	48.6	51.2	47.4	48.2	49.4	49.4	48.8	49.8	49.8
SciQ		92.6	91.2	91.8	92.6	92.4	91.8	92.0	92.0	91.4	<b>93.0</b>	91.6
Social-i-QA		75.6	77.0	76.2	74.4	76.6	76.2	<b>77.4</b>	77.4	75.0	77.0	75.6
Average		74.8	75.6	74.8	75.3	<b>76.0</b>	72.8	74.2	75.0	74.9	75.8	74.7

Table 4: Accuracy for in-task combined with cross-task prompting setup; prompt context is created using one source example and one human-labelled target example. Cross-task prompting compliments in-task prompting, leading to better performance (Com-QA: Commonsense-QA, C-POS: Conll2003-POS, C-NER: Conll2003-NER).




Figure 2: Variation of accuracy with the number of cross-task demonstrations in the context C for 8-shot prompt, where C consists of a mixture of cross-task and in-task examples, using (a) LLaMA-2 7B, (b) LLaMA-2 13B and (c) GPT3.5. Except for Financial-Phrasebank, an increase in cross-task examples give a stable performance.

use the annotated version of  $D_{pl}$  in the context; 2. *Pseudo-label (ZS)*: We use zero-shot prompting to label  $D_{pl}$  for the context examples; 3. *Pseudo-label (CT)*: We use the cross-task method as proposed in Section 4.

The results of this experiment are shown in Table 5. For LLaMA models, as evident, pseudo labels generated via cross-task prompting are substantially better than zero-shot pseudo labels; they are of higher quality and lead to comparable performance as Gold-label.

In the case of GPT 3.5, the scale of improvement is much smaller though. Interestingly, with only two datasets, namely, SciQ and Social-i-QA, we observe the pseudo examples from cross-task prompting to perform worse than gold-label examples, in all three models. Given the comparable performance of cross-task prompt-generated pseudo examples, we expect this to become a viable alternative to traditional ICL in data-scarce scenarios.

	Gold-label	Pseudo-examples(ZS)	Pseudo-examples(CT)
LLaMA-2 7B			
ARC-Challenge	44.40.37	7.24.52	44.40.37
Financial-Phrasebank	48.49.56	56.52.01	56.82.37
MedMCQA	30.60.84	21.014.61	30.61.03
SciQ	67.70.65	3.31.88	64.61.51
Social-i-QA	46.81.69	46.71.57	46.41.94
Average	47.62.62	26.94.91	48.51.44
LLaMA-2 13B			
ARC-Challenge	56.80.47	59.70.84	58.41.79
Financial-Phrasebank	67.95.03	67.95.03	67.99.50.3
MedMCQA	38.21.31	2.61.8	39.40.37
SciQ	84.10.75	29.339.5	84.10.75
Social-i-QA	63.62.78	43.627.7	61.43.66
Average	62.12.06	40.6214.97	62.22.32
GPT3.5			
ARC-Challenge	76.70.89	77.41.13	77.460.94
Financial-Phrasebank	80.44.35	86.22.84	83.436.41
MedMCQA	43.25.65	43.43.30	47.82.97
SciQ	92.81.04	91.20.56	91.50.25
Social-i-QA	77.81.06	76.00.65	76.60.43
Average	74.22.59	74.91.69	75.32.2

Table 5: Accuracy in pseudo-label generated prompting (ZS: zero-shot, CT: cross-task). Pseudo-label (CT) performs comparably to gold-label examples.

**Increasing number of pseudo-examples for in-task prompting.** Figure 3 shows the relation between the number of pseudo-demonstrations and

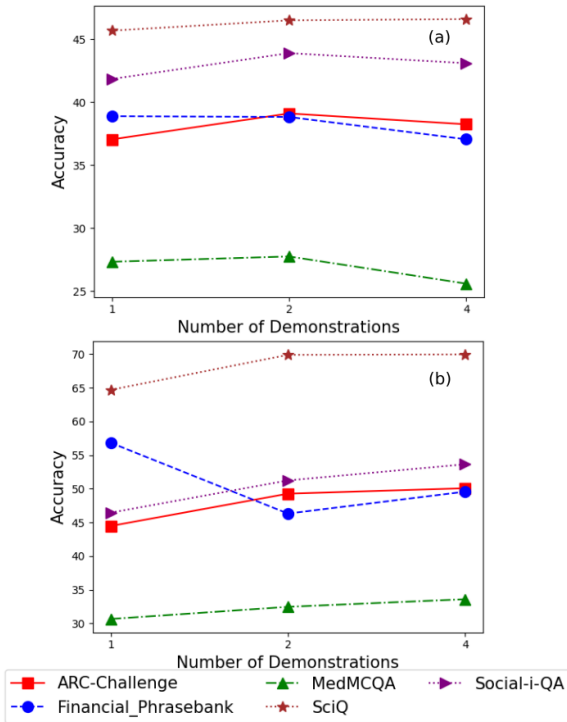


Figure 3: Variation of accuracy with change in the number of – (a) source demonstrations from  $D_s$  (the performance largely stagnates at  $k = 1$ ), (b) pseudo-label sampled from  $D_{pt}$  (the performance seems to increase with an increase in  $k$ ). All experiments are done on LLaMA-2 7B model.

accuracy. The general trend shows a rise in performance when more demonstrations are used in creating context  $C$ , except for Financial-Phrasebank, whose performance decreases with an increase in demonstrations.

## 5 Interpretability analysis

In this section, we focus on the internal workings of the LLMs using the hidden states of the model and build an understanding of why and how cross-task prompting works.

**What is an ideal source task?** Intuitively thinking of a source task that is similar to the target will result in better transfer of generalised signal from context to target inputs. To test this, we first compare the final layer outputs of the model corresponding to the source and target task definitions,  $d_s$  and  $d_t$ , respectively. Figure 4 shows the cosine similarity between the final layer hidden states corresponding to  $d_s$  and  $d_t$ . We note that for 80% of the time, the source definition that is the most semantically similar to the target definition also serves as its best suitor, leading to the best performance in Table 1.

**What does internal activation indicate?** To get

an even more in-depth picture, we analyse the layerwise activations for the LLaMA-2 7B model. For each layer, we compute the cosine similarities between the mean activations corresponding to the source and target task definitions. For a given target task, we then compute the rank correlation (in terms of Spearman correlation coefficient, Pearson coefficient, and, Kendall’s tau) between the cosine similarities and the absolute point performance change from zero-shot prediction for each source task (see Table 19 in Appendix). Precisely, this gives us an approximate idea about the underlying mechanism of cross-task signal transfer. For all the tasks, we can see a U-shaped pattern in the correlation values (c.f. Fig 5) – at the starting layers, there is a high correlation between the cross-task activation similarity and the cross-task improvement (likely due to semantically similar example selection), that quickly drops in the middle layers, and increases again in the later layers (only exception being the MedMCQA target task). One can intuitively claim that in the initial layers of the model (Layers 2 to 5), there is more task-specific computation going on where the cross-task transfer of information is the least. For Financial-Phrasebank, these task-specific layers cover more than 80% of the layers, with a gradual increase in correlation observable only after Layer 28. Hendel et al. (2023) provided a similar finding that the influence of the context kicks in only after a certain number of layers via *task vectors*. We see that the exact layer after which cross-task demonstrations will start signal to transfer is much more dependent on the target task and can vary widely.

## 6 Error analysis

We observe four types of error occurring in cross-task prompting (c.f. Table 6), as follows:

**Label space replication.** In example #1, the source example is from the *Conll2003-POS* dataset, a POS tagging task. In contrast, the target task is to predict the financial sentiment analysis. Therefore, the generation should be either negative, positive or neutral; however, we observe that the output is a sequence of POS tags instead. Here, the LLM is replicating the label space of the source task and the definition of the target task is not able to guide it to the correct target label space.

**Junk prediction.** In some cases, we observe that the output is neither from the label space of the source task nor from that of the target task. As




Figure 4: Heat map of cosine similarities between the mean activations of the task definitions for source-target pairs, extracted from the final layer of (a) LLaMA-2 7B (b) LLaMA-2 13B. The values in the cells show the absolute performance improvement over zero-shot prompting. For a target task, the source-target pair with the highest cosine similarity shows the most improvement in 80% cases. (ARC-C: ARC-Challenge, F-Phr: Financial-Phrasebank, MMCQ: MedMCQA, S-QA: Social-i-QA, AGN: AG-news, ARCE: ARC-Easy, Com-QA: Commonsense-QA, C-POS: Conll2003-POS, C-NER: Conll2003-NER).

#	Prompt	Output
1.	<b>Definition:</b> Given a sentence do token classification by doing Part-of-speech (POS) tagging, ... categories. <b>Sentence:</b> Net profit 6.08 vs 3.98 <b>Label:</b> JJ NN CD NNP CD <b>Definition:</b> Given a sentence ... label the sentence as "negative", "positive" or "neutral" <b>Sentence:</b> Consolidated net sales increased 16 % ... to a loss of EUR0.7 m in the prior year period. <b>Label:</b>	NNP CD NNP CD
2.	<b>Definition:</b> Given a sentence do token classification ... till the entity ends. <b>Sentence:</b> Total shares to be offered 2.5 million <b>Label:</b> O O O O O O O <b>Definition:</b> Given a sentence ... label the sentence as "negative", "positive" or "neutral" <b>Sentence:</b> The offer of some 30 million shares ... expected to be completed by Oct. 9 , Outokumpu said. <b>Label:</b>	N N N N N N N
3.	<b>Definition:</b> Given a question answering task ... most appropriate choice that answers the question <b>Question:</b> Which object in our solar system reflects light and is a satellite that orbits around one planet? A. Earth ... D. the Moon <b>Answer:</b> D <b>Definition:</b> Given a question answering task ... most appropriate choice that answers the question <b>Question:</b> Which of the following sets contains only objects that shine as a result of reflected light? A. moons, planets, and comets ... D. planets, stars, and moons <b>Answer:</b>	D
4.	<b>Definition:</b> Given a context and a question ... answer in a "True" and "False". <b>Context:</b> A blacklight (or often black light), also referred to as a UV-A light, ... much visible light. <b>Question:</b> is a black light and an ultraviolet light the same thing <b>Answer:</b> True <b>Definition:</b> Given a question from a scientific exam ... provide the label "A", "B", "C" and "D" as answer <b>Question:</b> When exposed to ultraviolet, ... characteristic visible wavelengths, a process called this? A. pigment ... D. fluorescence <b>Answer:</b>	fluorescence

Table 6: Error analysis of cross-task prompting. Four examples represent the major error characteristics (discussed in Section 6). The outputs shown are generated by LLaMA-2 13B. The examples shown here are shortened, complete examples are provided in Table 18 of Appendix.

in example #2, the source task is to classify each token into one named entity category (NER), and the target task is to predict the financial sentiment. The output in this case is found to be junk — it is a sequence of *N*'s but *N* is not a pre-defined named entity category, neither is it a label for the target task. Further, we observe that the correct label in this example is 'neutral' – the LLM could not follow the target task definition provided and gets confused between the correct prediction label of the target task and the token classification task where the output is expected to be a sequence of labels, and instead outputs a sequence with the initial letter of the correct label.

**Copying effect:** We also notice the copying effect, which has been observed by Lyu et al. (2023), where the predicted label is the same as the label of the context example which is semantically very similar to the target input. In example #3, the target input is semantically very similar to the context ex-

ample from the source task provided in the prompt, and consequently, the LLM incorrectly outputs the same label as that of the context example.

**Definition not followed:** We observe that for some instances the LLM does not adhere to the definition given in the form of task definition for the target task. In example #4, the LLM is supposed to output one among the four options *A*, *B*, *C*, *D*, instead it outputs the text corresponding to the option *D* – though *D* is the correct answer in this case, the LLM is not able to follow the definition properly.

## 7 Related work

In-context learning without gradient updates to the model was introduced by Brown et al. (2020) using the GPT-3 model. Multiple recent works sought robust ICL setup via different techniques: selecting the examples that are semantically most similar to the input (Liu et al., 2022), choosing low perplexity examples (Gonen et al., 2022), training a






Figure 5: Variation of (a) Spearman’s correlation coefficient, (b) Pearson correlation coefficient and (c) Kendall’s tau, with layers as reported in Table 19 for LLaMA-2 7B.

two-stage retrieval system to select the best examples (Rubin et al., 2022), etc. However, these works primarily aim to construct better in-task prompts where the examples and the input come from the same task. Tanwar et al. (2023) showed that cross-lingual ICL can be elicited with proper alignment between the source and target language examples. Raffel et al. (2020) introduced the *T5 transformer* model which has been trained to perform different text-processing tasks. Zhang et al. (2022) proposed a task prefix guided multi-task pre-training framework to solve this problem. More recently, a new prompting method called *Meta-CoT* has been proposed by Zou et al. (2023) which generalizes chain-of-thought prompting to include multi-task examples in the prompt and it has shown improvement in a number of tasks.

## 8 Conclusion

In this paper, we addressed LLMs’ adaptability to novel tasks. Exploiting the inherent ICL capabilities of these models, we established that LLMs can substantially enhance their performance in novel task scenarios, even when direct examples are lacking. This encouraging outcome unveils fresh possibilities for the practical integration of LLMs across a broader spectrum of applications. Furthermore, our study demonstrated the significance of generating pseudo-labels using cross-task prompting, presenting a potential solution for situations where annotated data is scarce. The observed correlation between the impact of cross-task examples and the similarity in model activations between source and target input tokens offers valuable insights into the underlying mechanisms of this phenomenon.

## Limitations

Despite presenting a potential future direction towards training-free task generalization using LLMs, this study has some important limitations. It is evident that the similarity between the source and the target tasks plays an important role in the performance. Hence, in real-world scenarios where the task novelty is extreme, such a method may fail to provide suitable performance. This is directly related to the fact that ours is the first study in this direction, and we have primarily focused on the empirical viability. A deeper understanding of generalizable task information captured inside the LLM circuits would help to come up with sophisticated solutions. Task novelty in our discussion does not presuppose access to novel knowledge. Hence, one can not mitigate the gap if a novel task requires the model to access newer information not present in the pretraining data or the in-context examples.

## References

- Sören Auer, Dante A. C. Barone, Cassiano Bartz, Eduardo G. Cortes, Mohamad Yaser Jaradeh, Oliver Karras, Manolis Koubarakis, Dmitry Mourmstev, Dmitrii Pliukhin, Daniil Radyush, Ivan Shilin, Markus Stocker, and Eleni Tsalapati. 2023. [The sciqa scientific question answering benchmark for scholarly knowledge](#). *Scientific Reports*, 13(1):7240.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss,

- Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). *CoRR*, abs/2005.14165.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). 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 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#). *ArXiv*, abs/1803.05457.
- Arthur Conmy, Augustine N Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. 2023. [Towards automated circuit discovery for mechanistic interpretability](#). *arXiv preprint arXiv:2304.14997*.
- Nelson Elhage, Neel Nanda, Catherine Olsson, Tom Henighan, Nicholas Joseph, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, Tom Conerly, et al. 2021. [A mathematical framework for transformer circuits](#). *Transformer Circuits Thread*, 1.
- Hila Gonen, Srini Iyer, Terra Blevins, Noah A. Smith, and Luke Zettlemoyer. 2022. [Demystifying prompts in language models via perplexity estimation](#).
- Xuanli He, Islam Nassar, Jamie Kiros, Gholamreza Hafari, and Mohammad Norouzi. 2022. [Generate, Annotate, and Learn: NLP with Synthetic Text](#). *Transactions of the Association for Computational Linguistics*, 10:826–842.
- Roe Hendel, Mor Geva, and Amir Globerson. 2023. [In-context learning creates task vectors](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9318–9333, Singapore. Association for Computational Linguistics.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale Reading comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3?](#) In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.
- Xinxi Lyu, Sewon Min, Iz Beltagy, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Z-ICL: Zero-shot in-context learning with pseudo-demonstrations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2304–2317, Toronto, Canada. Association for Computational Linguistics.
- P. Malo, A. Sinha, P. Korhonen, J. Wallenius, and P. Takala. 2014. [Good debt or bad debt: Detecting semantic orientations in economic texts](#). *Journal of the Association for Information Science and Technology*, 65.
- Christian David Márton, Léo Gagnon, Guillaume Lajoie, and Kanaka Rajan. 2021. [Efficient and robust multi-task learning in the brain with modular latent primitives](#). *arXiv preprint arXiv:2105.14108*.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Saurabh Pahune and Manoj Chandrasekharan. 2023. [Several categories of large language models \(llms\): A short survey](#). *arXiv preprint arXiv:2307.10188*.
- Ankit Pal, Logesh Kumar Umaphathi, and Malaikannan Sankarasubbu. 2022. [Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering](#). In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.
- Wen Qin, Chunshui Yu, et al. 2013. [Neural pathways conveying nonvisual information to the visual cortex](#). *Neural plasticity*, 2013.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Lance A. Ramshaw and Mitchell P. Marcus. 1995. [Text chunking using transformation-based learning](#). *ArXiv*, cmp-lg/9505040.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- 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.

- Kathy Ruddy and Richard Carson. 2013. [Neural pathways mediating cross education of motor function](#). *Frontiers in Human Neuroscience*, 7.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Socialliqa: Commonsense reasoning about social interactions](#). *CoRR*, abs/1904.09728.
- Lakshay Sharma, Laura Graesser, Nikita Nangia, and Utku Evci. 2019. [Natural language understanding with the quora question pairs dataset](#). *CoRR*, abs/1907.01041.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). 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 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Eshaan Tanwar, Subhabrata Dutta, Manish Borthakur, and Tanmoy Chakraborty. 2023. [Multilingual LLMs are better cross-lingual in-context learners with alignment](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6292–6307, Toronto, Canada. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Tu Vu, Minh-Thang Luong, Quoc V. Le, Grady Simon, and Mohit Iyyer. 2021. [Strata: Self-training with task augmentation for better few-shot learning](#). *CoRR*, abs/2109.06270.
- Kevin Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. 2022a. [Interpretability in the wild: a circuit for indirect object identification in gpt-2 small](#). *arXiv preprint arXiv:2211.00593*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022b. [Text embeddings by weakly-supervised contrastive pre-training](#). *ArXiv*, abs/2212.03533.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Jerry Wei, Le Hou, Andrew Lampinen, Xiangning Chen, Da Huang, Yi Tay, Xinyun Chen, Yifeng Lu, Denny Zhou, Tengyu Ma, and Quoc V. Le. 2023. [Symbol tuning improves in-context learning in language models](#).
- 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. Association for Computational Linguistics.
- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *NIPS*.
- Zhuosheng Zhang, Shuohang Wang, Yichong Xu, Yuwei Fang, Wenhao Yu, Yang Liu, Hai Zhao, Chengang Zhu, and Michael Zeng. 2022. [Task compass: Scaling multi-task pre-training with task prefix](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5671–5685, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Anni Zou, Zhuosheng Zhang, Hai Zhao, and Xiangru Tang. 2023. [Meta-cot: Generalizable chain-of-thought prompting in mixed-task scenarios with large language models](#).

## A Dataset details

### A.1 Source datasets

We have used the following datasets as source datasets:

**ARC-Easy:** ARC-Easy (Clark et al., 2018) is a part of the ARC (AI2 Reasoning Challenge) dataset consisting of easy natural science questions targeted for students of 3rd to 9th grade. The questions are of multiple-choice format, where one of the four given options is correct. The training set consists of 2251 questions that we use for selecting our source examples.

**AG-news:** AG-news (Zhang et al., 2015) is a text classification dataset containing news

articles categorized into four classes - *sports*, *business*, *technology*, and *world*. It has a training set size of 120K, from which we have randomly sampled 10K news articles to construct the source dataset for our experiments.

**BoolQ:** BoolQ (Clark et al., 2019) is a reading comprehension dataset consisting of yes/no questions asked from a passage given for each question. The questions are mostly non-factoid, and considerable inference ability is required to answer them based on the passages provided. In our usage, each question is labeled as either *true* or *false*. The training set consists of 9427 labeled question-passage pairs from which we select the source examples.

**Commonsense-QA:** Commonsense-QA (Talmor et al., 2019) is a commonsense question-answering dataset that consists of multiple-choice questions where one of the five options provided is correct. To answer the questions, logical reasoning abilities and in some cases prior knowledge are required. The training set consists of 9740 labeled questions from which source examples are chosen by us.

**Conll2003-POS:** Conll2003-POS (Tjong Kim Sang and De Meulder, 2003) is a collection of the data which is a part of the CoNLL-2003 shared task. In this dataset, each sentence is labeled as a sequence of part-of-speech (POS) tags (each token is assigned a POS tag). We construct our source dataset by sampling 10K sentences from the 14,041 sentences in the training set.

**Conll2003-NER:** Conll2003-NER (Tjong Kim Sang and De Meulder, 2003) is also a part of the CoNLL-2003 shared task. Here, each sentence is labeled as a sequence of named entity tags. The task in this case is to perform named-entity recognition (NER). IOB tagging scheme (Ramshaw and Marcus, 1995) is used for assigning the tags. Four types of entities are assumed to be present in the data – persons (PER), organizations (ORG), locations (LOC) and miscellaneous names (MISC). The training set consists of 14,041 sentences of which 10K are sampled by us as our source data.

**MNLI:** Multi-Genre Natural Language Inference (MNLI) (Williams et al., 2018) corpus is one of the largest available resources for natural language inference. The corpus consists of 393K labeled examples, where each example consists of a pair of sentences and the label is one among *neutral*, *contradiction* or *entailment* based on the relationship between their meanings. Our source dataset is constructed by sampling 10K examples from the 393K examples available.

**QQP:** Quora Question Pairs (QQP) (Sharma et al., 2019) dataset is curated for the task of natural language understanding. This dataset consists of question pairs collected from the popular question-answering website *Quora*, and the task is to determine if the questions are duplicates of each other. The training set consists of 364K question pairs, each labeled as *duplicate* or *not duplicate*. We sample 10K labeled question pairs from the training set for our source data.

**RACE:** RACE (Lai et al., 2017) is a reading comprehension dataset consisting of passages along with questions asked from them. The passages and questions are collected from English exams of school students aged between 12 to 18. The questions are of multiple-choice format where one of the four options is correct. Our source dataset is gathered by picking 10K passage-question pairs from 87.9K such pairs available in the training set.

**SST2:** SST2 (Socher et al., 2013) dataset is a part of the Stanford Sentiment Treebank corpus. It contains movie review snippets collected from the *Rotten Tomatoes* website. Each review is labeled as *positive* or *negative*. The 10K reviews for our source dataset are sampled from the training set consisting of 67.3K labeled reviews.

## A.2 Target datasets

We have used the following datasets as target datasets:

**ARC-Challenge:** ARC-Challenge (Clark et al., 2018) is also a part of the ARC (AI2 Reasoning Challenge) dataset and it consists of hard natural science questions targeted for students of 3rd to 9th grade. Those

questions that were answered incorrectly by both a retrieval-based algorithm and a word co-occurrence method are considered to be ‘hard’ enough for inclusion in this dataset. The questions are of multiple-choice format, where one of the four options is correct. The test set consists of 1172 questions out of which 500 were randomly selected for our target dataset.

**Social-i-QA:** Social-i-QA (Sap et al., 2019) is a commonsense reasoning dataset focusing on social situations. This dataset contains examples consisting of a social situation or action given as context and a multiple-choice question asked based on the context aimed at testing emotional and social intelligence. Each question has three options, out of which one is correct. We select 500 examples for our target data from the 1954 examples available in the validation set.

**SciQ:** SciQ (Auer et al., 2023) (or, SciQA) is a scientific question-answering dataset consisting of questions from different areas of Physics, Chemistry and Biology. The questions are of multiple-choice format, where one of the four options is correct. The test set consists of 1000 questions, of which 500 are sampled to prepare our target dataset.

**MedMCQA:** MedMCQA (Pal et al., 2022) is a multiple-choice question-answering dataset consisting of questions from post-graduate medical entrance exams in India. Each question has four options of which one is correct. Our target dataset is constructed by selecting 500 questions from the 4183 questions in the validation set.

**Financial-Phrasebank:** Financial-Phrasebank (Malo et al., 2014) is a financial sentiment analysis dataset containing sentences mined from a corpus of English news on all listed companies in OMX Helsinki. Each sentence is labeled as one of the three categories – *positive*, *negative*, or *neutral*, based on the influence the news snippet may have on the stock price. The training set consists of 2264 labeled sentences, from which 500 are sampled for our target dataset.

In each case, for preparing our target dataset we have selected 500 examples. The selection,

though random, is done in such a way that our target datasets are balanced, i.e. the number of examples with each of the different possible labels is almost equal.

## B Semantic context creation

One might assume that taking internal embeddings of the models, instead of an external model like Sentence-BERT, for extracting semantically similar examples might be better suited. However, we found that doing so for LLaMA-2 7B is computationally more expensive, with no significant gain in performance, as reported in Table 7.

We also experimented with Sentence-BERT and E5 (Wang et al., 2022b) to test which external model produces superior semantically similar context  $C$  for the LLMs to use. Table 8 shows the results for LLaMA-2 7B and GPT3.5 performance using  $C$  from the E5 model. We proceeded with Sentence-BERT instead of E5 as our external model, as the former has already been used for a similar role in prior works (Liu et al., 2022; Tanwar et al., 2023), and its performance is also comparable to that of E5.

## C Cross-task prompting in instruction-tuned models

The performance gain with cross-task prompting for instruction-tuned models is lesser compared to that for the corresponding base model, due to the improvement in zero-shot performance of instruction-tuned models. For instance, Table 9 shows the performance for LLaMA-2 7B Chat model. We observe that its zero-shot performance is much better compared to the LLaMA-2 7B base model.

## D Force decoding

Unlike greedy decoding, where the most probable token from the entire vocabulary of the model is the output, in force decoding the vocabulary is restricted to a set of tokens and the output is assigned out of these restricted tokens. Formally, for context  $C$  and a target task  $T$ , the output in the case of force decoding is:

$$\hat{y}_t = \operatorname{argmax}_{y \in L_T} p(y|C)$$

where,  $L_T$  is the label space of target task  $T$ . The label space of all five target tasks is shown in Table 11.

SRC		Zero-shot	ARC-Easy	AG-news	BoolQ	Com-QA	C-POS	C-NER	MNLI	QQP	RACE	SST2
TAR												
LLaMA-2 7B												
ARC-Challenge		4.6	<b>44</b>	32.2	36.6	43.2	36.8	35.2	34.6	33.8	37.4	33.2
Financial-Phrasebank		34.1	51.1	58.28	39.9	11.2	0.2	<b>65.6</b>	42.31	40.3	60.4	
MedMCQA		4.2	30.2	25.8	27	<b>32.4</b>	24.6	25.2	24.8	25.6	30.6	25.2
SciQ		8.0	60.2	44.8	46.6	<b>66.8</b>	33.6	28	47.8	25.6	61.8	42
Social-i-QA		41.1	41.7	35.5	37.3	<b>47.3</b>	39.3	37.9	37.9	36.3	<b>47.3</b>	39.7

Table 7: Accuracy for cross-task setup using one source example which is most similar to the target input based on their final hidden layer representation from LLaMA-2 7B. (Abbreviations: Com-QA → Commonsense-QA, C-POS → Conll2003-POS, C-NER → Conll2003-NER).

SRC		Zero-shot	ARC-Easy	AG-news	BoolQ	Com-QA	C-POS	C-NER	MNLI	QQP	RACE	SST2
TAR												
LLaMA-2 7B												
ARC-Challenge		4.6	44.80	34.20	36.00	44.00	32.00	31.80	35.80	35.20	<b>48.80</b>	33.60
Financial-Phrasebank		34.1	43.91	60.08	38.92	38.32	4.39	0.00	<b>67.86</b>	56.09	50.30	46.71
MedMCQA		4.2	30.40	24.40	28.20	<b>32.80</b>	24.20	24.60	23.60	25.60	29.20	23.40
SciQ		8.0	62.40	38.80	46.00	<b>68.40</b>	31.80	32.40	36.60	27.60	65.40	37.20
Social-i-QA		41.1	42.51	37.33	37.92	46.91	37.72	40.52	41.72	39.52	<b>47.90</b>	41.52
GPT3.5												
ARC-Challenge		74.6	77.00	74.80	75.80	76.60	70.20	69.80	72.40	74.40	<b>77.40</b>	73.00
Financial-Phrasebank		57.5	78.44	<b>82.63</b>	75.25	77.45	63.67	70.46	67.66	74.25	76.65	63.67
MedMCQA		49.6	47.00	49.60	49.20	<b>49.80</b>	48.40	46.20	47.20	46.60	49.60	46.80
SciQ		91.2	91.20	89.00	91.00	91.00	85.80	84.60	89.20	89.20	<b>93.40</b>	90.60
Social-i-QA		76.0	77.05	74.65	74.85	77.05	73.85	73.05	74.05	73.25	<b>77.25</b>	71.86

Table 8: Accuracy for cross-task setup using one source example which is most similar to the target input based on their embeddings generated by the E5 model. (Abbreviations: Com-QA → Commonsense-QA, C-POS → Conll2003-POS, C-NER → Conll2003-NER).

SRC		Zero-shot	ARC-Easy	AG-news	BoolQ	Com-QA	C-POS	C-NER	MNLI	QQP	RACE	SST2
TAR												
LLaMA-2 7B Chat												
ARC-Challenge		42.00	51.00	39.20	45.00	<b>51.40</b>	37.00	30.20	37.00	37.20	<b>51.40</b>	39.00
Financial-Phrasebank		79.64	76.05	<b>83.63</b>	77.84	68.66	43.91	12.18	77.25	80.84	78.84	83.43
MedMCQA		30.60	31.00	30.20	31.60	<b>34.60</b>	30.00	26.80	30.20	27.40	31.00	30.20
SciQ		65.60	66.60	59.00	67.00	<b>71.80</b>	58.20	45.00	53.80	60.80	69.20	59.60
Social-i-QA		52.10	48.10	47.70	47.70	49.50	47.90	36.33	42.32	41.92	<b>53.69</b>	45.91

Table 9: Accuracy for cross-task setup, using one source example, with the LLaMA-2 7B Chat model. (Abbreviations: Com-QA → Commonsense-QA, C-POS → Conll2003-POS, C-NER → Conll2003-NER).

Table 10 reports the performance of cross-task prompting for all source-target pairs using force decoding in LLaMA-2 7B and LLaMA-2 13B models. We noted that force decoding improves the performance of zero-shot prediction by a great margin. However, the same is not true for every source-target pair.

## E Random labeling

Recently, Wei et al. (2023) proposed using a random label space for pseudo-labeling examples of (target) task and utilise these to generate the context  $C$ . We experimented with this setup as it is a better alternative to zero-shot prompting, but our results (Table 12) showed that the model output is random with such  $C$ .

## F Prompt details

We show a few examples of cross-task prompts in Table 14 and in-task combined with cross-task prompts in Table 15. Additionally, task definitions of source and target tasks are provided in Table 16 and Table 17 respectively.

## G Error analysis

Table 18 contains detailed examples for erroneous predictions in cross-task prompting setup.

## H Activation analysis on LLaMA-2 7B

Full correlation analysis is presented in Table 19.

TAR \ SRC	SRC										
	Zero-shot	ARC-Easy	AG-news	BoolQ	Com-QA	C-POS	C-NER	MNLI	QQP	RACE	SST2
LLaMA-2 7B											
ARC-Challenge	38.80	45.40	33.60	36.00	45.00	35.20	32.80	35.40	35.20	<b>46.60</b>	35.00
Financial-Phrasebank	36.33	47.11	<b>63.27</b>	60.28	45.91	50.50	53.49	62.87	57.29	44.91	53.09
MedMCQA	26.00	28.00	26.60	28.80	<b>32.60</b>	28.20	26.80	27.00	27.00	29.60	27.40
SciQ	59.60	61.60	37.00	49.80	<b>65.60</b>	37.00	31.60	33.20	27.20	62.00	32.40
Social-i-QA	47.11	43.31	41.32	39.12	<b>48.90</b>	40.92	40.92	43.11	39.72	48.10	40.52
LLaMA-2 13B											
ARC-Challenge	52.80	<b>58.80</b>	54.60	56.00	56.80	51.20	50.20	53.40	54.00	57.20	52.00
Financial-Phrasebank	62.08	73.25	79.64	70.06	66.87	66.27	67.47	75.65	<b>80.24</b>	70.86	77.64
MedMCQA	36.60	38.20	36.00	37.60	<b>38.80</b>	36.00	32.20	33.80	31.20	39.60	33.00
SciQ	82.00	<b>83.60</b>	81.40	82.80	82.20	80.60	79.00	81.20	79.80	82.80	77.80
Social-i-QA	57.68	62.08	58.08	56.49	62.28	57.49	57.49	55.29	56.49	<b>65.07</b>	57.88

Table 10: Accuracy for cross-task setup using one source example by force decoding, where the label from the label space to which the LLM assigns the highest probability is considered as the output. (Abbreviations: Com-QA  $\rightarrow$  Commonsense-QA, C-POS  $\rightarrow$  Conll2003-POS, C-NER  $\rightarrow$  Conll2003-NER).

Target task	Label space
ARC-Challenge	{A, B, C, D}
Financial-Phrasebank	{positive, neutral, negative}
MedMCQA	{A, B, C, D}
SciQ	{A, B, C, D}
Social-i-QA	{A, B, C}

Table 11: Label space of target tasks

Shots \ TRC	TRC				
	ARC-Challenge	Financial-Phrasebank	MedMCQA	SciQ	Social-i-QA
LLaMA-2 7B					
1	25.07	33.33	25.00	25.00	33.27
8	25.00	34.13	24.80	25.40	32.73
LLaMA-2 13B					
1	25.13	31.00	25.00	25.07	33.33
8	24.40	33.33	26.00	25.20	33.53
GPT3.5					
1	24.00	49.50	22.60	26.60	32.58
8	24.00	56.89	24.40	37.00	33.13

Table 12: Accuracy for random label space labelling. We experimented with random labeled context  $C$  of size one and eight.

	With source definitions	W/O source definitions
LLaMA-2 7B		
ARC-Challenge	37.0	33.0
Financial-Phrasebank	38.8	27.0
MedMCQA	27.3	17.06
SciQ	45.7	48.1
Social-i-QA	41.8	42.0
Average	38.1	33.8
LLaMA-2 13B		
ARC-Challenge	53.2	52.42
Financial-Phrasebank	58.9	48.0
MedMCQA	34.2	24.36
SciQ	78.2	75.3
Social-i-QA	57.5	58.1
Average	56.4	51.6

Table 13: Accuracy With and Without (W/O) source definition. We note a drop in accuracy when source definitions are removed.

Target task	Source task	Prompt	Output
Financial-Phrasebank	QQP	<p>Definition: Given two question pairs do text classification based on whether they are duplicates or not. The questions are mined from the popular online discussion forum Quora. As duplicate questions might be present on Quora, the task is to label two identical questions as "duplicate" if they ask the same query else label the pair as "not duplicate".</p> <p>Question 1: My yearly income went from \$0 to \$55,000. By how many times did it increase?</p> <p>Question 2:What can I do with approximately 10.000 Baht per month to increase my income?</p> <p>Label: not duplicate</p> <p>Definition: Given a sentence mined from a financial news article, you are to determine the sentiment polarity of the sentence. The task deals with financial sentiment analysis. Based on the sentiment conveyed by the sentence, label the sentence as "negative", "positive" or "neutral"</p> <p>Sentence: Net income from life insurance doubled to EUR 6.8 mn from EUR 3.2 mn , and net income from non-life insurance rose to EUR 5.2 mn from EUR 1.5 mn in the corresponding period in 2009 .</p> <p>Label:</p>	positive
SciQ	Commonsense-QA	<p>Definition: The following task relates to commonsense reasoning. It consists of a question that can be easily solved using logical abilities and reasoning, a set of five options "A.", "B.", "C.", "D." and "E." are also provided along with the question, one of these options answers the question logically. Use your reasoning ability to select the most appropriate answer from the provided choices "A.", "B.", "C.", "D." and "E." and assign these choices (i.e "A.", "B.", "C.", "D." and "E.") as the label</p> <p>Question:In what substance do clouds float?</p> <p>A. sky B. top of mountain C. air D. ground level E. outer space</p> <p>Answer: C</p> <p>Definition: Given a question from a scientific exam about Physics, Chemistry, and Biology, among others. The question is in multiple choice format with four answer options "A.", "B.", "C." and "D.". Using your knowledge about the scientific fields answer the question and provide the label "A", "B", "C" and "D" as answer</p> <p>Question: What term means the amount of water vapor in the air?</p> <p>A. pressure B. humidity C. temperature D. ambient</p> <p>Answer:</p>	B
Social-i-QA	RACE	<p>Definition: Given a reading comprehension type question-answering from an english exam for school students. You are given a context and multiple choice question containing four options "A.", "B.", "C." and "D.". The question is answerable from the comprehension. Based on the question, the option and the context select the most appropriate answer from the provided choices "A.", "B.", "C." and "D.".</p> <p>Context: Mike is a factory worker. He is often very tired after a day's work. His wife, Jenny, has no job, so she stays at home to cook the meals. Every day he can have his dinner when he gets home from his factory. One day, Mike came home very late because he was very busy in the factory. He was very hungry when he got home. He was not happy when he found his dinner was not ready. He was very angry with his wife. He shouted at her, "I'm going out to eat in a restaurant." "Wait for five minutes," said his wife. "Why? Do you think that dinner will be ready in five minutes?" asked Mike. "Of course not," she answered. "But I can be ready to go with you in five minutes."</p> <p>Question: Mike works in _ .</p> <p>A. a factory B. an office C. a school D. a hospital</p> <p>Answer: A</p> <p>Definition: Given an action as the context and a related question, you are to answer the question based on the context using your social intelligence. The question is of multiple choice form with three options "A", "B" and "C". Select the most appropriate answer from the provided choices "A", "B" and "C".</p> <p>Context: Jesse is patient and hardworking and hungry in the morning.</p> <p>Question: How would you describe Jesse?</p> <p>A. ill prepared B. Exhausted and starved C. thoughtful</p> <p>Answer:</p>	B

Table 14: Examples of prompts for cross-task prompting technique. The number of source examples in each case is one. The outputs shown are generated by LLaMA-2 13B.



Target task	Source task	Prompt	Output
Financial-Phrasebank	SST2	<p>Definition: Given a movie review do text classification, based on the sentiment conveyed by the review label it as "positive" or "negative"</p> <p>Sentence: results is the best performance from either in years</p> <p>Label: positive</p> <p>Definition: Given a sentence mined from a financial news article, you are to determine the sentiment polarity of the sentence. The task deals with financial sentiment analysis. Based on the sentiment conveyed by the sentence, label the sentence as "negative", "positive" or "neutral"</p> <p>Sentence: For the last quarter of 2010 , Componenta 's net sales doubled to EUR131m from EUR76m for the same period a year earlier , while it moved to a zero pre-tax profit from a pre-tax loss of EUR7m .</p> <p>Label: positive</p> <p>Sentence: Both operating profit and net sales for the six-month period increased , respectively from EUR18 .1 m and EUR127 .6 m , as compared to the corresponding period in 2006 .</p> <p>Label:</p>	positive
MedMCQA	RACE	<p>Definition: Given a reading comprehension type question-answering from an english exam for school students. You are given a context and multiple choice question containing four options "A.", "B.", "C." and "D.". The question is answerable from the comprehension. Based on the question, the option and the context select the most appropriate answer from the provided choices "A.", "B.", "C." and "D.".</p> <p>Context: Plants need green leaves to make food. A plant needs sunlight and carbon dioxide from the air for making food and it also needs water and salts from the soil to make food too. There are certain cells in the leaves which change carbon dioxide and water into sugar. To do this the cells needs energy, which they get from the sunlight. Green leaves make food for the whole plant. A red leaf can make food too because under the red color ling of the leaf there are food—making cells. There are no leaves which are completely yellow, for they can't make food. The plant makes sugar for its food. In sunlight green leaves make a lot of sugar. The veins can't carry all this sugar away, so the leaves change the sugar into starch, which is kept and so stored in the leaves. At night, the starch changes back to sugar. It is then carried away from the leaves. Some of the sugar is used as food by the plant while the rest is stored as starch. In some plants, food is stored in the roots, in others it is stored in the stem and in leaves, fruits and seeds.</p> <p>Question: Sugar is made for its food by -</p> <p>A. sunlight B. veins C. stems D. green leaves</p> <p>Answer: D</p> <p>Definition: Given a multiple choice question containing four options "A.", "B.", "C." and "D." from a medical entrance exam. The question is related to a sub-field of medical science like Microbiology, Radiology, Ophthalmology, Surgery, Human anatomy, etc. Based on the question, the option and your knowledge of the medical field select the most appropriate answer from the provided choices "A.", "B.", "C." and "D.".</p> <p>Question: Growth hormone has its effect on growth through?</p> <p>A. Directly B. IG1-1 C. Tyroxine D. Intranuclear receptors</p> <p>Answer: B</p> <p>Question: Which of the following has intracellular receptor -</p> <p>A. Glucagon B. Insulin C. Epinephrine D. Thyroxine</p> <p>Answer:</p>	D
SciQ	ARC-Easy	<p>Definition: Given a question answering task from the 3rd to 9th-grade science exam. The question contains four options "A.", "B.", "C." and "D." Select the most appropriate choice that answers the question</p> <p>Question: From which part of the plant does a bee get food?</p> <p>A. flower B. seed C. stem D. root</p> <p>Answer: A</p> <p>Definition: Given a question from a scientific exam about Physics, Chemistry, and Biology, among others. The question is in multiple choice format with four answer options "A.", "B.", "C." and "D.". Using your knowledge about the scientific fields answer the question and provide the label "A", "B", "C" and "D" as answer</p> <p>Question: What type of organism is commonly used in preparation of foods such as cheese and yogurt?</p> <p>A. Viruses B. Protozoa C. Gymnosperms D. Mesophilic organisms</p> <p>Answer: D</p> <p>Question: A bee will sometimes do a dance to tell other bees in the hive where to find what?</p> <p>A. water B. food C. honey D. enemies</p> <p>Answer:</p>	B

Table 15: Examples of prompts for in-task prompting combined with cross-task prompting technique. The outputs shown are generated by LLaMA-2 13B.

Source task	Task definition
AG-news	Given a sentence do text classification, the sentence is a clipping from a news article that may be either related to sports, business, technology, or world news. You are to recognize the category of the sentence and label them as "sports", "business", "technology" or "world" news
ARC-Easy	Given a question answering task from the 3rd to 9th-grade science exam. The question contains four options "A.", "B.", "C." and "D." Select the most appropriate choice that answers the question
BoolQ	Given a context and a question do binary true and false type text classification. You are given a passage as context and a question related to the passage that can be answered as "True" or "False". Based on the context, question and your reasoning ability answer in a "True" and "False".
Commonsense-QA	The following task relates to commonsense reasoning. It consists of a question that can be easily solved using logical abilities and reasoning, a set of five options "A.", "B.", "C.", "D." and "E." are also provided along with the question, one of these options answers the question logically. Use your reasoning ability to select the most appropriate answer from the provided choices "A.", "B.", "C.", "D." and "E." and assign these choices (i.e "A.", "B.", "C.", "D." and "E.") as the label
Conll2003-NER	Given a sentence do token classification on it seek to locate and classify named entities mentioned in the sentence provided. The pre-defined named entity categories along with there labeles are Person (PER), Location (LOC), Organization (ORG) and Miscellaneous (MIS). If the token is not an entity mark it as None. As the entity is more than two tokens long use the prefix B with the named entity token to represent the beginning and use the prefix I till the entity ends.
Conll2003-POS	Given a sentence do token classification by doing Part-of-speech (POS) tagging, which is a process in natural language processing (NLP) where each word in a text is labeled with its corresponding part of speech. This can include nouns, verbs, adjectives, and other grammatical categories.
MNLI	Given Sentence 1 which is a premise and Sentence 2 which is a hypothesis do natural language inference on the pair. In natural language inference we mark whether the premise and hypothesis are "neutral", "contradiction" or "entailment". The pair are said to be "entailed" if the premise justifies/supports the hypothesis, if the pair contradict each other we label them as "contradiction" and label them "neutral" in all other cases
QQP	Given two question pairs do text classification based on whether they are duplicates or not. The questions are mined from the popular online discussion forum Quora. As duplicate quetion might be present on Quora, the task is to label two identical questions as "duplicate" if they ask the same query else label the pair as "not duplicate".
RACE	Given a reading comprehension type question-answering from an english exam for school students. You are given a context and multiple choice question containing four options "A.", "B.", "C." and "D.". The question is answerable from the comprehension. Based on the question, the option and the context select the most appropriate answer from the provided choices "A.", "B.", "C." and "D.".
SST2	Given a movie review do text classification, based on the sentiment conveyed by the review label it as "positive" or "negative"

Table 16: Task definitions of source tasks

Target task	Task definition
ARC-Challenge	Given a question answering task from the 3rd to 9th-grade science exam. The question contains four options "A.", "B.", "C." and "D." Select the most appropriate choice that answers the question
Financial-Phrasebank	Given a sentence mined from a financial news article, you are to determine the sentiment polarity of the sentence. The task deals with financial sentiment analysis. Based on the sentiment conveyed by the sentence, label the sentence as "negative", "positive" or "neutral"
MedMCQA	Given a multiple choice question containing four options "A.", "B.", "C." and "D." from a medical entrance exam. The question is related to a sub-field of medical science like Microbiology, Radiology, Ophthalmology, Surgery, Human anatomy, etc. Based on the question, the option and your knowledge of the medical field select the most appropriate answer from the provided choices "A.", "B.", "C." and "D.".
SciQ	Given a question from a scientific exam about Physics, Chemistry, and Biology, among others. The question is in multiple choice format with four answer options "A.", "B.", "C." and "D.". Using your knowledge about the scientific fields answer the question and provide the label "A", "B", "C" and "D" as answer
Social-i-QA	Given an action as the context and a related question, you are to answer the question based on the context using your social intelligence. The question is of multiple choice form with three options "A", "B" and "C". Select the most appropriate answer from the provided choices "A", "B" and "C".

Table 17: Task definitions of target tasks

#	Prompt	Output
1.	<p>Definition: Given a sentence do token classification by doing Part-of-speech (POS) tagging, which is a process in natural language processing (NLP) where each word in a text is labeled with its corresponding part of speech. This can include nouns, verbs, adjectives, and other grammatical categories. Sentence: Net profit 6.08 vs 3.98 Label: JJ NN CD NNP CD</p> <p>Definition: Given a sentence mined from a financial news article, you are to determine the sentiment polarity of the sentence. The task deals with financial sentiment analysis. Based on the sentiment conveyed by the sentence, label the sentence as "negative", "positive" or "neutral" Sentence: Consolidated net sales increased 16 % to reach EUR74 .8 m , while operating profit amounted to EUR0 .9 m compared to a loss of EUR0 .7 m in the prior year period . Label:</p>	<p>NNP CD NNP CD</p>
2.	<p>Definition: Given a sentence do token classification on it seek to locate and classify named entities mentioned in the sentence provided. The pre-defined named entity categories along with there labels are Person (PER), Location (LOC), Organization (ORG) and Miscellaneous (MIS). If the token is not an entity mark it as None. As the entity is more than two tokens long use the prefix B with the named entity token to represent the beginning and use the prefix I till the entity ends. Sentence: Total shares to be offered 2.5 million Label: O O O O O O</p> <p>Definition: Given a sentence mined from a financial news article, you are to determine the sentiment polarity of the sentence. The task deals with financial sentiment analysis. Based on the sentiment conveyed by the sentence, label the sentence as "negative", "positive" or "neutral" Sentence: The offer of some 30 million shares aimed to raise more than x20ac 500 million US\$ 640 million , was expected to be completed by Oct. 9 , Outokumpu said . Label:</p>	<p>NNNNNNN</p>
3.	<p>Definition: Given a question answering task from the 3rd to 9th-grade science exam. The question contains four options "A.", "B.", "C." and "D." Select the most appropriate choice that answers the question Question: Which object in our solar system reflects light and is a satellite that orbits around one planet? A. Earth B. Mercury C. the Sun D. the Moon Answer: D</p> <p>Definition: Given a question answering task from the 3rd to 9th-grade science exam. The question contains four options "A.", "B.", "C." and "D." Select the most appropriate choice that answers the question Question: Which of the following sets contains only objects that shine as a result of reflected light? A. moons, planets, and comets B. moons, comets, and stars C. planets, stars, and comets D. planets, stars, and moons Answer:</p>	<p>D</p>
4.	<p>Definition: Given a context and a question do binary true and false type text classification. You are given a passage as context and a question related to the passage that can be answered as "True" or "False". Based on the context, question and your reasoning ability answer in a "True" and "False". Context: A blacklight (or often black light), also referred to as a UV-A light, Wood's lamp, or simply ultraviolet light, is a lamp that emits long-wave (UV-A) ultraviolet light and not much visible light. Question: is a black light and an ultraviolet light the same thing Answer: True</p> <p>Definition: Given a question from a scientific exam about Physics, Chemistry, and Biology, among others. The question is in multiple choice format with four answer options "A.", "B.", "C." and "D.". Using your knowledge about the scientific fields answer the question and provide the label "A", "B", "C" and "D" as answer Question: When exposed to ultraviolet, some substances, such as minerals, glow in characteristic visible wavelengths, a process called this? A. pigment B. plasma C. chemical reaction D. fluorescence Answer:</p>	<p>fluorescence</p>

Table 18: Error analysis of cross-task prompting. Four examples represent the major error characteristics. The outputs shown are generated by LLaMA-2 13B.

Layer #	ARC-Challenge			Financial-Phrasebank			MedMCQA			SciQ			Social-i-QA		
	$\rho_s$	$\rho$	$\tau$	$\rho_s$	$\rho$	$\tau$	$\rho_s$	$\rho$	$\tau$	$\rho_s$	$\rho$	$\tau$	$\rho_s$	$\rho$	$\tau$
Layer-1	0.70	0.80	0.51	0.47	0.50	<b>0.37</b>	0.75	0.81	0.60	0.87	0.87	0.69	0.78	0.67	0.61
Layer-2	0.15	0.30	0.08	-0.11	0.05	-0.08	0.59	0.59	0.51	0.26	0.27	0.24	0.46	0.43	0.39
Layer-3	-0.09	-0.20	-0.08	0.08	0.28	0.05	0.01	0.14	0.03	0.13	0.04	0.10	0.05	-0.09	0.00
Layer-4	0.02	-0.17	-0.03	-0.36	-0.10	-0.34	0.14	0.18	0.09	-0.13	-0.12	-0.13	-0.09	-0.28	-0.05
Layer-5	0.07	-0.11	0.05	0.07	0.14	0.05	0.11	0.25	0.08	0.11	0.02	0.07	0.42	-0.04	0.34
Layer-6	0.26	0.19	0.16	-0.08	0.07	-0.07	0.76	0.67	0.61	0.36	0.29	0.29	0.41	0.18	0.30
Layer-7	0.19	0.06	0.12	-0.04	0.17	-0.07	0.85	0.69	0.69	0.25	0.26	0.14	0.39	0.13	0.30
Layer-8	0.40	0.25	0.29	-0.10	0.13	-0.16	0.93	0.77	0.81	0.46	0.34	0.34	0.48	0.23	0.37
Layer-9	0.28	0.28	0.21	0.01	0.10	-0.02	<b>0.95</b>	0.72	<b>0.84</b>	0.52	0.41	0.43	0.45	0.27	0.37
Layer-10	0.31	0.24	0.23	0.02	0.12	-0.02	0.83	0.66	0.69	0.49	0.40	0.39	0.52	0.28	0.39
Layer-11	0.27	0.20	0.19	-0.09	0.09	-0.04	0.91	0.71	0.76	0.49	0.37	0.43	0.45	0.26	0.36
Layer-12	0.32	0.15	0.23	-0.10	0.13	-0.07	0.84	0.71	0.70	0.34	0.34	0.31	0.45	0.26	0.34
Layer-13	0.27	0.13	0.19	-0.16	0.15	-0.16	0.86	0.73	0.75	0.32	0.33	0.30	0.44	0.21	0.31
Layer-14	0.32	0.13	0.20	-0.11	0.13	-0.09	0.81	0.73	0.63	0.37	0.32	0.33	0.47	0.27	0.39
Layer-15	0.26	0.15	0.16	-0.12	0.16	-0.11	0.85	0.77	0.72	0.41	0.37	0.34	0.47	0.29	0.40
Layer-16	0.26	0.13	0.16	-0.08	0.19	-0.09	0.81	0.76	0.63	0.37	0.39	0.33	0.48	0.31	0.40
Layer-17	0.26	0.19	0.16	-0.04	0.20	-0.07	0.94	0.83	0.81	0.50	0.48	0.42	0.66	0.39	0.52
Layer-18	0.26	0.21	0.16	-0.04	0.22	-0.07	<b>0.95</b>	0.88	0.82	0.70	0.61	0.51	0.65	0.44	0.49
Layer-19	0.26	0.22	0.16	-0.09	0.21	-0.04	0.94	<b>0.90</b>	0.81	0.70	0.69	0.51	0.77	0.54	0.58
Layer-20	0.26	0.27	0.16	-0.07	0.19	-0.02	0.88	0.89	0.73	0.75	0.70	0.56	0.79	0.54	0.58
Layer-21	0.24	0.32	0.14	-0.15	0.18	-0.11	0.88	0.89	0.73	0.75	0.72	0.56	0.83	0.57	0.64
Layer-22	0.34	0.38	0.25	-0.12	0.20	-0.07	0.88	0.89	0.73	0.81	0.77	0.64	0.85	0.63	0.67
Layer-23	0.45	0.42	0.34	-0.01	0.23	-0.02	0.82	0.86	0.64	0.81	0.81	0.64	0.85	0.67	0.67
Layer-24	0.45	0.44	0.34	-0.01	0.24	-0.02	0.82	0.85	0.64	0.85	0.83	0.73	0.85	0.68	0.68
Layer-25	0.45	0.45	0.34	0.02	0.25	0.00	0.81	0.83	0.60	0.87	0.84	0.78	0.85	0.71	0.67
Layer-26	0.45	0.45	0.34	0.04	0.26	0.02	0.81	0.84	0.60	0.87	0.85	0.78	0.85	0.71	0.67
Layer-27	0.45	0.47	0.34	0.04	0.28	0.02	0.81	0.81	0.60	0.87	0.86	0.78	0.86	0.73	0.72
Layer-28	0.45	0.48	0.34	0.04	0.30	0.02	0.81	0.82	0.60	0.87	0.86	0.78	0.86	0.73	0.72
Layer-29	0.45	0.50	0.34	0.19	0.31	0.11	0.81	0.82	0.60	0.87	0.87	0.78	0.86	0.75	0.70
Layer-30	0.45	0.52	0.34	0.19	0.33	0.11	0.81	0.82	0.60	0.87	0.88	0.78	0.85	0.76	0.67
Layer-31	0.72	<b>0.84</b>	<b>0.57</b>	0.26	0.49	0.20	0.83	0.87	0.69	<b>0.92</b>	<b>0.91</b>	<b>0.82</b>	<b>0.89</b>	0.84	<b>0.81</b>
Layer-32	<b>0.74</b>	0.80	<b>0.57</b>	<b>0.49</b>	<b>0.56</b>	0.33	0.78	0.88	0.69	0.90	<b>0.91</b>	0.78	0.83	<b>0.87</b>	0.76

Table 19: The correlation between the ranking of the source tasks for each layer of the LLM with their actual ranking obtained based on performance for each target task. The ranking for each layer is obtained based on the cosine similarity between the mean activations of the source and target task definitions derived from that layer.  $\rho_s$ ,  $\rho$  and  $\tau$  are the Spearman rank-order correlation coefficient, Pearson correlation coefficient and Kendall’s tau measures respectively. The results shown in this table are for LLaMA-2 7B.

SRC	p-value	t-statistic	Significant/Not
LLaMA-2 7B			
ARC-Easy	1.19E-102	22.5127	Significant
AG-news	1.62E-103	22.6187	Significant
BoolQ	1.07E-62	17.1466	Significant
Com-QA	2.20E-78	19.3822	Significant
C-POS	2.30E-14	7.5864	Significant
C-NER	2.30E-09	5.8815	Significant
MNLI	3.52E-84	20.1607	Significant
QQP	2.06E-53	15.7097	Significant
RACE	2.47E-151	28.079	Significant
SST2	2.48E-90	20.9642	Significant
LLaMA-2 13B			
ARC-Easy	2.79E-48	14.8683	Significant
AG-news	1.62E-31	11.7801	Significant
BoolQ	4.19E-29	11.2767	Significant
Com-QA	3.18E-36	12.7125	Significant
C-POS	6.77E-05	-3.8223	NOT Significant
C-NER	7.71E-02	-1.4253	NOT Significant
MNLI	3.12E-27	10.8716	Significant
QQP	6.46E-30	11.4484	Significant
RACE	2.62E-47	14.7043	Significant
SST2	2.94E-21	9.4765	Significant
GPT3.5			
ARC-Easy	2.73E-13	7.2515	Significant
AG-news	9.83E-12	6.7395	Significant
BoolQ	5.04E-11	6.4936	Significant
Com-QA	1.82E-10	6.2939	Significant
C-POS	6.12E-09	-5.7154	NOT Significant
C-NER	2.07E-04	-3.5359	NOT Significant
MNLI	2.39E-03	2.8244	Significant
QQP	4.05E-04	3.3535	Significant
RACE	8.00E-10	6.0566	Significant
SST2	7.04E-04	3.1968	Significant

Table 20: p-value and t-statistic obtained on performing a one-tailed T-test for each source task over all samples of the five target tasks. Significance is decided by taking  $\alpha = 0.05$  (Abbreviations: Com-QA  $\rightarrow$  Commonsense-QA, C-POS  $\rightarrow$  Conll2003-POS, C-NER  $\rightarrow$  Conll2003-NER).