

Rectifying Demonstration Shortcut in In-Context Learning

Joonwon Jang¹ Sanghwan Jang² Wonbin Kweon³ Minjin Jeon¹ Hwanjo Yu^{1,2,*}

Graduate School of AI, POSTECH¹

Department of Computer Science and Engineering, POSTECH²

Institute of Artificial Intelligence, POSTECH³

{kaoara, s.jang, kwb4453, minjinj, hwanjoyu}@postech.ac.kr

Abstract

Large language models (LLMs) are able to solve various tasks with only a few demonstrations utilizing their in-context learning (ICL) abilities. However, LLMs often rely on their pre-trained semantic priors of demonstrations rather than on the input-label relationships to proceed with ICL prediction. In this work, we term this phenomenon as the ‘Demonstration Shortcut’. While previous works have primarily focused on improving ICL prediction results for predefined tasks, we aim to rectify the Demonstration Shortcut, thereby enabling the LLM to effectively learn new input-label relationships from demonstrations. To achieve this, we introduce In-Context Calibration, a demonstration-aware calibration method. We evaluate the effectiveness of the proposed method in two settings: (1) the Original ICL Task using the standard label space and (2) the Task Learning setting, where the label space is replaced with semantically unrelated tokens. In both settings, In-Context Calibration demonstrates substantial improvements, with results generalized across three LLM families (OPT, GPT, and Llama2) under various configurations. ¹

1 Introduction

Large language models (LLMs) have demonstrated their effectiveness on a wide range of tasks through in-context learning (ICL), where models learn to perform a task from demonstrations (Brown et al., 2020). Leveraging their pre-trained knowledge, LLMs can associate various words in the demonstration with specific semantics (e.g., associating ‘extremely painful’ with ‘negative’), thereby performing new tasks using only a small set of input-label examples, without requiring parameter updates (Dong et al., 2022; Wei et al., 2023).

However, LLMs often rely on the semantics from their pre-trained knowledge of given demonstrations, resulting in insufficient task learning for the patterns of the provided input-label pairs. (Reynolds and McDonell, 2021; Min et al., 2022; Wei et al., 2023; Pan et al., 2023). This issue intensifies as the model size decreases (Wei et al., 2023). Kossen et al. (2023) suggest that smaller LLMs show promise in learning new mappings from demonstrations in some tasks, yet they still struggle to override semantic priors acquired during pre-training. Therefore, it is necessary to develop a method that enables LLMs of various sizes to effectively mitigate semantic priors preferences and learn to perform unseen tasks from demonstrations.

Prior works have primarily focused on the instabilities of LLMs in ICL prediction (Holtzman et al., 2021; Fei et al., 2023). To mitigate these instabilities, these studies introduced content-free tokens or utilized the entire test set to calibrate prediction probabilities (Holtzman et al., 2021; Fei et al., 2023; Zhou et al., 2023). However, they lack consideration of the semantic priors of LLMs on the demonstration and do not verify whether their approach enhances LLMs to learn new tasks from the demonstrations.

In this work, we investigate how the LLMs’ pre-trained knowledge on the demonstrations affects ICL. We define the following phenomenon as a *Demonstration Shortcut*: the reliance of LLMs on their pre-trained semantic priors of demonstrations in ICL prediction, rather than learning from the input-label relationships presented in these demonstrations. Due to the Demonstration Shortcut, LLMs’ ICL predictions may be overly dependent on the semantics of the given demonstrations even when the label distribution is uniform and the order is identical (Figure 1).

To tackle this problem, we propose In-Context Calibration, a method designed to rectify the

* Corresponding author
¹<https://github.com/Lainshower/In-Context-Calibration.git>

Demonstration Shortcut in ICL. In-Context Calibration estimates the semantic prior of LLMs on each demonstration sample with the in-context examples. Formally, for each example in the demonstration, we estimate its semantic prior relative to the remaining examples and calculate the expected semantic priors of the demonstrations. At test time, we use this term to rectify LLMs’ dependency on semantic priors and enable the model to learn the intended input-label relationships from the demonstrations.

We evaluate the effectiveness of In-Context Calibration on 27 classification datasets from two perspectives: (1) Original ICL Task and (2) Task Learning settings. In the Original ICL Task, we use the standard label space. In the Task Learning setting, the label space is replaced with semantically unrelated tokens. This requires LLMs to learn the novel input-label relationships to achieve high performance, as these relationships are never seen in pre-training. Our proposed method not only demonstrated enhanced performance across various tasks but also showed improvement in task learning abilities. Specifically, In-Context Calibration outperforms other ICL methods in Natural Language Inference (NLI) tasks, which demand high task learning ability. We also demonstrate that In-Context Calibration enhances ICL performance across various model types and sizes, effectively rectifying the ‘Demonstration Shortcut’ problem.

2 Backgrounds

How LLMs utilize demonstrations in ICL Following ICL’s accomplishments, extensive prior works have sought to understand how LLMs use demonstrations, yet there is still no consensus on the following two contradictory perspectives. One line of research claims that LLMs do not learn new input-label relationships from the demonstrations, with the evidence that ICL performance only marginally drops when labels in the demonstrations are replaced with random labels (Min et al., 2022). Instead, LLMs independently recognize the semantics of input and label of in-context demonstrations using their pre-trained knowledge and perform ICL prediction with its language modeling objective (Reynolds and McDonnell, 2021; Min et al., 2022). On the other hand, while some studies suggest that LLMs can learn novel tasks through demonstrations, there is a notable lack of

concrete experimental proof in real-world LLM applications (Xie et al., 2021; Zhang et al., 2023). Addressing this gap, Wei et al. (2023) provides evidence that larger LLMs can learn the new input-label mappings from demonstrations. Summarizing these perspectives, Pan et al. (2023) show that applying pre-trained knowledge to demonstrations for task recognizing is a broad capability across scales while learning new input-label mappings becomes more feasible as the scale increases. This indicates that as LLMs decrease in size, they rely more on pre-trained knowledge of demonstrations in ICL prediction. Furthermore, when labels in the demonstrations are flipped with different labels (e.g., labeling ‘positive’ as ‘negative’ and vice versa), these models struggle to override semantic priors obtained during pre-training. (Wei et al., 2023; Kossen et al., 2023).

Improving ICL through Calibration Various studies have focused on the instability of ICL prediction in LLMs. (Zhao et al., 2021; Jiang et al., 2023) reveal the instability of ICL prediction arises from the majority label bias and recency label bias, and Fei et al. (2023) identifies domain as a factor contributing to label bias behind this instability. These studies have attempted to estimate the instability of ICL prediction by introducing content-free tokens (Zhao et al., 2021) or using the entire test set to calibrate ICL prediction probabilities (Fei et al., 2023; Zhou et al., 2023). Although their approaches show an improvement in ICL prediction, they do not address the reliance of LLMs on the semantic priors of the demonstration. In other words, the primary objective of these methods is to enhance ICL prediction performance for pre-defined tasks, rather than enabling the model to learn input-label mappings from the demonstrations. Furthermore, they fail to demonstrate whether their calibration method allows LLMs to learn new input-label mappings through demonstrations. Moreover, it is unreasonable to assume that the entire test set is available when learning new input-label relationships. Based on the discussions above, our analysis focuses on the reliance of LLMs on their semantic priors of demonstrations, offering a novel perspective on the calibration method.

3 Demonstration Shortcut


In this section, we first introduce a new typology termed *Demonstration Shortcut*. This concept

Zero-shot

Classify the news articles into the categories of world, sports, business, and technology.

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. ~
 Answer:

LLM



In-Context Learning

Classify the news articles into the categories of world, sports, business, and technology.

Article: Vietnam Hosts Investment Conference - Hoping to boost business ties with Singapore ~

Answer: *world*

Article: Oracle Sales Data Seen Being Released - Oracle Corp. sales documents detailing highly ~

Answer: *technology*

Article: Holiday-Shopping Season Remains Sluggish - U.S. shoppers have kept a tight grip on ~

Answer: *business*


Article: Coulthard signs for Red Bull team The deal, thought to be worth 2 million, could prove ~

Answer: *sports*

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. ~

Answer:

Demo # 1



Classify the news articles into the categories of World, Sports, Business, and Technology.

Article: U.S. Gymnasts Win 3 Medals; Hamm Angry ATHENS, Greece - Three got medals, and ~

Answer: *world*

Article: Madden' Week No. 1 ~ the latest version of Electronic Arts Inc.'s pro football video ~

Answer: *technology*

Article: Man Utd warns of lower TV revenues - Manchester United has warned that lower TV ~

Answer: *business*

Article: Champions League matches in Group A Liverpool: 22-Chris Kirkland; 17-Josemi, ~

Answer: *sports*

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. ~

Answer:

Demo # 2




Figure 1: The overall illustration of the *Demonstration Shortcut*. In a zero-shot setting, an LLM predicts the test article to the world label. With the first demonstration set, the LLM predicts the business label through ICL. However, with the second demonstration set which has the same label order but different semantics in the examples the LLM predicts the sports label. GPT-J is used for these experiments. See Appendix A for a full description of the demonstrations.

refers to the reliance of LLMs on their pre-trained semantic priors of demonstrations for making ICL predictions, rather than learning the input-label relationships presented in the demonstrations. Figure 1 illustrates the underlying concept of the Demonstration Shortcut. In a zero-shot setting, the LLM predicts the test article to the world label (the ground-truth label is technology). Next, we constructed two demonstration sets from the same training dataset, each with a uniform label distribution. While these sets have identical label orders, they differ in the semantics of the examples. The first demonstration set mostly features business-related semantics across all examples, while the second set leans more toward sports-related semantics. After ICL with the first demonstration set, the LLM predicted the article to the business label. However, LLM predicted the article to the sports label with the second demonstration set. Despite both sets having uniform label distributions and identical orders, LLM relies on the semantic prior from each demonstration set, exhibiting

	Demo #1	Demo #2	Demo #3	Demo #4
World	0.20	0.22	0.26	0.17
Sports	0.25	0.30	0.25	0.25
Business	0.34	0.28	0.25	0.27
Technology	0.19	0.18	0.21	0.28

Table 1: Prediction distributions of GPT-J with different demonstration semantics sets. All four demonstration sets have uniform and identical label distribution.

a Demonstration Shortcut and failing to predict the correct answer. This indicates that an over-dependence on semantic prior interrupts the ability of LLMs to learn the new input-label mapping relationships from the demonstrations.

To deeply describe how semantic priors acquired from pre-training affect ICL prediction, consider the first example in Demo #1 of Figure 1, titled ‘*Vietnam Hosts Investment Conference - Hoping to Boost Business Ties with Singapore ~*’ (see Appendix A for the full example). The overall context of the article allows the LLM’s pre-training knowledge to associate its semantics with

the business label. Meanwhile, word-by-word examination (e.g., Vietnam or Singapore) also reveals potential associations of its pre-trained semantics with the world label (Tang et al., 2023). By iterating this process for all examples, regardless of the assigned ground-truth label, the LLM may proceed with ICL prediction based on pre-trained semantic distributions of the demonstrations, leading to the Demonstration Shortcut.

Substantiating our intuition, we conducted additional experiments based on Figure 1. We constructed Demo #3 to be characterized by world semantics across all examples, while Demo #4 features across technology-related semantics. All demonstration sets were designed to have a uniform label distribution and an identical sequence, as depicted in Figure 1. We ensured the test set also followed a uniform label distribution (25 samples for each label) and reported the average label prediction probabilities for these examples in Table 1. Despite having uniform and identical label distributions in the demonstrations, the LLM predictions for all demonstration sets exhibit different behaviors; this variance aligns with the overall semantics of each demonstration.

4 In-Context Calibration

This section introduces a novel calibration method to rectify the Demonstration Shortcut. We first revisit existing calibration methods designed to improve ICL predictions and analyze their limitations, particularly regarding Demonstration Shortcut. We then propose In-Context Calibration, our approach to rectifying the Demonstration Shortcut in ICL.

Revisiting Previous Methods Prior works on calibrating LLMs attempt to adjust ICL predictions by estimating the prompt prior with a content-free token ‘N/A’ (Zhao et al., 2021) or by estimating the task prior by utilizing the entire test distribution (Fei et al., 2023; Zhou et al., 2023). Specifically, Contextual Calibration (CC) estimates the content-free prediction prior as $P_{LM}(y|‘N/A’, [(x_i, y_i)]_{i \in [K]})$, where x_i represents the text input, y_i corresponds to a verbalized label name, and K denotes the total number of examples. Meanwhile, Domain Calibration (DC) estimates task prior with $\frac{1}{M} \sum_{r=1}^M P_{LM}(y|[r/w]_r, [(x_i, y_i)]_{i \in [K]})$, where $[r/w]$ represents random words drawn from the entire test set. However, the introduced terms are

not entirely content-free; their neutrality depends on the task type and the demonstrations (Fei et al., 2023; Zhou et al., 2023). Therefore, if these terms’ pre-trained semantics mismatch with the semantic priors of demonstrations, they are limited in rectifying the Demonstration Shortcut. Moreover, relying on the entire test distribution is impractical in real-world settings.

In-Context Calibration To overcome these limitations, we propose In-Context Calibration, which rectifies the Demonstration Shortcut of the model in the ICL setting.

For each x_i in the demonstration, we estimate the semantic prior of LLMs on each demonstration sample with remaining in-context demonstrations:

$$P_i = P_{LM}(y | x_i, [(x_j, y_j)]_{j \in [K] \setminus \{i\}}), \forall i, \quad (1)$$

where P_i represents the semantic distribution of x_i given the remaining $K - 1$ demonstrations. This allows us to estimate the contextual semantic prior of each demonstration sample while preserving the remaining in-context demonstrations order and conditions present in the original ICL setting.

Additionally, we estimate each demonstration sample’s word-by-word semantic distribution by applying the random shuffling function $R()$ to x_i , which shuffles the order of the words as follows:

$$P_{R(i)} = P_{LM}(y | R(x_i), [(x_j, y_j)]_{j \in [K] \setminus \{i\}}), \forall i. \quad (2)$$

The resulting random order of $R(x_i)$ is not grammatically meaningful, yet it enables context-agnostic estimation of the LLMs’ pre-trained semantics for each demonstration sample. In other words, $R(x_i)$ is stripped of context and retains only the words with their meanings, thereby preventing the LLMs from making predictions based solely on the semantics of individual words (Fei et al., 2023; Tang et al., 2023)²

We iterate this process across all K demonstrations and compute the average to estimate the semantic priors of the demonstrations:

$$\frac{1}{K} \sum_{i=1}^K (\lambda \cdot P_i + (1 - \lambda) \cdot P_{R(i)}), \quad (3)$$

where λ is the hyperparameter that controls the balance between the two terms. Considering every dependency from the K demonstrations and taking

²Please refer to Appendix E for additional details for the random shuffling function.

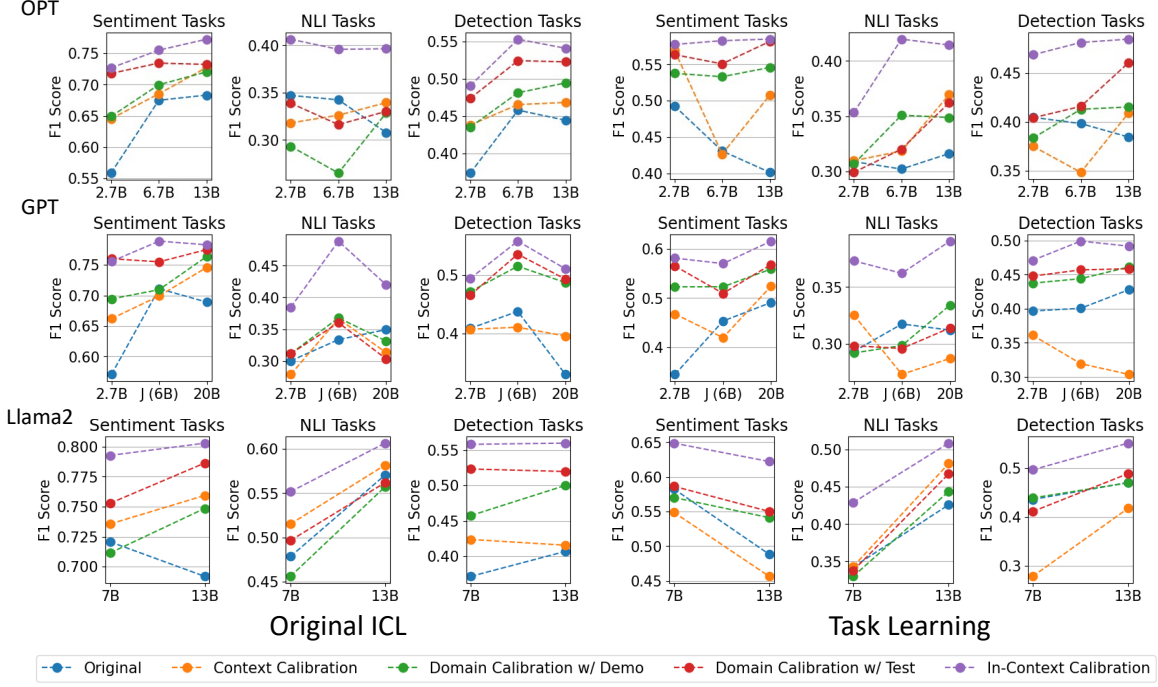


Figure 2: Averaged Macro F1 scores for OPT (Top), GPT (Medium), and Llama2 (Bottom) across Sentiment, NLI, and Detection Tasks. The left the left three columns depict the performance on the Original ICL Task. The right three columns plot the Task Learning scores. In both graphs, the x-axis represents the model size.

the average, we can estimate the expected semantic priors of demonstrations, enabling more effective *demonstration-aware calibration*. The model then makes ICL predictions based on the following estimates:

$$\hat{y}_{pred} = \arg \max_{y \in \mathcal{L}} \frac{P_{LM}(y | x_{test}, [(x_i, y_i)]_{i \in [K]})}{\frac{1}{K} \sum_{i=1}^K (\lambda \cdot P_i + (1 - \lambda) \cdot P_{R(i)}),} \quad (4)$$

where $P_{LM}(y | x_{test}, [(x_i, y_i)]_{i \in [K]})$ is the original ICL prediction.

5 Experimental Setups

We investigate the effectiveness of In-Context Calibration in two aspects: (1) the model’s performance on the Original ICL Task (using standard label space from the dataset) and (2) the Task Learning setting (label space is randomly mapped to semantically unrelated tokens), following the experimental settings of Pan et al. (2023). Since the newly introduced input-label mappings have never been parameterized during pre-training, the model must utilize its task learning abilities to handle the problem. In our main experiment, unless stated otherwise, we conduct task learning experi-

ments using string numbers³, and to avoid any pre-trained bias, each label is randomly assigned to a unique string number for every seed.

5.1 Datasets

We conducted experiments on 27 classification datasets across three types of tasks: Sentiment, NLI, and Detection classification task. Our dataset selection and prompts largely follow the methodologies of prior ICL works (Zhao et al., 2021; Min et al., 2022; Fei et al., 2023), and more details are described in Appendix B.

5.2 Base Models

To validate our method across a diverse set of models, we use three state-of-the-art LLM families: GPT (2.7B, J (6B), 20B) (Brown et al., 2020), OPT (2.7B, 6.7B, 13B) (Zhang et al., 2022), and Llama2 (7B, 13B) (Touvron et al., 2023). For the GPT models, we use the open-sourced versions provided by EleutherAI (Gao et al., 2020; Wang, 2021; Black et al., 2022) as Fei et al. (2023). Consequently, we utilize checkpoints from the Transformers library (Wolf et al., 2020) for all the aforementioned models.

³String numbers demonstrate better task learning ability than other tokens in Pan et al. (2023).

5.3 Implementation Details

Adopting a sampling-based evaluation approach, we sample different sets of demonstrations from the training set for each seed and report the mean and standard deviation of the results. We use $K = 8$ examples and conduct five evaluations using different random seeds, per the methodology described by Fei et al. (2023). Unless stated otherwise, we set λ to 0.5. For the baselines, we selected Contextual Calibration (CC) (Zhao et al., 2021) and Domain Calibration (DC) (Fei et al., 2023) to assess the performance on the Original ICL Task and Task Learning setting. For Domain Calibration, the original method involves constructing a bag of words from the entire test set, which is impractical for real-world inference. To facilitate a fair comparison, we adapted this method. The adapted version samples random words from the demonstration and is labeled as “w/ Demo” (with Demo), while we refer to the original method as “w/ Test” (with Test).

6 Experimental Results

6.1 Main Results

Figure 2 shows the Macro F1-scores of OPT, GPT, and Llama2 on three categories of datasets in the Original ICL Task and Task Learning setting. The results demonstrate that In-Context Calibration effectively rectifies the Demonstration Shortcut, enhancing performance in the Original ICL Task and improving learning capabilities for new input-label tasks. Specifically, for Llama2, In-Context Calibration resulted in an average F1 score improvement of 23% compared to the original inference in the Original ICL Task. Regarding CC using ‘N/A’ token and DC w/ Test sampling random words from the test set, if the newly introduced tokens fail to accurately estimate the neutrality of the demonstration’s semantic distribution, the model remains constrained by the Demonstration Shortcut, limiting performance improvements. This is particularly pronounced in NLI tasks, where previous works (Pan et al., 2023; Kossen et al., 2023) have shown to be challenging for LLMs to learn input-label pairs due to strong reliance on semantic priors in the ICL setting, while In-Context Calibration effectively rectifies the Demonstration Shortcut for all LLMs, resulting in notable performance increases.

This improvement is also evident in the Task Learning setting, where the label space is replaced

with string numbers. For GPT, In-Context Calibration achieved an average F1 Score increase of 27% over the original inference. Particularly in the NLI task, other methodologies struggled to mitigate the dependency of the semantic priors of the models on demonstration, leading to decreased performance compared to original inference in some cases. Across various dataset categories and model types, In-Context Calibration consistently outperformed baseline methods on tasks with novel input-label pairs by effectively reducing the model’s reliance on the demonstration’s semantic prior (see Appendix G for comprehensive results).

6.2 Ablation Study

Method	Original	TL
Original Inference	0.48	0.40
$R(x_i) \rightarrow \text{‘N/A’}$	0.49	0.38
$R(x_i) \rightarrow$ random words	0.57	0.46
In Context Calibration ($\lambda = 0.5$)	0.60	0.49
In Context Calibration ($\lambda = 0$)	0.56	0.46
In Context Calibration ($\lambda = 1$)	0.58	0.50

Table 2: Analysis of In-Context Calibration: Performance comparing $R(x_i)$ replacement with ‘N/A’ (CC) or randomly sampled test set words (DC), and effects of varying λ values, shown through average Macro F1 scores across 27 datasets in Original ICL Task (Original) and Task Learning (TL) settings using GPT-J.

We conducted an ablation study of the proposed method. First, we replaced the $R(x_i)$ term with either a ‘N/A’ token or words randomly sampled from the bag of words constructed from the entire test set. The replacement with the ‘N/A’ token resulted in lower performance compared to the original inference in the Task Learning setting, likely due to the instability of the ‘N/A’ token in estimating prompt neutrality (Fei et al., 2023). Furthermore, words randomly sampled from the test set underperformed in both tasks compared to In-Context Calibration with $\lambda = 0.5$. This underperformance could stem from the semantic distribution mismatch between the sampled words and the demonstrations, limiting the model’s ability to rectify the Demonstration Shortcut. However, the demonstration-aware calibrating of In-Context Calibration leads to performance improvements.

Ablation studies on λ demonstrate that using only $R(x_i)$ for calibration ($\lambda = 0$), which shuffles the token order for each demonstration, leads to limited performance gains in both the Original ICL




Figure 3: Averaged Macro F1 scores for Llama2-Chat across Sentiment, NLI, and Detection Tasks. The left three columns depict the performance of the Original ICL Task. The right three columns plot the Task Learning scores. In both graphs, the x-axis represents the model size.

Task and Task Learning setting, attributing to the loss of contextual information. On the other hand, not utilizing $R(x_i)$ for calibration ($\lambda = 1$) results in better learning of new input-label relationships. However, this approach is less effective in mitigating the model’s dependency on word-wise pre-trained semantics of demonstration when considering labels in the Original ICL Task. We conduct a detailed task-wise analysis of the λ value’s effect in Section 6.5 and Appendix F. Therefore, after comprehensive analysis for λ values, we set λ to 0.5 in the main experiment.⁴

6.3 Analysis with Enhanced Models

Instruction-Tuned Model In previous studies (Min et al., 2022; Wei et al., 2023), instruction tuning has been demonstrated to strengthen the usage of semantic priors in the demonstration. Therefore, we conducted experiments to determine whether In-Context Calibration consistently rectifies the Demonstration Shortcut and improves task learning ability in instruction-tuned LLMs. As depicted in Figure 3, In-Context Calibration increases the model’s F1 score across all Original ICL Tasks and particularly shows substantial improvement over other calibration methods in the Task Learning setting. Especially in the Task Learning setting, other calibration methods often underperformed the original inference. These experiments demonstrate that In-Context Calibration remains effective and enhances the model’s task learning abilities, even after instruction tuning has strengthened the LLM’s reliance on semantic priors.

Over 50B Scale Models We conducted experiments to verify that the proposed method consistently improves performance in larger models. As

demonstrated by Pan et al. (2023), in the Task Learning setting, the performance of some models (e.g., OPT) starts to match their performance in the Original ICL Task for models larger than 50B. These models can utilize the mapping information provided in the demonstrations. Therefore, we conducted experiments on OPT 66B, Llama2 70B, and Llama2-chat 70B, reporting average F1 scores for 27 datasets in both the Original ICL Task and Task Learning settings, as shown in Figure 4. Particularly for Llama2 70B, In-Context Calibration improved F1-score performance by 15% over the original inference in the Original ICL Task and by 9% in the Task Learning setting, while some methods hurt the models performance. These findings suggest that In-Context Calibration not only boosts performance in smaller-scale models but also consistently improves the learning ability of larger-scale models, by effectively rectifying the Demonstration Shortcut through its demonstration-aware calibration.

6.4 Analysis with Other Input-Label Mappings

Overriding Semantic Priors (Permutating Label Space) Wei et al. (2023) and Kossen et al. (2023) reveal that LLMs struggle to override semantic priors preference with input-label mappings in ICL setting. To test whether our proposed method helps the models override semantic priors preference by using its task learning abilities, we evaluate performance by randomly permuting the label space. For instance, in the AGNews dataset, the original label ‘sports’ is permuted to ‘world,’ ‘business’ data to ‘technology,’ ‘technology’ data to ‘sports,’ and ‘world’ data to ‘business’ to construct a demonstration set. The test labels are similarly permuted for evaluation. Due to semantic priors preferences, the models tend to rely more

⁴Please refer to the Appendix F for the comprehensive analysis for λ value.




Figure 4: Averaged Macro F1 scores across 27 classification tasks for over 50B scale LLMs. The left graphs depict performance in the Original ICL Task, while the right graphs plot task learning scores. In both sets of graphs, the x-axis denotes the model type. Full details of the data-type scores are provided in Appendix G.




Figure 5: Averaged Macro F1 scores for the GPT model are presented across 27 classification tasks, each featuring a permuted label space. The x-axis represents the model size. Results for the OPT and Llama2 models are provided in Appendix G.

on their pre-trained knowledge than on learning new relationships from input-label pairs, resulting in lower performance on permuted datasets. We report the average F1 score across 27 datasets. Figure 5 illustrates that the GPT model outperforms other calibration methods using In-Context Calibration. For results related to other models, please refer to Figure 7 in Appendix G. This indicates that In-Context Calibration’s demonstration-aware calibrating is needed to diminish the model’s semantic priors preferences and let it learn new tasks from demonstrations, especially those that contradict pre-trained knowledge.

Task Learning with different label mapping (Symbol Token) Pan et al. (2023) demonstrate that replacing the label space with symbols in the Task Learning setting leads to underperformance, attributing to their unnaturalness in pre-training stages. To verify whether our method enhances learning ability in a more general task learning en-

Figure 6: Averaged Macro F1 scores for the 6-7B scale model families are presented across 27 datasets with each label space replaced by symbol tokens. The x-axis represents the model type. Results for the 13-20B scale models are available in Figure 8.

vironment, we conducted experiments by mapping the label space to symbols. In other words, we randomly replace the label space with one of [$@$, $\#$, $\$$, ...] at every seed⁵. We report the average F1-score across 27 datasets as experimental results. Figure 6 shows that our method outperforms other calibration methods, particularly demonstrating a significant improvement over the original inference. This indicates that our method enhances LLM’s task learning ability in a broader context.

6.5 Analysis of λ across Different Task Categories

We investigated the utility of $R()$ across the different task categories, by calculating the task-wise average F1 score as shown in Table 3. For the Sentiment and Detection task, where the semantics of each word are crucial in associating a specific label in pre-trained knowledge, applying $R()$

⁵Please refer to the Appendix D for the detailed implementation.

Task	Orig.	ICC ($\lambda = 1$)	ICC ($\lambda = 0.5$)
Sentiment	0.69	0.75	0.78
NLI	0.33	0.51	0.49
Detection	0.41	0.52	0.54

Table 3: GPT-J’s averaged F1 scores across different task categories with and without applying $R()$. Orig. denotes original inference, while ICC stands for In-Context Calibration.

for calibration proves more effective. In contrast, for NLI tasks, where the LLMs must discern the logical relationship between two input sentences, In-Context Calibration demonstrates its effectiveness over the baseline in mitigating the reliance of LLMs on semantic priors of the demonstration (as show in Figure 2). However, applying $R()$ in NLI tasks disrupts the order of the sentences, necessitating contextual awareness calibrating for better performance. We reserved for additional results in Appendix F.

7 Conclusion

In this paper, we introduced the term ‘Demonstration Shortcut’, which refers to the reliance of LLMs on their pre-trained semantic priors of demonstrations for making ICL predictions. To rectify this Demonstration Shortcut and enable the model to learn the input-label relationships from the demonstrations, we proposed a novel method, In-Context Calibration, based on the provided demonstrations. This demonstration-aware calibration consistently yields the improved performance, regardless of model sizes or types, across various settings. With the introduction of In-Context Calibration, we anticipate more reliable applications of Large Language Models.

Limitations

In this work, we investigate the reliance of Large Language Models on their pre-trained semantic priors on demonstrations in in-context learning prediction. While In-Context Calibration demonstrates its effectiveness across various tasks and enhances the LLMs’ task learning abilities, our experiments primarily focus on classification tasks. However, the effect of the Demonstration Shortcut might manifest differently in generation tasks. Further analysis and adaptation of our In-Context Calibration method for these tasks are left for future research. Due to budget constraints, experiments with larger models (e.g., GPT4 API) and in multilingual settings were not feasible. Future studies with diverse settings and sufficient resources could provide a more comprehensive understanding.

Due to computational constraints, it was impractical to explore every possible λ value for each model. While there may be variations across different models that remain unexplored, we believe that the comprehensive analysis provided in the appendix will offer practical guidelines for selecting the λ value.

Ethical Considerations

Our work focuses on how Large Language Models utilize demonstrations in in-context learning. To enhance the ability of LLMs to learn input-label relationships from demonstrations, we conducted several additional inferences, requiring only minimal computational resources compared to updating model parameters. Additionally, we used only open-source LLMs and publicly available text classification datasets. Therefore, we do not concern about significant ethical issues arising from our work. On the contrary, we anticipate that future works could utilize our analysis to rectify harmful biases inherent in pre-trained model knowledge through demonstration-based methods.

Acknowledgements

We appreciate Keonwoo Kim, Jaehee Kim, Yukyung Lee, and Hyowon Cho for their invaluable comments. We also thank POSTECH DI LAB members and anonymous reviewers for their comments on the paper. This work was supported by Institute of Information communications Technology Planning Evaluation (IITP) grant funded

by the Korea government (MSIT) (No.2019-0-01906, Artificial Intelligence Graduate School Program (POSTECH)), Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2018-0-00584, (SW starlab) Development of Decision Support System Software based on Next-Generation Machine Learning), National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00217286), Digital Innovation Hub project supervised by the Daegu Digital Innovation Promotion Agency (DIP) grant funded by the Korea government (MSIT and Daegu Metropolitan City) in 2024 (No. DBSD1-07). MSIT: Ministry of Science and ICT, and Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (2022R1A6A1A03052954).

References

- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. [SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Sid Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace He, Connor Leahy, Kyle McDonell, Jason Phang, et al. 2022. Gpt-neox-20b: An open-source autoregressive language model. *arXiv preprint arXiv:2204.06745*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine learning challenges workshop*, pages 177–190. Springer.
- Ona de Gibert, Naiara Perez, Aitor García-Pablos, and Montse Cuadros. 2018. [Hate Speech Dataset from a White Supremacy Forum](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 11–20, Brussels, Belgium. Association for Computational Linguistics.
- Marie-Catherine De Marneffe, Mandy Simons, and Judith Tonhauser. 2019. The commitmentbank: Investigating projection in naturally occurring discourse. In *proceedings of Sinn und Bedeutung*, volume 23, pages 107–124.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. Mitigating label biases for in-context learning. *arXiv preprint arXiv:2305.19148*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. [Surface form competition: Why the highest probability answer isn't always right](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7038–7051, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Eduard Hovy, Laurie Gerber, Ulf Hermjakob, Chin-Yew Lin, and Deepak Ravichandran. 2001. [Toward semantics-based answer pinpointing](#). In *Proceedings of the First International Conference on Human Language Technology Research*.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 168–177.
- Zhongtao Jiang, Yuanzhe Zhang, Cao Liu, Jun Zhao, and Kang Liu. 2023. Generative calibration for in-context learning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2312–2333.
- Jannik Kossen, Tom Rainforth, and Yarin Gal. 2023. In-context learning in large language models learns label relationships but is not conventional learning. *arXiv preprint arXiv:2307.12375*.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.
- 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.
- Marco Marelli, Stefano Menini, Marco Baroni, Luisa Bentivogli, Raffaella Bernardi, and Roberto Zamparelli. 2014. [A SICK cure for the evaluation of compositional distributional semantic models](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages

- 216–223, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. Semeval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41.
- Ioannis Mollas, Zoe Chrysopoulou, Stamatis Karlos, and Grigorios Tsoumakas. 2020. Ethos: an online hate speech detection dataset. *arXiv preprint arXiv:2006.08328*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. Adversarial nli: A new benchmark for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.
- Jane Pan, Tianyu Gao, Howard Chen, and Danqi Chen. 2023. [What in-context learning “learns” in-context: Disentangling task recognition and task learning](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8298–8319, Toronto, Canada. Association for Computational Linguistics.
- Bo Pang and Lillian Lee. 2004. A sentimental education: sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, pages 271–es.
- L. PaNgB. 2005. Exploiting class relationships for sentiment categorization with respect to ratings. *IN: Proceedings of ACL r05*.
- Laria Reynolds and Kyle McDonell. 2021. Prompt programming for large language models: Beyond the few-shot paradigm. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7.
- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. [Social bias frames: Reasoning about social and power implications of language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.
- Emily Sheng and David Uthus. 2020. [Investigating societal biases in a poetry composition system](#). In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pages 93–106, Barcelona, Spain (Online). Association for Computational Linguistics.
- 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.
- Ruixiang Tang, Dehan Kong, Lo li Huang, and Hui Xue. 2023. [Large language models can be lazy learners: Analyze shortcuts in in-context learning](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. Semeval-2018 task 3: Irony detection in english tweets. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 39–50.
- Ben Wang. 2021. Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. 2023. Larger language models do in-context learning differently, 2023. URL <https://arxiv.org/abs/2303.03846>.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Huggingface’s transformers: State-of-the-art natural language processing](#).
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of in-context learning as implicit bayesian inference. In *International Conference on Learning Representations*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. Semeval-2019 task 6: Identifying and categorizing offensive language in social media (offenseval). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 75–86.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

- Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *NIPS*.
- Yufeng Zhang, Fengzhuo Zhang, Zhuoran Yang, and Zhaoran Wang. 2023. What and how does in-context learning learn? bayesian model averaging, parameterization, and generalization. *arXiv preprint arXiv:2305.19420*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 12697–12706. PMLR.
- Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2023. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. *arXiv preprint arXiv:2309.17249*.

A Full Description of the Articles

A full description of the articles in Figure 1 is provided in Figure 10. We used the AGNews dataset (Zhang et al., 2015) for this experiment. In a zero-shot setting, the LLM (GPT-J) predicts the test article the world label, although the ground-truth label is technology. We then sampled two demonstration sets from the same training set, each having a uniform label distribution and the same label order. The first demonstration set predominantly features business-related semantics across all examples, while the second set leans more toward sports-related semantics. Despite both demonstration sets having the same uniform label distribution and identical order, the LLM fails to learn the input-label relationships, instead relying on semantic priors of demonstration to make ICL predictions.

B Datasets

We use 27 text classification datasets for our experiments, most of which are widely used in existing ICL works (Min et al., 2022; Fei et al., 2023; Pan et al., 2023). **Sentiment** task datasets include AGNews (Zhang et al., 2015), CR (Hu and Liu, 2004), financial_phrasebank (Malo et al., 2014), poem_sentiment (Sheng and Uthus, 2020), MR (PaNgB, 2005), sst2 (Socher et al., 2013), Subj (Pang and Lee, 2004), and TREC (Hovy et al., 2001); **Natural Language Inference** task datasets include ANLI (Nie et al., 2020), WNLI (Levesque et al., 2012), RTE (Dagan et al., 2005), CB (De Marneffe et al., 2019), and SICK (Marelli et al., 2014) For the **Detection** task we use social_bias_frames (Sap et al., 2020), tweet_eval_stance_athesim, tweet_eval_stance_feminist (Mohammad et al., 2016), tweet_eval_hate (Basile et al., 2019), tweet_eval_irony (Van Hee et al., 2018), tweet_eval_offensive (Zampieri et al., 2019), hate_speech18 (de Gibert et al., 2018), ethos_binary, ethos_disability, ethos_gender, ethos_national_origin, ethos_race, ethos_religion, and ethos_violence (Mollas et al., 2020).

We constructed demonstrations by sampling from the training set and using the validation set for evaluation. In cases where a validation set does not exist, we utilized the test set. For evaluation, we sampled either a maximum of 500 examples or the entire dataset size, whichever is larger.

C Prompt Templates

Our natural language prompts are largely based on Zhao et al. (2021) and Fei et al. (2023). We have adjusted the templates as needed to better align with each dataset’s specific intent. The complete list of prompts is provided in Table 7.

D Implementation Details

We set $M = 20$ for the Domain Calibration implementation, following the original settings used by Fei et al. (2023). We mainly followed the Pan et al. (2023)’s setting for symbol token selection and incorporated additional symbol tokens to enhance generalization. Consequently, in each seed, every label was randomly mapped to one of the following symbols: [“@”, “#”, “\$”, “%”, “*”, “^”, “##”, “\$\$”, “%%”, “**”].

E Random Shuffling Function

In our main experiment, we applied the random shuffling function only once to each x_i to calculate $P_{R(i)}$. To validate whether a single shuffling process introduces randomness, we conducted additional experiments, varying the number of shuffles to 1, 5, and 10. We present supplementary ablation studies using GPT-J with default settings in our main experiments.

Method	Original ICL	TL
Original Inference	0.48	0.40
Context Calibration	0.47	0.34
Domain Calibration (w/ Test)	0.55	0.43
ICC (1)	0.60	0.49
ICC (5)	0.60	0.48
ICC (10)	0.60	0.49

Table 4: ICC stands for In-Context Calibration, and the number in (N) means shuffling number. We calculated the expectation of the randomly shuffled term when the shuffling number exceeded 1.

The results illustrate that a single shuffle does not incur significant randomness in the estimations. These observations are similar to those in Fei et al. (2023), which used different random seeds to demonstrate the rapid stabilization and convergence of the sampling process.

F Comprehensive analysis for λ Values

We present the complete results for various λ values of GPT-J in Table 5. GPT-J performs best with

Method	Original	TL
Original Inference	0.48	0.40
Context Calibration	0.47	0.34
Domain Calibration (w/ Test)	0.55	0.43
ICC ($\lambda = 0$)	0.56	0.46
ICC ($\lambda = 0.25$)	0.56	0.45
ICC ($\lambda = 0.5$)	0.60	0.49
ICC ($\lambda = 0.75$)	0.57	0.47
ICC ($\lambda = 1$)	0.58	0.50

Table 5: Average Macro F1 scores across 27 datasets for both Original ICL Task (Original) and Task Learning (TL) with different λ values using GPT-J. ICC stands for In-Context Calibration.

an Original ICL Task at 0.5 and the higher λ values lead to higher task learning ability.

	Original		TL	
	OPT-6.7B	Llama2-7B	OPT-6.7B	Llama2-7B
Original Inference	0.48	0.48	0.38	0.45
Context Calibration	0.48	0.48	0.35	0.43
Domain Calibration (w/ Test)	0.52	0.55	0.43	0.49
ICC ($\lambda = 0$)	0.55	0.56	0.46	0.51
ICC ($\lambda = 0.25$)	0.55	0.58	0.47	0.51
ICC ($\lambda = 0.5$)	0.57	0.61	0.49	0.52
ICC ($\lambda = 0.75$)	0.56	0.60	0.47	0.51
ICC ($\lambda = 1.0$)	0.57	0.59	0.49	0.53

Table 6: Average Macro F1 scores across 27 datasets for both Original ICL Task (Original) and Task Learning (TL) with different λ values for OPT-6.7B and Llama2-7B. ICC stands for In Context Calibration.

To further analyze the comprehensive impact of λ , we conducted supplementary studies using OPT-6.7B and Llama2-7B, which have similar sizes to GPT-J, exploring broader λ values (Table 6). As with GPT-Js λ value, OPT-6.7B and Llama2-7B also perform best with an Original ICL Task at 0.5 (OPT-6.7B also performs best at the value of 1.0). In line with our primary findings (Table 3), a higher λ value enhances task learning capabilities, leading to less grammatical information corruption in the original input sentences. We also analyze the task-wise value as in Table 8. Consistent across all models, a higher λ improves performance in the NLI task, corroborating our initial findings. Similar to the results seen in Fei et al. (2023), the semantics of words play a critical role in Sentiment and Detection tasks (where a lower λ still shows comparative performance, despite the lost of grammatical information), as the language model’s dependence on the label’s semantics is significant in these tasks. Therefore, (1) we recommend starting with a λ value of 0.5 in In-Context

Calibration and adjusting it based on the task-wise experimental results obtained from the validation set in the Original ICL Task and (2) search for a value with 0.5 or higher in the Task Learning setting. However, due to computational constraints, performing a grid search across all models for every λ value was impractical, so we opted for λ of 0.5 in the main experiment for all models for efficiency.

G Detailed Results

We demonstrate the results of the OPT, GPT, Llama2, and Llama2-Chat models in both the Original ICL Task and the Task Learning setting across Table 9 to 16. For models above 50B scale, please refer to Table 17 and 18. We average the performance across a total of 5 seeds, recording mean values and standard deviations. ‘Orig.’ denotes original inference, ‘CC’ refers to Context Calibration, ‘DC’ stands for Domain Calibration, and ‘ICC’ represents In-Context Calibration. Furthermore, our results indicate that In-Context Calibration outperforms the baselines in most tasks.




Figure 7: Averaged Macro F1 scores for the OPT (Top Graph) and Llama2 (Bottom graph) model are presented across 27 classification tasks, each featuring a permuted label space. The x-axis represents the model size.

Figure 7 presents experimental results for OPT and Llama2 under the same setting as figure 5. We demonstrate that our method is effective across various model types and sizes. Specifically, OPT

shows a 19% performance increase compared to the original inference, while Llama2 exhibits a 15% improvement.

In Figure 8, we present the performance of LLMs when labels are mapped to symbols. In-Context Calibration significantly improves performance across different model types and sizes, highlighting the consistent enhancement of task learning ability.

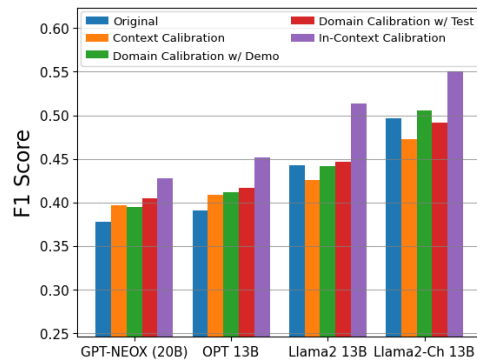


Figure 8: Averaged Macro F1 scores for the 13-20B scale model families are presented across 27 datasets with each label space replaced by symbol tokens. The x-axis represents the model type.

H Adding More In-Context Examples (K=8/12/16)

We study the effect of adding more in-context examples by evaluating the GPT-J. As shown in Figure 9, In-Context Calibration achieves approximately a 6% higher performance than other calibration methods in the Original ICL Task. Similar to Pan et al. (2023), Task Learning performance further improves as the number of demonstrations increases. In-Context Calibration consistently outperforms the baselines, whereas Context Calibration lags behind the original inference. In conclusion, In-Context Calibration consistently demonstrates superior performance in both the Original ICL Task and Task Learning settings, irrespective of the number of demonstrations.




Figure 9: The top graph depicts the average Macro F1 score for the Original ICL Task across 27 datasets for GPT-J. The bottom graph plots the average Macro F1 score for the Task Learning setting. In both graphs, the x-axis represents the number of demonstrations.


Zero-shot

Classify the news articles into the categories of world, sports, business, and technology.

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. It's profane and politically incorrect. It's packed wall to wall with tough thugs doing terrible things.

Answer:

LLM



In-Context Learning

Classify the news articles into the categories of world, sports, business, and technology.

Article: Vietnam Hosts Investment Conference - Hoping to boost business ties with Singapore, Vietnam's prime minister opened a forum here Tuesday with promises of economic reform and continued growth.

Answer: world

Article: Oracle Sales Data Seen Being Released - Oracle Corp. sales documents detailing highly confidential information, such as which companies receive discounts on Oracle's business software products and the size of the discounts, are likely to be made public, a federal judge said on Friday.

Answer: technology

Article: Holiday-Shopping Season Remains Sluggish - U.S. shoppers have kept a tight grip on their wallets this holiday season with indices on Tuesday showing sluggish sales in the second week of the season.

Answer: business

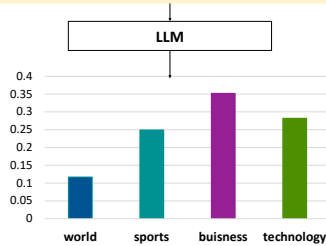
Article: Coulthard signs for Red Bull team The deal, thought to be worth 2 million, could prove more lucrative if Coulthard hits performance targets. That, however, is a big quote;if quote;, given the return of only seven points managed by Mark Webber last term.

Answer: sports

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. It's profane and politically incorrect. It's packed wall to wall with tough thugs doing terrible things.

Answer:

Demo # 1



Classify the news articles into the categories of world, sports, business, and technology.

Article: U.S. Gymnasts Win 3 Medals; Hamm Angry ATHENS, Greece - Three got medals, and Paul Hamm got mad. The United States upped its gymnastics medal haul to seven Sunday night, the most since the Americans won 16 at the boycotted Los Angeles Games in 1984...

Answer: world

Article: 'Madden' Week No. 1 Game Sales Top 1.3 Million LOS ANGELES - "Madden NFL 2005," the latest version of Electronic Arts Inc.'s pro football video game franchise, sold more than 1.3 million copies in its first week, giving a strong kick-off to one of EA's top sellers, the company said on Thursday.

Answer: technology

Article: Man Utd warns of lower TV revenues - Manchester United has warned that lower TV revenues would hit profits next year, but said it would meet its targets for this year.

Answer: business

Article: Champions League matches in Group A Liverpool: 22-Chris Kirkland; 17-Josemi, 23-Jamie Carragher, 4-Sami Hyypia, 6-John Arne Riise; 18-Antonio Nunez, 8-Steven Gerrard, 14-Xabi Alonso, 10-Luis Garcia; 7-Harry Kewell, 5-Milan Baros.

Answer: sports

Article: Game Daze; Grand Theft Auto: San Andreas; Midway Arcade. It's violent. It's profane and politically incorrect. It's packed wall to wall with tough thugs doing terrible things.

Answer:

Demo # 2




Figure 10: Full description of the *Demonstration Shortcut* illustrated in Figure 1. All articles and labels for the demonstrations and the test article were selected from the AGNews (Zhang et al., 2015) dataset. GPT-J was used in this experiment.

Dataset	Natural Language Template	Label Space
Sentiment Task		
sst2, MR, CR	Review: [INPUT] Sentiment: [LABEL]	positive, negative
financial_phrasebank	Sentence: [INPUT] Sentiment: [LABEL]	positive, negative, neutral
poem_sentiment	Verse text: [INPUT] Sentiment: [LABEL]	positive, negative, neutral
Subj	Input: [INPUT] Label: [LABEL]	objective subjective/personal
AG News	Classify the news articles into the categories of Label Space . Article: [INPUT] Answer: [LABEL]	world, technology sports, business
TREC	Classify the questions based on whether their answer type is a Label Space . Question: [INPUT] Answer: [LABEL]	number, location, description, entity, abbrev, person
NLI Task		
ANLI, SICK	[PREMISE] question: [HYPOTHESIS] true, neutral, or false? answer: [LABEL]	true, neutral, false
RTE, WNLI	[PREMISE] question: [HYPOTHESIS] True or False? answer: [LABEL]	True, False
CB	[PREMISE] question: [HYPOTHESIS] true, false, or neither? answer: [LABEL]	true, false, neither
Detection Task		
social_bias_frames	Post: [INPUT] Label: [LABEL]	neutral offensive/hate
tweet_eval_hate	Tweet: [INPUT] Label: [LABEL]	neutral, hate
tweet_eval_irony	Tweet: [INPUT] Label: [LABEL]	neutral ironic/contradict
tweet_eval_offensive	Tweet: [INPUT] Label: [LABEL]	neutral offensive/hate
tweet_eval_stance_athesim, tweet_eval_stance_feminist	Tweet: [INPUT] Label: [LABEL]	none, against, favor
ethos_binary, ethos_disability, ethos-religion, ethos_national_origin, ethos_race, ethos_religion, ethos_violence, hate_speech18	Text: [INPUT] Label: [LABEL]	neutral, hate

Table 7: Prompt templates used in our experiments are primarily adapted from the work of [Zhao et al. \(2021\)](#) and [Fei et al. \(2023\)](#). In the Task Learning setting, each label is randomly mapped to a string number corresponding to the size of the Label Space (i.e., from 0 to len(Label Space) - 1) for each seed.

	Sentiment		NLI		Detection	
	OPT-6.7B	Llama2-7B	OPT-6.7B	Llama2-7B	OPT-6.7B	Llama2-7B
Original Inference	0.674	0.720	0.342	0.478	0.457	0.371
Contextual Calibration	0.685	0.735	0.326	0.515	0.465	0.423
Domain Calibration w/ Test	0.734	0.752	0.316	0.496	0.524	0.523
ICC ($\lambda = 0$)	0.764	0.774	0.338	0.454	0.543	0.523
ICC ($\lambda = 0.25$)	0.764	0.796	0.337	0.511	0.522	0.555
ICC ($\lambda = 0.5$)	0.749	0.792	0.398	0.551	0.552	0.558
ICC ($\lambda = 0.75$)	0.751	0.782	0.428	0.557	0.532	0.564
ICC ($\lambda = 1.0$)	0.749	0.775	0.445	0.547	0.530	0.532

Table 8: OPT-6.7B and Llama2-7B’s averaged F1 scores across different task categories using various λ value. ICC stands for In Context Calibration.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task						
AGNews	2.7B	0.298 (0.072)	0.669 (0.117)	0.676 (0.676)	0.803 (0.018)	0.811 (0.024)
	6.7B	0.582 (0.104)	0.716 (0.108)	0.677 (0.677)	0.765 (0.061)	0.811 (0.023)
	13B	0.722 (0.070)	0.717 (0.104)	0.735 (0.735)	0.796 (0.045)	0.828 (0.051)
CR	2.7B	0.888 (0.031)	0.892 (0.021)	0.911 (0.911)	0.915 (0.008)	0.908 (0.018)
	6.7B	0.860 (0.042)	0.883 (0.024)	0.911 (0.911)	0.921 (0.007)	0.921 (0.005)
	13B	0.860 (0.082)	0.912 (0.013)	0.921 (0.921)	0.920 (0.013)	0.926 (0.006)
financial_phrasebank	2.7B	0.623 (0.191)	0.459 (0.062)	0.607 (0.607)	0.719 (0.115)	0.701 (0.088)
	6.7B	0.693 (0.075)	0.526 (0.106)	0.522 (0.522)	0.441 (0.041)	0.723 (0.069)
	13B	0.471 (0.116)	0.481 (0.085)	0.637 (0.637)	0.676 (0.169)	0.690 (0.098)
poem_sentiment	2.7B	0.410 (0.077)	0.367 (0.112)	0.414 (0.414)	0.542 (0.090)	0.545 (0.083)
	6.7B	0.516 (0.136)	0.423 (0.028)	0.500 (0.500)	0.511 (0.069)	0.618 (0.004)
	13B	0.231 (0.067)	0.485 (0.097)	0.414 (0.414)	0.472 (0.044)	0.492 (0.061)
MR	2.7B	0.875 (0.029)	0.878 (0.034)	0.804 (0.804)	0.900 (0.018)	0.900 (0.022)
	6.7B	0.895 (0.017)	0.908 (0.013)	0.913 (0.913)	0.917 (0.002)	0.918 (0.004)
	13B	0.893 (0.039)	0.921 (0.011)	0.926 (0.926)	0.924 (0.008)	0.925 (0.008)
sst2	2.7B	0.846 (0.130)	0.919 (0.048)	0.910 (0.910)	0.917 (0.034)	0.923 (0.035)
	6.7B	0.934 (0.030)	0.952 (0.007)	0.947 (0.947)	0.953 (0.010)	0.954 (0.009)
	13B	0.900 (0.084)	0.941 (0.011)	0.943 (0.943)	0.945 (0.017)	0.945 (0.022)
Subj	2.7B	0.360 (0.034)	0.546 (0.150)	0.489 (0.489)	0.566 (0.050)	0.614 (0.067)
	6.7B	0.581 (0.182)	0.628 (0.161)	0.642 (0.642)	0.714 (0.061)	0.702 (0.130)
	13B	0.794 (0.079)	0.772 (0.116)	0.760 (0.760)	0.726 (0.191)	0.852 (0.046)
TREC	2.7B	0.232 (0.027)	0.246 (0.035)	0.339 (0.339)	0.382 (0.033)	0.385 (0.039)
	6.7B	0.356 (0.049)	0.287 (0.044)	0.308 (0.308)	0.360 (0.026)	0.361 (0.020)
	13B	0.382 (0.112)	0.342 (0.041)	0.342 (0.342)	0.342 (0.041)	0.439 (0.051)
NLI Task						
ANLI	2.7B	0.224 (0.059)	0.269 (0.043)	0.229 (0.229)	0.222 (0.047)	0.294 (0.024)
	6.7B	0.215 (0.055)	0.263 (0.053)	0.212 (0.212)	0.258 (0.035)	0.316 (0.020)
	13B	0.216 (0.044)	0.268 (0.029)	0.243 (0.243)	0.281 (0.037)	0.307 (0.016)
WNLI	2.7B	0.317 (0.027)	0.302 (0.003)	0.330 (0.330)	0.350 (0.039)	0.410 (0.012)
	6.7B	0.333 (0.056)	0.337 (0.033)	0.380 (0.380)	0.330 (0.046)	0.461 (0.053)
	13B	0.360 (0.012)	0.374 (0.021)	0.398 (0.398)	0.378 (0.033)	0.414 (0.049)
RTE	2.7B	0.489 (0.044)	0.368 (0.034)	0.372 (0.372)	0.515 (0.020)	0.529 (0.035)
	6.7B	0.530 (0.035)	0.465 (0.052)	0.355 (0.355)	0.514 (0.057)	0.540 (0.024)
	13B	0.376 (0.064)	0.499 (0.076)	0.455 (0.455)	0.434 (0.095)	0.547 (0.037)
CB	2.7B	0.358 (0.076)	0.333 (0.133)	0.242 (0.242)	0.269 (0.014)	0.392 (0.065)
	6.7B	0.291 (0.047)	0.240 (0.029)	0.113 (0.113)	0.163 (0.057)	0.267 (0.064)
	13B	0.280 (0.076)	0.217 (0.076)	0.218 (0.218)	0.226 (0.008)	0.317 (0.067)
SICK	2.7B	0.285 (0.063)	0.270 (0.024)	0.317 (0.317)	0.299 (0.067)	0.328 (0.072)
	6.7B	0.267 (0.088)	0.361 (0.070)	0.393 (0.393)	0.359 (0.065)	0.407 (0.073)
	13B	0.225 (0.066)	0.225 (0.066)	0.225 (0.225)	0.352 (0.171)	0.443 (0.107)
Detection Task						
social_bias_frames	2.7B	0.466 (0.075)	0.644 (0.078)	0.615 (0.615)	0.665 (0.025)	0.667 (0.036)
	6.7B	0.610 (0.169)	0.651 (0.077)	0.665 (0.665)	0.674 (0.034)	0.685 (0.062)
	13B	0.556 (0.099)	0.687 (0.066)	0.600 (0.600)	0.639 (0.089)	0.711 (0.016)
tweet_eval_stance_athesim	2.7B	0.089 (0.007)	0.086 (0.002)	0.107 (0.107)	0.109 (0.014)	0.138 (0.015)
	6.7B	0.279 (0.094)	0.150 (0.063)	0.151 (0.151)	0.162 (0.044)	0.291 (0.048)
	13B	0.178 (0.044)	0.147 (0.063)	0.181 (0.181)	0.191 (0.055)	0.227 (0.044)
tweet_eval_stance_feminist	2.7B	0.175 (0.115)	0.185 (0.096)	0.273 (0.273)	0.298 (0.063)	0.347 (0.043)
	6.7B	0.398 (0.051)	0.306 (0.065)	0.267 (0.267)	0.312 (0.075)	0.395 (0.080)
	13B	0.135 (0.023)	0.247 (0.145)	0.325 (0.325)	0.337 (0.102)	0.390 (0.048)
tweet_eval_hate	2.7B	0.366 (0.062)	0.513 (0.084)	0.548 (0.548)	0.545 (0.029)	0.530 (0.052)
	6.7B	0.344 (0.039)	0.379 (0.089)	0.550 (0.550)	0.599 (0.034)	0.559 (0.072)
	13B	0.548 (0.114)	0.599 (0.047)	0.630 (0.630)	0.633 (0.036)	0.623 (0.044)
tweet_eval_irony	2.7B	0.387 (0.059)	0.517 (0.155)	0.601 (0.601)	0.653 (0.070)	0.647 (0.079)
	6.7B	0.451 (0.113)	0.581 (0.101)	0.597 (0.597)	0.625 (0.083)	0.650 (0.085)
	13B	0.434 (0.128)	0.394 (0.084)	0.622 (0.622)	0.602 (0.049)	0.602 (0.058)
tweet_eval_offensive	2.7B	0.538 (0.059)	0.603 (0.045)	0.588 (0.588)	0.613 (0.031)	0.621 (0.046)
	6.7B	0.472 (0.072)	0.602 (0.021)	0.626 (0.626)	0.627 (0.019)	0.632 (0.020)
	13B	0.585 (0.063)	0.605 (0.054)	0.604 (0.604)	0.655 (0.010)	0.643 (0.022)
ethos_binary	2.7B	0.490 (0.151)	0.581 (0.051)	0.625 (0.625)	0.681 (0.029)	0.674 (0.025)
	6.7B	0.619 (0.135)	0.740 (0.049)	0.656 (0.656)	0.754 (0.036)	0.762 (0.026)
	13B	0.674 (0.052)	0.658 (0.035)	0.705 (0.705)	0.722 (0.032)	0.719 (0.031)
ethos_disability	2.7B	0.402 (0.134)	0.442 (0.044)	0.232 (0.232)	0.334 (0.071)	0.421 (0.048)
	6.7B	0.418 (0.116)	0.367 (0.125)	0.261 (0.261)	0.421 (0.064)	0.460 (0.053)
	13B	0.446 (0.097)	0.383 (0.115)	0.327 (0.327)	0.408 (0.050)	0.457 (0.103)
ethos_gender	2.7B	0.368 (0.148)	0.412 (0.108)	0.350 (0.350)	0.380 (0.042)	0.406 (0.057)
	6.7B	0.507 (0.071)	0.329 (0.092)	0.400 (0.400)	0.440 (0.070)	0.483 (0.046)
	13B	0.397 (0.058)	0.376 (0.088)	0.381 (0.381)	0.449 (0.022)	0.453 (0.029)
ethos_national_origin	2.7B	0.293 (0.154)	0.326 (0.098)	0.373 (0.373)	0.375 (0.046)	0.410 (0.074)
	6.7B	0.376 (0.088)	0.350 (0.123)	0.411 (0.411)	0.412 (0.056)	0.429 (0.096)
	13B	0.427 (0.067)	0.433 (0.041)	0.306 (0.306)	0.397 (0.068)	0.434 (0.061)
ethos_race	2.7B	0.394 (0.114)	0.475 (0.066)	0.428 (0.428)	0.406 (0.058)	0.579 (0.066)
	6.7B	0.441 (0.042)	0.361 (0.079)	0.503 (0.503)	0.516 (0.077)	0.569 (0.043)
	13B	0.331 (0.141)	0.374 (0.117)	0.479 (0.479)	0.482 (0.039)	0.490 (0.081)
ethos_religion	2.7B	0.372 (0.173)	0.411 (0.150)	0.390 (0.390)	0.419 (0.060)	0.420 (0.037)
	6.7B	0.326 (0.131)	0.243 (0.104)	0.350 (0.350)	0.385 (0.053)	0.394 (0.090)
	13B	0.361 (0.108)	0.361 (0.108)	0.355 (0.355)	0.395 (0.108)	0.399 (0.134)
ethos_violence	2.7B	0.310 (0.084)	0.381 (0.115)	0.373 (0.373)	0.426 (0.062)	0.423 (0.108)
	6.7B	0.426 (0.100)	0.380 (0.104)	0.505 (0.505)	0.490 (0.060)	0.513 (0.026)
	13B	0.379 (0.124)	0.437 (0.052)	0.431 (0.431)	0.475 (0.066)	0.492 (0.081)
hate_speech18	2.7B	0.454 (0.193)	0.396 (0.086)	0.413 (0.413)	0.460 (0.046)	0.455 (0.027)
	6.7B	0.481 (0.022)	0.547 (0.028)	0.638 (0.638)	0.628 (0.040)	0.612 (0.050)
	13B	0.488 (0.047)	0.585 (0.066)	0.571 (0.571)	0.590 (0.048)	0.582 (0.045)

Table 9: Macro F1-score across 27 datasets for OPT model family under the Original ICL Task.

Dataset	Size	Method				
Sentiment Task		Orig.	CC	DC w/Demo	DC w/Test	ICC
AGNews	2.7B	0.635 (0.092)	0.724 (0.041)	0.596 (0.596)	0.771 (0.043)	0.727 (0.039)
	J (6B)	0.755 (0.129)	0.787 (0.073)	0.726 (0.726)	0.755 (0.046)	0.867 (0.012)
	20B	0.640 (0.066)	0.826 (0.031)	0.798 (0.798)	0.813 (0.014)	0.818 (0.030)
CR	2.7B	0.744 (0.106)	0.797 (0.095)	0.857 (0.857)	0.878 (0.032)	0.887 (0.019)
	J (6B)	0.869 (0.052)	0.885 (0.026)	0.904 (0.904)	0.911 (0.002)	0.912 (0.004)
	20B	0.895 (0.031)	0.919 (0.010)	0.914 (0.914)	0.918 (0.008)	0.919 (0.005)
financial_phrasebank	2.7B	0.561 (0.204)	0.481 (0.067)	0.673 (0.673)	0.688 (0.052)	0.696 (0.036)
	J (6B)	0.593 (0.110)	0.418 (0.115)	0.540 (0.540)	0.617 (0.073)	0.681 (0.056)
	20B	0.649 (0.194)	0.477 (0.062)	0.522 (0.522)	0.590 (0.077)	0.688 (0.075)
poem_sentiment	2.7B	0.410 (0.077)	0.367 (0.112)	0.414 (0.414)	0.542 (0.090)	0.545 (0.083)
	J (6B)	0.512 (0.157)	0.323 (0.148)	0.359 (0.359)	0.524 (0.098)	0.639 (0.082)
	20B	0.410 (0.146)	0.435 (0.107)	0.468 (0.468)	0.494 (0.057)	0.517 (0.064)
MR	2.7B	0.847 (0.067)	0.856 (0.027)	0.847 (0.847)	0.897 (0.005)	0.897 (0.003)
	J (6B)	0.893 (0.028)	0.900 (0.007)	0.903 (0.903)	0.902 (0.013)	0.915 (0.006)
	20B	0.875 (0.037)	0.894 (0.019)	0.917 (0.917)	0.915 (0.009)	0.917 (0.004)
sst2	2.7B	0.602 (0.269)	0.791 (0.101)	0.867 (0.867)	0.920 (0.017)	0.884 (0.029)
	J (6B)	0.927 (0.044)	0.940 (0.014)	0.942 (0.942)	0.954 (0.012)	0.950 (0.007)
	20B	0.944 (0.015)	0.956 (0.011)	0.958 (0.958)	0.956 (0.012)	0.960 (0.009)
Subj	2.7B	0.327 (0.005)	0.710 (0.177)	0.712 (0.712)	0.778 (0.051)	0.801 (0.042)
	J (6B)	0.606 (0.091)	0.556 (0.143)	0.598 (0.598)	0.675 (0.116)	0.696 (0.100)
	20B	0.578 (0.119)	0.578 (0.147)	0.721 (0.721)	0.709 (0.050)	0.740 (0.055)
TREC	2.7B	0.431 (0.057)	0.387 (0.037)	0.566 (0.566)	0.535 (0.077)	0.547 (0.075)
	J (6B)	0.416 (0.078)	0.503 (0.066)	0.533 (0.533)	0.564 (0.030)	0.564 (0.013)
	20B	0.481 (0.045)	0.615 (0.054)	0.572 (0.572)	0.620 (0.047)	0.608 (0.057)
NLI Task		Orig.	CC	DC w/Demo	DC w/Test	ICC
ANLI	2.7B	0.221 (0.035)	0.208 (0.052)	0.260 (0.260)	0.255 (0.061)	0.294 (0.050)
	J (6B)	0.256 (0.112)	0.235 (0.028)	0.267 (0.267)	0.252 (0.044)	0.394 (0.109)
	20B	0.237 (0.039)	0.229 (0.050)	0.198 (0.198)	0.184 (0.023)	0.296 (0.017)
WNLI	2.7B	0.334 (0.042)	0.325 (0.025)	0.373 (0.373)	0.400 (0.066)	0.417 (0.070)
	J (6B)	0.304 (0.000)	0.394 (0.065)	0.414 (0.414)	0.444 (0.055)	0.497 (0.044)
	20B	0.344 (0.032)	0.353 (0.085)	0.358 (0.358)	0.391 (0.040)	0.465 (0.079)
RTE	2.7B	0.385 (0.110)	0.392 (0.074)	0.407 (0.407)	0.384 (0.070)	0.492 (0.060)
	J (6B)	0.495 (0.084)	0.492 (0.085)	0.553 (0.553)	0.489 (0.090)	0.580 (0.049)
	20B	0.505 (0.050)	0.483 (0.034)	0.477 (0.477)	0.408 (0.055)	0.566 (0.021)
CB	2.7B	0.260 (0.059)	0.189 (0.116)	0.205 (0.205)	0.206 (0.032)	0.334 (0.107)
	J (6B)	0.278 (0.071)	0.345 (0.118)	0.237 (0.237)	0.256 (0.079)	0.484 (0.095)
	20B	0.312 (0.030)	0.188 (0.121)	0.291 (0.291)	0.229 (0.159)	0.353 (0.076)
SICK	2.7B	0.262 (0.120)	0.313 (0.062)	0.328 (0.328)	0.305 (0.107)	0.425 (0.102)
	J (6B)	0.356 (0.056)	0.392 (0.047)	0.264 (0.264)	0.304 (0.108)	0.515 (0.082)
	20B	0.301 (0.115)	0.493 (0.077)	0.233 (0.233)	0.222 (0.091)	0.556 (0.039)
Detection Task		Orig.	CC	DC w/Demo	DC w/Test	ICC
social_bias_frames	2.7B	0.419 (0.096)	0.467 (0.064)	0.528 (0.528)	0.537 (0.064)	0.537 (0.069)
	J (6B)	0.506 (0.066)	0.585 (0.029)	0.653 (0.653)	0.684 (0.023)	0.692 (0.045)
	20B	0.375 (0.015)	0.630 (0.033)	0.684 (0.684)	0.720 (0.026)	0.706 (0.028)
tweet_eval_stance_athesim	2.7B	0.291 (0.026)	0.272 (0.087)	0.314 (0.314)	0.345 (0.036)	0.344 (0.021)
	J (6B)	0.240 (0.119)	0.256 (0.036)	0.277 (0.277)	0.267 (0.013)	0.360 (0.018)
	20B	0.162 (0.063)	0.271 (0.073)	0.229 (0.229)	0.242 (0.067)	0.259 (0.031)
tweet_eval_stance_feminist	2.7B	0.136 (0.086)	0.304 (0.033)	0.312 (0.312)	0.281 (0.078)	0.362 (0.087)
	J (6B)	0.403 (0.100)	0.348 (0.025)	0.400 (0.400)	0.374 (0.030)	0.474 (0.046)
	20B	0.152 (0.062)	0.351 (0.022)	0.388 (0.388)	0.358 (0.048)	0.403 (0.036)
tweet_eval_hate	2.7B	0.403 (0.061)	0.468 (0.047)	0.492 (0.492)	0.523 (0.030)	0.524 (0.030)
	J (6B)	0.326 (0.034)	0.365 (0.041)	0.541 (0.541)	0.575 (0.030)	0.551 (0.050)
	20B	0.403 (0.061)	0.468 (0.047)	0.492 (0.492)	0.523 (0.030)	0.524 (0.030)
tweet_eval_irony	2.7B	0.432 (0.055)	0.333 (0.045)	0.419 (0.419)	0.417 (0.059)	0.436 (0.057)
	J (6B)	0.600 (0.108)	0.501 (0.077)	0.556 (0.556)	0.625 (0.074)	0.602 (0.076)
	20B	0.492 (0.123)	0.480 (0.085)	0.652 (0.652)	0.657 (0.034)	0.637 (0.037)
tweet_eval_offensive	2.7B	0.442 (0.017)	0.411 (0.103)	0.558 (0.558)	0.535 (0.072)	0.551 (0.070)
	J (6B)	0.595 (0.115)	0.509 (0.064)	0.662 (0.662)	0.689 (0.015)	0.691 (0.009)
	20B	0.464 (0.063)	0.402 (0.070)	0.512 (0.512)	0.512 (0.041)	0.532 (0.023)
ethos_binary	2.7B	0.485 (0.053)	0.543 (0.066)	0.608 (0.608)	0.556 (0.050)	0.628 (0.052)
	J (6B)	0.616 (0.167)	0.541 (0.062)	0.708 (0.708)	0.753 (0.018)	0.756 (0.020)
	20B	0.412 (0.064)	0.647 (0.076)	0.763 (0.763)	0.774 (0.018)	0.775 (0.009)
ethos_disability	2.7B	0.469 (0.000)	0.331 (0.107)	0.459 (0.459)	0.395 (0.069)	0.478 (0.072)
	J (6B)	0.405 (0.129)	0.382 (0.139)	0.352 (0.352)	0.390 (0.064)	0.415 (0.105)
	20B	0.257 (0.173)	0.208 (0.102)	0.229 (0.229)	0.255 (0.068)	0.319 (0.103)
ethos_gender	2.7B	0.494 (0.049)	0.432 (0.112)	0.483 (0.483)	0.473 (0.061)	0.489 (0.099)
	J (6B)	0.283 (0.119)	0.270 (0.095)	0.404 (0.404)	0.418 (0.045)	0.443 (0.019)
	20B	0.232 (0.107)	0.227 (0.090)	0.353 (0.353)	0.328 (0.021)	0.401 (0.078)
ethos_national_origin	2.7B	0.452 (0.004)	0.395 (0.061)	0.367 (0.367)	0.452 (0.040)	0.444 (0.034)
	J (6B)	0.332 (0.109)	0.310 (0.125)	0.424 (0.424)	0.445 (0.038)	0.449 (0.051)
	20B	0.240 (0.110)	0.253 (0.149)	0.430 (0.430)	0.351 (0.019)	0.446 (0.088)
ethos_race	2.7B	0.509 (0.067)	0.504 (0.098)	0.528 (0.528)	0.539 (0.045)	0.589 (0.066)
	J (6B)	0.358 (0.120)	0.416 (0.199)	0.510 (0.510)	0.553 (0.091)	0.568 (0.050)
	20B	0.293 (0.130)	0.309 (0.154)	0.341 (0.341)	0.363 (0.101)	0.495 (0.049)
ethos_religion	2.7B	0.454 (0.015)	0.503 (0.035)	0.425 (0.425)	0.512 (0.028)	0.517 (0.069)
	J (6B)	0.332 (0.125)	0.298 (0.148)	0.400 (0.400)	0.421 (0.056)	0.476 (0.083)
	20B	0.240 (0.111)	0.217 (0.117)	0.334 (0.334)	0.350 (0.064)	0.401 (0.063)
ethos_violence	2.7B	0.420 (0.037)	0.455 (0.082)	0.469 (0.469)	0.480 (0.077)	0.526 (0.066)
	J (6B)	0.369 (0.097)	0.269 (0.027)	0.546 (0.546)	0.531 (0.024)	0.568 (0.037)
	20B	0.279 (0.061)	0.334 (0.126)	0.497 (0.497)	0.491 (0.044)	0.499 (0.072)
hate_speech18	2.7B	0.502 (0.038)	0.460 (0.098)	0.544 (0.544)	0.574 (0.035)	0.558 (0.051)
	J (6B)	0.469 (0.009)	0.487 (0.114)	0.568 (0.568)	0.588 (0.022)	0.587 (0.021)
	20B	0.397 (0.151)	0.329 (0.114)	0.559 (0.559)	0.561 (0.062)	0.561 (0.069)

Table 10: Macro F1-score across 27 datasets for GPT model family under the Original ICL Task.

Dataset	Size	Method				
Sentiment Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
AGNews	7B	0.803 (0.010)	0.770 (0.013)	0.753 (0.753)	0.759 (0.016)	0.822 (0.018)
	13B	0.741 (0.046)	0.733 (0.035)	0.730 (0.730)	0.730 (0.026)	0.764 (0.047)
CR	7B	0.892 (0.039)	0.926 (0.010)	0.929 (0.929)	0.927 (0.007)	0.938 (0.013)
	13B	0.928 (0.023)	0.934 (0.013)	0.932 (0.932)	0.931 (0.015)	0.936 (0.015)
financial_phrasebank	7B	0.665 (0.192)	0.426 (0.066)	0.515 (0.515)	0.588 (0.102)	0.720 (0.082)
	13B	0.631 (0.141)	0.598 (0.066)	0.744 (0.744)	0.752 (0.037)	0.772 (0.029)
poem_sentiment	7B	0.408 (0.089)	0.369 (0.045)	0.368 (0.368)	0.509 (0.044)	0.544 (0.061)
	13B	0.431 (0.138)	0.415 (0.073)	0.404 (0.404)	0.531 (0.084)	0.575 (0.066)
MR	7B	0.928 (0.013)	0.942 (0.009)	0.939 (0.939)	0.940 (0.010)	0.943 (0.008)
	13B	0.940 (0.011)	0.947 (0.007)	0.946 (0.946)	0.945 (0.007)	0.947 (0.004)
sst2	7B	0.950 (0.018)	0.971 (0.009)	0.960 (0.960)	0.967 (0.009)	0.965 (0.007)
	13B	0.964 (0.007)	0.970 (0.005)	0.967 (0.967)	0.968 (0.004)	0.971 (0.004)
Subj	7B	0.608 (0.146)	0.699 (0.159)	0.628 (0.628)	0.674 (0.065)	0.800 (0.086)
	13B	0.318 (0.006)	0.744 (0.109)	0.722 (0.722)	0.827 (0.107)	0.837 (0.064)
TREC	7B	0.457 (0.086)	0.471 (0.056)	0.401 (0.401)	0.494 (0.052)	0.537 (0.029)
	13B	0.521 (0.074)	0.573 (0.075)	0.539 (0.539)	0.573 (0.074)	0.592 (0.053)
NLI Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
ANLI	7B	0.291 (0.053)	0.344 (0.057)	0.324 (0.324)	0.337 (0.026)	0.350 (0.059)
	13B	0.379 (0.032)	0.385 (0.037)	0.324 (0.324)	0.327 (0.033)	0.419 (0.024)
WNLI	7B	0.456 (0.081)	0.543 (0.112)	0.450 (0.450)	0.481 (0.099)	0.570 (0.046)
	13B	0.574 (0.111)	0.578 (0.114)	0.531 (0.531)	0.566 (0.107)	0.585 (0.062)
RTE	7B	0.714 (0.063)	0.642 (0.056)	0.659 (0.659)	0.721 (0.037)	0.729 (0.015)
	13B	0.749 (0.019)	0.740 (0.022)	0.759 (0.759)	0.756 (0.011)	0.770 (0.020)
CB	7B	0.454 (0.102)	0.531 (0.051)	0.392 (0.392)	0.448 (0.115)	0.557 (0.020)
	13B	0.580 (0.025)	0.624 (0.049)	0.614 (0.614)	0.598 (0.012)	0.652 (0.073)
SICK	7B	0.315 (0.138)	0.458 (0.113)	0.349 (0.349)	0.430 (0.111)	0.537 (0.068)
	13B	0.534 (0.091)	0.587 (0.118)	0.436 (0.436)	0.463 (0.067)	0.669 (0.052)
Detection Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
social_bias_frames	7B	0.365 (0.004)	0.734 (0.029)	0.687 (0.687)	0.736 (0.018)	0.741 (0.018)
	13B	0.359 (0.008)	0.620 (0.093)	0.753 (0.753)	0.751 (0.011)	0.753 (0.012)
tweet_eval_stance_athesim	7B	0.144 (0.080)	0.178 (0.072)	0.140 (0.140)	0.146 (0.070)	0.226 (0.032)
	13B	0.204 (0.056)	0.194 (0.089)	0.195 (0.195)	0.186 (0.081)	0.302 (0.053)
tweet_eval_stance_feminist	7B	0.374 (0.059)	0.401 (0.011)	0.395 (0.395)	0.411 (0.039)	0.499 (0.035)
	13B	0.560 (0.050)	0.412 (0.028)	0.417 (0.417)	0.433 (0.050)	0.560 (0.037)
tweet_eval_hate	7B	0.412 (0.113)	0.330 (0.033)	0.515 (0.515)	0.601 (0.064)	0.575 (0.034)
	13B	0.455 (0.078)	0.446 (0.105)	0.517 (0.517)	0.557 (0.026)	0.536 (0.068)
tweet_eval_irony	7B	0.283 (0.008)	0.440 (0.046)	0.526 (0.526)	0.613 (0.072)	0.667 (0.084)
	13B	0.283 (0.008)	0.356 (0.042)	0.512 (0.512)	0.609 (0.091)	0.633 (0.140)
tweet_eval_offensive	7B	0.214 (0.004)	0.575 (0.070)	0.471 (0.471)	0.644 (0.022)	0.646 (0.025)
	13B	0.231 (0.019)	0.555 (0.033)	0.614 (0.614)	0.629 (0.029)	0.636 (0.040)
ethos_binary	7B	0.759 (0.070)	0.690 (0.046)	0.809 (0.809)	0.823 (0.014)	0.824 (0.019)
	13B	0.683 (0.164)	0.664 (0.065)	0.784 (0.784)	0.789 (0.062)	0.803 (0.031)
ethos_disability	7B	0.317 (0.139)	0.152 (0.033)	0.189 (0.189)	0.291 (0.050)	0.354 (0.077)
	13B	0.381 (0.165)	0.218 (0.141)	0.334 (0.334)	0.352 (0.058)	0.395 (0.050)
ethos_gender	7B	0.341 (0.074)	0.210 (0.047)	0.297 (0.297)	0.380 (0.031)	0.438 (0.033)
	13B	0.325 (0.112)	0.227 (0.088)	0.315 (0.315)	0.325 (0.066)	0.390 (0.064)
ethos_national_origin	7B	0.321 (0.109)	0.193 (0.067)	0.230 (0.230)	0.369 (0.060)	0.396 (0.080)
	13B	0.319 (0.152)	0.256 (0.121)	0.380 (0.380)	0.411 (0.082)	0.413 (0.052)
ethos_race	7B	0.398 (0.074)	0.185 (0.055)	0.279 (0.279)	0.414 (0.054)	0.429 (0.080)
	13B	0.418 (0.139)	0.280 (0.133)	0.378 (0.378)	0.438 (0.083)	0.492 (0.098)
ethos_religion	7B	0.391 (0.093)	0.188 (0.022)	0.249 (0.249)	0.314 (0.065)	0.417 (0.080)
	13B	0.352 (0.116)	0.239 (0.114)	0.370 (0.370)	0.390 (0.125)	0.411 (0.133)
ethos_violence	7B	0.422 (0.137)	0.259 (0.010)	0.372 (0.372)	0.473 (0.053)	0.519 (0.121)
	13B	0.429 (0.113)	0.248 (0.000)	0.410 (0.410)	0.435 (0.087)	0.512 (0.108)
hate_speech18	7B	0.506 (0.075)	0.526 (0.137)	0.545 (0.545)	0.587 (0.080)	0.612 (0.051)
	13B	0.588 (0.098)	0.463 (0.176)	0.560 (0.560)	0.566 (0.049)	0.590 (0.062)

Table 11: Macro F1-score across 27 datasets for Llama2 model family under the Original ICL Task.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task	7B	0.734 (0.059)	0.774 (0.041)	0.773 (0.773)	0.767 (0.040)	0.784 (0.042)
	13B	0.765 (0.018)	0.745 (0.019)	0.778 (0.778)	0.772 (0.029)	0.774 (0.030)
AGNews	7B	0.926 (0.009)	0.920 (0.007)	0.924 (0.924)	0.927 (0.011)	0.930 (0.008)
	13B	0.919 (0.021)	0.928 (0.012)	0.928 (0.928)	0.928 (0.012)	0.937 (0.010)
CR	7B	0.869 (0.052)	0.574 (0.058)	0.759 (0.759)	0.774 (0.049)	0.877 (0.040)
	13B	0.720 (0.119)	0.708 (0.075)	0.806 (0.806)	0.790 (0.047)	0.829 (0.041)
financial_phrasebank	7B	0.562 (0.104)	0.354 (0.017)	0.527 (0.527)	0.562 (0.060)	0.575 (0.079)
	13B	0.464 (0.054)	0.354 (0.019)	0.472 (0.472)	0.521 (0.021)	0.521 (0.042)
poem_sentiment	7B	0.923 (0.024)	0.930 (0.009)	0.931 (0.931)	0.928 (0.014)	0.931 (0.009)
	13B	0.926 (0.011)	0.931 (0.009)	0.928 (0.928)	0.928 (0.011)	0.929 (0.005)
MR	7B	0.951 (0.004)	0.954 (0.015)	0.957 (0.957)	0.962 (0.005)	0.965 (0.007)
	13B	0.955 (0.006)	0.954 (0.004)	0.954 (0.954)	0.955 (0.006)	0.965 (0.007)
sst2	7B	0.334 (0.011)	0.620 (0.174)	0.671 (0.671)	0.671 (0.136)	0.689 (0.184)
	13B	0.330 (0.003)	0.727 (0.071)	0.735 (0.735)	0.745 (0.060)	0.741 (0.064)
Subj	7B	0.425 (0.037)	0.431 (0.036)	0.486 (0.486)	0.494 (0.037)	0.460 (0.022)
	13B	0.669 (0.052)	0.741 (0.036)	0.709 (0.709)	0.738 (0.042)	0.712 (0.040)
TREC						
NLI Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
	7B	0.429 (0.036)	0.446 (0.022)	0.439 (0.439)	0.440 (0.035)	0.469 (0.013)
ANLI	13B	0.416 (0.033)	0.411 (0.038)	0.385 (0.385)	0.385 (0.011)	0.443 (0.021)
	7B	0.652 (0.046)	0.716 (0.039)	0.675 (0.675)	0.686 (0.043)	0.721 (0.027)
WNLI	13B	0.594 (0.096)	0.630 (0.085)	0.630 (0.630)	0.594 (0.076)	0.671 (0.033)
	7B	0.758 (0.006)	0.756 (0.012)	0.751 (0.751)	0.758 (0.009)	0.770 (0.007)
RTE	13B	0.750 (0.009)	0.751 (0.015)	0.752 (0.752)	0.723 (0.018)	0.757 (0.007)
	7B	0.585 (0.074)	0.610 (0.095)	0.597 (0.597)	0.586 (0.091)	0.691 (0.063)
CB	13B	0.704 (0.032)	0.699 (0.040)	0.714 (0.714)	0.671 (0.065)	0.715 (0.044)
	7B	0.460 (0.053)	0.487 (0.044)	0.491 (0.491)	0.481 (0.036)	0.533 (0.033)
SICK	13B	0.510 (0.065)	0.445 (0.039)	0.397 (0.397)	0.408 (0.041)	0.537 (0.041)
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Detection Task	7B	0.364 (0.011)	0.743 (0.024)	0.777 (0.777)	0.776 (0.014)	0.777 (0.017)
	13B	0.369 (0.013)	0.713 (0.038)	0.758 (0.758)	0.765 (0.026)	0.771 (0.021)
social_bias_frames	7B	0.219 (0.043)	0.168 (0.063)	0.192 (0.192)	0.194 (0.048)	0.233 (0.080)
	13B	0.224 (0.030)	0.190 (0.082)	0.197 (0.197)	0.189 (0.079)	0.264 (0.068)
tweet_eval_stance_athesim	7B	0.496 (0.072)	0.402 (0.022)	0.412 (0.412)	0.418 (0.031)	0.551 (0.021)
	13B	0.578 (0.033)	0.479 (0.040)	0.524 (0.524)	0.496 (0.046)	0.596 (0.032)
tweet_eval_stance_feminist	7B	0.481 (0.077)	0.476 (0.093)	0.521 (0.521)	0.631 (0.047)	0.561 (0.079)
	13B	0.431 (0.081)	0.368 (0.062)	0.419 (0.419)	0.513 (0.060)	0.504 (0.023)
tweet_eval_hate	7B	0.283 (0.008)	0.524 (0.065)	0.607 (0.607)	0.651 (0.052)	0.663 (0.035)
	13B	0.283 (0.008)	0.373 (0.025)	0.560 (0.560)	0.568 (0.061)	0.635 (0.064)
tweet_eval_irony	7B	0.218 (0.006)	0.626 (0.032)	0.610 (0.610)	0.657 (0.065)	0.681 (0.071)
	13B	0.231 (0.019)	0.555 (0.033)	0.614 (0.614)	0.629 (0.029)	0.636 (0.040)
tweet_eval_offensive	7B	0.739 (0.053)	0.740 (0.040)	0.817 (0.817)	0.805 (0.025)	0.828 (0.049)
	13B	0.772 (0.050)	0.742 (0.055)	0.801 (0.801)	0.797 (0.029)	0.816 (0.033)
ethos_binary	7B	0.248 (0.118)	0.222 (0.114)	0.148 (0.148)	0.249 (0.058)	0.314 (0.080)
	13B	0.276 (0.111)	0.176 (0.094)	0.162 (0.162)	0.187 (0.083)	0.283 (0.071)
ethos_disability	7B	0.351 (0.093)	0.337 (0.067)	0.336 (0.336)	0.371 (0.028)	0.406 (0.020)
	13B	0.338 (0.061)	0.265 (0.088)	0.290 (0.290)	0.285 (0.059)	0.398 (0.041)
ethos_gender	7B	0.301 (0.092)	0.273 (0.125)	0.239 (0.239)	0.270 (0.031)	0.363 (0.088)
	13B	0.328 (0.086)	0.257 (0.067)	0.285 (0.285)	0.312 (0.077)	0.377 (0.081)
ethos_national_origin	7B	0.341 (0.154)	0.314 (0.166)	0.292 (0.292)	0.358 (0.076)	0.435 (0.128)
	13B	0.367 (0.112)	0.288 (0.104)	0.338 (0.338)	0.362 (0.123)	0.431 (0.140)
ethos_race	7B	0.331 (0.111)	0.293 (0.070)	0.278 (0.278)	0.305 (0.035)	0.360 (0.082)
	13B	0.322 (0.089)	0.274 (0.040)	0.280 (0.280)	0.305 (0.074)	0.325 (0.088)
ethos_religion	7B	0.394 (0.078)	0.345 (0.053)	0.379 (0.379)	0.408 (0.048)	0.457 (0.118)
	13B	0.439 (0.080)	0.348 (0.051)	0.394 (0.394)	0.419 (0.039)	0.470 (0.060)
ethos_violence	7B	0.604 (0.088)	0.642 (0.067)	0.629 (0.629)	0.664 (0.047)	0.672 (0.067)
	13B	0.554 (0.065)	0.569 (0.101)	0.628 (0.628)	0.603 (0.093)	0.623 (0.055)
hate_speech18	7B					
	13B					

Table 12: Macro F1-score across 27 datasets for Llama2-chat model family under the Original ICL Task.

Dataset	Size	Method				
Sentiment Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
AGNews	2.7B	0.193 (0.058)	0.167 (0.037)	0.187 (0.187)	0.195 (0.043)	0.261 (0.032)
	6.7B	0.184 (0.068)	0.143 (0.043)	0.228 (0.228)	0.226 (0.045)	0.305 (0.030)
	13B	0.205 (0.071)	0.155 (0.048)	0.193 (0.193)	0.294 (0.029)	0.311 (0.047)
CR	2.7B	0.838 (0.055)	0.738 (0.090)	0.865 (0.865)	0.902 (0.016)	0.905 (0.009)
	6.7B	0.587 (0.255)	0.507 (0.080)	0.844 (0.844)	0.867 (0.026)	0.847 (0.069)
	13B	0.687 (0.180)	0.841 (0.097)	0.875 (0.875)	0.900 (0.019)	0.884 (0.037)
financial_phrasebank	2.7B	0.339 (0.084)	0.390 (0.077)	0.298 (0.298)	0.392 (0.085)	0.377 (0.069)
	6.7B	0.216 (0.058)	0.320 (0.085)	0.212 (0.212)	0.403 (0.095)	0.393 (0.091)
	13B	0.396 (0.093)	0.376 (0.115)	0.265 (0.265)	0.405 (0.073)	0.403 (0.066)
poem_sentiment	2.7B	0.295 (0.036)	0.294 (0.025)	0.242 (0.242)	0.301 (0.017)	0.309 (0.036)
	6.7B	0.224 (0.094)	0.297 (0.039)	0.279 (0.279)	0.330 (0.052)	0.362 (0.037)
	13B	0.195 (0.102)	0.260 (0.035)	0.259 (0.259)	0.321 (0.042)	0.319 (0.062)
MR	2.7B	0.825 (0.122)	0.830 (0.049)	0.885 (0.885)	0.909 (0.019)	0.893 (0.033)
	6.7B	0.809 (0.143)	0.794 (0.069)	0.883 (0.883)	0.918 (0.011)	0.912 (0.016)
	13B	0.622 (0.252)	0.847 (0.112)	0.880 (0.880)	0.922 (0.014)	0.880 (0.058)
sst2	2.7B	0.565 (0.257)	0.583 (0.382)	0.650 (0.650)	0.624 (0.387)	0.609 (0.410)
	6.7B	0.503 (0.241)	0.458 (0.216)	0.596 (0.596)	0.590 (0.197)	0.624 (0.174)
	13B	0.578 (0.256)	0.659 (0.262)	0.676 (0.676)	0.695 (0.296)	0.707 (0.283)
Subj	2.7B	0.621 (0.241)	0.627 (0.203)	0.728 (0.728)	0.768 (0.152)	0.801 (0.081)
	6.7B	0.520 (0.176)	0.696 (0.128)	0.694 (0.694)	0.699 (0.092)	0.789 (0.098)
	13B	0.423 (0.146)	0.632 (0.166)	0.761 (0.761)	0.754 (0.106)	0.767 (0.138)
TREC	2.7B	0.114 (0.075)	0.748 (0.196)	0.206 (0.206)	0.244 (0.062)	0.264 (0.061)
	6.7B	0.188 (0.063)	0.088 (0.051)	0.207 (0.207)	0.226 (0.063)	0.240 (0.041)
	13B	0.100 (0.052)	0.159 (0.081)	0.176 (0.176)	0.185 (0.128)	0.226 (0.101)
NLI Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
ANLI	2.7B	0.226 (0.046)	0.247 (0.061)	0.247 (0.247)	0.219 (0.029)	0.303 (0.039)
	6.7B	0.199 (0.033)	0.204 (0.031)	0.267 (0.267)	0.242 (0.045)	0.296 (0.026)
	13B	0.210 (0.063)	0.240 (0.050)	0.259 (0.259)	0.245 (0.028)	0.265 (0.041)
WNLI	2.7B	0.337 (0.077)	0.299 (0.005)	0.331 (0.331)	0.309 (0.022)	0.362 (0.059)
	6.7B	0.318 (0.023)	0.302 (0.003)	0.341 (0.341)	0.306 (0.019)	0.433 (0.045)
	13B	0.388 (0.084)	0.392 (0.078)	0.373 (0.373)	0.381 (0.064)	0.434 (0.034)
RTE	2.7B	0.415 (0.087)	0.411 (0.075)	0.413 (0.413)	0.389 (0.047)	0.474 (0.023)
	6.7B	0.435 (0.070)	0.350 (0.006)	0.472 (0.472)	0.412 (0.064)	0.532 (0.043)
	13B	0.397 (0.098)	0.481 (0.092)	0.469 (0.469)	0.457 (0.099)	0.529 (0.035)
CB	2.7B	0.257 (0.103)	0.281 (0.074)	0.236 (0.236)	0.279 (0.148)	0.273 (0.099)
	6.7B	0.257 (0.086)	0.417 (0.026)	0.323 (0.323)	0.322 (0.073)	0.418 (0.077)
	13B	0.270 (0.094)	0.365 (0.115)	0.294 (0.294)	0.366 (0.097)	0.428 (0.054)
SICK	2.7B	0.244 (0.004)	0.271 (0.035)	0.269 (0.269)	0.314 (0.050)	0.334 (0.033)
	6.7B	0.252 (0.015)	0.309 (0.031)	0.296 (0.296)	0.301 (0.047)	0.323 (0.051)
	13B	0.281 (0.037)	0.286 (0.069)	0.298 (0.298)	0.301 (0.064)	0.304 (0.062)
Detection Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
social_bias_frames	2.7B	0.477 (0.095)	0.469 (0.064)	0.461 (0.461)	0.500 (0.091)	0.528 (0.069)
	6.7B	0.416 (0.111)	0.520 (0.105)	0.531 (0.531)	0.563 (0.074)	0.612 (0.034)
	13B	0.337 (0.052)	0.432 (0.028)	0.466 (0.466)	0.504 (0.073)	0.550 (0.085)
tweet_eval_stance_athesim	2.7B	0.244 (0.111)	0.199 (0.056)	0.342 (0.342)	0.269 (0.040)	0.366 (0.036)
	6.7B	0.258 (0.073)	0.176 (0.067)	0.349 (0.349)	0.265 (0.080)	0.353 (0.064)
	13B	0.278 (0.107)	0.237 (0.078)	0.300 (0.300)	0.327 (0.099)	0.333 (0.051)
tweet_eval_stance_feminist	2.7B	0.249 (0.085)	0.265 (0.077)	0.323 (0.323)	0.380 (0.034)	0.337 (0.094)
	6.7B	0.250 (0.078)	0.211 (0.121)	0.345 (0.345)	0.241 (0.111)	0.347 (0.101)
	13B	0.192 (0.070)	0.184 (0.088)	0.254 (0.254)	0.296 (0.060)	0.297 (0.070)
tweet_eval_hate	2.7B	0.422 (0.064)	0.383 (0.081)	0.425 (0.425)	0.437 (0.059)	0.474 (0.036)
	6.7B	0.433 (0.080)	0.339 (0.047)	0.478 (0.478)	0.473 (0.030)	0.491 (0.010)
	13B	0.406 (0.058)	0.423 (0.097)	0.489 (0.489)	0.494 (0.024)	0.489 (0.022)
tweet_eval_irony	2.7B	0.417 (0.076)	0.391 (0.037)	0.417 (0.417)	0.432 (0.085)	0.448 (0.074)
	6.7B	0.364 (0.043)	0.412 (0.092)	0.428 (0.428)	0.485 (0.085)	0.489 (0.088)
	13B	0.390 (0.028)	0.442 (0.096)	0.505 (0.505)	0.513 (0.048)	0.508 (0.045)
tweet_eval_offensive	2.7B	0.439 (0.066)	0.282 (0.080)	0.496 (0.496)	0.456 (0.061)	0.499 (0.036)
	6.7B	0.439 (0.035)	0.244 (0.048)	0.490 (0.490)	0.445 (0.080)	0.508 (0.033)
	13B	0.419 (0.004)	0.428 (0.049)	0.405 (0.405)	0.424 (0.073)	0.428 (0.053)
ethos_binary	2.7B	0.392 (0.061)	0.392 (0.061)	0.495 (0.495)	0.503 (0.072)	0.521 (0.052)
	6.7B	0.401 (0.058)	0.489 (0.114)	0.469 (0.469)	0.515 (0.070)	0.542 (0.043)
	13B	0.425 (0.095)	0.439 (0.078)	0.514 (0.514)	0.592 (0.079)	0.604 (0.112)
ethos_disability	2.7B	0.469 (0.000)	0.467 (0.176)	0.189 (0.189)	0.266 (0.056)	0.516 (0.082)
	6.7B	0.469 (0.000)	0.364 (0.209)	0.345 (0.345)	0.362 (0.084)	0.475 (0.073)
	13B	0.469 (0.000)	0.467 (0.004)	0.333 (0.333)	0.422 (0.092)	0.533 (0.030)
ethos_gender	2.7B	0.452 (0.016)	0.421 (0.111)	0.343 (0.343)	0.344 (0.075)	0.461 (0.033)
	6.7B	0.444 (0.000)	0.334 (0.137)	0.374 (0.374)	0.375 (0.071)	0.452 (0.070)
	13B	0.444 (0.000)	0.398 (0.095)	0.409 (0.409)	0.441 (0.035)	0.501 (0.062)
ethos_national_origin	2.7B	0.472 (0.033)	0.277 (0.145)	0.447 (0.447)	0.484 (0.057)	0.473 (0.076)
	6.7B	0.454 (0.001)	0.286 (0.099)	0.344 (0.344)	0.417 (0.117)	0.455 (0.049)
	13B	0.455 (0.000)	0.480 (0.030)	0.334 (0.334)	0.444 (0.066)	0.496 (0.037)
ethos_race	2.7B	0.456 (0.009)	0.385 (0.136)	0.294 (0.294)	0.444 (0.118)	0.536 (0.090)
	6.7B	0.487 (0.070)	0.285 (0.126)	0.442 (0.442)	0.504 (0.146)	0.552 (0.023)
	13B	0.480 (0.058)	0.507 (0.070)	0.374 (0.374)	0.424 (0.065)	0.560 (0.140)
ethos_religion	2.7B	0.456 (0.017)	0.411 (0.105)	0.326 (0.326)	0.345 (0.122)	0.470 (0.060)
	6.7B	0.448 (0.000)	0.349 (0.096)	0.324 (0.324)	0.407 (0.133)	0.520 (0.059)
	13B	0.448 (0.000)	0.508 (0.052)	0.423 (0.423)	0.462 (0.119)	0.580 (0.051)
ethos_violence	2.7B	0.420 (0.037)	0.455 (0.082)	0.469 (0.469)	0.480 (0.077)	0.526 (0.066)
	6.7B	0.415 (0.028)	0.412 (0.082)	0.324 (0.324)	0.421 (0.135)	0.498 (0.104)
	13B	0.401 (0.000)	0.559 (0.113)	0.537 (0.537)	0.594 (0.073)	0.602 (0.058)
hate_speech18	2.7B	0.469 (0.003)	0.400 (0.098)	0.258 (0.258)	0.379 (0.123)	0.476 (0.046)
	6.7B	0.490 (0.041)	0.330 (0.157)	0.408 (0.408)	0.426 (0.118)	0.523 (0.053)
	13B	0.470 (0.004)	0.487 (0.018)	0.354 (0.354)	0.455 (0.053)	0.482 (0.020)

Table 13: Macro F1-score across 27 datasets for OPT model family under the Task Learning setting.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task						
AGNews	2.7B	0.215 (0.044)	0.151 (0.039)	0.299 (0.299)	0.349 (0.063)	0.396 (0.055)
	J (6B)	0.205 (0.032)	0.197 (0.048)	0.238 (0.238)	0.267 (0.046)	0.314 (0.036)
	20B	0.229 (0.030)	0.206 (0.041)	0.239 (0.239)	0.281 (0.048)	0.314 (0.034)
CR	2.7B	0.603 (0.176)	0.770 (0.077)	0.772 (0.772)	0.849 (0.050)	0.835 (0.043)
	J (6B)	0.695 (0.138)	0.618 (0.125)	0.879 (0.879)	0.868 (0.057)	0.919 (0.003)
	20B	0.727 (0.159)	0.869 (0.044)	0.860 (0.860)	0.886 (0.022)	0.889 (0.025)
financial_phrasebank	2.7B	0.298 (0.111)	0.356 (0.091)	0.406 (0.406)	0.476 (0.032)	0.432 (0.102)
	J (6B)	0.358 (0.097)	0.423 (0.071)	0.405 (0.405)	0.439 (0.039)	0.506 (0.083)
	20B	0.319 (0.062)	0.350 (0.080)	0.361 (0.361)	0.386 (0.074)	0.410 (0.041)
poem_sentiment	2.7B	0.228 (0.075)	0.274 (0.067)	0.238 (0.238)	0.290 (0.022)	0.301 (0.035)
	J (6B)	0.223 (0.061)	0.260 (0.006)	0.228 (0.228)	0.245 (0.036)	0.286 (0.024)
	20B	0.299 (0.023)	0.266 (0.021)	0.302 (0.302)	0.354 (0.047)	0.367 (0.017)
MR	2.7B	0.507 (0.180)	0.778 (0.084)	0.743 (0.743)	0.820 (0.034)	0.805 (0.044)
	J (6B)	0.811 (0.188)	0.837 (0.068)	0.899 (0.899)	0.893 (0.023)	0.905 (0.034)
	20B	0.639 (0.247)	0.860 (0.085)	0.818 (0.818)	0.887 (0.035)	0.908 (0.009)
sst2	2.7B	0.339 (0.016)	0.537 (0.173)	0.604 (0.604)	0.605 (0.212)	0.618 (0.101)
	J (6B)	0.628 (0.264)	0.551 (0.262)	0.674 (0.674)	0.623 (0.262)	0.686 (0.271)
	20B	0.808 (0.233)	0.903 (0.081)	0.909 (0.909)	0.908 (0.082)	0.915 (0.077)
Subj	2.7B	0.337 (0.016)	0.557 (0.196)	0.745 (0.745)	0.754 (0.108)	0.795 (0.079)
	J (6B)	0.417 (0.137)	0.369 (0.104)	0.597 (0.597)	0.501 (0.137)	0.655 (0.074)
	20B	0.523 (0.191)	0.386 (0.127)	0.572 (0.572)	0.423 (0.097)	0.613 (0.115)
TREC	2.7B	0.186 (0.091)	0.205 (0.034)	0.259 (0.259)	0.288 (0.073)	0.318 (0.068)
	J (6B)	0.194 (0.084)	0.107 (0.055)	0.145 (0.145)	0.169 (0.099)	0.226 (0.050)
	20B	0.211 (0.071)	0.179 (0.059)	0.212 (0.212)	0.238 (0.080)	0.301 (0.079)
NLI Task						
ANLI	2.7B	0.230 (0.037)	0.276 (0.053)	0.245 (0.245)	0.281 (0.030)	0.328 (0.028)
	J (6B)	0.204 (0.031)	0.240 (0.045)	0.220 (0.220)	0.251 (0.042)	0.271 (0.032)
	20B	0.215 (0.036)	0.190 (0.043)	0.284 (0.284)	0.247 (0.030)	0.294 (0.019)
WNLI	2.7B	0.313 (0.026)	0.318 (0.026)	0.327 (0.327)	0.295 (0.008)	0.402 (0.051)
	J (6B)	0.333 (0.052)	0.308 (0.008)	0.322 (0.322)	0.309 (0.010)	0.386 (0.085)
	20B	0.304 (0.000)	0.304 (0.000)	0.333 (0.333)	0.315 (0.030)	0.415 (0.084)
RTE	2.7B	0.414 (0.076)	0.435 (0.068)	0.370 (0.370)	0.366 (0.013)	0.467 (0.046)
	J (6B)	0.415 (0.085)	0.345 (0.000)	0.385 (0.385)	0.373 (0.048)	0.460 (0.077)
	20B	0.456 (0.091)	0.381 (0.037)	0.473 (0.473)	0.411 (0.065)	0.537 (0.069)
CB	2.7B	0.216 (0.031)	0.272 (0.063)	0.226 (0.226)	0.251 (0.059)	0.295 (0.056)
	J (6B)	0.318 (0.143)	0.199 (0.010)	0.267 (0.267)	0.250 (0.112)	0.330 (0.113)
	20B	0.274 (0.100)	0.274 (0.101)	0.246 (0.246)	0.282 (0.087)	0.315 (0.073)
SICK	2.7B	0.313 (0.089)	0.307 (0.053)	0.309 (0.309)	0.323 (0.070)	0.349 (0.090)
	J (6B)	0.277 (0.040)	0.338 (0.096)	0.308 (0.308)	0.197 (0.068)	0.349 (0.108)
	20B	0.252 (0.020)	0.229 (0.044)	0.276 (0.276)	0.202 (0.080)	0.285 (0.037)
Detection Task						
social_bias_frames	2.7B	0.388 (0.110)	0.510 (0.053)	0.485 (0.485)	0.529 (0.078)	0.526 (0.047)
	J (6B)	0.337 (0.049)	0.387 (0.020)	0.474 (0.474)	0.486 (0.057)	0.533 (0.055)
	20B	0.448 (0.077)	0.385 (0.031)	0.621 (0.621)	0.591 (0.035)	0.608 (0.041)
tweet_eval_stance_athesim	2.7B	0.177 (0.086)	0.214 (0.082)	0.267 (0.267)	0.260 (0.077)	0.300 (0.071)
	J (6B)	0.307 (0.051)	0.249 (0.094)	0.377 (0.377)	0.345 (0.066)	0.397 (0.071)
	20B	0.306 (0.082)	0.279 (0.061)	0.429 (0.429)	0.417 (0.075)	0.427 (0.091)
tweet_eval_stance_feminist	2.7B	0.348 (0.067)	0.177 (0.073)	0.289 (0.289)	0.242 (0.071)	0.366 (0.052)
	J (6B)	0.283 (0.044)	0.150 (0.049)	0.317 (0.317)	0.210 (0.094)	0.367 (0.036)
	20B	0.276 (0.020)	0.167 (0.070)	0.248 (0.248)	0.232 (0.079)	0.306 (0.066)
tweet_eval_hate	2.7B	0.404 (0.063)	0.393 (0.078)	0.471 (0.471)	0.510 (0.025)	0.488 (0.011)
	J (6B)	0.368 (0.017)	0.371 (0.086)	0.494 (0.494)	0.493 (0.060)	0.498 (0.048)
	20B	0.461 (0.026)	0.343 (0.066)	0.476 (0.476)	0.469 (0.046)	0.486 (0.039)
tweet_eval_irony	2.7B	0.405 (0.047)	0.393 (0.119)	0.444 (0.444)	0.517 (0.086)	0.484 (0.080)
	J (6B)	0.397 (0.027)	0.284 (0.021)	0.402 (0.402)	0.442 (0.073)	0.451 (0.085)
	20B	0.438 (0.058)	0.315 (0.027)	0.459 (0.459)	0.455 (0.104)	0.463 (0.102)
tweet_eval_offensive	2.7B	0.434 (0.016)	0.325 (0.113)	0.451 (0.451)	0.503 (0.081)	0.507 (0.062)
	J (6B)	0.430 (0.026)	0.326 (0.135)	0.491 (0.491)	0.478 (0.066)	0.508 (0.028)
	20B	0.455 (0.045)	0.224 (0.017)	0.528 (0.528)	0.509 (0.060)	0.533 (0.071)
ethos_binary	2.7B	0.412 (0.082)	0.502 (0.105)	0.507 (0.507)	0.545 (0.051)	0.543 (0.090)
	J (6B)	0.498 (0.117)	0.324 (0.017)	0.648 (0.648)	0.637 (0.063)	0.675 (0.042)
	20B	0.464 (0.125)	0.356 (0.036)	0.597 (0.597)	0.618 (0.053)	0.629 (0.062)
ethos_disability	2.7B	0.469 (0.000)	0.331 (0.107)	0.459 (0.459)	0.395 (0.069)	0.478 (0.072)
	J (6B)	0.469 (0.000)	0.431 (0.033)	0.431 (0.431)	0.462 (0.046)	0.496 (0.053)
	20B	0.469 (0.000)	0.340 (0.170)	0.379 (0.379)	0.392 (0.149)	0.472 (0.051)
ethos_gender	2.7B	0.419 (0.052)	0.351 (0.145)	0.485 (0.485)	0.484 (0.031)	0.497 (0.029)
	J (6B)	0.444 (0.000)	0.254 (0.105)	0.254 (0.254)	0.463 (0.064)	0.522 (0.044)
	20B	0.451 (0.013)	0.293 (0.089)	0.358 (0.358)	0.393 (0.069)	0.462 (0.057)
ethos_national_origin	2.7B	0.455 (0.000)	0.302 (0.136)	0.443 (0.443)	0.424 (0.065)	0.473 (0.036)
	J (6B)	0.459 (0.007)	0.277 (0.167)	0.471 (0.471)	0.413 (0.079)	0.478 (0.031)
	20B	0.455 (0.000)	0.320 (0.151)	0.413 (0.413)	0.413 (0.100)	0.497 (0.051)
ethos_race	2.7B	0.509 (0.067)	0.504 (0.098)	0.528 (0.528)	0.539 (0.045)	0.589 (0.086)
	J (6B)	0.452 (0.000)	0.446 (0.133)	0.369 (0.369)	0.403 (0.149)	0.489 (0.096)
	20B	0.500 (0.069)	0.335 (0.120)	0.462 (0.462)	0.457 (0.071)	0.506 (0.050)
ethos_religion	2.7B	0.448 (0.000)	0.348 (0.111)	0.518 (0.518)	0.530 (0.050)	0.526 (0.052)
	J (6B)	0.470 (0.044)	0.383 (0.139)	0.501 (0.501)	0.472 (0.074)	0.532 (0.069)
	20B	0.473 (0.032)	0.327 (0.126)	0.472 (0.472)	0.460 (0.105)	0.486 (0.086)
ethos_violence	2.7B	0.431 (0.059)	0.394 (0.115)	0.453 (0.453)	0.440 (0.010)	0.480 (0.049)
	J (6B)	0.401 (0.000)	0.395 (0.135)	0.507 (0.507)	0.522 (0.114)	0.535 (0.076)
	20B	0.437 (0.112)	0.305 (0.115)	0.514 (0.514)	0.473 (0.062)	0.508 (0.094)
hate_speech18	2.7B	0.477 (0.014)	0.383 (0.171)	0.501 (0.501)	0.501 (0.038)	0.512 (0.069)
	J (6B)	0.472 (0.002)	0.343 (0.154)	0.490 (0.490)	0.488 (0.075)	0.509 (0.075)
	20B	0.501 (0.063)	0.332 (0.169)	0.465 (0.465)	0.496 (0.068)	0.515 (0.078)

Table 14: Macro F1-score across 27 datasets for GPT model family under the Task Learning setting.

Dataset	Size	Method				
Sentiment Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
AGNews	7B	0.377 (0.045)	0.317 (0.068)	0.344 (0.344)	0.334 (0.088)	0.462 (0.019)
	13B	0.381 (0.053)	0.315 (0.054)	0.320 (0.320)	0.297 (0.064)	0.488 (0.026)
CR	7B	0.920 (0.016)	0.802 (0.116)	0.906 (0.906)	0.928 (0.010)	0.932 (0.005)
	13B	0.623 (0.163)	0.569 (0.204)	0.774 (0.774)	0.821 (0.041)	0.823 (0.061)
financial_phrasebank	7B	0.346 (0.109)	0.408 (0.083)	0.398 (0.398)	0.383 (0.062)	0.418 (0.093)
	13B	0.367 (0.095)	0.384 (0.071)	0.395 (0.395)	0.368 (0.073)	0.434 (0.101)
poem_sentiment	7B	0.230 (0.066)	0.254 (0.054)	0.221 (0.221)	0.263 (0.037)	0.304 (0.044)
	13B	0.268 (0.005)	0.312 (0.068)	0.249 (0.249)	0.272 (0.039)	0.324 (0.021)
MR	7B	0.893 (0.067)	0.919 (0.028)	0.919 (0.919)	0.933 (0.019)	0.933 (0.013)
	13B	0.610 (0.257)	0.471 (0.212)	0.633 (0.633)	0.678 (0.236)	0.683 (0.268)
sst2	7B	0.839 (0.234)	0.917 (0.057)	0.876 (0.876)	0.900 (0.101)	0.916 (0.069)
	13B	0.657 (0.267)	0.784 (0.111)	0.893 (0.893)	0.930 (0.034)	0.938 (0.025)
Subj	7B	0.577 (0.178)	0.433 (0.172)	0.498 (0.498)	0.482 (0.110)	0.706 (0.072)
	13B	0.549 (0.186)	0.414 (0.118)	0.604 (0.604)	0.522 (0.181)	0.754 (0.085)
TREC	7B	0.242 (0.085)	0.200 (0.039)	0.226 (0.226)	0.263 (0.075)	0.286 (0.046)
	13B	0.329 (0.093)	0.329 (0.081)	0.314 (0.314)	0.331 (0.077)	0.345 (0.114)
NLI Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
ANLI	7B	0.196 (0.005)	0.225 (0.042)	0.225 (0.225)	0.232 (0.039)	0.269 (0.041)
	13B	0.196 (0.039)	0.255 (0.055)	0.229 (0.229)	0.250 (0.053)	0.260 (0.029)
WNLI	7B	0.352 (0.096)	0.302 (0.003)	0.303 (0.303)	0.302 (0.003)	0.440 (0.106)
	13B	0.408 (0.111)	0.419 (0.125)	0.438 (0.438)	0.420 (0.120)	0.467 (0.118)
RTE	7B	0.571 (0.150)	0.558 (0.066)	0.565 (0.565)	0.585 (0.097)	0.645 (0.070)
	13B	0.622 (0.178)	0.724 (0.043)	0.585 (0.585)	0.677 (0.134)	0.749 (0.024)
CB	7B	0.250 (0.111)	0.286 (0.066)	0.226 (0.226)	0.231 (0.062)	0.362 (0.027)
	13B	0.477 (0.064)	0.526 (0.050)	0.522 (0.522)	0.521 (0.079)	0.558 (0.028)
SICK	7B	0.337 (0.114)	0.284 (0.073)	0.270 (0.270)	0.269 (0.080)	0.348 (0.086)
	13B	0.336 (0.127)	0.350 (0.088)	0.354 (0.354)	0.348 (0.092)	0.361 (0.084)
Detection Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
social_bias_frames	7B	0.419 (0.085)	0.367 (0.019)	0.479 (0.479)	0.498 (0.103)	0.537 (0.077)
	13B	0.532 (0.134)	0.568 (0.046)	0.623 (0.623)	0.640 (0.064)	0.670 (0.070)
tweet_eval_stance_athesim	7B	0.366 (0.081)	0.269 (0.074)	0.411 (0.411)	0.375 (0.088)	0.449 (0.032)
	13B	0.403 (0.110)	0.336 (0.078)	0.387 (0.387)	0.385 (0.098)	0.514 (0.067)
tweet_eval_stance_feminist	7B	0.302 (0.049)	0.289 (0.087)	0.354 (0.354)	0.334 (0.060)	0.425 (0.030)
	13B	0.421 (0.089)	0.356 (0.052)	0.404 (0.404)	0.354 (0.125)	0.481 (0.081)
tweet_eval_hate	7B	0.457 (0.111)	0.302 (0.017)	0.446 (0.446)	0.483 (0.089)	0.494 (0.067)
	13B	0.461 (0.085)	0.313 (0.032)	0.404 (0.404)	0.475 (0.071)	0.499 (0.078)
tweet_eval_irony	7B	0.478 (0.078)	0.285 (0.009)	0.436 (0.436)	0.424 (0.095)	0.489 (0.069)
	13B	0.397 (0.055)	0.396 (0.089)	0.412 (0.412)	0.447 (0.077)	0.489 (0.045)
tweet_eval_offensive	7B	0.475 (0.059)	0.215 (0.006)	0.450 (0.450)	0.320 (0.064)	0.491 (0.046)
	13B	0.461 (0.063)	0.411 (0.161)	0.495 (0.495)	0.510 (0.140)	0.568 (0.062)
ethos_binary	7B	0.500 (0.161)	0.316 (0.019)	0.500 (0.500)	0.598 (0.098)	0.622 (0.097)
	13B	0.645 (0.192)	0.479 (0.063)	0.580 (0.580)	0.656 (0.064)	0.736 (0.038)
ethos_disability	7B	0.476 (0.015)	0.289 (0.117)	0.423 (0.423)	0.281 (0.086)	0.484 (0.097)
	13B	0.499 (0.047)	0.385 (0.084)	0.509 (0.509)	0.495 (0.055)	0.511 (0.038)
ethos_gender	7B	0.434 (0.021)	0.315 (0.117)	0.399 (0.399)	0.363 (0.098)	0.461 (0.058)
	13B	0.404 (0.080)	0.446 (0.128)	0.376 (0.376)	0.417 (0.144)	0.522 (0.064)
ethos_national_origin	7B	0.454 (0.001)	0.337 (0.133)	0.391 (0.391)	0.393 (0.077)	0.484 (0.031)
	13B	0.420 (0.069)	0.347 (0.130)	0.433 (0.433)	0.381 (0.104)	0.500 (0.031)
ethos_race	7B	0.459 (0.015)	0.265 (0.154)	0.499 (0.499)	0.454 (0.066)	0.557 (0.106)
	13B	0.463 (0.023)	0.503 (0.070)	0.513 (0.513)	0.452 (0.105)	0.586 (0.116)
ethos_religion	7B	0.468 (0.039)	0.327 (0.125)	0.495 (0.495)	0.420 (0.088)	0.544 (0.067)
	13B	0.480 (0.111)	0.359 (0.141)	0.458 (0.458)	0.448 (0.064)	0.504 (0.027)
ethos_violence	7B	0.386 (0.133)	0.252 (0.008)	0.317 (0.317)	0.350 (0.081)	0.483 (0.054)
	13B	0.495 (0.076)	0.364 (0.135)	0.398 (0.398)	0.388 (0.116)	0.522 (0.095)
hate_speech18	7B	0.449 (0.091)	0.141 (0.058)	0.488 (0.488)	0.433 (0.133)	0.509 (0.042)
	13B	0.477 (0.009)	0.486 (0.096)	0.507 (0.507)	0.503 (0.049)	0.514 (0.064)

Table 15: Macro F1-score across 27 datasets for Llama2 model family under the Task Learning setting.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task	7B	0.315 (0.108)	0.330 (0.044)	0.263 (0.263)	0.320 (0.062)	0.391 (0.042)
	13B	0.496 (0.134)	0.413 (0.095)	0.479 (0.479)	0.374 (0.087)	0.528 (0.138)
AGNews	7B	0.690 (0.160)	0.525 (0.190)	0.609 (0.609)	0.690 (0.144)	0.723 (0.114)
	13B	0.841 (0.080)	0.713 (0.202)	0.848 (0.848)	0.840 (0.093)	0.853 (0.083)
CR	7B	0.218 (0.072)	0.241 (0.112)	0.344 (0.344)	0.380 (0.103)	0.388 (0.064)
	13B	0.467 (0.110)	0.462 (0.080)	0.474 (0.474)	0.531 (0.085)	0.488 (0.068)
financial_phrasebank	7B	0.282 (0.073)	0.245 (0.097)	0.299 (0.299)	0.343 (0.053)	0.376 (0.070)
	13B	0.310 (0.036)	0.212 (0.090)	0.307 (0.307)	0.351 (0.057)	0.373 (0.045)
poem_sentiment	7B	0.906 (0.032)	0.867 (0.058)	0.895 (0.895)	0.892 (0.040)	0.927 (0.005)
	13B	0.800 (0.122)	0.496 (0.136)	0.779 (0.779)	0.788 (0.176)	0.833 (0.109)
MR	7B	0.951 (0.009)	0.910 (0.050)	0.952 (0.952)	0.952 (0.013)	0.959 (0.010)
	13B	0.744 (0.139)	0.626 (0.217)	0.828 (0.828)	0.868 (0.035)	0.844 (0.062)
sst2	7B	0.644 (0.120)	0.497 (0.194)	0.586 (0.586)	0.522 (0.184)	0.655 (0.105)
	13B	0.555 (0.103)	0.527 (0.092)	0.702 (0.702)	0.522 (0.114)	0.698 (0.142)
Subj	7B	0.191 (0.040)	0.151 (0.046)	0.173 (0.173)	0.207 (0.054)	0.237 (0.070)
	13B	0.380 (0.117)	0.389 (0.143)	0.360 (0.360)	0.417 (0.116)	0.442 (0.104)
TREC						
NLI Task	Orig.	CC	DC w/Demo	DC w/Test	ICC	
	7B	0.268 (0.026)	0.268 (0.020)	0.268 (0.268)	0.266 (0.025)	0.271 (0.022)
ANLI	13B	0.281 (0.048)	0.264 (0.026)	0.280 (0.280)	0.280 (0.037)	0.281 (0.048)
	7B	0.528 (0.153)	0.477 (0.078)	0.531 (0.531)	0.494 (0.166)	0.563 (0.133)
WNLI	13B	0.513 (0.133)	0.548 (0.140)	0.562 (0.562)	0.536 (0.128)	0.596 (0.120)
	7B	0.629 (0.082)	0.633 (0.062)	0.635 (0.635)	0.636 (0.073)	0.663 (0.040)
RTE	13B	0.729 (0.051)	0.731 (0.050)	0.722 (0.722)	0.726 (0.038)	0.731 (0.050)
	7B	0.419 (0.058)	0.337 (0.032)	0.429 (0.429)	0.368 (0.042)	0.452 (0.052)
CB	13B	0.565 (0.023)	0.561 (0.101)	0.523 (0.523)	0.515 (0.051)	0.597 (0.091)
	7B	0.256 (0.059)	0.272 (0.048)	0.255 (0.255)	0.256 (0.045)	0.280 (0.076)
SICK	13B	0.255 (0.082)	0.234 (0.053)	0.237 (0.237)	0.236 (0.072)	0.264 (0.070)
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Detection Task	7B	0.437 (0.099)	0.471 (0.073)	0.448 (0.448)	0.451 (0.085)	0.485 (0.091)
	13B	0.706 (0.091)	0.677 (0.106)	0.694 (0.694)	0.665 (0.078)	0.719 (0.048)
social_bias_frames	7B	0.394 (0.050)	0.489 (0.069)	0.484 (0.484)	0.465 (0.038)	0.530 (0.072)
	13B	0.552 (0.127)	0.471 (0.042)	0.536 (0.536)	0.485 (0.078)	0.561 (0.095)
tweet_eval_stance_athesim	7B	0.371 (0.103)	0.393 (0.083)	0.483 (0.483)	0.417 (0.070)	0.501 (0.055)
	13B	0.540 (0.037)	0.443 (0.031)	0.506 (0.506)	0.435 (0.090)	0.550 (0.063)
tweet_eval_stance_feminist	7B	0.468 (0.091)	0.409 (0.076)	0.434 (0.434)	0.475 (0.088)	0.496 (0.087)
	13B	0.496 (0.144)	0.408 (0.075)	0.478 (0.478)	0.505 (0.069)	0.511 (0.108)
tweet_eval_hate	7B	0.403 (0.059)	0.363 (0.071)	0.395 (0.395)	0.389 (0.065)	0.422 (0.035)
	13B	0.417 (0.041)	0.308 (0.016)	0.393 (0.393)	0.385 (0.060)	0.484 (0.034)
tweet_eval_irony	7B	0.442 (0.070)	0.290 (0.067)	0.344 (0.344)	0.393 (0.065)	0.441 (0.098)
	13B	0.538 (0.071)	0.455 (0.138)	0.559 (0.559)	0.467 (0.161)	0.581 (0.120)
tweet_eval_offensive	7B	0.528 (0.116)	0.448 (0.086)	0.515 (0.515)	0.534 (0.117)	0.548 (0.090)
	13B	0.700 (0.130)	0.654 (0.106)	0.732 (0.732)	0.716 (0.124)	0.751 (0.090)
ethos_binary	7B	0.559 (0.088)	0.526 (0.051)	0.566 (0.566)	0.625 (0.076)	0.615 (0.102)
	13B	0.542 (0.068)	0.359 (0.118)	0.295 (0.295)	0.243 (0.126)	0.556 (0.081)
ethos_disability	7B	0.466 (0.045)	0.461 (0.029)	0.430 (0.430)	0.475 (0.123)	0.502 (0.025)
	13B	0.443 (0.028)	0.436 (0.138)	0.447 (0.447)	0.330 (0.113)	0.481 (0.041)
ethos_gender	7B	0.455 (0.003)	0.500 (0.046)	0.444 (0.444)	0.420 (0.089)	0.543 (0.048)
	13B	0.365 (0.144)	0.485 (0.031)	0.420 (0.420)	0.310 (0.104)	0.454 (0.075)
ethos_national_origin	7B	0.549 (0.096)	0.551 (0.035)	0.448 (0.448)	0.460 (0.080)	0.588 (0.050)
	13B	0.398 (0.139)	0.369 (0.070)	0.323 (0.323)	0.304 (0.101)	0.494 (0.139)
ethos_race	7B	0.522 (0.077)	0.533 (0.052)	0.447 (0.447)	0.444 (0.099)	0.558 (0.073)
	13B	0.436 (0.192)	0.372 (0.135)	0.463 (0.463)	0.374 (0.144)	0.475 (0.027)
ethos_religion	7B	0.465 (0.056)	0.481 (0.069)	0.511 (0.511)	0.523 (0.045)	0.525 (0.080)
	13B	0.580 (0.067)	0.501 (0.056)	0.501 (0.501)	0.464 (0.083)	0.582 (0.072)
ethos_violence	7B	0.536 (0.060)	0.619 (0.072)	0.612 (0.612)	0.660 (0.069)	0.621 (0.045)
	13B	0.556 (0.086)	0.526 (0.102)	0.591 (0.591)	0.502 (0.107)	0.571 (0.106)
hate_speech18	7B					
	13B					

Table 16: Macro F1-score across 27 datasets for Llama2-chat model family under the Task Learning setting.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task						
AGNews	OPT 66B	0.793 (0.058)	0.820 (0.049)	0.761 (0.761)	0.794 (0.098)	0.827 (0.073)
	Llama2 70B	0.885 (0.016)	0.893 (0.014)	0.897 (0.897)	0.900 (0.008)	0.891 (0.014)
	Llama2-ch 70B	0.806 (0.035)	0.782 (0.051)	0.804 (0.804)	0.819 (0.042)	0.813 (0.034)
CR	OPT 66B	0.852 (0.045)	0.916 (0.012)	0.917 (0.917)	0.922 (0.006)	0.922 (0.009)
	Llama2 70B	0.918 (0.039)	0.926 (0.025)	0.928 (0.928)	0.933 (0.011)	0.933 (0.007)
	Llama2-ch 70B	0.927 (0.003)	0.927 (0.004)	0.927 (0.927)	0.927 (0.004)	0.928 (0.003)
financial_phrasebank	OPT 66B	0.503 (0.134)	0.506 (0.105)	0.426 (0.426)	0.490 (0.106)	0.542 (0.140)
	Llama2 70B	0.745 (0.190)	0.461 (0.073)	0.685 (0.685)	0.683 (0.119)	0.765 (0.061)
	Llama2-ch 70B	0.780 (0.188)	0.691 (0.082)	0.739 (0.739)	0.723 (0.073)	0.811 (0.120)
poem_sentiment	OPT 66B	0.470 (0.134)	0.389 (0.031)	0.350 (0.350)	0.436 (0.070)	0.502 (0.073)
	Llama2 70B	0.570 (0.104)	0.420 (0.058)	0.473 (0.473)	0.548 (0.020)	0.618 (0.034)
	Llama2-ch 70B	0.652 (0.077)	0.345 (0.015)	0.569 (0.569)	0.622 (0.071)	0.663 (0.072)
MR	OPT 66B	0.873 (0.052)	0.931 (0.009)	0.931 (0.931)	0.931 (0.006)	0.931 (0.009)
	Llama2 70B	0.926 (0.007)	0.927 (0.013)	0.928 (0.928)	0.939 (0.010)	0.933 (0.009)
	Llama2-ch 70B	0.930 (0.012)	0.929 (0.007)	0.930 (0.930)	0.929 (0.015)	0.930 (0.014)
sst2	OPT 66B	0.931 (0.031)	0.961 (0.005)	0.963 (0.963)	0.962 (0.003)	0.967 (0.002)
	Llama2 70B	0.958 (0.009)	0.961 (0.009)	0.956 (0.956)	0.967 (0.006)	0.962 (0.007)
	Llama2-ch 70B	0.961 (0.009)	0.963 (0.009)	0.962 (0.962)	0.960 (0.009)	0.964 (0.007)
Subj	OPT 66B	0.684 (0.170)	0.729 (0.036)	0.745 (0.745)	0.787 (0.086)	0.819 (0.056)
	Llama2 70B	0.634 (0.185)	0.742 (0.161)	0.504 (0.504)	0.694 (0.172)	0.823 (0.076)
	Llama2-ch 70B	0.749 (0.067)	0.719 (0.147)	0.745 (0.745)	0.739 (0.103)	0.790 (0.071)
TREC	OPT 66B	0.416 (0.068)	0.419 (0.032)	0.540 (0.540)	0.599 (0.050)	0.616 (0.021)
	Llama2 70B	0.594 (0.109)	0.614 (0.102)	0.632 (0.632)	0.617 (0.113)	0.650 (0.089)
	Llama2-ch 70B	0.606 (0.095)	0.579 (0.059)	0.588 (0.588)	0.607 (0.053)	0.627 (0.071)
NLI Task						
ANLI	Orig.	0.245 (0.054)	0.264 (0.040)	0.233 (0.233)	0.261 (0.016)	0.297 (0.029)
	CC	0.497 (0.048)	0.471 (0.058)	0.444 (0.444)	0.452 (0.030)	0.518 (0.050)
	DC w/Demo	0.456 (0.021)	0.440 (0.015)	0.447 (0.447)	0.459 (0.012)	0.490 (0.012)
WNLI	DC w/Test	0.402 (0.030)	0.411 (0.037)	0.411 (0.411)	0.407 (0.036)	0.440 (0.025)
	ICC	0.861 (0.022)	0.870 (0.016)	0.857 (0.857)	0.860 (0.022)	0.870 (0.016)
	Llama2 70B	0.628 (0.070)	0.546 (0.070)	0.627 (0.627)	0.627 (0.057)	0.694 (0.055)
RTE	OPT 66B	0.595 (0.054)	0.548 (0.071)	0.556 (0.556)	0.613 (0.044)	0.647 (0.024)
	Llama2 70B	0.706 (0.193)	0.714 (0.165)	0.721 (0.721)	0.714 (0.159)	0.753 (0.105)
	Llama2-ch 70B	0.798 (0.010)	0.797 (0.013)	0.797 (0.797)	0.808 (0.008)	0.808 (0.008)
CB	OPT 66B	0.402 (0.038)	0.425 (0.079)	0.073 (0.073)	0.329 (0.137)	0.503 (0.062)
	Llama2 70B	0.617 (0.011)	0.768 (0.054)	0.750 (0.750)	0.742 (0.075)	0.791 (0.028)
	Llama2-ch 70B	0.661 (0.092)	0.670 (0.068)	0.664 (0.664)	0.679 (0.052)	0.769 (0.061)
SICK	OPT 66B	0.387 (0.104)	0.405 (0.172)	0.358 (0.358)	0.359 (0.016)	0.503 (0.057)
	Llama2 70B	0.495 (0.140)	0.463 (0.092)	0.393 (0.393)	0.446 (0.078)	0.583 (0.099)
	Llama2-ch 70B	0.458 (0.113)	0.402 (0.029)	0.419 (0.419)	0.405 (0.036)	0.512 (0.069)
Detection Task						
social_bias_frames	Orig.	0.573 (0.109)	0.676 (0.037)	0.707 (0.707)	0.700 (0.023)	0.710 (0.024)
	CC	0.363 (0.003)	0.797 (0.031)	0.709 (0.709)	0.747 (0.044)	0.798 (0.018)
	DC w/Demo	0.382 (0.019)	0.774 (0.025)	0.782 (0.782)	0.786 (0.031)	0.783 (0.031)
tweet_eval_stance_athesim	DC w/Test	0.276 (0.005)	0.277 (0.002)	0.287 (0.287)	0.250 (0.076)	0.297 (0.055)
	ICC	0.370 (0.091)	0.335 (0.113)	0.361 (0.361)	0.356 (0.088)	0.402 (0.102)
	Llama2 70B	0.315 (0.120)	0.207 (0.098)	0.167 (0.167)	0.192 (0.064)	0.355 (0.099)
tweet_eval_stance_feminist	OPT 66B	0.341 (0.092)	0.361 (0.039)	0.393 (0.393)	0.381 (0.112)	0.450 (0.069)
	Llama2 70B	0.534 (0.081)	0.470 (0.019)	0.481 (0.481)	0.477 (0.049)	0.585 (0.057)
	Llama2-ch 70B	0.594 (0.087)	0.481 (0.027)	0.475 (0.475)	0.483 (0.062)	0.600 (0.057)
tweet_eval_hate	OPT 66B	0.462 (0.123)	0.563 (0.095)	0.600 (0.600)	0.610 (0.026)	0.610 (0.040)
	Llama2 70B	0.515 (0.122)	0.373 (0.023)	0.322 (0.322)	0.394 (0.031)	0.540 (0.055)
	Llama2-ch 70B	0.557 (0.092)	0.455 (0.079)	0.525 (0.525)	0.634 (0.051)	0.605 (0.088)
tweet_eval_irony	OPT 66B	0.549 (0.128)	0.624 (0.108)	0.622 (0.622)	0.660 (0.023)	0.664 (0.012)
	Llama2 70B	0.284 (0.008)	0.516 (0.105)	0.494 (0.494)	0.520 (0.083)	0.744 (0.033)
	Llama2-ch 70B	0.285 (0.007)	0.416 (0.043)	0.567 (0.567)	0.617 (0.026)	0.643 (0.037)
tweet_eval_offensive	OPT 66B	0.600 (0.068)	0.618 (0.076)	0.638 (0.638)	0.650 (0.035)	0.661 (0.015)
	Llama2 70B	0.216 (0.007)	0.675 (0.042)	0.638 (0.638)	0.645 (0.058)	0.681 (0.028)
	Llama2-ch 70B	0.389 (0.213)	0.642 (0.037)	0.591 (0.591)	0.663 (0.021)	0.663 (0.026)
ethos_binary	OPT 66B	0.554 (0.117)	0.638 (0.065)	0.631 (0.631)	0.686 (0.028)	0.689 (0.039)
	Llama2 70B	0.816 (0.024)	0.683 (0.035)	0.557 (0.557)	0.606 (0.101)	0.819 (0.016)
	Llama2-ch 70B	0.870 (0.014)	0.820 (0.049)	0.871 (0.871)	0.861 (0.019)	0.872 (0.011)
ethos_disability	OPT 66B	0.358 (0.141)	0.363 (0.146)	0.228 (0.228)	0.342 (0.052)	0.413 (0.065)
	Llama2 70B	0.464 (0.006)	0.310 (0.194)	0.148 (0.148)	0.242 (0.133)	0.471 (0.046)
	Llama2-ch 70B	0.318 (0.082)	0.220 (0.125)	0.189 (0.189)	0.261 (0.081)	0.358 (0.055)
ethos_gender	OPT 66B	0.378 (0.091)	0.414 (0.063)	0.283 (0.283)	0.377 (0.050)	0.427 (0.072)
	Llama2 70B	0.374 (0.130)	0.265 (0.127)	0.194 (0.194)	0.231 (0.089)	0.397 (0.090)
	Llama2-ch 70B	0.342 (0.047)	0.275 (0.062)	0.311 (0.311)	0.325 (0.043)	0.381 (0.039)
ethos_national_origin	OPT 66B	0.413 (0.138)	0.398 (0.119)	0.421 (0.421)	0.463 (0.096)	0.491 (0.065)
	Llama2 70B	0.427 (0.064)	0.258 (0.120)	0.191 (0.191)	0.194 (0.049)	0.461 (0.080)
	Llama2-ch 70B	0.399 (0.062)	0.333 (0.115)	0.347 (0.347)	0.334 (0.019)	0.416 (0.064)
ethos_race	OPT 66B	0.406 (0.170)	0.407 (0.141)	0.489 (0.489)	0.487 (0.111)	0.545 (0.051)
	Llama2 70B	0.462 (0.045)	0.288 (0.118)	0.184 (0.184)	0.219 (0.097)	0.501 (0.041)
	Llama2-ch 70B	0.450 (0.100)	0.435 (0.128)	0.415 (0.415)	0.427 (0.085)	0.511 (0.100)
ethos_religion	OPT 66B	0.389 (0.142)	0.413 (0.102)	0.403 (0.403)	0.434 (0.095)	0.459 (0.075)
	Llama2 70B	0.374 (0.092)	0.280 (0.175)	0.162 (0.162)	0.214 (0.104)	0.416 (0.122)
	Llama2-ch 70B	0.319 (0.093)	0.271 (0.072)	0.206 (0.206)	0.274 (0.025)	0.327 (0.067)
ethos_violence	OPT 66B	0.414 (0.094)	0.469 (0.087)	0.444 (0.444)	0.470 (0.040)	0.517 (0.053)
	Llama2 70B	0.498 (0.080)	0.343 (0.027)	0.254 (0.254)	0.293 (0.030)	0.539 (0.106)
	Llama2-ch 70B	0.563 (0.051)	0.484 (0.106)	0.529 (0.529)	0.514 (0.068)	0.628 (0.075)
hate_speech18	OPT 66B	0.497 (0.033)	0.602 (0.046)	0.584 (0.584)	0.591 (0.063)	0.615 (0.020)
	Llama2 70B	0.617 (0.129)	0.552 (0.160)	0.385 (0.385)	0.456 (0.226)	0.642 (0.063)
	Llama2-ch 70B	0.614 (0.065)	0.591 (0.106)	0.609 (0.609)	0.646 (0.077)	0.616 (0.081)

Table 17: Macro F1-score across 27 datasets for over 50B scale LLMs under the Original ICL Task setting.

Dataset	Size	Method				
		Orig.	CC	DC w/Demo	DC w/Test	ICC
Sentiment Task						
AGNews	OPT 66B	0.172 (0.068)	0.198 (0.071)	0.157 (0.157)	0.159 (0.033)	0.336 (0.027)
	Llama2 70B	0.745 (0.113)	0.744 (0.107)	0.751 (0.751)	0.739 (0.107)	0.785 (0.065)
	Llama2-ch 70B	0.666 (0.098)	0.655 (0.080)	0.683 (0.683)	0.667 (0.089)	0.705 (0.047)
CR	OPT 66B	0.768 (0.156)	0.853 (0.047)	0.891 (0.891)	0.897 (0.037)	0.913 (0.019)
	Llama2 70B	0.720 (0.261)	0.722 (0.238)	0.803 (0.803)	0.811 (0.151)	0.805 (0.147)
	Llama2-ch 70B	0.912 (0.019)	0.864 (0.044)	0.917 (0.917)	0.916 (0.014)	0.913 (0.016)
financial_phrasebank	OPT 66B	0.371 (0.093)	0.471 (0.163)	0.542 (0.542)	0.510 (0.186)	0.566 (0.122)
	Llama2 70B	0.393 (0.131)	0.352 (0.066)	0.386 (0.386)	0.379 (0.118)	0.434 (0.107)
	Llama2-ch 70B	0.497 (0.192)	0.321 (0.125)	0.482 (0.482)	0.465 (0.162)	0.514 (0.201)
poem_sentiment	OPT 66B	0.245 (0.046)	0.211 (0.048)	0.288 (0.288)	0.272 (0.056)	0.329 (0.067)
	Llama2 70B	0.379 (0.080)	0.326 (0.104)	0.330 (0.330)	0.398 (0.024)	0.397 (0.070)
	Llama2-ch 70B	0.495 (0.127)	0.291 (0.129)	0.420 (0.420)	0.464 (0.065)	0.520 (0.083)
MR	OPT 66B	0.588 (0.270)	0.681 (0.252)	0.754 (0.754)	0.722 (0.202)	0.765 (0.187)
	Llama2 70B	0.927 (0.014)	0.897 (0.088)	0.925 (0.925)	0.940 (0.010)	0.930 (0.016)
	Llama2-ch 70B	0.910 (0.038)	0.813 (0.080)	0.906 (0.906)	0.911 (0.021)	0.918 (0.028)
sst2	OPT 66B	0.670 (0.237)	0.798 (0.121)	0.886 (0.886)	0.860 (0.110)	0.892 (0.081)
	Llama2 70B	0.910 (0.039)	0.743 (0.111)	0.840 (0.840)	0.921 (0.055)	0.916 (0.053)
	Llama2-ch 70B	0.883 (0.118)	0.794 (0.135)	0.861 (0.861)	0.906 (0.056)	0.902 (0.066)
Subj	OPT 66B	0.575 (0.131)	0.662 (0.186)	0.689 (0.689)	0.678 (0.186)	0.804 (0.085)
	Llama2 70B	0.689 (0.048)	0.699 (0.117)	0.722 (0.722)	0.640 (0.186)	0.725 (0.117)
	Llama2-ch 70B	0.677 (0.038)	0.626 (0.127)	0.678 (0.678)	0.645 (0.106)	0.690 (0.079)
TREC	OPT 66B	0.102 (0.061)	0.178 (0.035)	0.146 (0.146)	0.185 (0.015)	0.265 (0.067)
	Llama2 70B	0.376 (0.078)	0.429 (0.103)	0.447 (0.447)	0.455 (0.065)	0.470 (0.049)
	Llama2-ch 70B	0.369 (0.061)	0.395 (0.081)	0.362 (0.362)	0.407 (0.049)	0.475 (0.025)
NLI Task						
ANLI	OPT 66B	0.215 (0.046)	0.229 (0.032)	0.245 (0.245)	0.225 (0.046)	0.303 (0.046)
	Llama2 70B	0.368 (0.089)	0.367 (0.028)	0.403 (0.403)	0.389 (0.029)	0.432 (0.035)
	Llama2-ch 70B	0.381 (0.110)	0.397 (0.040)	0.443 (0.443)	0.453 (0.035)	0.463 (0.032)
WNLI	OPT 66B	0.321 (0.025)	0.451 (0.088)	0.433 (0.433)	0.450 (0.080)	0.494 (0.053)
	Llama2 70B	0.661 (0.180)	0.704 (0.176)	0.695 (0.695)	0.711 (0.169)	0.772 (0.076)
	Llama2-ch 70B	0.627 (0.108)	0.591 (0.126)	0.593 (0.593)	0.601 (0.162)	0.705 (0.061)
RTE	OPT 66B	0.421 (0.083)	0.423 (0.077)	0.422 (0.422)	0.445 (0.081)	0.464 (0.080)
	Llama2 70B	0.720 (0.140)	0.680 (0.148)	0.706 (0.706)	0.697 (0.141)	0.736 (0.135)
	Llama2-ch 70B	0.683 (0.194)	0.677 (0.187)	0.672 (0.672)	0.688 (0.194)	0.699 (0.194)
CB	OPT 66B	0.283 (0.065)	0.337 (0.052)	0.349 (0.349)	0.322 (0.096)	0.379 (0.050)
	Llama2 70B	0.664 (0.123)	0.673 (0.122)	0.679 (0.679)	0.674 (0.117)	0.721 (0.124)
	Llama2-ch 70B	0.459 (0.135)	0.393 (0.164)	0.453 (0.453)	0.485 (0.127)	0.493 (0.111)
SICK	OPT 66B	0.304 (0.072)	0.338 (0.144)	0.333 (0.333)	0.308 (0.109)	0.371 (0.120)
	Llama2 70B	0.519 (0.115)	0.519 (0.157)	0.531 (0.531)	0.515 (0.130)	0.566 (0.080)
	Llama2-ch 70B	0.392 (0.197)	0.448 (0.211)	0.452 (0.452)	0.464 (0.177)	0.497 (0.105)
Detection Task						
social_bias_frames	OPT 66B	0.477 (0.109)	0.548 (0.117)	0.544 (0.544)	0.586 (0.079)	0.599 (0.065)
	Llama2 70B	0.555 (0.062)	0.626 (0.153)	0.666 (0.666)	0.613 (0.113)	0.675 (0.083)
	Llama2-ch 70B	0.712 (0.071)	0.726 (0.042)	0.674 (0.674)	0.690 (0.053)	0.747 (0.030)
tweet_eval_stance_athesim	OPT 66B	0.252 (0.111)	0.378 (0.116)	0.424 (0.424)	0.344 (0.128)	0.413 (0.105)
	Llama2 70B	0.530 (0.135)	0.531 (0.137)	0.530 (0.530)	0.524 (0.129)	0.562 (0.145)
	Llama2-ch 70B	0.522 (0.098)	0.440 (0.056)	0.500 (0.500)	0.427 (0.050)	0.562 (0.097)
tweet_eval_stance_feminist	OPT 66B	0.317 (0.069)	0.291 (0.060)	0.343 (0.343)	0.286 (0.072)	0.353 (0.040)
	Llama2 70B	0.462 (0.089)	0.477 (0.064)	0.498 (0.498)	0.478 (0.046)	0.528 (0.082)
	Llama2-ch 70B	0.528 (0.078)	0.424 (0.035)	0.522 (0.522)	0.495 (0.067)	0.550 (0.072)
tweet_eval_hate	OPT 66B	0.448 (0.085)	0.516 (0.056)	0.509 (0.509)	0.517 (0.038)	0.534 (0.021)
	Llama2 70B	0.498 (0.044)	0.499 (0.091)	0.511 (0.511)	0.528 (0.062)	0.546 (0.023)
	Llama2-ch 70B	0.528 (0.080)	0.524 (0.101)	0.570 (0.570)	0.598 (0.040)	0.572 (0.036)
tweet_eval_irony	OPT 66B	0.365 (0.130)	0.434 (0.061)	0.421 (0.421)	0.451 (0.064)	0.476 (0.089)
	Llama2 70B	0.419 (0.102)	0.422 (0.138)	0.419 (0.419)	0.414 (0.138)	0.445 (0.126)
	Llama2-ch 70B	0.548 (0.059)	0.459 (0.107)	0.556 (0.556)	0.513 (0.113)	0.560 (0.097)
tweet_eval_offensive	OPT 66B	0.400 (0.118)	0.495 (0.048)	0.469 (0.469)	0.453 (0.054)	0.514 (0.053)
	Llama2 70B	0.516 (0.084)	0.506 (0.065)	0.512 (0.512)	0.509 (0.068)	0.525 (0.061)
	Llama2-ch 70B	0.552 (0.107)	0.507 (0.101)	0.566 (0.566)	0.542 (0.126)	0.578 (0.101)
ethos_binary	OPT 66B	0.382 (0.082)	0.542 (0.081)	0.508 (0.508)	0.540 (0.044)	0.569 (0.023)
	Llama2 70B	0.719 (0.094)	0.665 (0.075)	0.743 (0.743)	0.733 (0.035)	0.769 (0.051)
	Llama2-ch 70B	0.819 (0.060)	0.746 (0.124)	0.815 (0.815)	0.819 (0.043)	0.831 (0.059)
ethos_disability	OPT 66B	0.459 (0.077)	0.497 (0.194)	0.298 (0.298)	0.296 (0.137)	0.501 (0.039)
	Llama2 70B	0.472 (0.113)	0.436 (0.178)	0.434 (0.434)	0.406 (0.170)	0.513 (0.110)
	Llama2-ch 70B	0.615 (0.131)	0.639 (0.119)	0.556 (0.556)	0.607 (0.115)	0.640 (0.128)
ethos_gender	OPT 66B	0.481 (0.035)	0.490 (0.031)	0.423 (0.423)	0.413 (0.109)	0.501 (0.075)
	Llama2 70B	0.478 (0.027)	0.494 (0.067)	0.509 (0.509)	0.494 (0.026)	0.556 (0.043)
	Llama2-ch 70B	0.456 (0.032)	0.439 (0.061)	0.488 (0.488)	0.487 (0.072)	0.499 (0.065)
ethos_national_origin	OPT 66B	0.455 (0.000)	0.477 (0.036)	0.296 (0.296)	0.392 (0.117)	0.508 (0.045)
	Llama2 70B	0.444 (0.073)	0.344 (0.137)	0.359 (0.359)	0.375 (0.141)	0.471 (0.079)
	Llama2-ch 70B	0.512 (0.062)	0.471 (0.105)	0.493 (0.493)	0.507 (0.066)	0.522 (0.057)
ethos_race	OPT 66B	0.452 (0.000)	0.517 (0.054)	0.442 (0.442)	0.491 (0.145)	0.560 (0.050)
	Llama2 70B	0.451 (0.001)	0.476 (0.037)	0.473 (0.473)	0.520 (0.064)	0.543 (0.099)
	Llama2-ch 70B	0.602 (0.068)	0.583 (0.087)	0.590 (0.590)	0.590 (0.056)	0.612 (0.049)
ethos_religion	OPT 66B	0.448 (0.000)	0.522 (0.074)	0.422 (0.422)	0.505 (0.081)	0.540 (0.020)
	Llama2 70B	0.501 (0.069)	0.454 (0.025)	0.482 (0.482)	0.464 (0.053)	0.521 (0.096)
	Llama2-ch 70B	0.509 (0.070)	0.532 (0.095)	0.580 (0.580)	0.589 (0.112)	0.590 (0.096)
ethos_violence	OPT 66B	0.451 (0.107)	0.541 (0.098)	0.511 (0.511)	0.481 (0.057)	0.622 (0.056)
	Llama2 70B	0.492 (0.078)	0.483 (0.078)	0.464 (0.464)	0.461 (0.069)	0.510 (0.072)
	Llama2-ch 70B	0.571 (0.106)	0.596 (0.110)	0.574 (0.574)	0.604 (0.077)	0.630 (0.092)
hate_speech18	OPT 66B	0.472 (0.004)	0.504 (0.044)	0.450 (0.450)	0.491 (0.070)	0.512 (0.034)
	Llama2 70B	0.504 (0.056)	0.495 (0.029)	0.547 (0.547)	0.527 (0.048)	0.548 (0.056)
	Llama2-ch 70B	0.544 (0.104)	0.455 (0.214)	0.531 (0.531)	0.543 (0.142)	0.565 (0.118)

Table 18: Macro F1-score across 27 datasets for over 50B scale LLMs under the Task Learning setting.