

# BATCH CALIBRATION: RETHINKING CALIBRATION FOR IN-CONTEXT LEARNING AND PROMPT ENGINEERING

Han Zhou<sup>1,2,\*</sup> Xingchen Wan<sup>1</sup> Lev Prolev<sup>1</sup> Diana Mincu<sup>1</sup> Jilin Chen<sup>1</sup>

Katherine Heller<sup>1</sup> Subhrajit Roy<sup>1</sup>

<sup>1</sup>Google Research <sup>2</sup>University of Cambridge

{hzhouml, xingchenw, levp, dmincu, jilinc, kheller, subhrajitroy}@google.com

## ABSTRACT

Prompting and in-context learning (ICL) have become efficient learning paradigms for large language models (LLMs). However, LLMs suffer from prompt brittleness and various bias factors in the prompt, including but not limited to the formatting, the choice verbalizers, and the ICL examples. To address this problem that results in unexpected performance degradation, calibration methods have been developed to mitigate the effects of these biases while recovering LLM performance. In this work, we first conduct a systematic analysis of the existing calibration methods, where we both provide a unified view and reveal the failure cases. Inspired by these analyses, we propose *Batch Calibration* (BC), a simple yet intuitive method that controls the contextual bias from the batched input, unifies various prior approaches, and effectively addresses the aforementioned issues. BC is zero-shot, inference-only, and incurs negligible additional costs. In the few-shot setup, we further extend BC to allow it to *learn* the contextual bias from labeled data. We validate the effectiveness of BC with PaLM 2-(S, M, L) and CLIP models and demonstrate state-of-the-art performance over previous calibration baselines across more than 10 natural language understanding and image classification tasks.

## 1 INTRODUCTION

Prompting large language models (LLMs) (Chowdhery et al., 2022; Anil et al., 2023) has become an efficient learning paradigm for adapting LLMs to a new task by conditioning on human-designed instructions. The remarkable in-context learning (ICL) ability of LLMs also leads to efficient few-shot learners that can generalize from few-shot input-label pairs (Brown et al., 2020; Liu et al., 2023). However, the predictions of LLMs are highly sensitive and even biased to the choice of templates (Min et al., 2022b), verbalizers (Holtzman et al., 2021), and demonstrations (Liu et al., 2022a), resulting in barriers for pursuing efficiently adaptable and robust LLM applications.

Extensive research has been devoted to mitigating these biases, which we explicitly refer to the a-priori propensity of LLMs to predict certain classes over others unfairly. Lu et al. (2022) provide an analysis of the impacts of the order of ICL examples to LLMs and have explored the order selection mechanisms for ICL. On the other hand, Zhao et al. (2021) reveal the bias of language models toward certain answers and propose to calibrate the LLM given content-free tokens. More recently, Fei et al. (2023) detect the domain-label bias, and Han et al. (2023) treat the calibration of LLMs, a technique for mitigating the label bias, as learning a robust decision boundary. Though multiple calibration solutions have been provided, the field currently lacks a unified analysis that systematically distinguishes and explains the unique characteristics, merits, and downsides of each approach.

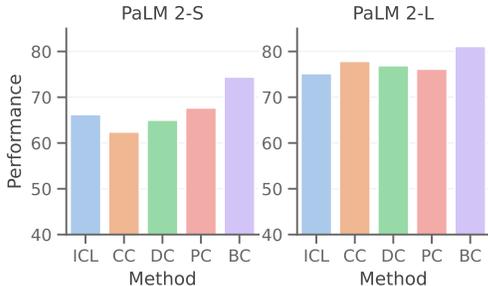


Figure 1: Batch Calibration (BC) achieves the best performance on 1-shot ICL over calibration baselines on an average of 13 classification tasks on PaLM 2-S and PaLM 2-L (Anil et al., 2023).

\*Work done as a Student Researcher at Google.

In this work, we first conduct a comprehensive analysis across existing calibration methods for LLMs. We approach the calibration problem from a distinctive point of view by interpreting the decision boundaries for each calibration method together with the ICL decision boundary. We start observing fatal failure cases for each method by extending them to more challenging and under-explored evaluation tasks. We then conclude the current limitation for each method with a novel interpretation from the decision boundary perspective, pointing to the need for a unified and widely applicable solution for conquering diverse bias sources in the field of LLM efficient learning.

Inspired by these findings, we propose *Batch Calibration* (BC), a zero-shot and inference-only calibration method for prompting and ICL. The central objective of BC is to accurately model the bias from the prompt context (referred to as *contextual bias* in this paper) by marginalizing the LLM scores in the batched input. The simplicity of the design of BC only brings negligible computation overhead at the output of the LLM. We further extend BC to the black-box few-shot learning (BCL), a practical case where labeled data is available, by introducing a *single* learnable parameter into BC, which enables it to adapt and *learn* the contextual bias from the available data resources.

We conducted extensive experiments on more than 10 natural language understanding tasks together with image classification tasks. BC stands as the most widely applicable calibration method while achieving state-of-the-art results. With the proposed black-box few-shot BCL framework, we show that further slight gains can be achieved by leveraging more labeled data. We provide further analysis with BC on robustness with templates, ICL choices and orders, and verbalizers, validating that BC can effectively alleviate prompt brittleness and make prompt engineering easier. To summarize, we provide the following contributions:

- We provide a unified and systematic analysis of existing calibration methods through their decision boundaries, investigate the common use of content-free tokens as an estimator of contextual bias, and identify their deficiency with individual case studies.
- We propose Batch Calibration (BC), a zero-shot and inference-only calibration method for ICL, that mitigates the bias from the batch. We further extend BC to learn from few-shot data.
- We show that while conceptually simple, BC attains state-of-the-art performance in both zero-shot and few-shot learning setups over widely selected tasks with PaLM-2 and CLIP models.

## 2 RELATED WORK

**Understanding and Improving ICL.** Lu et al. (2022) show the sensitivity of LLMs to ICL examples. This phenomenon is further explained through the effect from pretraining term frequencies (Razeghi et al., 2022) and corpora (Shin et al., 2022). Meanwhile, Xie et al. (2022) explain the ICL process through implicit Bayesian inference, and Wei et al. (2023) show the emergent ability of LLMs by learning new input-label mappings. Various methods have been proposed to optimally select better in-context templates (Sorensen et al., 2022; Pan et al., 2023; Yin et al., 2023) and examples (Rubin et al., 2022; Liu et al., 2022a; Wan et al., 2023b). Specifically, Wan et al. (2023a) introduce a selection criteria for in-context examples based on the confidence. Recently, noisy channel prompting (Min et al., 2022a), flipped learning (Ye et al., 2023), and self-ensembling (Li et al., 2023) have been proposed for robust ICL. Learning to assign labels by k-nearest neighbors (Xu et al., 2023) and training decoder networks (Cui et al., 2023) are also effective alternatives for few-shot ICL.

**Bias in ICL and Calibrating LLMs.** Zhao et al. (2021) reveal the instability of LLMs in few-shot learning and demonstrate three bias sources: majority label bias, recency bias, and common token bias, as the bias factors behind the instability. They propose contextual calibration (CC) to mitigate these biases by grounding the prediction based on a content-free token as sample inputs. Si et al. (2023) characterize the feature bias of LLMs, and Wang et al. (2023) introduce the positional bias in candidate choices. Fei et al. (2023) further observe the existence of domain-label bias and propose domain-context calibration (DC) that uses random in-domain tokens for estimating the bias. Meanwhile, Han et al. (2023) analyze the impact of decision boundary for text classification tasks and propose to estimate prototypical clusters by Gaussian mixture models, thereby learning a robust decision boundary. Concurrently with our work, Pezeshkpour & Hruschka (2023) spot the positional bias in multiple-choice questions, and Zheng et al. (2023) propose to debias the positional bias in multiple choices with permutation-based prior estimation. We will discuss quantitatively and provide

a unified analysis of these methods in Sec. 3. As we will show, our proposed method differentiates from these methods as a generalizable solution across challenging classification tasks and modalities.

### 3 A SYSTEMATIC ANALYSIS OF CALIBRATION

#### 3.1 BIAS IN PROMPTING AND IN-CONTEXT LEARNING (ICL)

Prompting is an efficient learning paradigm that allows LLMs to perform zero-shot inference by conditioning on a human-designed instruction. Formally, denoting a test query-target pair  $\{x_i, y_i\}$  and instruction as the context  $C$  for a classification task, LLMs make prediction by computing:  $\arg \max_{y \in \mathcal{Y}} \mathbf{p}(y|x_i, C)$ , where  $\mathbf{p} \in \mathbb{R}^J$  is the logits, and  $\mathcal{Y}$  denotes the verbalizers that define the label set for  $J$  classes. ICL further enables LLM to learn from  $k$  input-label pairs (i.e., few-shot setup),  $s^{(i)} = \text{Template}(x^{(i)}, y^{(i)}) \forall i \in \{1, \dots, k\}$ , by concatenating few-shot demonstrations in a pre-defined template as the context,  $C = \text{Concat}(s^{(1)}, \dots, s^{(k)})$ . Though ICL has demonstrated strong performance with easy implementations, the prediction of LLMs is shown to be biased towards certain answers due to different elements of  $\mathbf{p}(y|x_i, C)$  (Lu et al., 2022). In the ICL context  $C$ , majority label bias and recency label bias (Zhao et al., 2021) can bias the prediction of LLMs toward the most frequent label and the label towards the end of the demonstration, respectively. Among verbalizer tokens  $y_j \in \mathcal{Y}$ , LLMs are shown to be inherently biased towards predicting the label-tokens that appear more frequently from pretraining term statistics (Shin et al., 2022; Razeghi et al., 2022). These bias factors significantly degrade the performance of LLMs for robust ICL applications.

#### 3.2 OVERVIEW OF ICL CALIBRATION METHODS.

As introduced in Sec. 2, various *calibration* methods have been proposed to mitigate the issue of bias identified above. In this section, we provide an overview of the state-of-the-art calibration methods.

**Contextual Calibration (Zhao et al., 2021) (CC).** Motivated by a common calibration technique that applies affine transformation on the model outputs (Platt et al., 1999; Guo et al., 2017), Zhao et al. (2021) propose to calibrate the LLM prediction by first measuring the entire test-time distribution  $\hat{\mathbf{p}}$  by a content-free input. Using ‘‘N/A’’ as a content-free example, the model score distribution is generated by  $\hat{\mathbf{p}}_{cf} := \mathbf{p}(y|[N/A], C)$ . CC then generates the calibrated output by transforming the uncalibrated scores  $\mathbf{p}(y|x, C)$  with  $\mathbf{W} \in \mathbb{R}^{J \times J}$  via  $\mathbf{W}\mathbf{p}(y|x, C)$ , where  $\mathbf{W} = \text{diag}(\hat{\mathbf{p}}_{cf})^{-1}$  offsets the uncalibrated scores with the model score (a contextual prior) triggered by the content-free sample.

**Domain-Context Calibration (Fei et al., 2023) (DC).** Instead of using a single content-free token, Fei et al. (2023) propose DC that estimates a contextual prior  $\hat{\mathbf{p}}(y|C)$  by using a random in-domain sequence. It randomly sampled  $L$  tokens at an average sentence length from an unlabeled text set. Then, it estimates the content-free prediction prior by averaging the model score  $T$  times, such that:  $\hat{\mathbf{p}}_{\text{random}} = \frac{1}{T} \sum_{t=1}^T \mathbf{p}(y|[RANDOM\ TEXT]_t, C)$ . The final test-time prediction is then calibrated by dividing the estimated prior prediction, or equivalently in logits space,  $\mathbf{p}(y|x_i, C) - \hat{\mathbf{p}}_{\text{random}}$ .

**Prototypical Calibration (Han et al., 2023) (PC).** PC learns a decision boundary with Gaussian mixture models (GMMs). It estimates  $J$  prototypical clusters for the model output  $\mathbf{p}$  for  $J$  classes:  $P_{\text{GMM}}(\mathbf{p}) = \sum_{j=0}^{J-1} \alpha_j P_G(\mathbf{p}|\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)$ , where  $P_G$  denotes a multi-variate Gaussian distribution, and the parameters: mixing coefficient  $\alpha$ , mean vector  $\boldsymbol{\mu}$ , covariance matrix  $\boldsymbol{\Sigma}$  are estimated by the Expectation-Maximization (Moon, 1996). Followed by an automatic label assignment strategy, the predicted label is then computed by  $\arg \max_j P_G(\mathbf{p}_j|\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$  in the inference time. This EM-GMM process can require up to  $T$  repetitions to stabilize its estimation of clusters where  $T$  is a hyperparameter of the algorithm.

#### 3.3 DESIGN PRINCIPLES BEHIND CALIBRATIONS

Summarizing the calibration methods with distinctive design principles discussed so far, in Table 1, we present a unified view of the characteristics of each method with their mathematical formulation, computation cost, and strengths & weaknesses. Though each approach demonstrates a clear motivation for calibrating ICL, it is still unclear which method surpasses others in what scenarios.

Table 1: Calibration methods with their mathematical formulation and their equivalent decision boundary derivations in a two-dimensional problem. The cost for the number of API calls is denoted as #Forward, where 1 counts for the original ICL forward cost, and PC and BC incur no additional API costs. We refer derivations of decision boundaries to Sec. 3.3. The potential failure case for each calibration method in practical scenarios is marked as  $\times$ .

Method	Token	#Forward	Comp. Cost	Cali. Form	Learning Term	Decision Boundary $h(\mathbf{p})$	Multi-Sentence	Multi-Class
CC	N/A	1 + 1	Inverse	$\mathbf{W}\mathbf{p} + \mathbf{b}$	$\mathbf{W} = \text{diag}(\hat{\mathbf{p}})^{-1}, \mathbf{b} = \mathbf{0}$	$p_0 = \alpha p_1$	$\times$	$\checkmark$
DC	Random	20 + 1	Add	$\mathbf{W}\mathbf{p} + \mathbf{b}$	$\mathbf{W} = \mathbf{I}, \mathbf{b} = -\frac{1}{T} \sum_t \mathbf{p}(y \text{text}_t, C)$	$p_0 = p_1 + \alpha$	$\times$	$\checkmark$
PC	-	1	EM-GMM	-	$\sum_j \alpha_j P_G(\mathbf{p} \mu_j, \Sigma_j)$	$P_G(\mathbf{p} \mu_0, \Sigma_0) = P_G(\mathbf{p} \mu_1, \Sigma_1)$	$\checkmark$	$\times$
BC (Ours)	-	1	Add	$\mathbf{W}\mathbf{p} + \mathbf{b}$	$\mathbf{W} = \mathbf{I}, \mathbf{b} = -\mathbb{E}_x[\mathbf{p}(y x, C)]$	$p_0 = p_1 + \alpha$	$\checkmark$	$\checkmark$

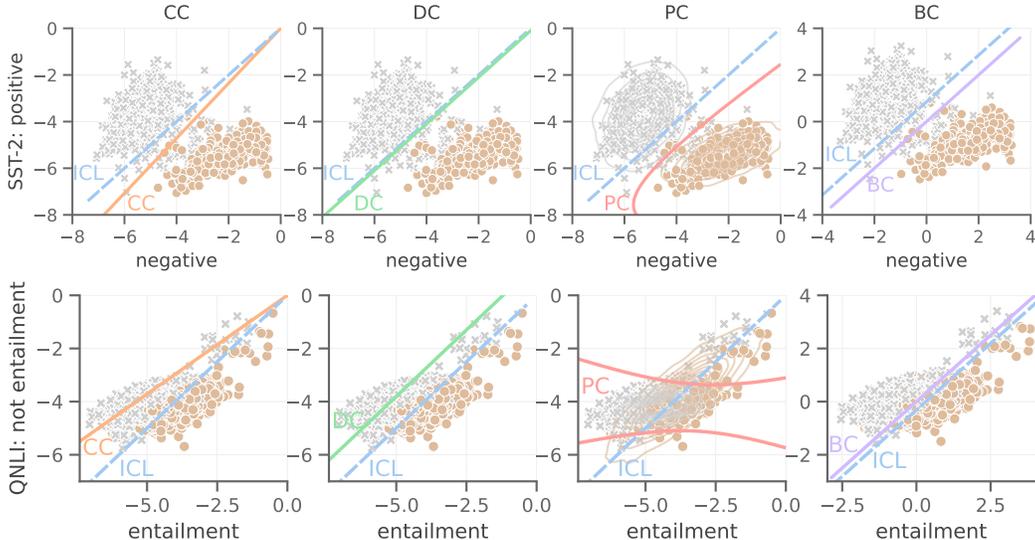


Figure 2: Visualization of the decision boundaries of uncalibrated ICL, and after applying existing calibration methods and the proposed BC (to be introduced in Sec 4) in representative binary classification tasks of SST-2 (top row) (Socher et al., 2013) and QNLI (bottom row) (Wang et al., 2018) on 1-shot PaLM 2-S. We show success and failure cases for each baseline method (CC, DC, and PC), whereas BC is consistently effective. Refer to Appendix §E for more examples.

We proceed with an in-depth analysis of existing methods in representative tasks. We provide a novel view of calibration methods from a multi-variate decision boundary perspective. In pursuit of practical guidelines for ICL calibration, we set out two important research questions behind their design principles: **1)** What constitutes a better decision boundary for calibrations? **2)** Is content-free prior a good estimator of contextual bias?

**What Constitutes a Better Decision Boundary for Calibrations?** To address this research question, we first derive the decision boundary for each category of calibration methods. We recall that the classification by a LLM is based on  $\arg \max_{j \in \{0, \dots, J-1\}} p_j$  where  $p_j$  denotes the  $j$ -th element of output vector  $\mathbf{p}$ . Consider binary classification problem for simplicity: the decision boundary  $h(\mathbf{p})$  for ICL is given by the line  $p_0 - p_1 = 0$ : the model predicts class 0,  $y_0$ , if  $p_0 - p_1 \geq 0$ , and class 1 otherwise. Consequently, CC and DC that apply an affine transformation at  $\mathbf{p}$  is equivalent to a linear transformation to the decision boundary. In CC with  $\mathbf{W} = \text{diag}(\hat{\mathbf{p}})^{-1}, \mathbf{b} = \mathbf{0}$ , the decision boundary can then be derived as  $p_0 \times \frac{1}{\hat{p}_0} = p_1 \times \frac{1}{\hat{p}_1} \rightarrow p_0 - p_1 \times \frac{\hat{p}_0}{\hat{p}_1} = 0$ , which is a *rotation* of the ICL’s linear decision boundary around the origin. Similarly, DC with  $\mathbf{W} = \mathbf{I}, \mathbf{b} = -\frac{1}{T} \sum_t \mathbf{p}(y|\text{RANDOM TEXT}_t, C) = -\hat{\mathbf{p}}$  is equivalent to a *shift* of ICL’s linear decision boundary away from the origin, such that  $p_0 - p_1 = (\hat{p}_0 - \hat{p}_1)$ . It is worth noting that both calibration choices lead to a linear decision boundary, indicating that the calibration problem can be framed as an unsupervised decision boundary learning problem. Based on this intuition, we further derive the decision boundary for PC as  $P_G(\mathbf{p}|\mu_0, \Sigma_0) - P_G(\mathbf{p}|\mu_1, \Sigma_1) = 0$ , which delivers a non-linear boundary between the estimated Gaussian mixtures. We conduct a preliminary experiment to visualize the derived decision bounds from existing calibration methods alongside the ICL baseline. In Fig. 2,

we observe that uncalibrated ICL is biased towards predicting `negative` in the SST-2 task. This biased prediction is then mitigated by each calibration method, where we observe a rotated decision boundary from CC, a shifted boundary from DC, and a non-linear boundary between the GMMs by PC. However, in the QNLI task (bottom row of Fig. 2), we observe failure cases in the calibration baselines, in particular, PC (third figure from the left), where it fails to capture the correct distribution for each class. From Fig. 2 and the additional results in Fig. 10 in Appendix §E, we find that while theoretically more flexible, the non-linear decision boundaries learned by PC tend to be susceptible to overfitting and may suffer from instability in EM-GMM. We hypothesize that the PC boundary is even more vulnerable to instability for more challenging multi-class tasks due to the increased difficulties of learning clusters and assigning classes correctly. Conversely, we find that linear decision boundaries, as evidenced by CC and DC, can be more robust and generalizable across tasks. We validate this hypothesis by proposing BC with extensive experiments in Sec. 5.2.

### Is Content-free Input a Good Estimator of the Contextual Prior?

CC and DC both use a linear decision boundary but differ from each other by leveraging different formats of a content-free input to estimate the contextual prior. However, as we observed in Fig. 2, they both exhibit failure cases in QNLI, a question-answering NLI task. We hypothesize that contrary to the proposals made by CC and DC, relying on content-free tokens for calibration is *not* always optimal and may even introduce additional bias, depending on the task type. For example, in a textual entailment task involving question-sentence pairs, we empirically observe that an ICL template employed with a content-free token ‘N/A’ such as ‘Question: N/A, Sentence: N/A, Answer:’ will result in a biased prediction towards ‘entailment’, because although ‘N/A’ is intended to be content-free, the LLM may nevertheless construe ‘N/A’ in the sentence to be substantively entailed to the ‘N/A’ in the question due to surface text equivalency. This phenomenon holds true for other multi-text classification tasks, such as paraphrasing and word disambiguation tasks. Consequently, the prior estimated via a single content-free token can lead to further bias. DC introduces multiple randomly sampled tokens to form a content-free input, e.g. ‘Question: that What old rubisco’s the did Which?’. We suspect a possible reason is that random sequences, when used in conjunction with in-context demonstrations, can be susceptible to spurious relations among them that often lead to unfair priors further skewing the predictions, which is also reflected in Fig. 3, where CC and DC fail to mitigate the contextual bias in the QNLI task. In sum, the empirical observation shows that content-free inputs can be inappropriate prior estimators, especially for multi-sentence classification tasks.

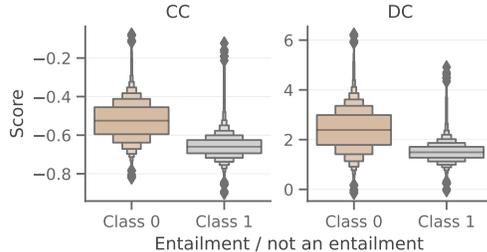


Figure 3: The distribution of ICL scores after applying CC and DC on QNLI. Due to an unfair content-free prior, the prediction by 1-shot PaLM-2 is biased towards entailment.

## 4 BATCH CALIBRATION

Inspired by the previous discussions, we now propose Batch Calibration (BC), a zero-shot, inference-only, and generalizable calibration technique with negligible computation cost. We further discuss how BC can be extended to *vision*-language models as well as the black-box *few-shot* learning setup.

**Batch Calibration (BC).** Following the discussion in Sec. 3.3, we argue that the most critical component for calibration is to accurately estimate the contextual bias term  $\mathbf{p}(y|C)$ . Both CC and DC, which use content-free and in-domain random tokens as trigger signals, respectively, have failure cases in multi-sentence classification when the estimation of the contextual bias is inaccurate. On the other hand, PC is vulnerable to overfitting and may incorrectly model the mixtures, especially in high-dimensional space. We, therefore, opt for a linear decision boundary for its robustness, and instead of relying on content-free tokens, we propose to estimate the contextual bias for each class  $\mathbf{p}(y = y_j|C)$  from a batch with  $M$  samples,  $\{x^1, \dots, x^M\}$ , in a *content-based* manner by marginalizing the output score over all samples  $x \sim P(x)$  within the batch:

$$\mathbf{p}(y|C)_j = \mathbb{E}_{x \sim P(x)} \left[ \mathbf{p}(y = y_j|x, C) \right] \approx \frac{1}{M} \sum_{i=1}^M \mathbf{p}(y = y_j|x^{(i)}, C) \forall y_j \in \mathcal{Y}. \quad (1)$$

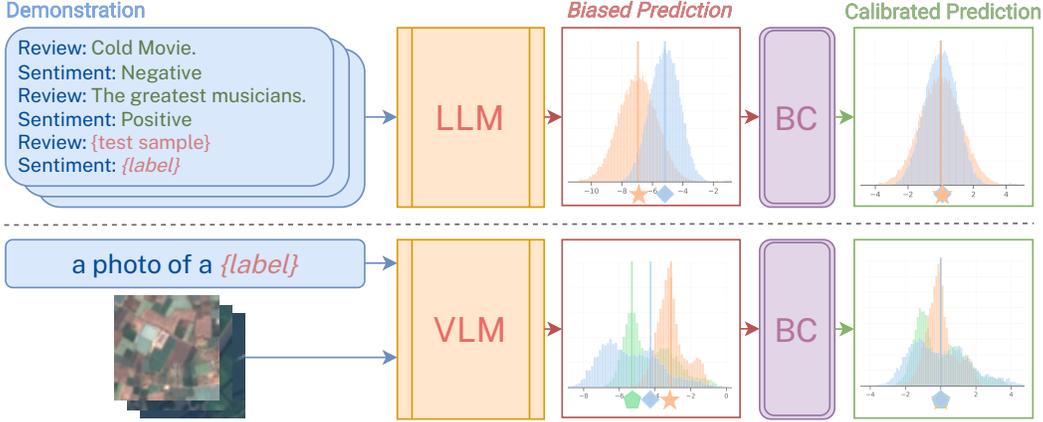


Figure 4: Illustration of Batch Calibration (BC). Batches of demonstrations with in-context examples and test samples are passed into the LLM. Due to implicit bias sources in the context, the score distribution from the LLM becomes highly biased. BC is a modular and adaptable layer option appended to the output of the LLM/VLM. BC generates calibrated scores according to Eq. 1 & 2. Highlighted symbols indicate the distribution means (visualized for illustration only).

We then obtain the calibrated probability by dividing the output probability over the contextual prior, which is equivalently by shifting the log-probability by the estimated mean of each class:

$$\hat{y}_i = \arg \max_{y \in \mathcal{Y}} \mathbf{p}_{\text{BC}}(y|x_i, C) = \arg \max_{y \in \mathcal{Y}} [\mathbf{p}(y|x_i, C) - \hat{\mathbf{p}}(y|C)]. \quad (2)$$

It is noteworthy that this BC procedure is zero-shot and only involves unlabeled test samples. BC incurs negligible computation costs. We may either compute the correction term  $\hat{\mathbf{p}}(y|C)$  once all test samples are seen or, alternatively, in an on-the-fly manner that dynamically processes the outputs. To do so, we may use a running estimate of the contextual bias for BC. At the  $n + 1$  mini-batch, the bias term is given by:

$$\mathbf{p}_r^{n+1}(y|C) = \frac{n}{n+1} \mathbf{p}_r^n(y|C) + \frac{1}{n+1} \hat{\mathbf{p}}^{n+1}(y|C), \quad (3)$$

thereby allowing BC’s calibration term to be estimated from a small number of mini-batches that is subsequently stabilized when more mini-batches arrive.

**Adjustable Batch Calibration Layer (BCL).**

While BC is designed to be zero-shot and inference-only, it is also common that some *labeled* data are available. In this section, we describe a simple, adapted variant of BC that may further refine the calibration and mitigate any estimation errors from the unlabeled data, which we term *BCL*. Specifically, instead of deducting the bias term  $\hat{\mathbf{p}}$  from the test data only, we introduce a single additional hyperparameter *strength*  $\gamma \in \mathbb{R}$ :

$$\mathbf{p}_{\text{BCL}}(y|x_i, C) = \mathbf{p}(y|x_i, C) - \gamma \hat{\mathbf{p}}(y|C), \quad (4)$$

where  $\gamma$  controls the strength of BC. To select the appropriate  $\gamma$ , we simply perform a grid search by uniformly sampling  $T$  different  $\gamma$  values in  $[a, b]$  (we set  $[a, b] := [-5, 5]$ , but any reasonable range may be used). The strength  $\gamma$  is then learned by  $\gamma^* = \arg \max_{\gamma \in [a, b]} R(\mathbf{p}_{\text{BC}}, \gamma)$ , where  $R(\cdot, \cdot)$  is the evaluation function (e.g., accuracy) on the set of *labeled* data, allowing the amount of calibration to be adjusted from evaluation metrics directly.

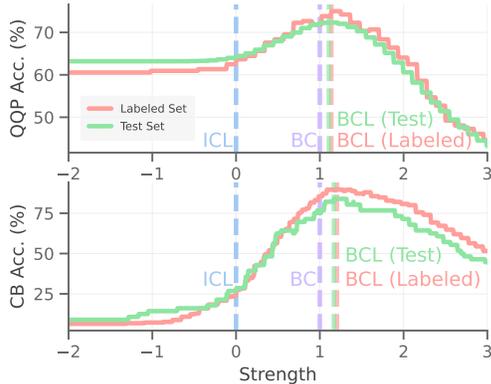


Figure 5: *BC benefits from labeled data:* The performance of the adjustable BCL compared to the zero-shot BC with a changing strength. The strength  $\gamma$  at 0 and 1 represent the uncalibrated ICL and BC, respectively. We highlight the optimal strength learned from a labeled set and the best test strength. Refer to Appendix §E for more examples.

Table 2: Accuracy (%) on natural language classification tasks with 1-shot PaLM 2-S and PaLM 2-L Models (Anil et al., 2023). We report the mean and standard deviation for all results for 5 different in-context examples. We reproduce all baselines, and the implementation details are described in Appendix §D. The **best** and **second-best** results are marked in bold fonts and ranked by color.

Model Method	PaLM 2-S					PaLM 2-L				
	ICL	CC	DC	PC	BC	ICL	CC	DC	PC	BC
SST-2	93.62 <sub>0.62</sub>	<b>95.50</b> <sub>0.25</sub>	94.29 <sub>0.32</sub>	<b>95.71</b> <sub>0.10</sub>	95.44 <sub>0.15</sub>	93.16 <sub>5.18</sub>	<b>95.82</b> <sub>0.62</sub>	94.91 <sub>2.01</sub>	95.64 <sub>0.47</sub>	<b>95.78</b> <sub>0.55</sub>
MNLI	<b>68.52</b> <sub>7.98</sub>	60.07 <sub>11.26</sub>	63.45 <sub>1.99</sub>	59.29 <sub>13.79</sub>	<b>75.12</b> <sub>2.76</sub>	72.77 <sub>3.65</sub>	<b>79.45</b> <sub>3.46</sub>	71.53 <sub>4.86</sub>	78.68 <sub>7.10</sub>	<b>81.34</b> <sub>2.29</sub>
QNLI	<b>81.20</b> <sub>1.90</sub>	56.86 <sub>3.29</sub>	65.62 <sub>3.53</sub>	69.82 <sub>17.73</sub>	<b>82.45</b> <sub>1.82</sub>	64.68 <sub>3.53</sub>	<b>69.71</b> <sub>4.89</sub>	68.97 <sub>3.27</sub>	61.01 <sub>15.26</sub>	<b>87.90</b> <sub>1.24</sub>
MRPC	66.42 <sub>10.15</sub>	<b>70.44</b> <sub>0.94</sub>	68.58 <sub>0.21</sub>	<b>71.86</b> <sub>1.29</sub>	70.05 <sub>2.40</sub>	<b>73.19</b> <sub>1.21</sub>	72.40 <sub>3.53</sub>	68.68 <sub>0.40</sub>	<b>75.39</b> <sub>2.60</sub>	70.39 <sub>2.56</sub>
QQP	63.91 <sub>0.66</sub>	<b>65.55</b> <sub>5.34</sub>	53.92 <sub>9.35</sub>	65.28 <sub>3.42</sub>	<b>71.48</b> <sub>1.46</sub>	<b>82.57</b> <sub>0.75</sub>	81.17 <sub>2.03</sub>	78.32 <sub>1.82</sub>	<b>81.42</b> <sub>0.24</sub>	79.56 <sub>1.40</sub>
BoolQ	83.99 <sub>3.90</sub>	87.14 <sub>1.60</sub>	87.64 <sub>1.10</sub>	<b>88.70</b> <sub>0.15</sub>	<b>87.83</b> <sub>0.10</sub>	90.02 <sub>0.60</sub>	<b>90.15</b> <sub>0.54</sub>	87.77 <sub>1.17</sub>	64.40 <sub>22.37</sub>	<b>90.10</b> <sub>0.22</sub>
CB	45.71 <sub>10.61</sub>	29.64 <sub>7.85</sub>	65.71 <sub>3.20</sub>	<b>81.07</b> <sub>9.42</sub>	<b>78.21</b> <sub>3.19</sub>	<b>92.86</b> <sub>2.19</sub>	85.72 <sub>7.78</sub>	<b>92.86</b> <sub>2.82</sub>	89.29 <sub>7.25</sub>	<b>93.21</b> <sub>1.49</sub>
COPA	<b>96.40</b> <sub>2.30</sub>	95.80 <sub>2.05</sub>	<b>96.40</b> <sub>2.88</sub>	<b>96.20</b> <sub>2.05</sub>	<b>96.40</b> <sub>2.07</sub>	<b>98.60</b> <sub>1.14</sub>	97.20 <sub>1.10</sub>	97.40 <sub>0.89</sub>	<b>99.00</b> <sub>0.71</sub>	97.00 <sub>1.00</sub>
RTE	<b>80.94</b> <sub>1.29</sub>	79.78 <sub>0.92</sub>	76.82 <sub>1.72</sub>	80.43 <sub>1.07</sub>	<b>83.47</b> <sub>1.10</sub>	75.09 <sub>2.11</sub>	80.00 <sub>2.48</sub>	79.21 <sub>1.95</sub>	<b>86.64</b> <sub>2.62</sub>	<b>85.42</b> <sub>2.48</sub>
WiC	50.69 <sub>0.59</sub>	50.56 <sub>0.50</sub>	49.97 <sub>0.13</sub>	<b>51.38</b> <sub>3.56</sub>	<b>61.10</b> <sub>2.07</sub>	51.35 <sub>1.90</sub>	55.58 <sub>6.38</sub>	54.67 <sub>6.02</sub>	<b>57.87</b> <sub>11.08</sub>	<b>64.83</b> <sub>8.59</sub>
ANLI-R1	<b>46.24</b> <sub>4.21</sub>	42.54 <sub>3.20</sub>	40.26 <sub>3.66</sub>	40.28 <sub>6.46</sub>	<b>59.82</b> <sub>0.51</sub>	63.06 <sub>2.63</sub>	71.92 <sub>3.71</sub>	<b>73.56</b> <sub>3.88</sub>	72.30 <sub>8.05</sub>	<b>75.00</b> <sub>3.03</sub>
ANLI-R2	40.44 <sub>0.90</sub>	38.36 <sub>0.82</sub>	38.44 <sub>3.46</sub>	<b>41.88</b> <sub>4.50</sub>	<b>50.16</b> <sub>0.82</sub>	58.40 <sub>1.19</sub>	65.36 <sub>3.75</sub>	<b>65.48</b> <sub>1.91</sub>	64.98 <sub>2.94</sub>	<b>67.30</b> <sub>2.34</sub>
ANLI-R3	42.53 <sub>0.99</sub>	38.78 <sub>1.04</sub>	<b>43.67</b> <sub>5.25</sub>	37.50 <sub>0.81</sub>	<b>55.75</b> <sub>1.66</sub>	61.35 <sub>3.14</sub>	<b>67.32</b> <sub>0.98</sub>	66.23 <sub>0.72</sub>	63.03 <sub>6.03</sub>	<b>66.38</b> <sub>0.74</sub>
Avg.	66.20	62.39	64.98	<b>67.65</b>	<b>74.41</b>	75.16	<b>77.83</b>	76.89	76.13	<b>81.09</b>

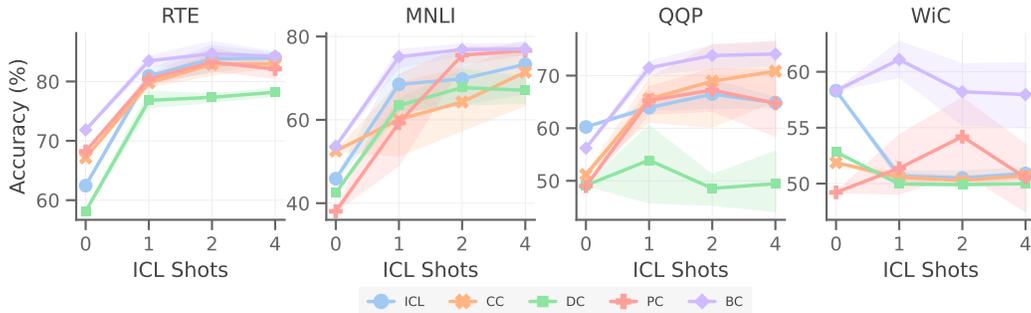


Figure 6: The ICL performance on various calibration techniques over the number of ICL shots on PaLM 2-S. Each shot indicates 1 example per class in the demonstration. Lines and shades denote the mean and standard deviation over 5 random seeds, respectively.

We give concrete examples in Fig. 5, which illustrates the effect of BCL where we plot the accuracy in QQP (Wang et al., 2018) and CB (De Marneffe et al., 2019) tasks over a range of  $\gamma$ . We observe that  $\gamma = 1$ , which corresponds to BC without adjustment (purple line), leads to a strong but not optimal performance. By using the  $\gamma$  learned from the labeled data (a 128-shot randomly sampled set in this case), BCL estimates the contextual bias more precisely by leveraging the labeled data and achieves a performance that is very close to the optimal. We refer readers to Table 3 for more results.

**Calibrating Vision-Language Models.** Recently, vision-language models (VLM) (Radford et al., 2021), which simultaneously encode visual and textual information, have demonstrated strong zero-shot generalization capability by rewriting class labels. However, the sources of bias as LLMs have also been observed in prompting VLMs (Alayrac et al., 2022) but have not been adequately addressed. In this work, we propose to apply BC to Zero-Shot (ZS) CLIP (Radford et al., 2021) and mitigate the biases in its zero-shot classifications. We follow the same notation from Sec. 3.1, where the test image is now  $x$ , and the prompt template becomes the context,  $C$ . Similarly, we append the BC layer at the output of the ZS CLIP and calibrate for each class following Eq. 1 & 2.

## 5 EXPERIMENTS

### 5.1 EXPERIMENTAL SETUP

**Evaluation Data.** For natural language tasks, in contrast to previous works that only report on relatively simple single-sentence classification tasks (Zhao et al., 2021; Fei et al., 2023; Han et al., 2023), we conduct experiments on 13 more diverse and challenging classification tasks, including the

Table 3: Accuracy (%) on natural language classification tasks with the zero-shot BC and the BCL that involves an optimal strength term learned from a labeled set. The experiments are evaluated with the same in-context example on 1-shot PaLM 2-S.

Method	SST-2	MNLI	QNLI	MRPC	QQP	BoolQ	CB	COPA	RTE	WiC	ANLI <sub>R1</sub>	ANLI <sub>R2</sub>	ANLI <sub>R3</sub>	Avg.
BC	95.4	75.0	83.5	68.6	70.3	87.9	75.0	98.0	<b>84.1</b>	<b>63.3</b>	<b>59.8</b>	<b>51.1</b>	<b>53.3</b>	74.3
BCL	<b>96.3</b>	<b>75.0</b>	<b>83.5</b>	<b>74.3</b>	<b>72.3</b>	<b>88.8</b>	<b>83.9</b>	<b>99.0</b>	82.7	63.2	58.0	49.7	52.2	<b>75.3</b>

standard GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019) datasets. Specifically, we consider commonsense reasoning: BoolQ (Clark et al., 2019), COPA (Roemmele et al., 2011); word disambiguation: WiC (Pilehvar & Camacho-Collados, 2019); sentiment classification: SST-2 (Socher et al., 2013); paraphrasing: QQP, MRPC (Dolan & Brockett, 2005); natural language inference and entailment: ANLI-R $\{1,2,3\}$  (Nie et al., 2020), CB (De Marneffe et al., 2019), RTE, QNLI (QA/NLI), MNLI (Williams et al., 2018). For image classification tasks, we include SVHN (Yuval, 2011), EuroSAT (Helber et al., 2019), and CLEVR (Johnson et al., 2017).

**Models.** We conduct experiments mainly on the state-of-the-art PaLM 2 (Anil et al., 2023) for its variants with different sizes, PaLM 2-S, PaLM 2-M, and PaLM 2-L. PaLM 2 has achieved results that are competitive with GPT-4 (OpenAI, 2023), and readers are referred to Anil et al. (2023) for more model details. For VLMs, we report the results on CLIP ViT-B/16 (Radford et al., 2021).

## 5.2 MAIN EXPERIMENTS

**Experiments on Natural Language Tasks.** We present the results across a diverse set of NLP tasks in Table 2. Notably, BC consistently outperforms ICL, yielding significant performance enhancement of 8% and 6% on PaLM 2-S and PaLM 2-L, respectively. This shows that the BC implementation successfully mitigates the contextual bias from the in-context examples and unleashes the full potential of LLM in efficient learning and quick adaptation to new tasks. In addition, BC improves over the state-of-the-art PC baseline by 6% on PaLM 2-S, and surpasses the competitive CC baseline by another 3% on average on PaLM 2-L. Specifically, BC is a generalizable and cheaper technique across all evaluated tasks, delivering stable performance improvement, whereas previous baselines exhibit varying degrees of instability across tasks: DC baseline is the least competitive; CC displays more failure cases in multi-sentence classification tasks, particularly for paraphrasing and NLI tasks, as we hypothesized in Sec 3.3; PC, while occasionally competitive, exhibits large performance fluctuations, as evidenced by its large standard deviation, resulting in frequent substantial performance degradation.

We further analyze the performance of BC by varying the ICL shots from 0 to 4 shots as shown in Fig. 6, and BC again outperforms all baseline methods. We also observe an overall trend for improved performance when more shots are available, and the performance disparities between BC and ICL converge on some tasks, which suggests that BC allows LLMs to more effectively take advantage of more in-context demonstrations. We also observe that PC exhibits the worst stability across the number of shots. In Table 3, we show that further slight gains, 1% on average, can be achieved by involving an adjustable strength parameter that refines the calibration and mitigates estimation errors. This alternative design not only makes BC applicable to both zero-shot and few-shot setups but also shows its capability to improve further from limited labeled data.

**Calibrating Vision-Language Models** We further extend the applicable scenarios of BC to multi-modal learning to handle the bias inherent in the prompt template designs in CLIP. We select three tasks in which the previous visual-prompting method shows significant improvement (Oh et al., 2023). As shown in Fig. 7, BC significantly im-

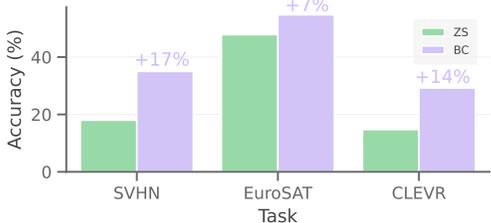


Figure 7: BC improves zero-shot (ZS) image classification: Accuracy (%) on image classification tasks with the zero-shot CLIP ViT-16/B. The BC implementation is zero-shot, and we apply BC together with the CLIP to demonstrate the effectiveness of BC in vision-language models. Refer to additional tasks in Appendix §E.

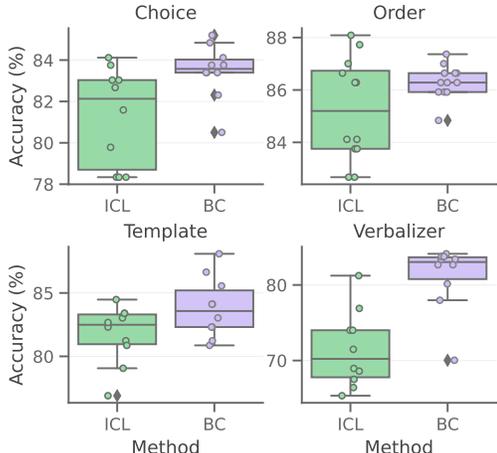


Figure 8: *BC makes prompt engineering easier:* Performance of BC with respect to ICL choices, ICL orders, prompt templates, and verbalizers.

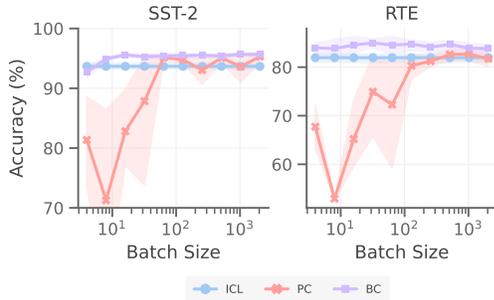


Figure 9: *BC is data-efficient and insensitive to the batch size:* Performance of BC across different sizes of an initial unlabeled set without using a running estimate of the contextual bias. We compare BC with the state-of-the-art PC baseline that also leverages unlabeled estimate set, and experiments are conducted on PaLM 2-S.

proves the zero-shot baseline by 12% on average. This observation further highlights the presence of contextual bias even within vision-language models, and BC can successfully restore the performance of VLM in image classification tasks, suggesting that BC may serve as a versatile and common technique for mitigating contextual biases across multiple modalities.

### 5.3 ROBUSTNESS AND ABLATION STUDIES

**Robustness.** We analyze the robustness of BC with respect to common prompt engineering design choices that were previously shown to significantly affect LLM performance (Lu et al., 2022; Liu et al., 2022a): choices and orders of in-context examples, the prompt template for ICL, and the verbalizers, as shown in Fig. 8 evaluated on RTE. Setup details are listed in Appendix §F. First, we find that BC is more robust to ICL choices and can mostly achieve the same performance with different ICL examples. Additionally, given a single set of ICL shots, altering the order between each ICL example has minimal impact on the BC performance. However, it is worth noting that an optimal order selection can still lead to promising ICL performance. Furthermore, we analyze the robustness of BC under 10 designs of prompt templates, where BC shows consistent improvement over the ICL baseline. Therefore, though BC makes further improvements, a well-designed template can further enhance the performance of BC. Lastly, we examine the robustness of BC to variations in verbalizer designs. Remarkably, even when employing unconventional choices such as emoji pairs as the verbalizers (listed in Tables 7 & 10) leading to dramatic oscillations of ICL performance, BC largely recovers performance. This observation shows BC robustifies LLM predictions under common prompt design choices and makes prompt engineering easier.

**Batch Size.** We study the impact of batch size on the performance of BC as shown in Fig. 9. In contrast to PC, which also leverages an unlabeled estimate set, BC is remarkably more sample efficient, achieving a strong performance with only around 10 unlabeled samples, whereas PC requires more than 500 unlabeled samples before its performance stabilizes.

## 6 CONCLUSION

We first revisit previous calibration methods while addressing two critical research questions from an interpretation of decision boundaries, revealing their failure cases and deficiencies. We then propose Batch Calibration, a zero-shot and inference-only calibration technique. We also introduce an adjustable extension, BCL, which offers more refined calibrations when labeled data is accessible. While methodologically simple and easy to implement with negligible computation cost, we show that BC scales from a language-only setup to the vision-language context, achieving state-of-the-art performance in both modalities. BC significantly improves the robustness of LLMs with respect to prompt designs, and we expect easy prompt engineering with BC while exploring the potential of BC to generative tasks in the future.

## ACKNOWLEDGEMENT

We thank Emily Salkey for her sincere project management support. We also thank Mohammad Havaei, Chirag Nagpal, Stephen Pfohl, Alexander D’Amour, and Ahmad Beirami for fruitful suggestions and feedbacks and the PaLM 2 team at Google for helping with infrastructure questions.

## REFERENCES

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L. Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022.
- Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.
- Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pp. 3606–3613, 2014.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. BoolQ: Exploring the surprising difficulty of natural yes/no questions. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 2924–2936, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- Ganqu Cui, Wentao Li, Ning Ding, Longtao Huang, Zhiyuan Liu, and Maosong Sun. Decoder tuning: Efficient language understanding as decoding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15072–15087, Toronto, Canada, July 2023. Association for Computational Linguistics.
- Marie-Catherine De Marneffe, Mandy Simons, and Judith Tonhauser. The commitmentbank: Investigating projection in naturally occurring discourse. In *proceedings of Sinn und Bedeutung*, volume 23, pp. 107–124, 2019.
- Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric Xing, and Zhiting Hu. RLPrompt: Optimizing discrete text prompts with reinforcement learning. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 3369–3391, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.
- Shizhe Diao, Zhichao Huang, Ruijia Xu, Xuechun Li, LIN Yong, Xiao Zhou, and Tong Zhang. Black-box prompt learning for pre-trained language models. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

- William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005.
- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. Mitigating label biases for in-context learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14014–14031, Toronto, Canada, July 2023. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pp. 1321–1330, 2017.
- Zhixiong Han, Yaru Hao, Li Dong, Yutao Sun, and Furu Wei. Prototypical calibration for few-shot learning of language models. In *The Eleventh International Conference on Learning Representations*, 2023.
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 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*, pp. 7038–7051, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pp. 2790–2799, 2019.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- Xuefeng Hu, Gokhan Uzunbas, Sirius Chen, Rui Wang, Ashish Shah, Ram Nevatia, and Ser-Nam Lim. Mixnorm: Test-time adaptation through online normalization estimation. *arXiv preprint arXiv:2110.11478*, 2021.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 1988–1997, 2017.
- Ansh Khurana, Sujoy Paul, Piyush Rai, Soma Biswas, and Gaurav Aggarwal. Sita: Single image test-time adaptation. *arXiv preprint arXiv:2112.02355*, 2021.
- Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- Chengzu Li, Han Zhou, Goran Glavaš, Anna Korhonen, and Ivan Vulić. On task performance and model calibration with supervised and self-ensembled in-context learning. *arXiv preprint arXiv:2312.13772*, 2023.
- Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 4582–4597, Online, August 2021. Association for Computational Linguistics.

- Hyesu Lim, Byeongeun Kim, Jaegul Choo, and Sungha Choi. TTN: A domain-shift aware batch normalization in test-time adaptation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*, 2023.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. What makes good in-context examples for GPT-3? In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pp. 100–114, Dublin, Ireland and Online, May 2022a. Association for Computational Linguistics.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35, 2023.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. P-tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 61–68, Dublin, Ireland, May 2022b. Association for Computational Linguistics.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 8086–8098, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*, 2013.
- Sewon Min, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Noisy channel language model prompting for few-shot text classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5316–5330, Dublin, Ireland, May 2022a. Association for Computational Linguistics.
- Sewon Min, Xinxin Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 11048–11064, Abu Dhabi, United Arab Emirates, December 2022b. Association for Computational Linguistics.
- Muhammad Jehanzeb Mirza, Jakub Micorek, Horst Possegger, and Horst Bischof. The norm must go on: Dynamic unsupervised domain adaptation by normalization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 14745–14755, 2022.
- Todd K Moon. The expectation-maximization algorithm. *IEEE Signal processing magazine*, 13(6): 47–60, 1996.
- Zachary Nado, Shreyas Padhy, D Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *arXiv preprint arXiv:2006.10963*, 2020.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. Adversarial NLI: A new benchmark for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4885–4901, Online, July 2020. Association for Computational Linguistics.
- Changdae Oh, Hyeji Hwang, Hee Young Lee, YongTaek Lim, Geunyoung Jung, Jiyoung Jung, Hosik Choi, and Kyungwoo Song. Blackvip: Black-box visual prompting for robust transfer learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pp. 24224–24235, 2023.
- OpenAI. Gpt-4 technical report, 2023.

- Jane Pan, Tianyu Gao, Howard Chen, and Danqi Chen. What in-context learning “learns” in-context: Disentangling task recognition and task learning. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 8298–8319, Toronto, Canada, July 2023. Association for Computational Linguistics.
- Pouya Pezeshkpour and Estevam Hruschka. Large language models sensitivity to the order of options in multiple-choice questions. *arXiv preprint arXiv:2308.11483*, 2023.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. WiC: the word-in-context dataset for evaluating context-sensitive meaning representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 1267–1273, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- John Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- Archiki Prasad, Peter Hase, Xiang Zhou, and Mohit Bansal. GrIPS: Gradient-free, edit-based instruction search for prompting large language models. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 3845–3864, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pp. 8748–8763, 2021.
- Yasaman Razeghi, Robert L Logan IV, Matt Gardner, and Sameer Singh. Impact of pretraining term frequencies on few-shot numerical reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 840–854, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S. Gordon. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, Technical Report SS-11-06, Stanford, California, USA, March 21-23, 2011*, 2011.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. Learning to retrieve prompts for in-context learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2655–2671, Seattle, United States, July 2022. Association for Computational Linguistics.
- Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Seongjin Shin, Sang-Woo Lee, Hwijeeen Ahn, Sungdong Kim, HyoungSeok Kim, Boseop Kim, Kyunghyun Cho, Gichang Lee, Woomyoung Park, Jung-Woo Ha, and Nako Sung. On the effect of pretraining corpora on in-context learning by a large-scale language model. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5168–5186, Seattle, United States, July 2022. Association for Computational Linguistics.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4222–4235, Online, November 2020. Association for Computational Linguistics.
- Chenglei Si, Dan Friedman, Nitish Joshi, Shi Feng, Danqi Chen, and He He. Measuring inductive biases of in-context learning with underspecified demonstrations. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 11289–11310, Toronto, Canada, July 2023. Association for Computational Linguistics.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics.
- Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.
- Taylor Sorensen, Joshua Robinson, Christopher Rytting, Alexander Shaw, Kyle Rogers, Alexia Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. An information-theoretic approach to prompt engineering without ground truth labels. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 819–862, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- Xingchen Wan, Ruoxi Sun, Hanjun Dai, Serkan Arik, and Tomas Pfister. Better zero-shot reasoning with self-adaptive prompting. In *Findings of the Association for Computational Linguistics: ACL 2023*, pp. 3493–3514, Toronto, Canada, July 2023a. Association for Computational Linguistics.
- Xingchen Wan, Ruoxi Sun, Hootan Nakhost, Hanjun Dai, Julian Eisenschlos, Serkan Arik, and Tomas Pfister. Universal self-adaptive prompting. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 7437–7462, Singapore, December 2023b. Association for Computational Linguistics.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, Brussels, Belgium, November 2018. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. Superglue: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 3261–3275, 2019.
- Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.
- Peiyi Wang, Lei Li, Liang Chen, Dawei Zhu, Binghuai Lin, Yunbo Cao, Qi Liu, Tianyu Liu, and Zhifang Sui. Large language models are not fair evaluators. *arXiv preprint arXiv:2305.17926*, 2023.
- Jerry Wei, Jason Wei, Yi Tay, Dustin Tran, Albert Webson, Yifeng Lu, Xinyun Chen, Hanxiao Liu, Da Huang, Denny Zhou, et al. Larger language models do in-context learning differently. *arXiv preprint arXiv:2303.03846*, 2023.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022, 2022*.
- Benfeng Xu, Quan Wang, Zhendong Mao, Yajuan Lyu, Qiaoqiao She, and Yongdong Zhang. \$k\$-nearest neighbor prompting: Beyond-context learning with calibration-free nearest neighbor inference. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023, 2023*.
- Seonghyeon Ye, Doyoung Kim, Joel Jang, Joongbo Shin, and Minjoon Seo. Guess the instruction! flipped learning makes language models stronger zero-shot learners. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023, 2023*.

- Fan Yin, Jesse Vig, Philippe Laban, Shafiq Joty, Caiming Xiong, and Chien-Sheng Wu. Did you read the instructions? rethinking the effectiveness of task definitions in instruction learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3063–3079, Toronto, Canada, July 2023. Association for Computational Linguistics.
- Fuming You, Jingjing Li, and Zhou Zhao. Test-time batch statistics calibration for covariate shift. *arXiv preprint arXiv:2110.04065*, 2021.
- Netzer Yuval. Reading digits in natural images with unsupervised feature learning. In *Proceedings of the NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- Tianjun Zhang, Xuezhi Wang, Denny Zhou, Dale Schuurmans, and Joseph E. Gonzalez. TEMPERA: Test-time prompt editing via reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pp. 12697–12706, 2021.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. On large language models’ selection bias in multi-choice questions. *arXiv preprint arXiv:2309.03882*, 2023.
- Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korhonen. Survival of the most influential prompts: Efficient black-box prompt search via clustering and pruning. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 13064–13077, Singapore, December 2023. Association for Computational Linguistics.
- Han Zhou, Xingchen Wan, Ivan Vulić, and Anna Korhonen. Autopeft: Automatic configuration search for parameter-efficient fine-tuning. *Transactions of the Association for Computational Linguistics*, 2024.
- Yuliang Zou, Zizhao Zhang, Chun-Liang Li, Han Zhang, Tomas Pfister, and Jia-Bin Huang. Learning instance-specific adaptation for cross-domain segmentation. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXIII*, pp. 459–476, 2022.

## A ADDITIONAL RELATED WORK

**Prompt Learning.** Prompt learning is an efficient learning pipeline for LLM as an alternative to traditional full-model fine-tuning (Liu et al., 2023). Soft prompting (Li & Liang, 2021; Lester et al., 2021; Liu et al., 2022b) enables fast adaptation of LLM by appending learnable continuous prompts in the embedding space while freezing the rest of weights. The recent development of parameter-efficient fine-tuning methods (Houlsby et al., 2019; Hu et al., 2022; Zhou et al., 2024), which learn additional modules, may also be interpreted as a form of soft prompting (He et al., 2022). However, these soft prompt learning methods inevitably require gradients and internal model access. On the other hand, hard prompting (Shin et al., 2020) is an appealing learning category for learning discrete prompts. Recent efforts have been devoted to black-box prompt search without accessing model gradients, and more interpretable prompts can be found by reinforcement learning (Deng et al., 2022; Zhang et al., 2023), gradient estimation (Diao et al., 2023), and other derivative-free search algorithms (Prasad et al., 2023; Zhou et al., 2023).

**Test-Time Adaptation.** Test-time adaptation aims to mitigate the covariate shift using the test-time statistics. Wang et al. propose to minimize the entropy by updating the affine parameters in the BN layer. Nado et al. (2020) and Schneider et al. (2020) introduce using test-time batch statistics for the standardization in BN and mixing it with source statistics to conquer covariate shift, respectively. Similarly, mixing the statistics with predefined hyperparameters (You et al., 2021; Khurana et al., 2021), interpolating source and target-domain statistics (Lim et al., 2023), or using a running average estimate (Mirza et al., 2022; Hu et al., 2021) have also been proposed to adapt the BN layer. Zou et al. (2022) introduce strength parameters in adapting the standardization statistics in semantic

segmentation tasks. We differentiate from test-time BN approaches by mitigating the bias in the novel context of LLM, and there is no source statistic similar to a BN layer in computer vision backbones.

## B LIMITATIONS AND FUTURE WORK

BC is a test-time method that relies on the target statistics from batched input. To mitigate any potential estimation errors from unlabelled data, we introduce the adjustable BCL extension to incorporate source statistics from labelled data. Though BC has shown remarkable sample efficiency in terms of batch sizes, it still requires a batch of inputs to estimate the contextual bias. We introduce a running estimation for BC from mini-batches, which subsequently stabilizes the prediction of LLMs when more mini-batches arrive.

In future work, we will endeavor to explore calibration for generative tasks while extending BC to more models across modalities. We suspect that the contextual bias may also exist in short-form generation tasks. Motivated by Zhao et al. (2021), one possible solution for generative tasks is to calibrate over the logits of the first output token since the following tokens are likely deterministic based on the first token. Overall, BC is zero-shot, inference-only, and incurs negligible additional costs. We expect easy prompt engineering with BC for users building towards their own robust and responsible LLM applications.

## C DATASET STATISTICS

Table 4: Details of the dataset used for evaluation in the Table 2.  $|\text{Test}|$  denotes the number of test samples, where we consistently use the validation split as the test split because labels are not publicly available for some datasets.

Dataset	Objective	#sentences	#classes	$ \text{Test} $
SST-2	Sentence Classification	1	2	872
MNLI	NLI	2	3	9815
QNLI	Question-Answering NLI	2	2	5463
MRPC	Paraphrasing	2	2	408
QQP	Paraphrasing	2	2	40430
BoolQ	Commonsense Reasoning	2	2	3270
CB	NLI	2	3	56
COPA	Commonsense Reasoning	3	2	100
RTE	NLI	2	2	277
WiC	Context Comprehension	3	2	638
ANLI-R1	NLI	2	3	1000
ANLI-R2	NLI	2	3	1000
ANLI-R3	NLI	2	3	1200

## D IMPLEMENTATION DETAILS

**Contextual Calibration (Zhao et al., 2021) (CC).** We follow the original implementation of CC and take the mean of the log-probability over three content-free tokens as the test sample in the predefined template: ‘N/A’, ‘’, ‘[MASK]’. It incurs 3 additional API costs from LLMs.

**Domain-Context Calibration (Fei et al., 2023) (DC).** We reproduce the DC baseline by using the same test set as the unlabeled text set to construct its bag-of-words. We then randomly sample tokens for an average length to form the content-free and in-domain input from the bag-of-words. This process is then repeated randomly for 20 times, and we take the mean of the log-probability following the original implementation. It incurs 20 additional API costs from LLMs.

**Prototypical Calibration (Han et al., 2023) (PC).** For a fair comparison, we use the same test set as the unlabeled estimate set for PC. We follow the same hyper-parameters reported by PC with 100 maximum iterations for EM and 100 times random initialization for the whole learning process to stabilize its estimation. It is noteworthy that this number of repetitions is costly and relatively slow, especially when the  $|\text{Test}|$  is large.

**Batch Calibration (BC).** In all reported experiments, we compute the correction log-probability term  $\hat{p}(y|C)$  once after all test samples are seen. In the  $n$ -shot ICL experiments reported in Table 2 and Fig. 6, the  $k$ -shot ICL is concatenating  $k$  random training sample per class. In the BCL experiment that uses labeled samples, we use  $J \times 128$  randomly selected training samples as the labeled data. In the robustness study, we use 1 randomly sampled example as the context to study the performance of BC with respect to the ICL choices. We then conduct the ICL order experiment by re-ordering 4 randomly sampled ICL examples. The rest experiments are conducted on the standard 1-shot ICL setup.

## E ADDITIONAL EXPERIMENTS

Table 5: Accuracy (%) on natural language classification tasks with 0-shot PaLM 2-S and 1-shot PaLM 2-M models in a single seed.

Model	PaLM 2-S					PaLM 2-M				
	0-shot					1-shot				
Method	ICL	CC	DC	PC	BC	ICL	CC	DC	PC	BC
SST-2	94.61	94.50	94.61	87.84	<b>95.18</b>	94.95	95.87	94.95	<b>96.22</b>	96.10
MNLI	45.87	52.54	42.50	38.04	<b>53.67</b>	45.50	54.43	56.26	43.81	<b>60.02</b>
QNLI	49.28	48.97	49.44	<b>50.28</b>	49.55	78.88	75.56	62.95	77.39	<b>78.91</b>
MRPC	69.12	61.76	<b>69.85</b>	<b>69.85</b>	64.95	57.11	<b>73.53</b>	68.87	69.85	65.93
QQP	<b>60.23</b>	51.16	49.12	48.98	56.20	66.18	<b>79.67</b>	74.32	70.27	75.13
BoolQ	86.51	<b>86.97</b>	76.88	55.41	84.04	87.37	88.53	87.28	<b>88.78</b>	87.31
CB	<b>85.71</b>	58.93	55.36	46.43	67.86	71.43	69.64	67.86	50.00	<b>80.36</b>
COPA	88.00	66.00	<b>90.00</b>	52.00	88.00	<b>97.00</b>	96.00	96.00	<b>97.00</b>	96.00
RTE	62.45	67.15	58.12	68.23	<b>71.84</b>	77.62	79.06	68.23	77.98	<b>80.51</b>
WiC	58.31	51.88	52.82	49.22	<b>58.30</b>	61.13	64.11	52.04	65.52	<b>68.03</b>
ANLI-R1	39.80	44.70	43.00	37.00	<b>50.00</b>	52.40	52.40	52.70	35.70	<b>54.00</b>
ANLI-R2	36.80	41.50	40.70	40.20	<b>45.10</b>	46.00	<b>50.70</b>	47.80	35.80	50.00
ANLI-R3	42.67	46.42	43.08	35.50	<b>48.50</b>	43.50	45.67	49.33	32.42	<b>50.50</b>
Avg.	63.03	59.42	58.88	52.23	<b>64.09</b>	67.62	71.17	67.58	64.67	<b>72.52</b>

Table 6: Accuracy (%) on image classification tasks with the zero-shot CLIP ViT-16/B. We additionally report on UCF101 (Soomro et al., 2012), FGVC Aircraft (Maji et al., 2013), and DTD (Cimpoi et al., 2014).

Method	SVHN	EuroSAT	UCF	CLEVR	Aircraft	DTD	Avg.
ZS	18.0	47.8	<b>66.7</b>	14.7	<b>24.8</b>	<b>44.4</b>	36.7
ZS+BC	<b>35.0</b>	<b>54.7</b>	66.0	<b>29.2</b>	22.3	41.7	<b>41.5</b>

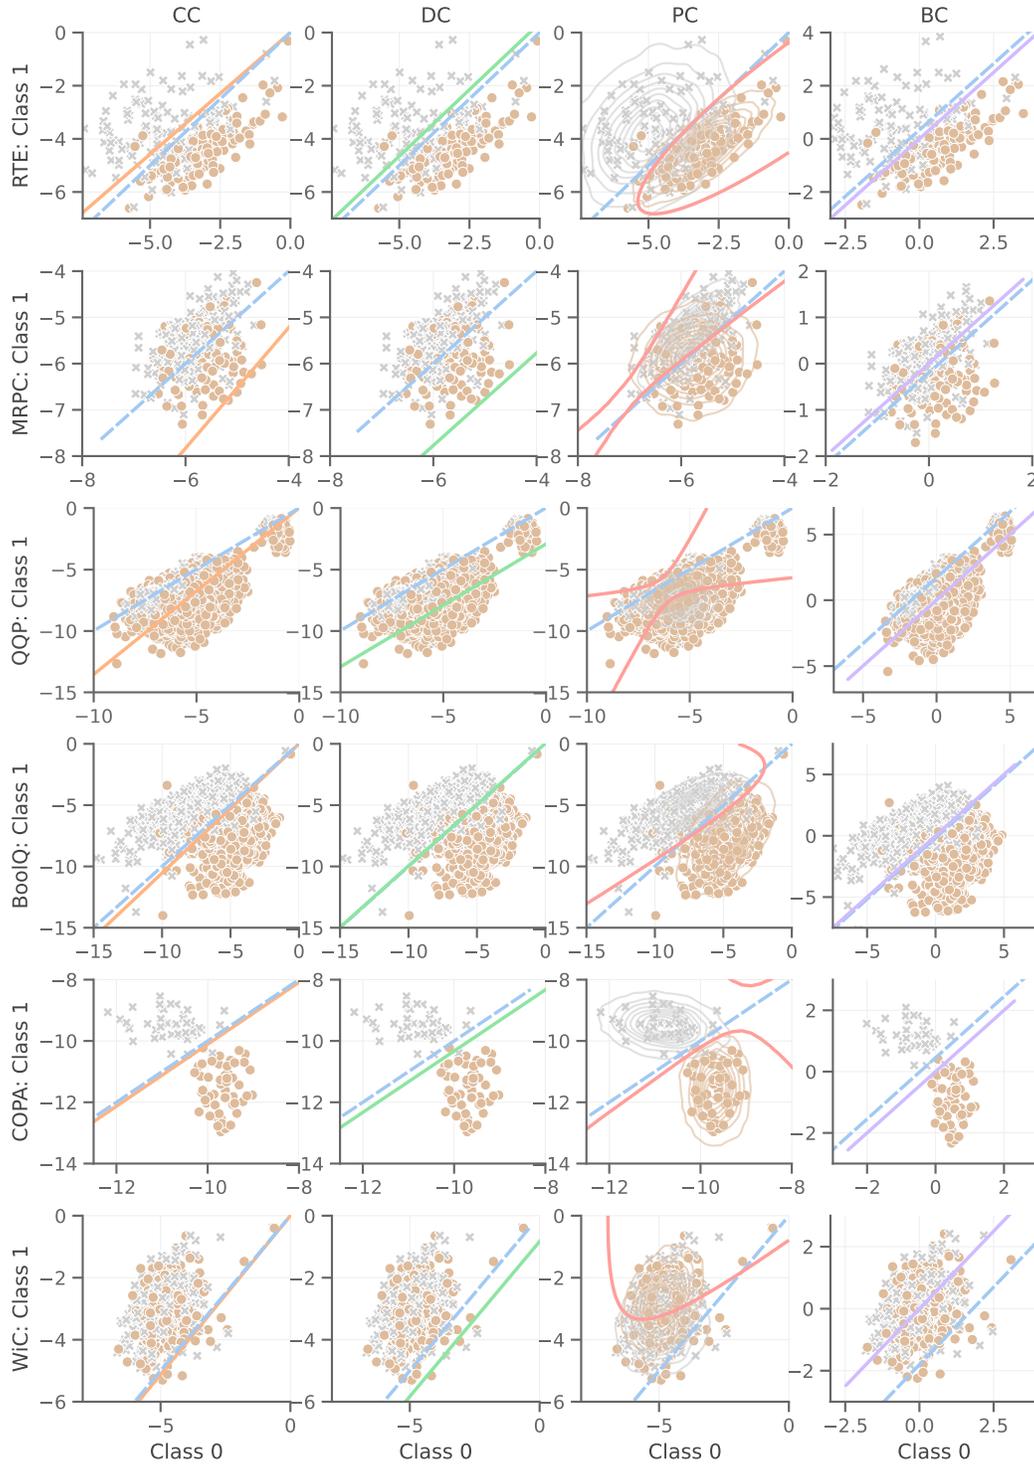


Figure 10: Visualization of the decision boundaries of uncalibrated ICL, and after applying existing calibration methods and the proposed BC. We list all binary classification tasks from the evaluation set.

Table 7: Performance of BC with respect to prompt templates and verbalizers per example in Fig. 8. We report accuracy (%) on the RTE task on 1-shot PaLM 2-S. We refer each example ID to the index of templates and verbalizers listed in Tables 9 & 10.

ID		1	2	3	4	5	6	7	8	9	10
Template	ICL	81.2	80.9	82.3	76.9	83.0	83.4	83.4	82.7	79.1	84.5
	BC	84.1	80.9	82.3	85.6	84.1	81.2	88.1	82.3	83.0	86.6
Verbalizer	ICL	81.2	76.9	74.0	74.0	67.5	69.0	71.5	68.6	65.3	66.4
	BC	84.1	83.4	83.4	80.1	78.0	70.0	82.7	82.7	83.8	83.8

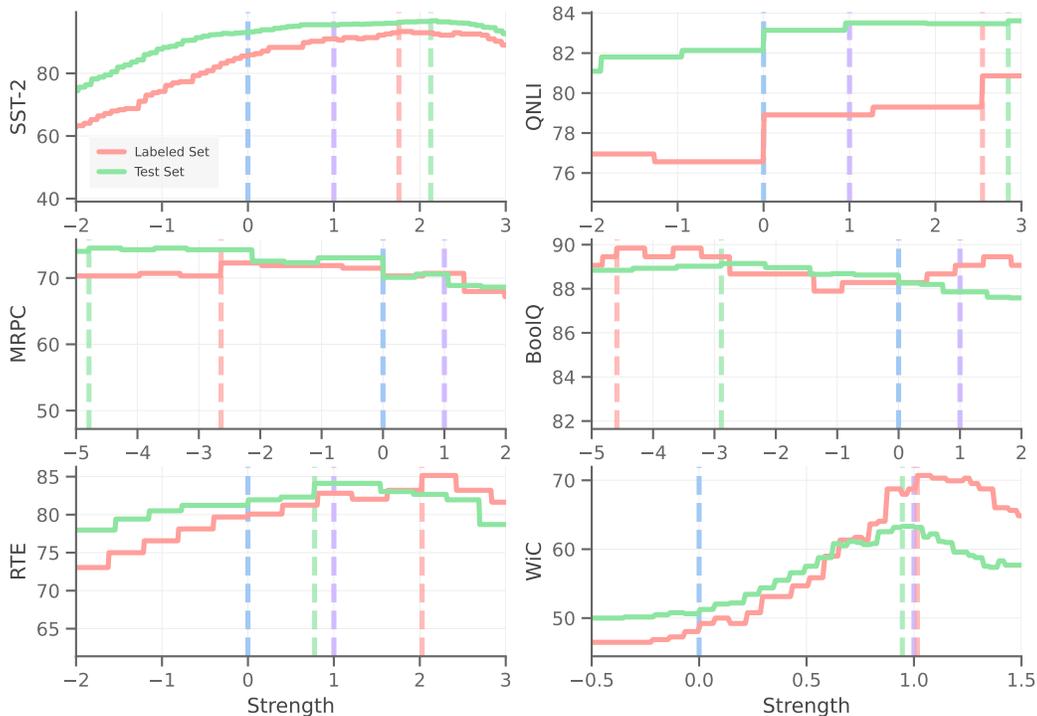


Figure 11: The performance of an adaptable batch calibration layer (BCL) compared to the zero-shot BC with a changing strength. The strength  $\gamma$  at 0 and 1 represent the uncalibrated ICL and BC, respectively. We highlight the optimal strength learned from a labeled set by a red vertical line and the best test strength by a green line.

## F PROMPT TEMPLATES

Table 8: Prompt templates for all  $k$ -shot ICL experiments. We follow the template styles from Han et al. (2023) and Brown et al. (2020).

Dataset	Template	Label Set
SST-2	Review: {sentence} Sentiment: {label}	negative / positive
MNLI CB ANLI	Premise: {premise} Hypothesis: {hypothesis} Answer: {label}	yes / maybe / no
QNLI	Question: {question} Sentence: {sentence} Label: {label}	yes / no
MRPC	Sentence 1: {sentence1} Sentence 2: {sentence2} Equivalence: {label}	no / yes
QQP	Question 1: {question1} Question 2: {question2} Duplicate: {label}	no / yes
BoolQ	{passage} Question: {question} Answer: {label}	no / yes
COPA	Premise: {premise} Choice1: {choice1} Choice2: {choice2} Answer: {label}	1 / 2
RTE	Premise: {sentence1} Hypothesis: {sentence2} Answer: {label}	yes / no
WiC	Sentence1: {sentence1} Sentence2: {sentence2} Word: {word} Answer: {label}	false / true

Table 9: Prompt templates for the robustness experiment conducted on RTE in Fig. 8.

ID	Template	Label Set
1	Premise: {sentence1} Hypothesis: {sentence2} Answer: {label}	yes / no
2	{sentence1} Hypothesis: {sentence2} Answer: {label}	
3	{sentence1} Question: {sentence2} Answer: {label}	
4	{sentence1} Question: {sentence2} {label}	
5	{sentence1} Question: {sentence2} yes or no? Answer: {label}	
6	Sentence 1: {sentence1} Sentence 2: {sentence2} Answer: {label}	
7	Premise: {sentence1} Hypothesis: {sentence2} Label: {label}	
8	Sentence 1: {sentence1} Sentence 2: {sentence2} Label: {label}	
9	Determine if the sentence 2 is true based on the Sentence 1 below Sentence 1: {sentence1} Sentence 2: {sentence2} Answer: {label}	
10	Determine if the sentence 2 is true or false based on the Sentence 1 below Sentence 1: {sentence1} Sentence 2: {sentence2} Answer: {label}	

Table 10: Verbalizer choices for the robustness experiment conducted on RTE in Fig. 8, where we include emoji pairs for ID 8, 9, 10.

ID	Label Set	Template
1	yes / no	Premise: {sentence1} Hypothesis: {sentence2} Answer: {label}
2	true / false	
3	correct / incorrect	
4	positive / negative	
5	good / bad	
6	great / terrible	
7	it was true / it was false	
8	:thumbs_up / :thumbs_down	
9	:man_gesturing_ok / :man_gesturing_no	
10	:check_mark / :cross_mark	

Table 11: Prompt templates for the 0-shot experiments.

Dataset	Template	Label Set
SST-2	Review: {sentence} Sentiment: {label}	negative / positive
MNLI CB ANLI	{premise} Question: {hypothesis} yes, no, or maybe? Answer: {label}	yes / maybe / no
QNLI	{question} Question: {sentence} yes or no? Answer: {label}	yes / no
MRPC	Sentence 1: {sentence1} Sentence 2: {sentence2} Equivalence: {label}	no / yes
QQP	Question 1: {question1} Question 2: {question2} Duplicate: {label}	no / yes
BoolQ	{passage} Question: {question} Answer: {label}	no / yes
COPA	Premise: {premise} Choice1: {choice1} Choice2: {choice2} Answer: {label}	1 / 2
RTE	{sentence1} Question: {sentence2} yes or no? Answer: {label}	yes / no
WiC	{sentence1} {sentence2} Question: Is the word '{word}' used in the same way in the two sentences above? Answer: {label}	no / yes