


Mitigating Label Biases for In-context Learning

Conference Paper**Author(s):**

Fei, Yu; [Hou, Yifan](#) ; Chen, Zeming; Bosselut, Antoine

Publication date:

2023-07

Permanent link:

<https://doi.org/10.3929/ethz-b-000669685>

Rights / license:

[Creative Commons Attribution 4.0 International](#)

Originally published in:

<https://doi.org/10.18653/v1/2023.acl-long.783>

Mitigating Label Biases for In-context Learning

Yu Fei^{†1}, Yifan Hou^{*2}, Zeming Chen^{*3}, Antoine Bosselut³

¹UC Irvine, ²ETH Zurich, ³NLP Lab, IC, EPFL, Switzerland

yu.fei@uci.edu, yifan.hou@inf.ethz.ch,
{zeming.chen, antoine.bosselut}@epfl.ch

Abstract

Various design settings for in-context learning (ICL), such as the choice and order of the in-context examples, can bias the model’s predictions. While many studies discuss these design choices, there have been few systematic investigations into categorizing them and mitigating their impact. In this work, we define a typology for three types of label biases in ICL for text classification: *vanilla-label bias*, *context-label bias*, and *domain-label bias* (which we conceptualize and detect for the first time).

Our analysis demonstrates that prior label bias calibration methods fall short of addressing all three types of biases. Specifically, domain-label bias restricts LLMs to random-level performance on many tasks regardless of the choice of in-context examples. To mitigate the effect of these biases, we propose a simple bias calibration method that estimates a language model’s *label bias* using random in-domain words from the task corpus. After controlling for this estimated bias when making predictions, our novel *domain-context calibration* significantly improves the ICL performance of GPT-J and GPT-3 on a wide range of tasks. The gain is substantial on tasks with large domain-label bias (up to 37% in Macro-F1). Furthermore, our results generalize to models with different scales, pretraining methods, and manually-designed task instructions, showing the prevalence of label biases in ICL.

1 Introduction

Large language models (LLMs) can perform unseen tasks by conditioning on a context prompt that consists of a few training example-label pairs (Brown et al., 2020). However, such in-context learning ability is highly sensitive to various design settings, such as the choice (Liu et al., 2021) and order (Lu et al., 2021) of the in-context samples.

*Equal contribution.

[†]Work done while interning at EPFL.

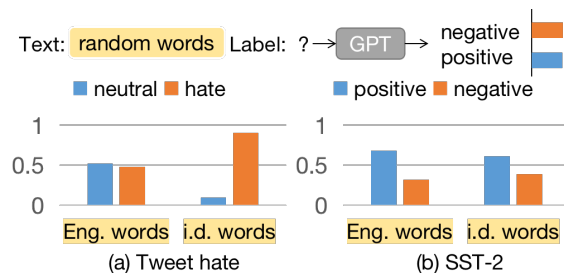


Figure 1: An illustration of domain-label bias. (a) On Tweet hate, seeing random in-domain (i.d.) words sampled from the dataset severely biases the model towards predicting label *hate*, while random English (Eng.) words show no such preference. (b) On a movie review dataset, SST-2, we do not observe such bias.

Recently, Zhao et al. (2021) showed that the instability of ICL largely arises from the fact that these design settings bias the model toward predicting certain answers (e.g., LLMs often predict the label of the last in-context example). As a result, the sensitivity of the results in ICL studies calls for a systematic discussion of biases in ICL and new methods to properly categorize, detect, and comprehensively mitigate various types of biases.

In this work, we conduct a thorough investigation of biases in ICL for text classification. We start by defining a typology of three types of *label biases* (the model’s undesirable preference toward certain label names): *vanilla label bias*, *context-label bias*, and *domain-label bias*. What we term *vanilla label bias* captures the model’s non-contextualized preference for the label names (e.g., the common token bias mentioned by Zhao et al. (2021) caused by different frequencies of label names in the pretraining corpus). *Context-label bias* summarizes the effects of the context prompt (e.g., LLMs tend to prefer the majority and last label of the in-context examples). Finally, *domain-label bias* captures the effects of the task corpus on the model’s predictions.

We show that domain-label biases significantly affect a model’s prediction in ICL. For example,

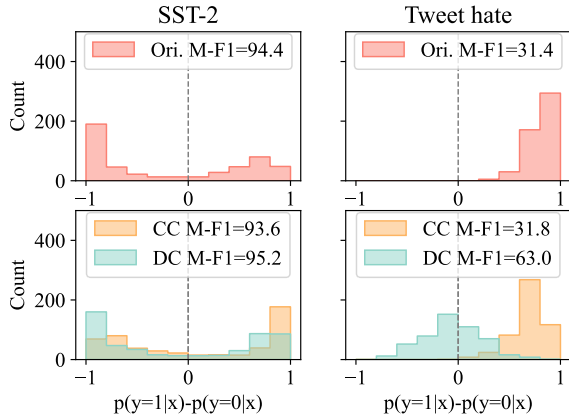


Figure 2: GPT-J has a balanced prediction distribution on SST-2. However, on Tweet hate, the original (Ori.) and calibrated model (with contextual calibration (CC)) predict most inputs as hateful. Our domain-context calibration (DC) largely mitigates such bias and substantially boosts the performance (Macro-F1).

on a hate detection task with two nearly balanced classes, simply seeing random words sampled from the dataset severely biases the model towards predicting the label *hate* (Fig. 1(a)), while seeing random English words does not show such an effect. More importantly, on many tasks with large domain-label bias, LLMs achieve no better than random performance, regardless of the choice of in-context examples (Fig. 2). Moreover, we find that existing bias mitigation methods, such as Contextual Calibration (CC; Zhao et al., 2021) do not combat this effect.

To this end, we propose Domain-context Calibration (DC) to mitigate label biases in ICL. DC first estimates the effects of different label biases holistically using random words sampled from the task corpus. Specifically, we compute the probabilities assigned by the model to each label using random in-domain words as the task input (with optional real in-context learning examples prepended). Using random words limits the semantic meaning of the input, allowing us to estimate the vanilla-label and context-label biases while using in-domain words accounts for the effect of the task corpus. Then, at inference time, we use this label bias estimate to calibrate the model’s output probabilities.

We evaluate the impact of DC on 24 classification datasets, showing that DC improves the average ICL performance of GPT-J (Wang and Komatsuzaki, 2021) and GPT-3 by 20% and 18%. We observe substantial gains on tasks with large domain-label bias (up to 37% in Macro-F1). DC also bene-

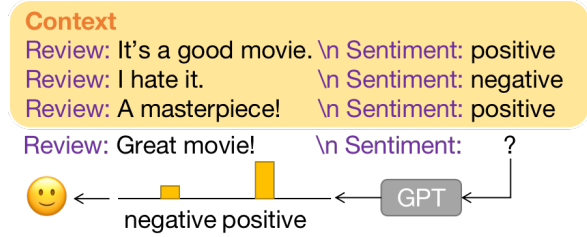


Figure 3: An illustration of few-shot in-context learning for sentiment classification. The context prompt consists of task-specific templates and a few example-label pairs.

fits models with different scales, instruction-tuning (e.g., Instruct-GPT, Ouyang et al., 2022), and provided with task instructions. Finally, we show that DC improves the zero-shot prompting performance of smaller models like RoBERTa (Liu et al., 2019), demonstrating that label bias can be mitigated in prompt-based frameworks beyond ICL.

Overall, our work proposes a new typology of label biases in prompt-based methods, and a simple method for mitigating them. When studying ICL on a diverse collection of datasets, the results on datasets with severe label bias can obfuscate the actual behaviors of the model. Thus, rigorous design for dataset selection (that accounts for confounders) and fine-grained analysis of individual datasets are essential for effectively studying ICL.

2 Categorizing Label Biases in ICL

In this paper, we focus on in-context learning (ICL; Fig. 3) for classification tasks. Formally, we consider a dataset of examples $\{x_i, y_i\}$ where x_i are text inputs and each y_i can be mapped to a verbalization in a label name set \mathcal{L} . We assume each class has one label name. For example, in a sentiment task, \mathcal{L} could be composed of *positive* and *negative* as label names. Given a context prompt C consisting of a few labeled examples and an input text x_i , the model M determines the label of x_i by computing: $\arg \max_{y \in \mathcal{L}} P_M(y|x_i, C)$. Using this notation, we define our typology of label biases based on the mathematical formulation of ICL.

2.1 A Typology of Label Biases

To perform a classification task, a model needs to learn the underlying text-label mapping, i.e., $P(y|x)$. In supervised learning, such mapping is learned by optimizing the model using the training data. In ICL, on the other hand, the model is fixed, and it determines the label of a text by computing the probabilities of predicting the label names

$P_M(y|x, C)$. Notice that there are three components involved in the inference: the label name y , the text x from a specific task corpus, and the context C . Accordingly, as shown in Fig. 4, we can define three types of label biases that lead to a discrepancy between $P_M(y|x, C)$ and $P(y|x)$.

Vanilla-label bias pertains to the uncontextual preference of the model towards predicting certain label names. One possible cause is the pre-training term frequencies of the label names. Zhao et al. (2021) reported a high correlation between the frequency of the DBPedia dataset label names and the rate at which GPT-3 predicts those labels.

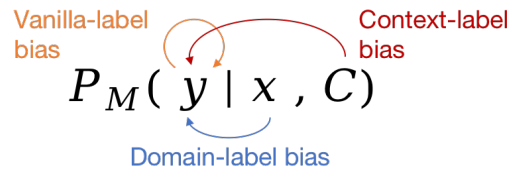
Context-label bias summarizes the effect of the context prompt. With in-context learning, the model “learns” from a few examples, and the learning is particularly sensitive to seemingly arbitrary decisions such as the order of the in-context examples (Lu et al., 2021) and the task template used to map the example to text that the model can process (Mishra et al., 2021; Holtzman et al., 2021).

Domain-label bias captures the effect of the task corpus. Beyond the text-label association demonstrated in the in-context examples, the model also relies on its prior knowledge of the task when making predictions. We show that the association of words to the label names learned from pre-training is a potential pitfall and discuss domain-label bias in more detail in the next section.

3 Domain Label Bias

To illustrate how the domain of a task can induce label bias, consider a case where an LLM predicts whether a patient is *sick* or *healthy* based on some medical descriptions. Because medical descriptions are associated more often with people having health problems in the natural corpus, frequently used words in such documents are likely to have a stronger correlation with *sick* than *healthy*, leading to a systematic bias in the model’s predictions.

Supporting this intuition, we find that for many datasets, conditioning on random words from the dataset examples biases the model toward predicting certain label names. For example, in the hate speech detection task depicted in Fig. 1, we compute the model’s preference (prior) on both label names given random words as the input. A model such as GPT-J has no preference for either of the classes (*neutral* v.s. *hate*) given random English words, but given random in-domain words sampled



	Vanilla-lab.	Context-lab.	Domain-lab.
CC	✓	✓	✗
DC	✓	✓	✓

Figure 4: Three types of biases in ICL. Contextual calibration (CC) considers only the first two types, while our domain-context calibration (DC) handles all of them.

from the dataset, the label priors shift dramatically, becoming 0.95 (*hate*) v.s. 0.05 (*neutral*).

Motivated by this experiment, we quantify the domain-label bias of a model for a particular task using the distance of the model’s priors estimated using random English words and in-domain words. To make the measure more comparable on tasks with different numbers of classes, we define the following metric:

$$bias = \frac{1}{2} \sum_{y \in \mathcal{L}} \left| P_M(y|x_{Eng.}) - P_M(y|x_{i.d.}) \right|, \quad (1)$$

where $x_{Eng.}$ and $x_{i.d.}$ correspond to L random English or random in-domain words and L is the average text length of the dataset.

We find that datasets¹ exhibit different levels of domain-label bias (see Fig. 17 in App. A). More importantly, LLMs behave distinctively on datasets with small and large domain-label bias. As shown in Fig. 5, while GPT-J performs competitively on datasets with small domain-label bias, it rarely outperforms the random baselines on large-bias datasets, indicating that domain-label bias significantly affects ICL. Contextual calibration, which only considers vanilla-label bias and context-label bias, fails to handle domain-label bias.

4 Domain-context Calibration

In this section, we propose Domain-context Calibration (DC), which mitigates the effects of the multiple label biases of our typology (§ 2.1). Following contextual calibration (CC; Zhao et al., 2021), we estimate the overall label bias of a model

¹We note that domain-label bias is model-dependent. However, we observe a high correlation of domain-label bias between LLMs in general (see App. B). Also, by definition, domain-label bias depends on the task formulation, particularly the choice of label names, which we discuss in § 7.

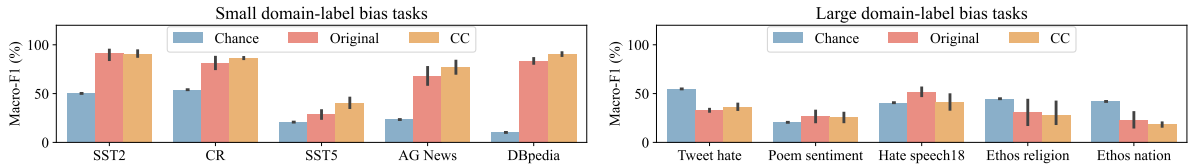


Figure 5: GPT-J in-context learning performance on datasets with varying levels of domain-label bias. On datasets with large domain-label bias, GPT-J usually under-performs random baseline. More importantly, on these datasets, applying contextual calibration (CC) is not helpful and can lead to even worse performance.

with respect to a task by estimating the label probabilities on a *content-free* example text. However, unlike CC, which uses a single, seemingly content-free token (*e.g.*, “N/A”) to approximate the label bias, we use random words sampled from the unlabeled evaluation dataset as the content-free text. Then, for all examples we classify for the task, we re-calibrate the model’s prediction probability using this estimated bias.

More formally, given a dataset, we first construct a bag-of-words \mathcal{B} using the unlabeled texts $\{x_i\}$. Assuming x_i ’s have average length L , we sample L words randomly from \mathcal{B} to form a content-free random text, which captures the word distribution of the dataset domain. However, the random text still remains nearly content-free as it is not grammatically meaningful and potentially contains words from all classes, making it suitable for calibration. We repeat this process M times and average the estimated priors:

$$\bar{P}(y|C) = \frac{1}{M} \sum_{j=1}^M P(y|[random\ text]_j, C). \quad (2)$$

The model then makes predictions according to the following estimate:

$$\hat{y}_i = \arg \max_{y \in \mathcal{L}} \frac{P(y|x_i, C)}{\bar{P}(y|C)}.$$

where $P(y|x_i, C)$ is the original probability assigned to label y for a particular example x_i .

5 Experimental Setup

We conduct comprehensive experiments to analyze the effectiveness of our domain-context calibration in mitigating label biases in ICL.

Datasets We conduct experiments on 24 text classification datasets that cover a wide range of tasks. Most of these datasets are recently used for studying ICL (Zhao et al., 2021; Min et al., 2022; Lu et al., 2021). To control the evaluation budget, we

use a subset of the 24 datasets for GPT-3 experiments following Min et al. (2022). More details can be found in Appendix C.

Model and implementation details We use GPT-J (6B) and GPT-3 (175B) as models in our study. For all experiments, unless stated otherwise, we use $k = 8$ examples sampled randomly from the training set to construct the context prompt and evaluate 5 times using different random seeds. Following Min et al. (2022), we use simple and unified templates for all datasets and do not use any task instructions to keep human engineering at a minimal level. We discuss the effect of task instructions in § 6.1. The templates and label names we used for all datasets can be found in App. E. To further control the budget for evaluating with GPT-3, we follow Lu et al. (2021) and sample a subset of size 500 for all datasets whose test sets have sizes exceeding this number. For domain-context calibration, we use the unlabeled test set for sampling random in-domain words and aggregate using $M = 20$ random texts (Eq. 2). We discuss the sampling of random words in more detail in Appendix F. We use Open-AI’s API for GPT-3 experiments and Tesla V100 GPUs for GPT-J inference.

Evaluation Details For each model, we compare the performance of domain-context calibration to the following baselines: random performance, uncalibrated performance, and contextual calibration performance. Following prior work, we use the Macro-F1 score as the evaluation metric.

6 Experimental Results

We report the average Macro-F1 scores of GPT-J (6B) and GPT-3 (175B) across the entire evaluation suite in Figure 6. Furthermore, we stratify our results into three equal-sized subsets according to their levels of domain-label bias.

Our main finding is that **domain-context calibration generally improves in-context learning, especially on tasks with large domain-label bias.**

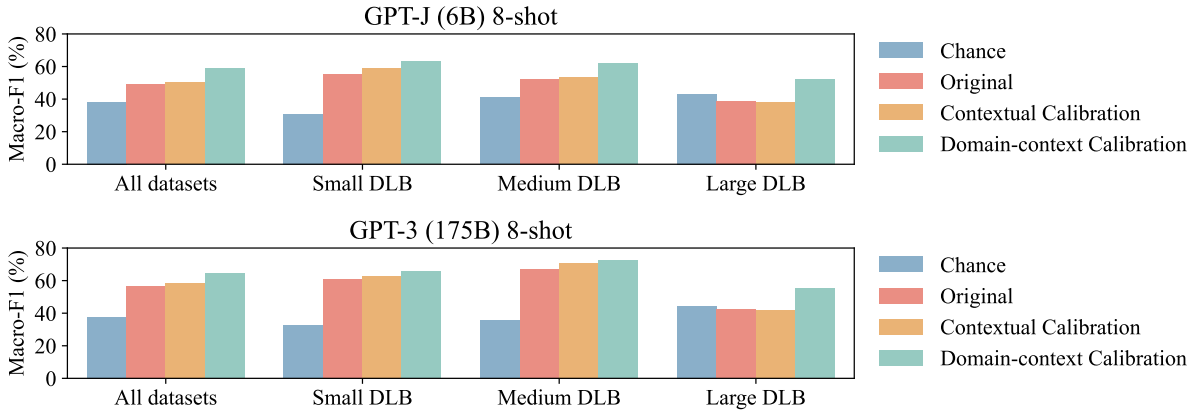


Figure 6: The 8-shot ICL performance of GPT-J and GPT-3. We report both the aggregated results on all datasets as well as the average scores on datasets with different levels of *domain-label bias* (DLB). **Domain-context calibration improves the performance of both models substantially, with larger gains on datasets having larger DLB.**

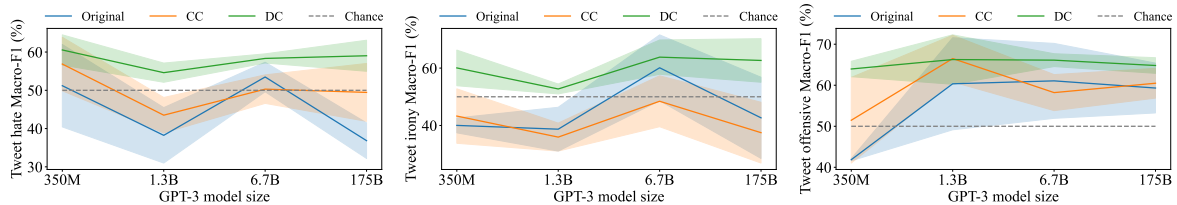


Figure 7: The average Macro-F1 (\pm one standard deviation) plots of GPT-3 with five different choices of in-context examples and different model sizes on three datasets with large domain-label bias. **Scaling up alone does not always increase the performance on tasks with severe domain-label bias. DC boosts the performance of GPT-3 of all different sizes while reducing the variance due to different choices of in-context examples.**

On all datasets, DC consistently boosts the performance for both models with an average improvement (Macro-F1) of 20% (GPT-J) and 18% (GPT-3).² As Fig. 6 shows, the performance gain of DC over baselines (original predictions and CC) largely increases when the degree of domain-label bias increases. On the tasks with the largest domain-label bias, DC is the only method that significantly outperforms the random baseline and achieves up to 37% (GPT-J) and 35% (GPT-3) performance improvement over other baselines, indicating that DC effectively mitigates domain-label bias.

6.1 Generalizability

Following the finding that DC improves ICL significantly on datasets with large domain-label bias, we analyze the robustness of DC under changes in model scale, number of in-context learning examples, and task instructions (all of which have been shown to improve ICL). We use three datasets that exhibit a high level of domain-label bias for GPT-J.

²See Tab. 6 in App. D for results on individual datasets.

Scaling up the model We evaluate GPT-3 models with sizes ranging from 350M to 175B. As Fig. 7 shows, larger models (both the original prediction and CC) do not exhibit better performance on tasks with large domain-label bias. However, DC consistently improves the performance of GPT-3 models of all sizes while reducing the variance due to different choices of in-context examples.

Adding more in-context examples In Table 1, we study the effect of adding more in-context examples by evaluating GPT-J and GPT-3 using 0, 8, and 16 in-context examples. For both models, adding more examples does not seem to benefit the model’s original and CC performance on tasks with large domain-label bias. However, in all settings, DC gives the best results, and for GPT-3, DC further improves the performance when provided with more in-context examples.

Task instructions and instruction-tuning Instruction-tuning and task instructions have been shown to be beneficial to ICL. As shown in Table 2, for GPT-3, providing task instructions³ improves the performance with DC much larger than the

Shot	GPT-J			GPT-3		
	Ori.	CC	DC	Ori.	CC	DC
0	56.0	30.9	65.9	51.5	36.2	61.4
8	50.8	48.4	64.2	47.4	49.8	61.8
16	53.7	45.9	62.7	48.0	49.8	63.1

Table 1: Average Macro-F1 scores on Tweet_eval-hate, Tweet_eval-irony, and Tweet_eval-offensive. **Adding more in-context examples does not resolve domain-label bias, while DC always gives the best results.**

original and CC performance, showing that DC enables GPT-3 to make better use of the task instruction. For Instruct-GPT3 (text-davinci-002), adding task instructions largely improves the model’s original performance. Still, DC yields significant improvement, while CC actually hurts the model’s performance.

Task Instruction	GPT-3			Text-davinci-002		
	Ori.	CC	DC	Ori.	CC	DC
✗	47.4	49.8	61.8	58.5	57.3	71.9
✓	47.7	54.6	68.6	64.1	56.7	68.1

Table 2: Average Macro-F1 scores on three tweet datasets. **DC benefits instruction-tuning models and works in conjunction with task instructions.**

6.2 Analysis

To understand why DC outperforms CC, we conduct a systematic analysis using GPT-J of three differences between DC and CC: 1) the effect of a pre-defined content-free token such as “N/A” compared to using random words; 2) the length of the random word sequence; 3) the source of random words. Below, we summarize our results from Fig. 8.

Content-free token can also be biased First, we find that replacing the pre-defined content-free token from CC (*i.e.*, “N/A”) with a single random English word improves GPT-J’s overall performance, indicating that specific content-free tokens may themselves can be biased toward particular labels. For example, as shown in Fig. 10, on sentiment tasks, calibration GPT-J using “N/A” leads to a systematic bias toward the positive class. Calibrating using random English words to estimate the label bias avoids this problem.

Calibrating with random texts of the average input length is beneficial As shown in Fig. 8, when

³The instructions we used can be found in App. E.

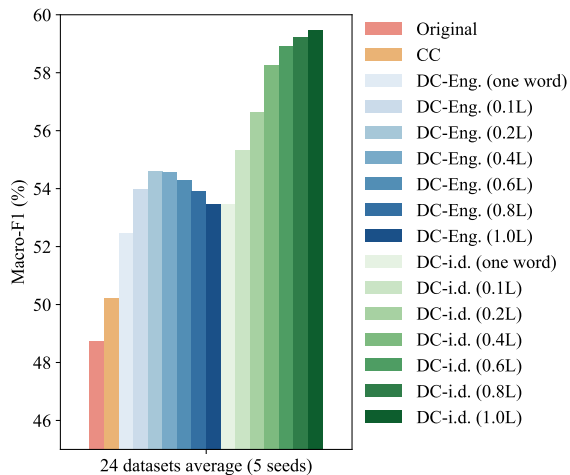


Figure 8: Analysis of differences between CC and DC. *DC-Eng.*: DC with random English words. *DC-i.d.*: DC with random in-domain words. The number of random words used in DC is shown in the bracket, where L is the average text length of the evaluation dataset. DC improves CC along three dimensions: 1) replacing “N/A” with random words to remove content-free token bias; 2) increasing the number of content-free tokens to estimate the effect of the context prompt more accurately; 3) mitigating domain-label bias by calibrating using in-domain words.

calibrating using random English words, increasing the number of words improves performance. Intuitively, using random texts of the average input length for calibration gives a more precise estimate of the effect of the context prompt. To test this, we select the longest and shortest 10% samples of all 24 datasets to construct a dataset with long and short inputs for a task. Then, we test the calibration performance using random English words of different lengths⁴. As shown in Fig. 11, longer (shorter) texts prefer longer (shorter) random texts as calibration sequences to estimate the label bias.

Calibrating using random in-domain words removes domain-label bias Finally, calibrating using random in-domain words yields a large improvement over calibrating using random English words. We plot the prediction distributions of GPT-J on Tweet hate after calibrating with both random English and in-domain words of various lengths in Fig. 9. We see that, when only calibrating using a few in-domain words, the word distribution of the dataset is not well-captured, and thus the domain-label bias is not effectively removed. When cali-

⁴We use random English words rather than random in-domain words to better study the effect of the length.

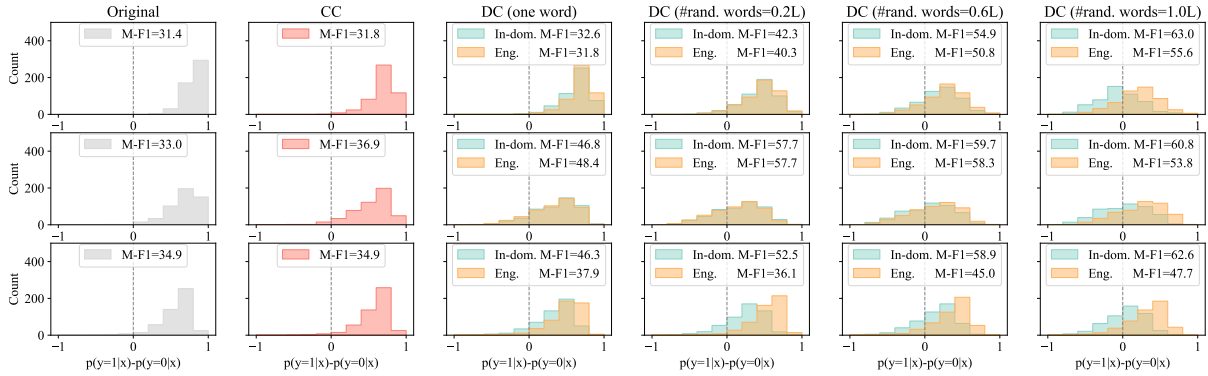


Figure 9: The 8-shot GPT-J prediction distributions of three random seeds on Tweet hate. Regardless of the choice of in-context examples, GPT-J predicts most inputs as *hate* (gray), even after CC (red). **By calibrating using more random in-domain words, the model’s bias towards the *hate* class is gradually mitigated (green), while calibrating using more random English words cannot effectively remove domain-label bias (orange).**

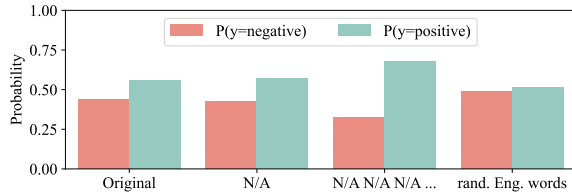


Figure 10: We compute the average original and calibrated marginal probabilities of GPT-J on three sentiment datasets (SST-2, CR, and MR). After calibrating with a single content-free token (e.g., “N/A”), GPT-J is slightly more biased towards positive. Such bias increases when we calibrate with more “N/A”s, while calibrating using the same number of random English words does not show such bias. This shows that **predefined content-free tokens also lead to biases.**

brating using more in-domain words, the prediction becomes more balanced, while after calibrating using more random English words, the model is still biased towards predicting label *hate*. Interestingly, we notice that the more DC mitigates the domain-label bias, the more task performance increases.

6.3 Zero-shot Prompting

Smaller LLMs pre-trained using masked language modeling can also be efficiently adapted to unseen tasks by reformulating the task into a cloze problem using a natural language prompt (i.e., zero-shot prompting (Schick and Schütze, 2020)). To demonstrate that DC can be used even with smaller models, we evaluate the zero-shot prompting ability of RoBERTa-large (Liu et al., 2019) when it is calibrated using DC.⁵ We report our results in Tab. 7 (App.G) and find that across the same 24 datasets,

⁵Implementation details can be found in App. G

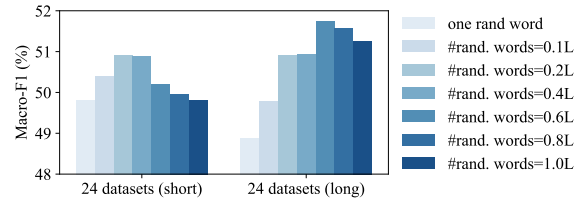


Figure 11: The GPT-J 8-shot (5 random DC seeds) average performance of calibrating using different numbers of random English words on the 10% longest and shortest texts of all 24 datasets. L is the average text length of the evaluation dataset. **Longer (shorter) texts prefer calibrating with longer (shorter) random texts.**

DC achieves a significant performance gain of 26%. Similar to ICL with GPT models, label biases affect RoBERTa’s zero-shot prompting priors. However, DC effectively mitigates these biases, leading to significant performance improvements.

7 Discussion

Label name selection as label bias mitigation

Our results outline how LLMs can be biased to certain label names for different tasks. Intuitively, because the task is formulated as *generating* the label name given an example, the mechanism elicits the model’s prior knowledge about the task. To better understand whether domain-label bias could be mitigated through more careful label name selection, we test GPT-J with three different pairs of label names on Tweet hate: 1) *neutral* v.s. *hate*, which is the most task-relevant set of label names but introduces severe domain-label bias; 2) *favor* v.s. *against*, a pair of less task-relevant antonyms used by Min et al. (2022); 3) X v.s. Y , which are meaningless placeholders.

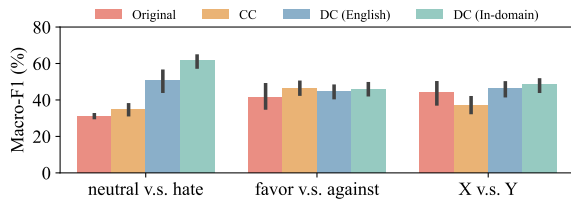


Figure 12: The Macro-F1 scores of GPT-J (8-shots) on Tweet hate with different label names. DC (English): DC with random English words; DC (In-domain): DC with random in-domain words. **Using less task-relevant label names mitigates domain-label bias but limits the model’s performance.**

As shown in Fig. 12, calibrating using random English words or in-domain words makes little difference when choosing (X, Y) or $(favor, against)$ as the label names, showing that they do not introduce domain-label bias. However, although GPT-J is able to achieve better original and CC performance on these label names, $(neutral, hate)$ yields the best performance after removing domain-label bias using DC. Thus, with proper calibration, using the most task-indicative words as label names is likely to be the best option. Surprisingly, the manually picked label names $(favor, against)$ under-perform the meaningless ones (X, Y) after applying DC, hinting that human-plausible label names are not necessarily good label names for LLMs.

Select datasets for ICL analysis The varying levels of domain-label bias in our studied datasets suggest a large variance in how ICL will perform on different datasets. Consequently, macro-averaging the performance on datasets with differing levels of label biases potentially obfuscates diverse results among different tasks. Our work encourages future studies to select datasets carefully to cover varying degrees of label bias among reported results, and to perform fine-grained analysis of individual datasets to when studying ICL performance.

Alternate causes of domain label bias When evaluating models for real-world applications such as hate speech detection, we usually use hard examples (e.g., non-hateful, but “hateful-looking” examples) to check the robustness of the model. However, LLMs trained on natural corpora are likely to be susceptible to adversarial word-level features (LLMs use word associations learned from pre-training to perform ICL). To some degree, adversarial examples could also be a source or large domain-label bias on many datasets.

8 Related Work

In-context learning (ICL) is the standard paradigm for adapting LLMs (Chowdhery et al., 2022; Wei et al., 2022; Zhang et al., 2022). Many recent works focus on understanding its mechanism to improve adaptation performance. For example, Lu et al. (2021) showed that ICL is sensitive to the order of in-context examples. Razeghi et al. (2022) demonstrated that the ICL performance on numerical reasoning tasks is strongly correlated with the pre-training term frequencies. Liu et al. (2021) found that using examples semantically close to the input texts is beneficial. Min et al. (2022) showed that for classification tasks, the input-label pairing format plays the most crucial role in ICL. Sorensen et al. (2022) found that structure of the prompt also significantly affected ICL performance, and that better prompts could be selected based on mutual information between the prompts and the model’s output. Complementary to these works, we comprehensively study the label bias problem in ICL. The existence of domain-label bias indicates that ICL is largely affected by the word-level associations LLMs learn during pre-training.

Other recent works discuss the bias problem in ICL. Zhao et al. (2021) proposed contextual calibration to mitigate three types of biases in ICL: the majority bias, recency bias, and common token bias. Holtzman et al. (2021) focused on the zero-shot setting and found that different surface forms of the answer can compete for probability mass given a question, leading to a bias when predicting with a single label name for each class. In contrast to these works, which consider a specific type of bias, we propose a typology of label biases and propose domain-context calibration that handles all biases in our typology.

The ability of the largest models, like PaLM (Chowdhery et al., 2022), to perform in-context learning under the flipped label or semantically-unrelated label settings (Wei et al., 2023) is very relevant to this work. As the largest models tend to have emergent abilities, it would be interesting to test how vulnerable these models are to label biases (especially domain-label bias) and how domain-context calibration would help. Unfortunately, we currently do not have access to them (e.g., PaLM and Flan-PaLM (Chung et al., 2022)). Nevertheless, the similar behavior of PaLM and Instruct-GPT (as shown in Wei et al. (2023)) and the fact that Instruct-GPT also suffers from domain-label bias

(Table 2) indicate that these more capable models may still be susceptible to label biases. Also, how scaling up or instruction-tuning would relieve label biases is an interesting direction to explore.

9 Conclusion

In this work, we define a typology of label biases that affect in-context learning (ICL). We categorize existing findings of label biases into two types: vanilla-label bias and context-label bias, and identify a new type of bias, domain-label bias, that significantly influences ICL performance. To mitigate these label biases, we propose domain-context calibration, which significantly improves ICL performance on a wide range of tasks, particularly on datasets with large domain-label bias. The various levels of domain-label bias in different datasets also suggest that when analyzing ICL, we need to select datasets with diverse types of label biases and report stratified that acknowledge this diversity beyond single aggregated scores.

Limitations

Data and Task Limitation In this work, we analyze domain-label bias and apply our domain-context calibration to English. We leave analysis and mitigation methods for multilingual tasks to future works. In experiments, we discuss calibration on classification tasks. The effect of domain-label bias could exist differently for open-end tasks like text generation. Our analysis of domain-label bias also emphasizes more on the word-level bias. Other types of biases associated with a domain, such as topics and genders, may also impact model prediction. We leave the diverse analysis to future works. Due to budget limitations, we conduct experiments on a subset of the 24 reported datasets for GPT-3. One can evaluate all 24 datasets to get a complete picture with enough budget.

Model Limitation For large language models, we only focus on the GPT models and only select RoBERTa as the small-scale language model in experiments. Future work could consider expanding to other model types, such as PaLM for large models and DeBERTa for small models. Access to the OpenAI API for GPT-3 is also necessary for parts of our experiments. Future work can consider experimenting with open-source LLMs like the OPT-175B or BLOOM-176B model.

Ethics Statement

Our work focuses on analyzing the general label bias problem of the in-context learning ability of LLMs and improving their performance with a tuning-free method, which involves no large neural model pre-training, re-training, or fine-tuning. As we only use LLMs for inference, developing and applying our approach requires only minimal computational resources compared to methods that require dataset-specific model fine-tuning or engineering. We do not anticipate significant ethical issues introduced by our approach, as we use only off-the-shelf LLMs, and the datasets involved are all publicly available text classification datasets. The discussion of biases in our work is general and not specific to any real word context. Still, our analysis and typology of the label biases of ICL may motivate future work to analyze the bias problem of ICL and LLMs in areas with larger social impacts, such as healthcare or legal scenarios.

References

- Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. [TweetEval: Unified benchmark and comparative evaluation for tweet classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- 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.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. Surface form competition: Why the highest probability answer isn’t always right. *arXiv preprint arXiv:2104.08315*.
- 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.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. The winograd schema challenge. In *Thirteenth international conference on the principles of knowledge representation and reasoning*.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, et al. 2021. Datasets: A community library for natural language processing. *arXiv preprint arXiv:2109.02846*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*.
- Pekka Malo, Ankur Sinha, Pekka Korhonen, Jyrki Wallenius, and Pyry Takala. 2014. Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology*, 65(4):782–796.
- 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.
- 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? *arXiv preprint arXiv:2202.12837*.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2021. Reframing instructional prompts to gptk’s language. *arXiv preprint arXiv:2109.07830*.
- Ioannis Mollas, Zoe Chrysopoulou, Stamatis Karlos, and Grigorios Tsoumakas. 2022. Ethos: a multi-label hate speech detection dataset. *Complex & Intelligent Systems*, pages 1–16.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. *arXiv preprint cs/0409058*.
- Bo Pang and Lillian Lee. 2005. Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales. *arXiv preprint cs/0506075*.
- Yasaman Razeghi, Robert L Logan IV, Matt Gardner, and Sameer Singh. 2022. Impact of pretraining term frequencies on few-shot reasoning. *arXiv preprint arXiv:2202.07206*.
- Timo Schick and Hinrich Schütze. 2020. Exploiting cloze questions for few shot text classification and natural language inference. *arXiv preprint arXiv:2001.07676*.
- 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 Y 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.
- Taylor Sorensen, Joshua Robinson, Christopher Michael Rytting, Alexander Glenn Shaw, Kyle Jeffrey Rogers, Alexia Pauline Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. 2022. An information-theoretic approach to prompt engineering without ground truth labels. *arXiv preprint arXiv:2203.11364*.

- Ellen M Voorhees and Dawn M Tice. 2000. Building a question answering test collection. In *Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 200–207.
- Ben Wang and Aran Komatsuzaki. 2021. GPT-J-6B: A 6 Billion Parameter Autoregressive Language Model. <https://github.com/kingoflolz/mesh-transformer-jax>.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*.
- 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. *arXiv preprint arXiv:2303.03846*.
- 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 Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*, pages 12697–12706. PMLR.

A Domain-label Bias of All Datasets

We compute and illustrate the domain-label bias of all datasets we used with different LLMs in Fig. 17. Different datasets exhibit different levels of domain-label bias. Regarding task types, The detection tasks (red) show the largest domain-label bias, while the NLI tasks (orange) have the least. On sentiment and topic tasks, the domain-label bias is mostly small but can vary depending on the domain of the dataset. For example, for sentiment classification datasets, movie review datasets like SST-2 have relatively small domain-label bias. While financial statements and poem sentiment, whose texts are from rare domains, have much larger biases. We discuss the model dependency of domain-label bias in the next section.

B Correlation of Domain-label Bias Estimated with Different LLMs

We compute the correlation of domain-label bias (defined by eq. (1)) computed with 5 different models on all 24 evaluation datasets. We use GPT-3 Ada (350M), GPT-3 Babbage (1.3B), GPT-3 curie (6.7B), GPT-3 DaVinci (175B), and GPT-J (6B). We show the correlation plot in Figure 13. Although domain-label bias is model-dependent by definition, the biases computed by different models are highly correlated.

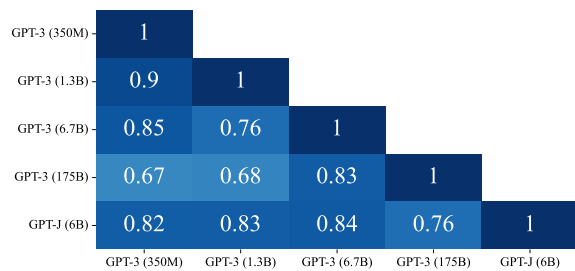


Figure 13: Although *domain-label bias* is model-dependent, we observe a high correlation of *domain-label bias* between LLMs on the evaluation datasets.

C Full Dataset Information

We use 24 datasets falling into three categories: sentiment and topic classification, NLI, and Detection. Most of the used datasets are from existing works. W(Min et al., 2022; Lu et al., 2021; Zhao et al., 2021)e added a few more detection datasets for better studying the domain-label bias as they tend to show the largest domain-label bias. We use the HuggingFace venison Lhoest et al. (2021) of all

datasets and use the test set, if available, for evaluation. Otherwise, we use the development set. We summarize the full dataset information in Table 3.

D Full Few-shot Results

We report the full 8-shot results on individual datasets with the standard deviations (5 random seeds) in Table 6. We further show the few-shot performance gain (GPT-J) of DC over CC on all datasets in Figure 14.

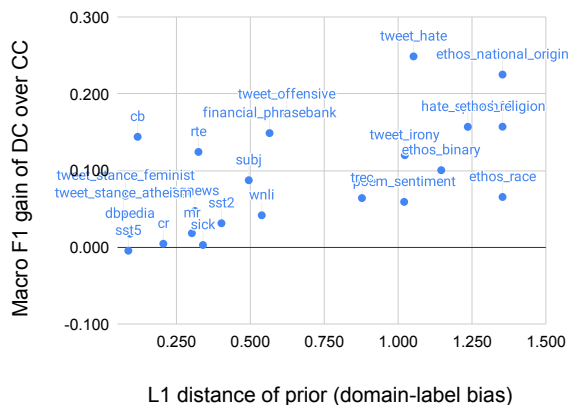


Figure 14: On tasks with more severe domain-label bias, we can observe more performance gain of DC over CC.

E Templates and Task Instructions

We show the templates and label names for all datasets in Table 4. The task instructions used in Table 2 are illustrated in Table 5. We always use exactly one word for every label name. To avoid label names being tokenized into subwords, we always use lower-cased label names except for tasks answering with True or False.

F Sampling Analysis

In this section, we analyze two factors related to the random word sampling process involved in DC: 1) the number of random texts to use for estimating the prior (M in eq. (2)), and 2) the size of the unlabeled dataset to sample random in-domain words from. We conduct experiments on two small-domain-label-bias datasets (SST-2 and AG News) and two large-domain-label-bias datasets (Tweet hate and Tweet irony).

How many random texts should we sample?

First, we sample different numbers of random texts M for estimating the model’s prior as in eq. (2). As

shown in Figure 15, DC is able to achieve a good estimate with a relatively small number of sampling. We choose $M = 20$ as it achieves a good balance between computational efficiency and stability of prior estimation.

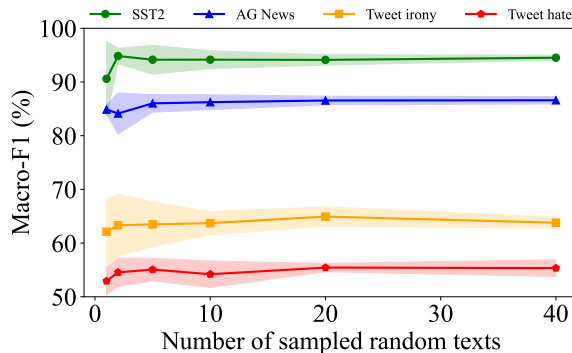


Figure 15: The few-shot performance of GPT-J after applying domain-context calibration with different numbers of sampled random texts for estimating the model’s prior. For each number, we sample random texts 5 times with different random seeds to show the stability of the sampling process. Using 20 random texts already provides a stable estimate of label biases.

How large should the unlabeled task corpus be?

In the main experiments, we use the whole unlabeled test set to construct a bag-of-words and sample random words from it. Here, we study the effect of the unlabeled dataset set size on the performance of DC. As shown in Figure 16, DC is able to achieve a good estimate with 50 unlabeled texts from the dataset.

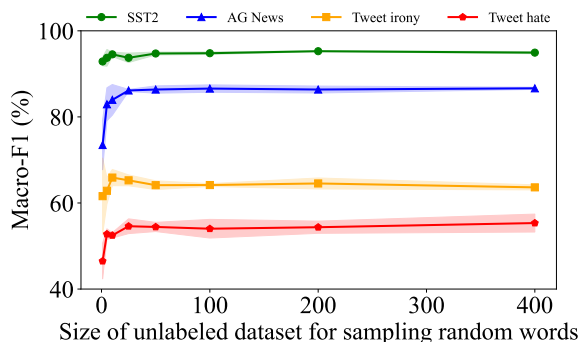


Figure 16: The few-shot performance of GPT-J with DC when sampling random words from a different size of the unlabeled dataset. For each size, we sample random texts 5 times with different random seeds to show the stability of the sampling process. Using 50 unlabeled texts from the dataset already provides a stable estimate of label biases.

G Zero-shot Prompting Experiment

Templates for zero-shot prompting We adapt templates from Gao et al. (2020) for our zero-shot prompting experiments.

Sentiment and detection tasks:

$\langle input \rangle$ It was $\langle mask \rangle$]

Subj:

$\langle input \rangle$ This is $\langle mask \rangle$

Topic tasks:

$\langle mask \rangle$: $\langle input \rangle$

NLI tasks:

$\langle sentence_1 \rangle$? $\langle mask \rangle$, $\langle sentence_2 \rangle$

Full Results We show the full zero-shot prompting results with RoBERTa-large on individual datasets in Table 7.

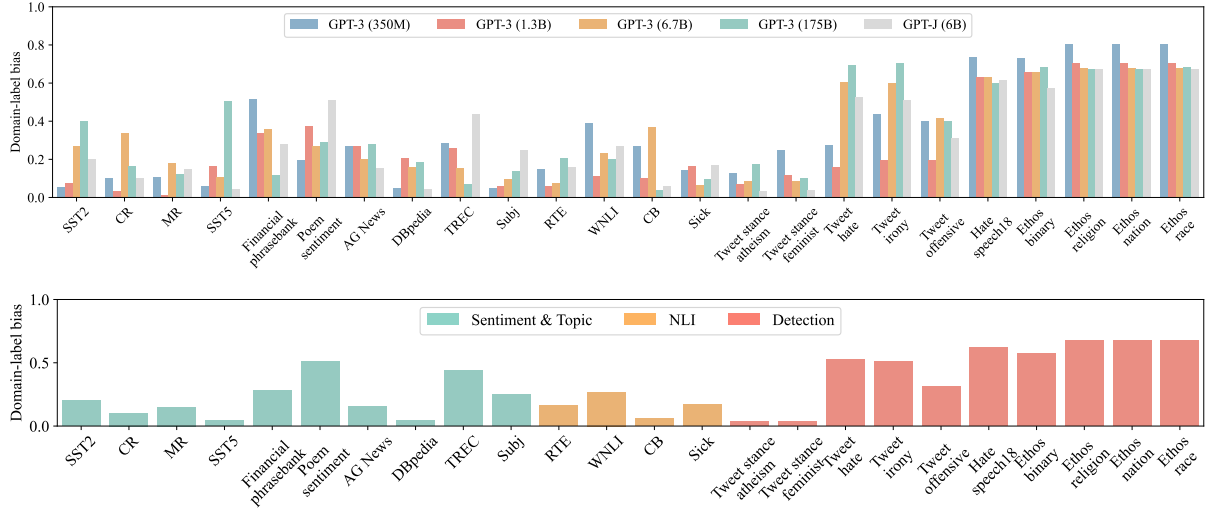


Figure 17: Top: Domain-label bias of different models on all evaluation datasets. Bottom: Domain-label bias of GPT-J on all evaluation datasets categorized based on the task types.

Dataset	# Class	Balanced	GPT-J	GPT-3
<i>Sentiment and topic classification</i>				
SST-2 (Socher et al., 2013)	2	✓	✓	✓
SST-5 (Socher et al., 2013)	5	✗	✓	✓
MR (Pang and Lee, 2005)	2	✓	✓	
CR (Hu and Liu, 2004)	2	✓	✓	
financial_phrasebank (Malo et al., 2014)	3	✗	✓	✓
poem_sentiment (Sheng and Uthus, 2020)	4	✗	✓	
Subj (Pang and Lee, 2004)	2	✗	✓	
AG News (Zhang et al., 2015)	4	✓	✓	✓
DBpedia (Zhang et al., 2015)	14	✓	✓	
TREC (Voorhees and Tice, 2000)	6	✗	✓	✓
<i>Natural language inference</i>				
glue-wnli (Levesque et al., 2012)	2	✗	✓	
RTE (Dagan et al., 2005)	2	✗	✓	✓
CB (De Marneffe et al., 2019)	3	✗	✓	✓
sick (Marelli et al., 2014)	3	✗	✓	
<i>Detection</i>				
tweet_eval-hate (Barbieri et al., 2020)	2	✗	✓	✓
tweet_eval-irony (Barbieri et al., 2020)	2	✗	✓	✓
tweet_eval-offensive (Barbieri et al., 2020)	2	✗	✓	✓
tweet_eval-stance_atheism (Barbieri et al., 2020)	3	✗	✓	
tweet_eval-stance_feminist (Barbieri et al., 2020)	3	✗	✓	
hate_speech18 (de Gibert et al., 2018)	2	✗	✓	✓
ethos-binary (Mollas et al., 2022)	2	✗	✓	
ethos-religion (Mollas et al., 2022)	2	✗	✓	
ethos-national_origin (Mollas et al., 2022)	2	✗	✓	
ethos-race (Mollas et al., 2022)	2	✗	✓	✓

Table 3: Full dataset information. To control the evaluation budget, we use a subset of the 24 datasets for GPT-3 experiments following Min et al. (2022). The GPT-J and GPT-3 columns indicate whether the corresponding datasets are used in Fig. 6.

Dataset	Template	Label Name
SST-2, SST-5, 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, mixed
Subj	Input: [INPUT] Label: [LABEL]	subjective, objective
AG News	Article: [INPUT] Answer: [LABEL]	world, sports, business, technology & science
DBpedia	Article: [INPUT] Article type: [LABEL]	company, school, artist, athlete, politics, transportation, building, nature, village, animal, plant, album, film, book
TREC	Question: [INPUT] Answer type: [LABEL]	number, location, person, description, entity, abbrev
RTE, glue-wnli	[PREMISE] question: [HYPOTHESIS] True or False? answer: [LABEL]	True, False
CB	[PREMISE] question: [HYPOTHESIS] true, false, or neither? answer: [LABEL]	true, false, neither
sick	[PREMISE] question: [HYPOTHESIS] entailment, neutral, or contradiction? answer: [LABEL]	entailment, neutral, contradiction
tweet_eval-hate	Tweet: [INPUT] Label: [LABEL]	neutral, hate
tweet_eval-irony	Tweet: [INPUT] Label: [LABEL]	neutral, ironic
tweet_eval-offensive	Tweet: [INPUT] Label: [LABEL]	neutral, offensive
tweet_eval-stance_atheism, tweet_eval-stance_feminist	Tweet: [INPUT] Label: [LABEL]	none, against, favor
hate_speech18, ethos-race, ethos-binary, ethos-religion, ethos-national_origin,	Text: [INPUT] Label: [LABEL]	neutral, hate

Table 4: **Templates** of all 24 datasets used in our experiments. We mainly adapt templates used in Zhao et al. (2021) and Min et al. (2022). We remove all task instructions and unify the format for similar tasks. We always use lower-case label names to avoid label names being tokenized into subwords except for “True or False” tasks.

Dataset	Instruction and Template
tweet_eval-hate	Classify tweets that are hateful against immigrants or women as hate and tweets that are not hateful against immigrants or women as neutral. Tweet: [INPUT] Label: [LABEL]
tweet_eval-irony	Classify tweets that are ironic as ironic, and tweets that are not ironic as neutral. Tweet: [INPUT] Label: [LABEL]
tweet_eval-offensive	Classify tweets that are offensive as offensive, and tweets that are not offensive as neutral. Tweet: [INPUT] Label: [LABEL]

Table 5: **Task instructions** used in Table 2.

Dataset	GPT-J			GPT-3		
	Original	CC	DC	Original	CC	DC
<i>Sentiment and topic classification</i>						
SST-2	91.0 _{6.0}	90.8 _{3.2}	94.0 _{1.3}	96.0 _{1.0}	96.4 _{0.2}	96.3 _{0.5}
CR	81.4 _{6.9}	86.5 _{0.8}	87.0 _{4.0}	-	-	-
MR	93.1 _{0.7}	91.3 _{1.1}	93.1 _{0.5}	-	-	-
SST-5	28.9 _{4.6}	40.8 _{5.4}	40.3 _{4.9}	32.5 _{3.3}	41.1 _{0.9}	42.4 _{2.9}
Financial phrasebank	46.4 _{6.9}	46.7 _{4.2}	61.6 _{3.3}	58.9 _{11.6}	60.6 _{4.6}	69.5 _{7.1}
Poem sentiment	26.6 _{6.1}	25.5 _{5.2}	31.4 _{3.0}	-	-	-
AG News	68.4 _{9.9}	76.8 _{7.2}	81.5 _{5.1}	79.9 _{7.0}	86.0 _{1.5}	85.9 _{1.1}
DBpedia	83.5 _{3.0}	90.6 _{1.7}	92.4 _{1.2}	-	-	-
TREC	55.3 _{7.0}	63.9 _{1.1}	70.3 _{3.1}	69.0 _{13.2}	76.5 _{7.9}	76.9 _{7.9}
Subj	65.2 _{12.2}	61.9 _{13.3}	70.7 _{4.3}	-	-	-
<i>Natural language inference</i>						
RTE	43.1 _{5.8}	37.8 _{4.3}	50.3 _{5.6}	61.8 _{10.6}	65.8 _{3.8}	64.5 _{5.2}
WNLI	33.5 _{0.0}	33.9 _{0.6}	38.1 _{3.1}	-	-	-
CB	24.8 _{6.1}	27.8 _{9.3}	42.3 _{3.6}	53.7 _{1.7}	49.0 _{8.4}	51.1 _{8.9}
Sick	25.6 _{6.7}	41.1 _{3.7}	41.4 _{10.9}	-	-	-
<i>Detection</i>						
Tweet hate	32.8 _{1.2}	36.4 _{3.4}	61.2 _{2.3}	36.8 _{4.7}	49.5 _{7.6}	59.0 _{4.1}
Tweet irony	60.0 _{10.4}	50.3 _{7.7}	62.4 _{5.3}	42.7 _{14.1}	37.5 _{10.6}	62.7 _{7.6}
Tweet offensive	59.4 _{11.4}	51.0 _{6.3}	68.3 _{2.9}	59.3 _{6.1}	60.5 _{3.6}	64.8 _{1.9}
Tweet stance atheism	23.2 _{5.1}	23.0 _{6.0}	27.4 _{3.2}	-	-	-
Tweet stance feminist	40.4 _{10.0}	34.6 _{2.6}	41.3 _{1.9}	-	-	-
Hate speech18	51.5 _{4.7}	41.6 _{8.6}	57.3 _{2.5}	49.9 _{11.5}	47.0 _{6.2}	52.0 _{3.9}
Ethos binary	48.4 _{16.1}	60.1 _{8.2}	70.2 _{2.5}	-	-	-
Ethos religion	30.7 _{14.3}	28.0 _{13.8}	43.8 _{6.7}	-	-	-
Ethos nation	23.1 _{8.7}	18.2 _{2.1}	40.7 _{7.8}	-	-	-
Ethos race	36.4 _{11.8}	44.8 _{17.4}	51.4 _{6.4}	40.0 _{12.2}	33.2 _{7.6}	47.0 _{6.9}

Table 6: Full 8-shot results. We report the average Macro-F1 scores over 5 random seeds with the standard deviations. To control the evaluation budget, we use a subset of the 24 datasets for GPT-3 experiments following Min et al. (2022).

Setting	SST-2	CR	MR	SST-5	FP	PS	AG	DB	TREC	Subj	RTE	WNLI
Ori.	82.5	82.2	78.5	28.3	32.5	24.1	52.0	65.5	7.0	37.9	39.2	32.9
CC	79.4	81.0	76.0	19.3	40.0	21.7	59.1	73.1	5.0	51.6	46.0	33.6
DC	87.2	82.2	82.5	35.1	43.3	26.2	62.7	74.2	22.6	44.3	44.6	33.8
Setting	Tw-H	Tw-I	Tw-O	Tw-A	Tw-F	HS18	Eth-B	Eth-Re	Eth-N	Eth-Ra	CB	Sick
Ori.	32.9	30.5	28.4	23.7	22.4	37.7	45.0	21.0	17.8	17.3	21.3	38.2
CC	30.5	30.8	32.8	19.5	28.0	26.5	37.0	17.6	15.5	15.8	34.9	39.7
DC	59.0	49.7	56.3	28.5	29.4	44.3	58.7	42.7	36.9	44.1	27.5	54.2

Table 7: Zero-shot prompting results with RoBERTa-large (Macro-F1).

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
At the end of the paper, after the conclusion section.
- A2. Did you discuss any potential risks of your work?
At the end of the paper, after the limitation section.
- A3. Do the abstract and introduction summarize the paper’s main claims?
We summarize our main contributions in the abstract and introduction sections.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

Section 6.

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 5.

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 5.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 6 and Appendix D.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 5 and Appendix C.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.