

Mitigating Word Bias in Zero-shot Prompt-based Classifiers

Adian Liusie, Potsawee Manakul, Mark J. F. Gales

ALTA Institute, Department of Engineering, University of Cambridge
a1826@cam.ac.uk, pm574@cam.ac.uk, mjfg@eng.cam.ac.uk

Abstract

Prompt-based classifiers are an attractive approach for zero-shot classification. However, the precise choice of the prompt template and label words can largely influence performance, with semantically equivalent settings often showing notable performance difference. This discrepancy can be partly attributed to word biases, where the classifier may be biased towards classes. To address this problem, it is possible to optimise classification thresholds on a labelled data set, however, this mitigates some of the advantages of prompt-based classifiers. This paper instead approaches this problem by examining the expected marginal probabilities of the classes. Here, probabilities are reweighted to have a uniform prior over classes, in an unsupervised fashion. Further, we draw a theoretical connection between the class priors and the language models’ word prior, and offer the ability to set a threshold in a zero-resource fashion. We show that matching class priors correlates strongly with the oracle upper bound performance and demonstrate large consistent performance gains for prompt settings over a range of NLP tasks.¹

1 Introduction

Large language models (LLM) have shown impressive general ability for natural language processing (NLP) tasks. LLMs can effectively handle a range of NLP tasks through ‘prompting’, where a natural language instruction is added to the input, conditioning the model to the task at hand. Prompting can either be an emergent ability learned through scaling up model size (Brown et al., 2020; Wei et al., 2022) or an ability learned through instruction tuning (Wei et al., 2021; Chung et al., 2022; Ouyang et al., 2022). Despite the recent popularity of prompting, there is a known sensitivity of prompt-based LLMs to elements such as prompt

¹code available on github at <https://github.com/adianliusie/robust-prompt-classifier>

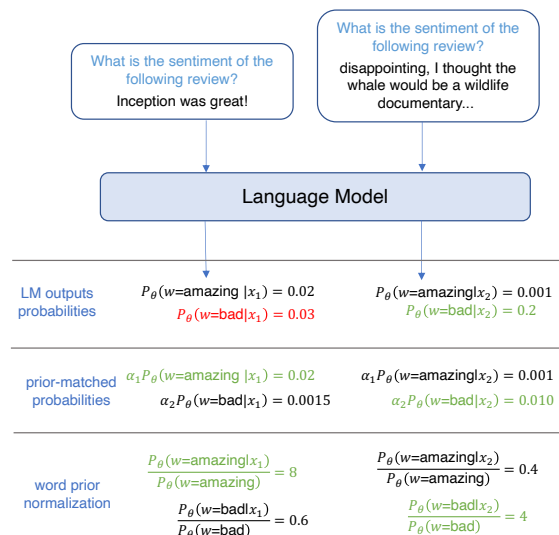


Figure 1: Instead of using the raw LM output probabilities of the label words, we consider mitigating bias by finding weights that make the classifier unbiased over classes. This is connected to normalising by word priors, which we use as a zero-resource de-biasing approach.

template and label words (Gao et al., 2021; Schick and Schütze, 2021). Previous works have demonstrated that prompt templates can significantly impact task performance (Shin et al., 2020; Zhou et al.) and that factors such as chosen label words can influence system performance for classification tasks (Zhao et al., 2021; Holtzman et al., 2021).

This work focuses on the influence of ‘word biases’ for prompt-based classifiers. i.e. the bias that prompts may have towards certain classes, independent of the input text. To account for this bias, one could use a labelled dataset to find optimal class decision thresholds. This, however, requires labelled task data, which may limit the zero-shot benefits of prompt-based classifiers. We propose a simple unsupervised solution of re-weighting probabilities, where we use unlabelled data to search for weight parameters that ensure a uniform prior over classes. We show that this prior matching leads to greater

robustness for diverse prompt settings and that the unsupervised weights which debias the classifier is highly correlated with the oracle weights that maximise accuracy. Further, we provide theoretical analysis that draws a connection between word priors and inherent class bias, which we use to motivate a zero-resource normalisation approach that is competitive with prior matching. Overall, we demonstrate that our unsupervised approach highly reduces sensitivity to the chosen prompt and label words, and that settings which initially fail can often be made effective through a simple probability re-weighting.

Our contributions are 1) We propose a simple unsupervised probability re-weighting method, and empirically demonstrate greater robustness to prompt and label word choice, with large accuracy gains across prompt settings for a range of standard NLP tasks. 2) We theoretically connect the weight parameters to word priors and use this to motivate a zero-resource re-weighting approach. 3) We show that the weights of prior matching are highly correlated with the optimal oracle weights that maximize accuracy, illustrating that our approach is a near-optimal use of a system’s output probabilities.

2 Mitigating Bias by Re-weighting

Prompt-based classifiers Given an input sequence $\mathbf{x} \in \mathcal{X}$, large language models (LLMs) model $P_\theta(\mathbf{w}|\mathbf{x})$, the output probability distribution over all possible sequences $\mathbf{w} \in \mathcal{X}$. For a classification task \mathcal{T} , a prompt-based classifier 1) reformats the input text \mathbf{x} to prompt $\mathbf{p} \in \mathcal{X}$ by including the task instruction, and 2) selects class words $\{w_k\}_{1:K}$ which are associated to each output class $\{y_k\}_{1:K}$. For example in sentiment classification, one can use prompt ‘*what is the sentiment of the following review? <x>*’, (where $\langle x \rangle$ is the current input \mathbf{x} , e.g. ‘*Inception was absolutely brilliant*’), and class words $w_0=bad$ and $w_1=good$ for the negative and positive classes respectively. For a prompt classifier, $Q = \{\mathbf{p}, \{w_k\}_{1:K}\}$, class probabilities can be set to be proportional to the probability of the associated class word, where the final decision \hat{y} is the class with the highest probability (Zhao et al., 2021; Jiang et al., 2020).

$$\tilde{P}_\theta(y_k|\mathbf{x}, Q) = \frac{P_\theta(w_k|\mathbf{p}(\mathbf{x}))}{\sum_{w_i} P_\theta(w_i|\mathbf{p}(\mathbf{x}))} \quad (1)$$

$$\hat{y} = \operatorname{argmax}_k \tilde{P}_\theta(y_k|\mathbf{x}, Q) \quad (2)$$

However, as a large language model, the prompt-based classifier may return probabilities that are influenced by distributional statistics of words (Gardner et al., 2021; Liusie et al., 2022). This may lead to inherent class bias, where label words may have high probability not because they better answer the prompt, but because they have a high LM prior.

Optimal Weights To account for this, one can define weight parameters $\alpha = \{\alpha_k\}_{1:K}$, where each $\alpha_k \in \mathbb{R}^+$ scales the probabilities of the classifier,

$$\hat{P}_\theta(y_k|\mathbf{x}, Q, \alpha) = \frac{\alpha_k \tilde{P}_\theta(y_k|\mathbf{x}, Q)}{\sum_i \alpha_i \tilde{P}_\theta(y_i|\mathbf{x}, Q)} \quad (3)$$

Given labelled task dataset $\mathcal{D} = \{(\mathbf{x}^{(j)}, y^{(j)})\}_{j=1}^N$, one can then find the optimal weights α^* that maximises the accuracy of the prompt classifier $\hat{P}_\theta(y_k|\mathbf{x}, Q, \alpha)$ over the dataset,

$$\alpha^* = \operatorname{argmax}_\alpha \operatorname{Accuracy}(Q, \alpha, \mathcal{D}) \quad (4)$$

Prior-Matching The previous approach requires labelled data, which may limit the benefit of using prompt-based classifiers. As an alternative, one can find the values $\bar{\alpha}$ that ensure that the classifier is unbiased, such that the class prior $\hat{P}(y_k|Q, \alpha)$ matches the true prior $P(y_k)$

$$\hat{P}_\theta(y_k|Q, \alpha) = \mathbb{E}_{\mathbf{x}}\{\hat{P}_\theta(y_k|\mathbf{x}, Q, \alpha)\} \quad (5)$$

$$\approx \frac{1}{N} \sum_{j=1}^N \hat{P}_\theta(y_k^{(j)}|\mathbf{x}^{(j)}, Q, \alpha) \quad (6)$$

$$\bar{\alpha} = \operatorname{argmin}_\alpha \sum_{\forall y_k} |\hat{P}_\theta(y_k|Q, \alpha) - P(y_k)| \quad (7)$$

A deterministic solution that exactly matches the distributions exists, which can be found with a search with 1 degree of freedom (that can be accounted for by setting $\alpha_1 = 1$). If there is no expected class bias, one can assume equal probabilities over all classes, $P(y_k) = \mathcal{U}(y_k) = \frac{1}{N}$. This approach is therefore unsupervised and only requires text inputs $\mathcal{D}_x = \{\mathbf{x}^{(j)}\}_{j=1}^M$, which therefore can be applied at inference to any test set.

Null-Input Approximation The dependence of prior-matching on unlabelled dataset \mathcal{D}_x is a drawback. In Appendix A, we show that one can make the analytical approximation

$$\bar{\alpha}_k \approx \frac{1}{\mathbb{E}_{\mathbf{x}}\{P_\theta(w_k|\mathbf{x}, Q)\}} = \frac{1}{P_\theta(w_k|Q)} \quad (8)$$

method	inputs	labels	imdb	rt	amazon	snli	mnli	qqp
baseline	✗	✗	85.4±12.7	78.8±14.0	86.0±13.8	45.2±13.7	43.5±11.3	65.4±14.0
null-input	✗	✗	92.1±3.2	89.1±3.8	95.0±1.8	75.2±10.4	66.1±9.7	77.4±6.6
prior-match	✓	✗	93.1±3.3	90.9±1.6	96.0±0.8	78.5±9.3	69.8±9.7	79.1±2.4
optimal	✓	✓	93.5±2.7	91.2±1.5	96.1±0.7	79.4±8.2	70.8±8.6	82.3±2.8

Table 1: Average dataset accuracy and standard deviations, over all prompts and label words. **baseline** and **null-input** are zero-resource classification methods, **prior matching** uses the text inputs but not labels, while **optimal** is an oracle approach that uses the labels to search for the best thresholds. Results for FlanT5 large

Inspired by Zhao et al. (2021), we consider a resource-free approximation of the word prior (equation 8) by considering the output word probabilities of the null input \emptyset (i.e. an empty string).

$$P_{\theta}(w_k|Q) \approx P_{\theta}(w_k|\mathbf{p}(\emptyset)) \quad (9)$$

This enables a zero-resource approximation of weight parameters $\bar{\alpha}$.

3 Experiments

3.1 Experimental Setup

Data Experimental results are run on standard NLP benchmarks, including sentiment classification (*IMDB* (Maas et al., 2011), *Rotten Tomatoes* (Pang and Lee, 2005) and *amazon*), natural language inference (*SNLI* (Bowman et al., 2015) and *MNLI* (Williams et al., 2018)) and paraphrase detection (*QQP* (Wang et al., 2018)). Evaluation is reported on standard test sets, except for amazon polarity where 5000 test examples were randomly sampled.

Models We use FlanT5 large² (Chung et al., 2022), a T5 with further instruction-tuning stage where the system was trained in a multi-task fashion over 1,836 tasks, each prepended with a natural instruction prompt. This work evaluates FlanT5 for different NLP tasks with arbitrary prompting set ups. For each task, we select 6 prompt templates and for binary classification tasks consider 25 possible class words pairs, while for NLI we have 64 class word triplets (where all permutations of valid class words are considered). All prompts and label words used are given in Appendix B. Further experiments for FlanT5 base and Llama-2-chat can be found in Appendix D.

Methods We consider 4 different methods to leverage LLM probabilities for classification. Class

word probability via equation 1 (**baseline**). Normalised probabilities calculated using null-inputs priors via equation 9 (**null-input**). Optimising α_k with a search to have unbiased class prior via equation 7 (**prior-match**). The oracle upper-bound performance, found by optimising the optimal accuracy threshold via equation 4 (**optimal**).

3.2 Experimental Results

Classification Robustness Table 1 shows the mean and standard deviation of accuracies among all prompt and class word settings for a given task. We observe large consistent gains from both re-weighting approaches, with prior-matching increasing baseline accuracy by between 6.7% to 12.1% for sentiment classification, 13.7% for qqp, and over 25% for natural language inference. Prior-matching also demonstrates performance very similar to the oracle upper-bound, often within 1%, showing that the unsupervised prior-match approach is competitive with the supervised threshold search. Prior-matching also performs better than null-input by a small margin in all tasks, where this small gap confirms that the word-prior normalisation is a very reasonable zero-shot approximation.

Prompt Robustness Figure 2 illustrates a box-plot of rotten tomatoes performance over all class-words for each considered method, over all 6 prompts. As observed in Table 1, naively using raw label word probabilities (dark blue) leads to considerable fluctuations in accuracy; some prompt and label word settings lead to reasonable accuracy (92%+ accuracy), however there is observed brittleness to label word choice, with many settings demonstrating poor performance. Prior matching (green) leads to significant robustness, with nearly all sensible settings above 85% accuracy. We further find that, as shown in Table 1, the unsupervised approach has accuracies very comparable to those

²<https://huggingface.co/google/flan-t5-large>

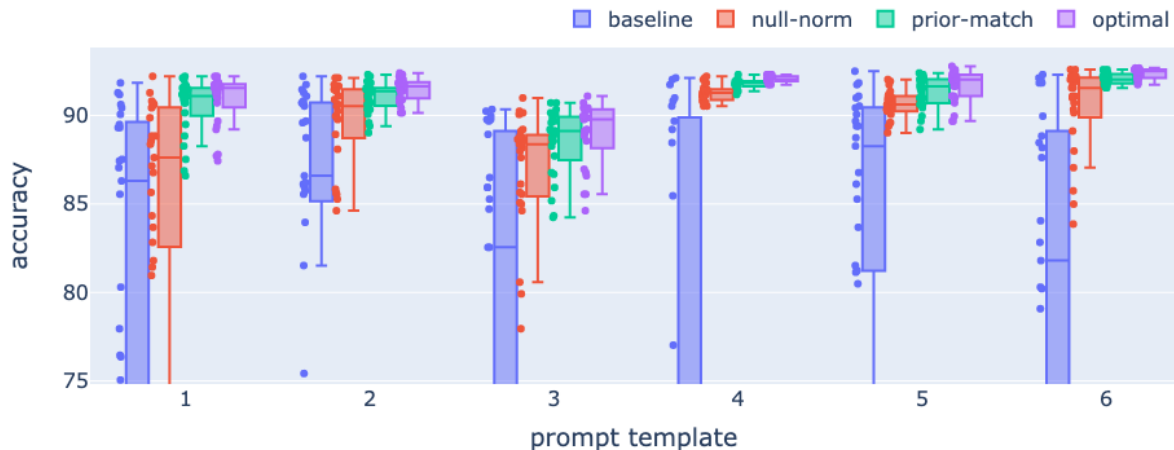


Figure 2: boxplots of the accuracy of all label-word pairs for **rotten tomatoes**, over all the considered prompts

when using optimal thresholds.

In Figure 3, we consider similar boxplots for SNLI and observe larger gains through reweighting. This was as higher probabilities are often assigned to the entailment and contradiction labels words, leading to under-classification of the neutral class. We observe greater sensitivity to prompt choice and label words for snli than as observed in rotten tomatoes, even with reweighting.

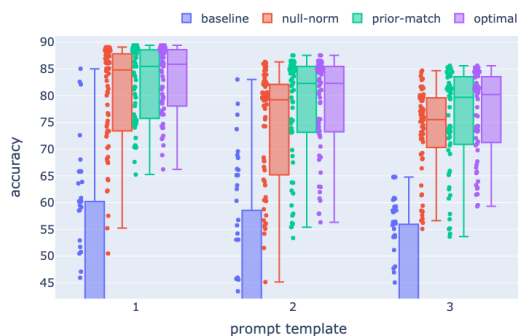


Figure 3: boxplots of the accuracy of all label-word sets for **snli**, for the first 3 prompts

Weight Alignment Figure 4 shows a scatter plot of the weights found by the optimal threshold search α^* (equation 4), with those found from the unsupervised prior matching method $\bar{\alpha}$ (equation 7) and the zero-resource word prior approximation (equation 9). We see a clear linear relationship between optimal and prior-match, illustrating that accounting for the marginal bias is almost equivalent with maximising accuracy, however, achieved in an unsupervised fashion. Null-input is also well correlated with the optimal thresholds,

but there is a less direct relationship. Similar linear relationships are observed also for other binary-classification tasks and prompts, as shown in Appendix C.

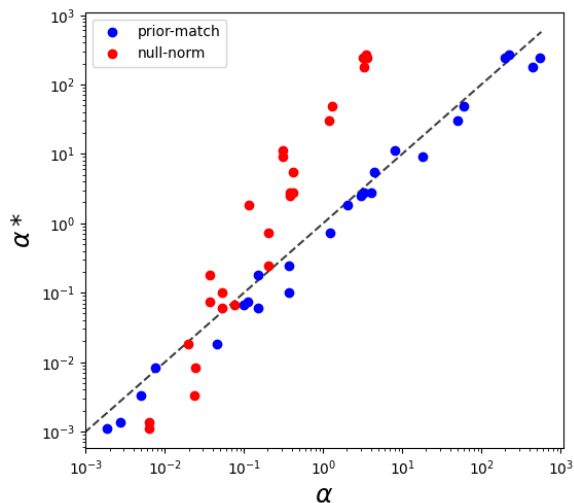


Figure 4: Scatter plot of the optimal weights α^* (equation 4) with the prior match weights $\bar{\alpha}$ (equation 7) and the approximation via null-input (equation 9), for all settings of prompt 1 on **amazon**

4 Conclusions

This paper analyzes prompt-based classifiers and demonstrates that inherent class bias is a significant factor that influences the sensitivity of the system to prompt and label words. We propose an unsupervised approach of prior matching, which we demonstrate performs competitively to the supervised alternative of searching for optimal thresholds, while avoiding the need for labelled data. We

relate prior matching with word biases, and motivate a zero-resource approach of debiasing model probabilities. We show that our methods lead to practical approaches that reduce the sensitivity to design choices such as prompts and label words.

Limitations

This work considered sentiment classification, natural language inference, and paraphrase detection, and could have been extended over a greater suite of tasks to guarantee its effectiveness. Further, this paper ran experiments on FlanT5 and Llama2, and this work has not yet explored a larger range of prompted language models. FlanT5 has also been instruction-tuned on similar tasks, so the findings may be limited in scenarios where known capabilities have to be elicited from models robustly.

Ethical Considerations

Though this work suggests methods to improve the robustness of prompt-based classifiers to prompts and label words, this does not imply that all design choices will work. In some set ups, the system may be ineffective and have poor generalisation over the task. Deploying machine learning classifiers in real-world classification settings has many associated risks, and careful analysis should be made before deploying such systems.

Acknowledgements

This work is supported by Cambridge University Press & Assessment (CUP&A), a department of The Chancellor, Masters, and Scholars of the University of Cambridge, and the Cambridge Commonwealth, European & International Trust.

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. 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.
- 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*.
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners. 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)*, pages 3816–3830.
- Matt Gardner, William Merrill, Jesse Dodge, Matthew E Peters, Alexis Ross, Sameer Singh, and Noah A Smith. 2021. Competency problems: On finding and removing artifacts in language data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1801–1813.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. Surface form competition: Why the highest probability answer isn’t always right. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7038–7051.
- Zhengbao Jiang, Frank F Xu, Jun Araki, and Graham Neubig. 2020. How can we know what language models know? *Transactions of the Association for Computational Linguistics*, 8:423–438.
- Adian Liusie, Vatsal Raina, Vyas Raina, and Mark Gales. 2022. [Analyzing biases to spurious correlations in text classification tasks](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 78–84, Online only. Association for Computational Linguistics.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. [Learning word vectors for sentiment analysis](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Bo Pang and Lillian Lee. 2005. [Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales](#). In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, pages 115–124, Ann Arbor, Michigan. Association for Computational Linguistics.

Timo Schick and Hinrich Schütze. 2021. Exploiting cloze-questions for few-shot text classification and natural language inference. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269.

Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. 2020. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. In *International Conference on Learning Representations*.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. **A broad-coverage challenge corpus for sentence understanding through inference**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

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.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.

A Derivation of Zero-Resource Equation

For a prompt classifier $Q = \{\mathbf{p}, \{w_k\}_{1:K}\}$, class probabilities are assumed to be proportional to the probability of the associated class word,

$$\tilde{P}_\theta(y_k|\mathbf{x}, Q) = \frac{P_\theta(w_k|\mathbf{p}(\mathbf{x}))}{\sum_{w_i} P_\theta(w_i|\mathbf{p}(\mathbf{x}))} \quad (10)$$

Given the task dataset $\mathcal{D} = \{\{\mathbf{x}^{(j)}, y^{(j)}\}\}_{j=1}^N$, one can calculate the assumed prior of the prompt classifier over the output classes,

$$\tilde{P}_\theta(y_k|Q) = \mathbb{E}_{\mathbf{x}}\{\tilde{P}_\theta(y_k|Q, \mathbf{x})\} \quad (11)$$

$$\approx \frac{1}{N} \sum_{j=1}^N \tilde{P}_\theta(y_k|Q, \mathbf{x}^{(j)}) \quad (12)$$

$$= \frac{1}{N} \sum_{j=1}^N \frac{P_\theta(w_k|\mathbf{p}(\mathbf{x}^{(j)}))}{\sum_{w_i} P_\theta(w_i|\mathbf{p}(\mathbf{x}^{(j)}))} \quad (13)$$

This can be compared to the actual prior of the task/domain,

$$P(y_k) \approx P(y_k|\mathcal{D}) = \frac{1}{N} \sum_{j=1}^N \mathbb{1}(y, y^{(j)}) \quad (14)$$

If \mathcal{D} is sufficiently large, then an unbiased classifier should have a class prior similar to that approximated via the labels. However, if they diverge, one may wish to debias the classifier by scaling class probabilities by factors α_k ,

$$\hat{P}_\theta(y_k|\mathbf{x}, Q, \boldsymbol{\alpha}) = \frac{\alpha_k \tilde{P}_\theta(y_k|\mathbf{x}, Q)}{\sum_i \alpha_i \tilde{P}_\theta(y_i|\mathbf{x}, Q)} \quad (15)$$

$$= \frac{\alpha_k P_\theta(w_k|\mathbf{x}, Q)}{Z(\mathbf{x}, Q, \boldsymbol{\alpha})} \quad (16)$$

Where $Z(\mathbf{x}, Q, \boldsymbol{\alpha}) = \sum_i \alpha_i P_\theta(w_i|\mathbf{x}, Q, \boldsymbol{\alpha})$ and $P_\theta(w_k|\mathbf{x}, Q) \equiv P_\theta(w_k|\mathbf{p}(\mathbf{x}))$. The parameters $\bar{\boldsymbol{\alpha}}$ that lead to an unbiased classifier can then be determined in a deterministic fashion.

$$\bar{\boldsymbol{\alpha}} = \underset{\boldsymbol{\alpha}}{\operatorname{argmin}} \sum_{\forall y_k} |\hat{P}_\theta(y_k|Q, \boldsymbol{\alpha}) - P(y_k)| \quad (17)$$

Note that by constraining $\alpha_1 = 1$, there will exist a deterministic solution that ensures that $\hat{P}_\theta(y_k|Q, \boldsymbol{\alpha}) = P(y_k)$. For given weight parameters $\boldsymbol{\alpha}$, Consider the prompt-classifier priors, $\hat{P}_\theta(y_k|Q, \boldsymbol{\alpha})$. One can approximate this using a Taylor series of the expectation of a ratio, yielding

$$\hat{P}_\theta(y_k|Q, \boldsymbol{\alpha}) = \mathbb{E}_{\mathbf{x}}[\hat{P}_\theta(y_k|\mathbf{x}, Q, \boldsymbol{\alpha})] \quad (18)$$

$$= \mathbb{E}_{\mathbf{x}}\left[\frac{\alpha_k P_\theta(w_k|\mathbf{x}, Q)}{Z(\mathbf{x}, Q, \boldsymbol{\alpha})}\right] \quad (19)$$

$$\approx \frac{\mathbb{E}_{\mathbf{x}}[\alpha_k P_\theta(w_k|\mathbf{x}, Q)]}{\mathbb{E}_{\mathbf{x}}[Z(\mathbf{x}, Q, \boldsymbol{\alpha})]} \quad (20)$$

$$= \frac{\alpha_k P_\theta(w_k|Q)}{Z(Q, \boldsymbol{\alpha})} \quad (21)$$

By equating the predicted prior with the true prior, we find an approximation for α_k

$$\hat{P}_\theta(y_k|Q) = P(y_k|\mathcal{D}) \quad (22)$$

$$\frac{\alpha_k P_\theta(w_k|Q)}{Z(Q)} = P(y_k|\mathcal{D}) \quad (23)$$

$$\Rightarrow \alpha_k = \frac{Z(Q) \cdot P(y_k|\mathcal{D})}{P_\theta(w_k|Q)} \quad (24)$$

A final insight is that in many cases it is assumed that there should be no inherent class bias, and so $P(y_k|\mathcal{D})$ can be assumed to be uniform and be included in the normalisation term.

B Prompts and Label Words

B.1 Sentiment Classification

prompt
classify the following review: how was the movie? which word best describes the text? what is the sentiment? what is the reviewer's verdict? is the following movie good or bad?

Table 2: sentiment classification prompts

positive	negative
good	bad
great	terrible
amazing	poor
fantastic	horrible
positive	negative

Table 3: label words for sentiment classification

B.2 Natural Language Inference

prompt
is the second text an entailment of the first text? does the second text directly follow from the first text? are the texts related? are the texts consistent? does text 1 imply text 2? can text 2 be logically derived from text 1? does the hypothesis logically follow the premise?

Table 4: NLI prompts

entailment	neutral	contradiction
yes	maybe	no
correct	unclear	incorrect
yeah	potentially	nope
follows	neutral	contradiction

Table 5: label words for NLI

B.3 Paraphrase Identification

prompt
is the second text a paraphrase of the first text? are the two texts semantically equivalent? are the texts paraphrases of each other? do the texts have the same meaning? is the meaning of text 1 the same as in text 2? would the two texts be classified as paraphrases?

Table 6: NLI prompts

paraphrase	not paraphrase
yes	no
correct	incorrect
yeah	not
positive	negative
true	false

Table 7: label words for sentiment classification

C Threshold Alignment Plots

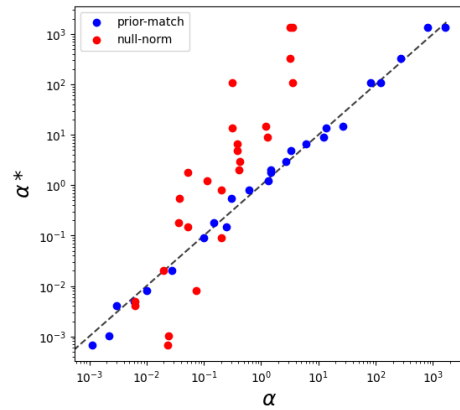


Figure 5: weights alignment plot for **rotten tomatoes**

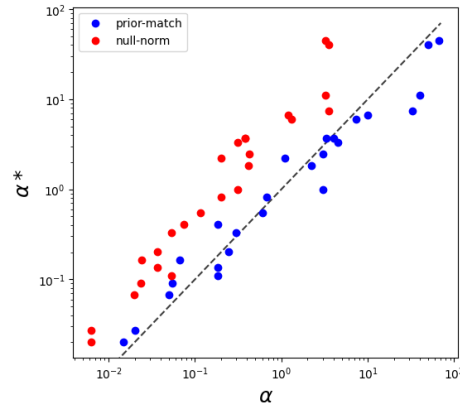


Figure 6: weights alignment plot for **imdb**

D Impact of LLM Choice

method	inputs	labels	imdb	rt	amazon	snli	mnli	qqp
baseline	✗	✗	82.1 \pm 11.1	70.1 \pm 11.7	83.4 \pm 12.8	37.4 \pm 6.1	37.0 \pm 4.3	52.5 \pm 11.4
null-input	✗	✗	87.5 \pm 3.2	78.5 \pm 4.8	91.1 \pm 2.3	41.8 \pm 3.8	40.2 \pm 4.1	53.9 \pm 10.2
prior-match	✓	✗	89.1 \pm 2.4	80.8 \pm 2.9	92.0 \pm 1.3	44.7 \pm 6.3	41.8 \pm 3.8	58.5 \pm 5.5
optimal	✓	✓	89.3 \pm 2.0	81.2 \pm 2.9	92.1 \pm 1.4	47.6 \pm 5.9	43.5 \pm 3.7	65.3 \pm 2.9

Table 8: Robustness performance when using FlanT5 base as the base LLM (set-up equivalent to Table 1).

method	inputs	labels	imdb	rt	amazon	snli	mnli	qqp
baseline	✗	✗	85.8 \pm 8.7	78.4 \pm 10.3	86.4 \pm 10.3	35.1 \pm 2.7	36.9 \pm 3.2	51.0 \pm 11.6
null-input	✗	✗	87.4 \pm 6.9	83.2 \pm 6.6	90.7 \pm 6.4	37.9 \pm 5.4	39.4 \pm 3.7	51.8 \pm 8.4
prior-match	✓	✗	90.5 \pm 3.1	86.3 \pm 3.7	93.1 \pm 2.4	39.5 \pm 5.4	41.2 \pm 3.1	52.6 \pm 1.9
optimal	✓	✓	90.8 \pm 2.8	86.7 \pm 3.6	93.2 \pm 2.4	42.8 \pm 4.6	42.8 \pm 2.3	66.8 \pm 0.4

Table 9: Robustness performance when using Llama-2-chat 7B as the base LLM.

Tables 8 and 9 show the prompt-based classifier performance for the different methods when using different FlanT5-base and Llama-2-chat 7B respectively. For sentiment classification and natural language tasks, we similarly observe that the various re-weighting methods lead to considerable boosts in accuracy. Both null-norm and prior-match again lead to performance near that to the optimal weights, with considerable performance boost over the baseline. However, for paraphrase detection, we only observe moderate performance boosts over the baseline setting with a larger performance discrepancy with the optimal weights.

E Boxplots

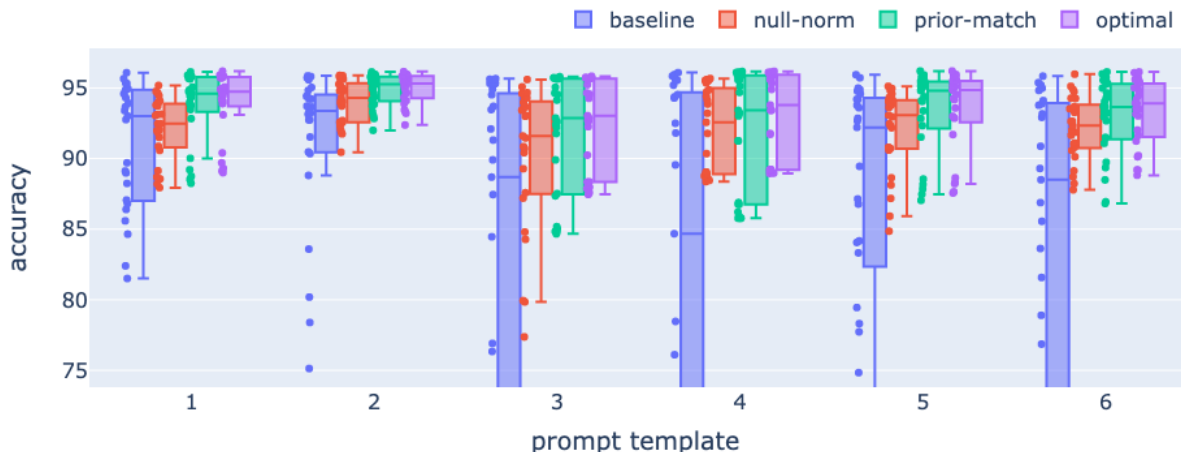


Figure 7: boxplots of the accuracy of all label-word pairs on **IMDB**, over all the considered prompts

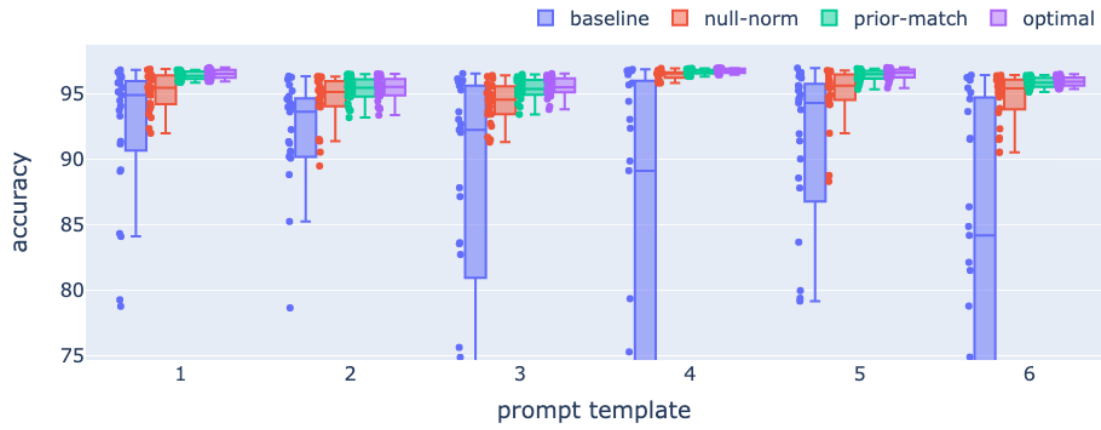


Figure 8: boxplots of the accuracy of all label-word pairs on **amazon**, over all the considered prompts

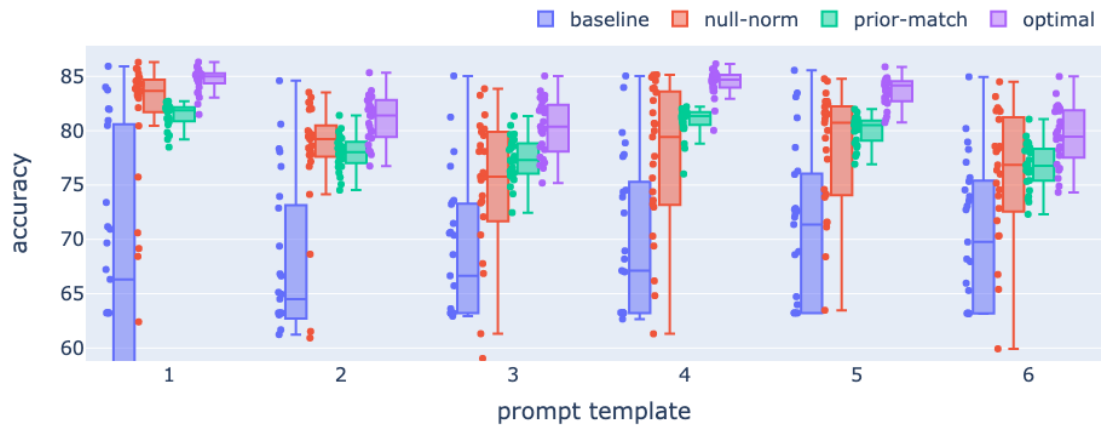


Figure 9: boxplots of the accuracy of all label-word pairs on **qqp**, over all the considered prompts