
A Study on the Calibration of In-context Learning

Hanlin Zhang¹ Yi-Fan Zhang² Yaodong Yu³
Dhruv Madeka⁴ Dean Foster⁴ Eric Xing^{5,6} Himabindu Lakkaraju⁴ Sham Kakade^{1,4}
¹Harvard University ²Chinese Academy of Sciences ³UC Berkeley ⁴Amazon
⁵Carnegie Mellon University ⁶Mohamed Bin Zayed University of Artificial Intelligence

Abstract

Accurate uncertainty quantification is crucial for the safe deployment of language models (LMs), and prior research has demonstrated improvements in the calibration of modern LMs. Our study focuses on in-context learning (ICL), a prevalent method for adapting static LMs through tailored prompts, and examines the balance between performance and calibration across a broad spectrum of natural language understanding and reasoning tasks. Through comprehensive experiments, we observe that, with an increasing number of ICL examples, models initially exhibit increased miscalibration before achieving better calibration and miscalibration tends to arise in low-shot settings. Moreover, we find that methods aimed at improving usability, such as fine-tuning and chain-of-thought (CoT), can lead to miscalibration and unreliable natural language explanations, suggesting that new methods may be required for scenarios where models are expected to be reliable.

1 Introduction

Language models (LMs) that encompass transformer-based architectures (Brown et al., 2020; Chowdhery et al., 2023; OpenAI, 2023) can generate coherent and contextually relevant texts for various use cases. Despite their impressive performance, these models occasionally produce erroneous or overconfident outputs, leading to concerns about their calibration (Dawid, 1982; DeGroot and Fienberg, 1983) which measures how faithful a model’s prediction uncertainty is. Such a problem is pressing when users adapt them using a recent paradigm called in-context learning (Brown et al., 2020) to construct performant predictors, especially for applications in safety-critical domains (Bhatt et al., 2021; Kadavath et al., 2022; Pan et al., 2023).

We provide an in-depth evaluation and analysis of how well these models are calibrated - that is, the alignment between the model’s confidence in its predictions and the actual correctness of those predictions. This token-level calibration assessment enables us to measure the discrepancy between the model’s perceived and actual performance through a Bayesian uncertainty lens, providing a valuable metric for assessing the model’s accuracy and reliability.

We find that LM such as LLaMA (Touvron et al., 2023a) is poorly calibrated in performant settings and there exists a calibration-accuracy trade-off (Fig.1) for low-shot settings ($k < 4$), i.e. as we increase the amount of in-context samples, both prediction accuracy and calibration error increase. Such a trade-off can be improved using more ICL examples ($k = 8$) and larger models. Crucially, this calibration degradation worsens when fine-tuning occurs using specialized data to improve usability, such as curated instructions (Dubois et al., 2023), dialogues (Zheng et al., 2023), or human preference data (Ziegler et al., 2019). Though previous work (Braverman et al., 2020) shows the entropy of each generation step is drifting and can be recalibrated via scaling techniques (Platt et al., 1999) such as temperature scaling (Guo et al., 2017), we show that in contrast to classic regimes, the miscalibration issue in ICL can not be easily addressed using such well-established recalibration approaches that rely on additional validation data.

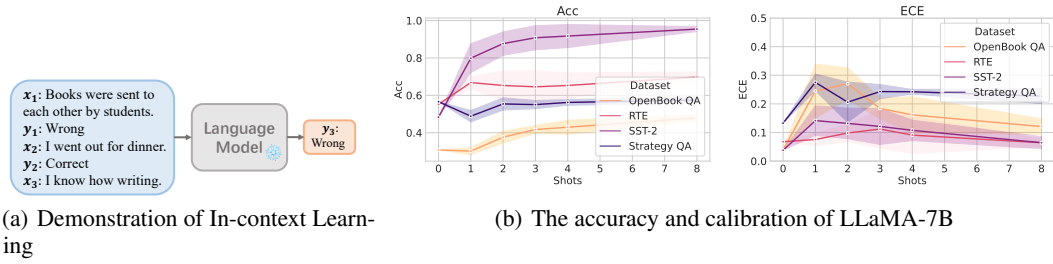


Figure 1: **The accuracy-calibration trade-off of in-context learning.** (a) ICL concerns taking task-specific examples as the prompt to adapt a frozen LLM to predict the answer. (b) Classification accuracy and expected calibration error of ICL. As the number of ICL samples increases, the prediction accuracy improves (**Left**); at the same time, the calibration first worsens ($k < 3$) and then becomes better (**Right**).

Table 1: **Accuracy and Calibration** of LLaMA-30B model with three sizes across four text classification datasets and four reasoning datasets. Results are excluded when the data exceeds the context length limit.

Dataset	LLaMA-30B											
	0-shot		1-shot		2-shot		3-shot		4-shot		8-shot	
	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc
Text Classification												
AGNews	0.261	0.37	0.043	0.830	0.049	0.817	0.067	0.810	0.049	0.821	0.047	0.855
RTE	0.023	0.672	0.051	0.742	0.060	0.747	0.050	0.738	0.048	0.748	0.058	0.752
CB	0.069	0.500	0.312	0.696	0.216	0.789	0.217	0.834	0.192	0.814	0.181	0.796
SST-2	0.083	0.607	0.163	0.930	0.139	0.940	0.126	0.961	0.112	0.964	0.080	0.964
Reasoning with Scratchpad												
Strategy QA	0.204	0.450	0.154	0.619	0.174	0.654	0.172	0.660	0.161	0.672	0.152	0.665
Commonsense QA	0.048	0.356	0.232	0.589	0.290	0.608	0.253	0.675	0.283	0.644	0.289	0.653
World Tree	0.112	0.534	0.211	0.570	0.251	0.621	0.185	0.680	0.206	0.646	-	-
OpenBook QA	0.036	0.386	0.231	0.561	0.255	0.604	0.207	0.644	0.206	0.648	0.191	0.662

Moreover, we study the trade-off in reasoning tasks that involve generation of explanations (Camburu et al., 2018; Nye et al., 2021; Wei et al., 2022) before the answer, showing that the model can produce confidently wrong answers (using confidence histograms and reliability plots) when prompted with explanations on Strategy QA (Geva et al., 2021), Commonsense QA (Talmor et al., 2018), OpenBook QA (Mihaylov et al., 2018), World Tree (Jansen et al., 2018). We carefully design our human evaluation and observe that, with the increase in model sizes and the quantity of ICL examples, there is a corresponding rise in the proportion of confidently predicted examples among those incorrectly forecasted. Moreover, we find that a high proportion of wrong predictions are of high confidence and showcase those typical confidently wrong examples of LMs.

In-context learning has been hypothesized as meta-learning by gradient descent in their forward pass (Von Oswald et al., 2023), which might hopefully yield calibrated predictions (Błasiok et al., 2023) if the models are getting close to local optimality with respect to test loss through meta-optimization. However, the fact that choosing ICL samples from the validation set does not naturally lead to calibrated predictions shows that ICL learns in a fairly different way than SGD. We design controlled experiments to illustrate the task learning properties of ICL, showing that when examples in the prompt demonstrate consistent task properties, the learning performance would be improved.

2 Background

Setting. Given a pre-trained language model $\mathcal{P}_\theta(w_t|w_{<t})$, we seek to adapt it using the prompt $w_0 = [x_1, y_1, x_2, y_2, \dots, x_{n-1}, y_{n-1}, x_n]$ to generate a predicted answer $y_n = \mathcal{P}_\theta(w_0)$. In the context of reasoning, a popular approach is to hand-craft some explanations/rationales/chain-of-thoughts e in the prompt $w_0 = [x_1, e_1, y_1, x_2, e_2, y_2, \dots, x_{n-1}, e_{n-1}, y_{n-1}, x_n]$ to generate explanation e_n and answer y_n , for the test sample: $\overbrace{w_1, w_2, \dots, w_k}^{e_n}, y_n = \mathcal{P}_\theta(w_0)$.




Figure 2: Reliability plots and confidence histograms of LLaMA models on 4-shot learning tasks. Results of different sizes 7B (left), 13B (middle), and 30B (right) are plotted.

We extract answer token probabilities of LMs, e.g. for binary classification tasks, we filter and extract probabilities $P(\text{"Yes"})$ and $P(\text{"No"})$, based on which we calculate the following statistics for studying the confidence and calibration of LMs:

Confidence and feature norm. We record the maximum probability of the answer token as its confidence $\text{Conf} = \mathcal{P}_\theta(y_n|w_{<n})$ and the feature norm z_n as the intermediate hidden state before the linear prediction layer.

Entropy rate. We denote the entropy of a token w_t at position t as $H(w_t|w_{<t}) = -\mathbb{E}_{w_t \sim \mathcal{P}_\theta(\cdot|w_{<t})}[\log \mathcal{P}_\theta(w_t|w_{<t})]$. We typically measure it based on the answer token via setting $w_t = y_n$. Note that auto-regressive LMs are trained via maximizing the negative log-likelihood objective $\mathcal{L} = -\mathbb{E}_t[\log \mathcal{P}_\theta(w_t|w_{<t})]$ on massive corpora.

Empirical estimate of the expected calibration error (ECE) In the realm of probabilistic classifiers, calibration is a crucial concept. A classifier, denoted as \mathcal{P}_θ with parameters θ and operating over C classes, is said to be "canonically calibrated" when, for every probability distribution p over the C classes and for every label y , the probability that the label is y given the classifier's prediction is p matches the component of p corresponding to y . This is mathematically represented as:

$$\forall p \in \Delta^{C-1}, \forall y \in Y : P(Y = y | \mathcal{P}_\theta(X) = p) = p_y. \quad (1)$$

Here, Δ^{C-1} symbolizes the $(C-1)$ -dimensional simplex, which encompasses all potential probability distributions over the C classes.

A simpler calibration criterion is the "top-label calibration." In this case, a classifier is deemed calibrated if, for every top predicted probability p^* , the probability that the true label belongs to the class with the highest predicted probability, given that this maximum predicted probability is p^* , equals p^* . Formally:

$$\forall p^* \in [0, 1] : P(Y \in \arg \max p | \max \mathcal{P}_\theta(X) = p^*) = p^*. \quad (2)$$

To gauge the calibration of a model, we adopt Expected Calibration Error (ECE) defined as:

$$\mathbb{E}[|p^* - \mathbb{E}[Y \in \arg \max \mathcal{P}_\theta(X) | \max \mathcal{P}_\theta(X) = p^*]|]. \quad (3)$$

In real-world applications, this quantity cannot be computed without quantization. So, the ECE is approximated by segmenting predicted confidences into M distinct bins, B_1, \dots, B_M . The approximation is then computed as:

$$\widehat{\text{ECE}} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)|.$$

Here, $\text{acc}(B_m)$ is the accuracy within bin B_m , and $\text{conf}(B_m)$ is the average confidence of predictions in bin B_m . The total number of samples is represented by n , and the dataset consists of n independent and identically distributed samples, $\{(x_i, y_i)\}_{i=1}^n$. In our work, we use this estimator to approximate the ECE.

Table 2: **Acc** and **ECE** of LLaMA-7B model on SST-2 with different prompt repetition strategies.

	1-shot		2-shot		4-shot		8-shot		Avg Acc	Avg ECE
	ACC	ECE	ACC	ECE	ACC	ECE	ACC	ECE		
Vanilla	0.740	0.098	0.877	0.132	0.917	0.108	0.954	0.064	0.872	0.100
Repeat prompt	0.740	0.098	0.693	0.155	0.801	0.117	0.820	0.111	0.764	0.120
Repeat context	0.740	0.098	0.668	0.208	0.657	0.220	0.607	0.219	0.668	0.186

3 Experiments

We briefly summarize our results and findings in Appendix B.1.

3.1 Experimental Settings

Models. We study decoder-only autoregressive LMs involving LLaMA (Touvron et al., 2023a), ranging from 7B to 30B, and its variants fine-tuned with instruction, dialog, or RLHF like Alpaca (Dubois et al., 2023), Vicuna (Zheng et al., 2023), and LLaMA2-Chat (Touvron et al., 2023b).

Datasets and tasks. We used both traditional NLU tasks such as AGNews (Zhang et al., 2015), TREC (Voorhees and Tice, 2000), CB (Schick and Schütze, 2021), SST-2 (Socher et al., 2013), DBPedia (Zhang et al., 2015), as well as reasoning question answering tasks like Strategy QA (Geva et al., 2021), Commonsense QA (Talmor et al., 2018), OpenBook QA (Mihaylov et al., 2018), World Tree (Jansen et al., 2018). Notably, the reasoning task performance can be greatly improved in general via prompting methods like scratchpad (Nye et al., 2021; Wei et al., 2022) that enables models to generate natural language explanations before predicting an answer.

In-context learning settings. For k -shot learning, we prompt the model via sampling k examples from the training set for each test example. Each experiment is repeated 10 times to reduce variance and we report the mean results. We use $M = 10$ bins for calculating calibration errors.

3.2 Numerical Results

Performance of LLaMA. We record the performance and calibration errors for k -shot learning ($k = 0, 1, 2, 3, 4, 8$), characterizing the calibration-accuracy trade-off in both classic and realistic settings (Tab. 1). Our findings are two-fold: as more in-context examples are included, we observe a concurrent rise in both accuracy and calibration error across most low-shot situations. Especially, when $k = 0$ increases to $k = 1$, there is a marked boost in both accuracy and calibration error, demonstrating the importance of in-context examples in learning performance while one single example may not be able to reduce aleatoric uncertainty. Moreover,

In particular, for reasoning tasks, we explore prompting approaches that explicitly include explanations in reasoning tasks, i.e. scratchpad (Nye et al., 2021) or chain-of-thought (Wei et al., 2022), showing that calibration significantly degrades after generating a long context for reasoning and explaining the final answer.

Post-hoc recalibration. We experiment with three strategies (Algorithm 1) in applying temperature scaling methods (Guo et al., 2017) to fix miscalibration:

1. (**0-shot**) Learn a temperature from the training split and apply it to all test samples with different shot numbers;
2. (**k -shot**) We learn one temperature for each k -shot ICL, i.e., we learn different temperatures for different shot numbers in ICL;
3. (**Fix w**) For each experiment, we fix the prompt and learn the temperature for the fixed prompt. That is, for every possible ICL prompt, we learn a corresponding temperature for calibration.

The training algorithm is shown in Appendix Alg. 1. Looking into Tab. 3, none of the above strategies achieves satisfactory calibration performance, which is in contrast to the well-studied supervised learning setting where scaling the confidence scores can effectively reduce calibration errors (Guo et al., 2017). The fact that applying a post-processing calibration method, such as temperature scaling,

Table 3: ECE for different calibration strategies of base models LLaMA-7B across various shot settings.

Dataset	Strategy	Shot						Avg
		0	1	2	3	4	8	
SST-2	Vanilla	0.038	0.098	0.133	0.136	0.124	0.064	0.099
	0-shot	0.026	0.157	0.265	0.275	0.282	0.222	0.205
	k -shot	0.026	0.104	0.103	0.101	0.087	0.035	0.076
	Fix w	0.027	0.079	0.056	0.051	0.051	0.030	0.049
CB	Vanilla	0.133	0.274	0.164	0.148	0.192	0.206	0.186
	0-shot	0.126	0.272	0.143	0.142	0.183	0.177	0.174
	k -shot	0.126	0.247	0.137	0.153	0.203	0.251	0.186
	Fix w	0.138	0.183	0.127	0.141	0.150	0.194	0.155
RTE	Vanilla	0.130	0.103	0.122	0.147	0.130	0.083	0.119
	0-shot	0.117	0.100	0.130	0.154	0.130	0.094	0.121
	k -shot	0.114	0.106	0.130	0.161	0.121	0.104	0.123
	Fix w	0.114	0.078	0.097	0.104	0.079	0.079	0.092
AGNews	Vanilla	0.086	0.077	0.166	0.203	0.088	0.063	0.114
	0-shot	0.101	0.087	0.149	0.198	0.117	0.078	0.122
	k -shot	0.114	0.085	0.161	0.203	0.075	0.067	0.117
	Fix w	0.132	0.081	0.113	0.126	0.089	0.079	0.103

cannot directly resolve the miscalibration issue suggests that ICL might have substantially different properties compared to making predictions via classical supervised learning models.




Figure 3: Accuracy and calibration errors of base models LLaMA and Mistral, as well as fine-tuned variants. The reported values for Acc and ECE represent averages across experiments conducted with 0, 1, ..., 8 shots.




Figure 4: **Illustration of confidence distribution:** the number of samples whose confidence is greater than a threshold on Commonsense QA.

The effect of fine-tuning. We show that vicuna, alpaca, and LLaMA2-Chat are all more accurate but less calibrated than their LLaMA counterpart backbones (Fig. 3), the margin is especially large for reasoning tasks and vicuna. Our finding indicates that fine-tuning might significantly degrade calibration, corroborating the evidence reported in GPT-4 (OpenAI, 2023), albeit it can improve the reasoning accuracy dramatically. Our results provide evidence that though fine-tuned on carefully curated datasets can greatly improve question-answering performance, especially for hard tasks like reasoning problems, attention may need to be paid when assessing the calibration of those models’ predictions. We also note that the pattern of zero-shot performance is fairly different for two fine-tuned models, i.e. vicuna, and alpaca. Moreover, we include results of Mistral-7B (Jiang et al., 2023), a sparse Mixture of Experts (MoEs) architecture with sliding window attention. As a base




Figure 5: The number of **wrongly classified** examples whose confidence is above a threshold with different numbers of shots on Commonsense QA.

model, it shows similar performance and calibration compared with LLaMA2-7B, indicating that our conclusion still holds for the model pre-trained with significantly different data and architectures.

The effect of prompt formats. In our study, we explore the effects of different prompt strategies using three distinct methods. We consider predicting the label y_n of test example x_n . First, the *Repeat-context* approach involves constructing prompts as $w_0 = [x_1, x_1, \dots, x_1, y_1, x_n]$, where the context x_1 is repeated $n-1$ times, but the label y_1 is not included in the repetition. Next, the *Repeat-prompt* strategy shapes the prompt as $w_0 = [x_1, y_1, \dots, x_1, y_1, x_n]$, where both the context x_1 and the label y_1 are repeated $n-1$ times. Finally, the *Normal* involves constructing the prompt as $w_0 = [x_1, y_1, x_2, y_2, \dots, x_{n-1}, y_{n-1}, x_n]$, systematically incorporating distinct context-label pairs.

The findings, as detailed in Tab. 2, reveal certain insights: firstly, integrating labels within prompts significantly decreases uncertainty and enhances learning performance. In contrast, simply repeating the context without labels does not lead to better outcomes. Secondly, the diversity of ICL examples in the prompt greatly affects performance, a potential explanation is it promotes better task learning (Pan, 2023). Those observations corroborate that ICL is performant when the number of ICL examples is large and they demonstrate consistent task properties. Importantly, the trade-off persists for different controlled scenarios, i.e. as we increase the number of ICL examples, models tend to be first more miscalibrated and then calibrated.

Table 4: **Norm of representation, entropy, and confidence** of LLaMA-30B model across three text classification datasets.

Dataset	LLaMA-30B											
	Norm				Entropy				Confidence			
	0-shot	1-shot	4-shot	8-shot	0-shot	1-shot	4-shot	8-shot	0-shot	1-shot	4-shot	8-shot
AGNews	78.8	92.3	92.1	92.2	3.920	0.650	0.595	0.444	0.214	0.821	0.819	0.865
CB	88.4	91.7	89.2	87.9	3.857	1.266	0.935	0.823	0.193	0.566	0.629	0.577
DBPdia	77.9	89.5	91.0	90.1	4.105	1.438	0.848	0.718	0.078	0.578	0.705	0.671

3.3 Qualitative Results

Reliability diagram and confidence histogram. A reliability diagram is a graphical tool used to evaluate the calibration of probabilistic predictions of a model across multiple classes; it compares the predicted probabilities of each class against the actual outcomes, with a perfectly calibrated model having its values lie on the diagonal $y = x$ line. A confidence histogram, on the other hand, displays the distribution of the model’s prediction confidences across all classes, showing how often the model predicts certain probabilities.

Recall that we found significant miscalibration for reasoning with CoT settings, therefore we closely examine the poorly calibrated reasoning cases using the above plots (Fig. 2 and Fig. 6). Our results on 4-shot settings show that for the reasoning problems (Strategy QA, Commonsense QA, OpenBook QA, World Tree) we consider, models are consistently over-confident with ECEs above 0.15. Larger models are better both in both ACC and ECE but for OpenBook QA, calibration worsens as the model size increases. Moreover, it’s observed that confidence scores tend to concentrate on high values as we enlarge the model size. Especially in Commonsense QA and OpenBook QA, the confidence level of nearly all predictions of 13B and 30B models predominantly exceeds 0.8.

References

- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2022. What learning algorithm is in-context learning? investigations with linear models. *arXiv preprint arXiv:2211.15661*.
- Anastasios N Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. 2022. Conformal risk control. *arXiv preprint arXiv:2208.02814*.
- Amos Azaria and Tom Mitchell. 2023. The internal state of an llm knows when its lying. *arXiv preprint arXiv:2304.13734*.
- Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. 2023. Transformers as statisticians: Provable in-context learning with in-context algorithm selection. *arXiv preprint arXiv:2306.04637*.
- Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16.
- Umang Bhatt, Javier Antorán, Yunfeng Zhang, Q Vera Liao, Prasanna Sattigeri, Riccardo Fogliato, Gabrielle Melançon, Ranganath Krishnan, Jason Stanley, Omesh Tickoo, et al. 2021. Uncertainty as a form of transparency: Measuring, communicating, and using uncertainty. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 401–413.
- Jarosław Błasiok, Parikshit Gopalan, Lunjia Hu, and Preetum Nakkiran. 2023. When does optimizing a proper loss yield calibration? *arXiv preprint arXiv:2305.18764*.
- Mark Braverman, Xinyi Chen, Sham Kakade, Karthik Narasimhan, Cyril Zhang, and Yi Zhang. 2020. Calibration, entropy rates, and memory in language models. In *International Conference on Machine Learning*, pages 1089–1099. PMLR.
- 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.
- Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z Gajos. 2021. To trust or to think: cognitive forcing functions can reduce overreliance on ai in ai-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–21.
- Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Jeremy R Cole, Michael JQ Zhang, Daniel Gillick, Julian Martin Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. Selectively answering ambiguous questions. *arXiv preprint arXiv:2305.14613*.
- A Philip Dawid. 1982. The well-calibrated bayesian. *Journal of the American Statistical Association*, 77(379):605–610.
- Morris H DeGroot and Stephen E Fienberg. 1983. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 32(1-2):12–22.
- Shrey Desai and Greg Durrett. 2020. Calibration of pre-trained transformers. *arXiv preprint arXiv:2003.07892*.
- Yann Dubois, Xuechen Li, Rohan Taori, Tianyi Zhang, Ishaan Gulrajani, Jimmy Ba, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Alpacafarm: A simulation framework for methods that learn from human feedback](#).

- Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. Mitigating label biases for in-context learning. *arXiv preprint arXiv:2305.19148*.
- Adam Fisch, Tal Schuster, Tommi Jaakkola, and Regina Barzilay. 2020. Efficient conformal prediction via cascaded inference with expanded admission. *arXiv preprint arXiv:2007.03114*.
- Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. 2022. What can transformers learn in-context? a case study of simple function classes. *Advances in Neural Information Processing Systems*, 35:30583–30598.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Ana Valeria González, Gagan Bansal, Angela Fan, Yashar Mehdad, Robin Jia, and Srinivasan Iyer. 2021. Do explanations help users detect errors in open-domain qa? an evaluation of spoken vs. visual explanations. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1103–1116.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.
- Zhixiong Han, Yaru Hao, Li Dong, Yutao Sun, and Furu Wei. 2023. Prototypical calibration for few-shot learning of language models. In *The Eleventh International Conference on Learning Representations*.
- Peter A Jansen, Elizabeth Wainwright, Steven Marmorstein, and Clayton T Morrison. 2018. Worldtree: A corpus of explanation graphs for elementary science questions supporting multi-hop inference. *arXiv preprint arXiv:1802.03052*.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. How can we know when language models know? on the calibration of language models for question answering. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Adam Tauman Kalai and Santosh S. Vempala. 2023. [Calibrated language models must hallucinate.](#)
- Amita Kamath, Robin Jia, and Percy Liang. 2020. Selective question answering under domain shift. *arXiv preprint arXiv:2006.09462*.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*.
- Philippe Laban, Tobias Schnabel, Paul N Bennett, and Marti A Hearst. 2022. Summac: Re-visiting nli-based models for inconsistency detection in summarization. *Transactions of the Association for Computational Linguistics*, 10:163–177.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.
- Kevin Liu, Stephen Casper, Dylan Hadfield-Menell, and Jacob Andreas. 2023. [Cognitive dissonance: Why do language model outputs disagree with internal representations of truthfulness?](#)
- Sabrina J Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. Reducing conversational agents’ overconfidence through linguistic calibration. *Transactions of the Association for Computational Linguistics*, 10:857–872.

- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*.
- Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. 2021. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.
- OpenAI. 2023. Gpt-4 technical report. <https://cdn.openai.com/papers/gpt-4.pdf>.
- Alexander Pan, Jun Shern Chan, Andy Zou, Nathaniel Li, Steven Basart, Thomas Woodside, Hanlin Zhang, Scott Emmons, and Dan Hendrycks. 2023. Do the rewards justify the means? measuring trade-offs between rewards and ethical behavior in the machiavelli benchmark. In *International Conference on Machine Learning*, pages 26837–26867. PMLR.
- Jane Pan. 2023. *What In-Context Learning “Learns” In-Context: Disentangling Task Recognition and Task Learning*. Ph.D. thesis, Princeton University.
- Suzanne Petryk, Spencer Whitehead, Joseph E Gonzalez, Trevor Darrell, Anna Rohrbach, and Marcus Rohrbach. 2023. Simple token-level confidence improves caption correctness. *arXiv preprint arXiv:2305.07021*.
- John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74.
- Allan Raventós, Mansheej Paul, Feng Chen, and Surya Ganguli. 2023. Pretraining task diversity and the emergence of non-bayesian in-context learning for regression. *arXiv preprint arXiv:2306.15063*.
- 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.
- Tal Schuster, Adam Fisch, Jai Gupta, Mostafa Dehghani, Dara Bahri, Vinh Q Tran, Yi Tay, and Donald Metzler. 2022. Confident adaptive language modeling. *arXiv preprint arXiv:2207.07061*.
- Tal Schuster, Adam Fisch, Tommi Jaakkola, and Regina Barzilay. 2021. Consistent accelerated inference via confident adaptive transformers. *arXiv preprint arXiv:2104.08803*.
- Andy Shih, Dorsa Sadigh, and Stefano Ermon. 2023. Long horizon temperature scaling. *arXiv preprint arXiv:2302.03686*.
- Chenglei Si, Navita Goyal, Sherry Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé III, and Jordan Boyd-Graber. 2023. Large language models help humans verify truthfulness—except when they are convincingly wrong. *arXiv preprint arXiv:2310.12558*.
- 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.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2018. Commonsenseqa: A question answering challenge targeting commonsense knowledge. *arXiv preprint arXiv:1811.00937*.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. *arXiv preprint arXiv:2305.14975*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, João Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. Transformers learn in-context by gradient descent. In *International Conference on Machine Learning*, pages 35151–35174. PMLR.
- 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.
- 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*.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 295–305.
- Theodore Zhao, Mu Wei, J Samuel Preston, and Hoifung Poon. 2023. Automatic calibration and error correction for large language models via pareto optimal self-supervision. *arXiv preprint arXiv:2306.16564*.
- Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *ICML*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging llm-as-a-judge with mt-bench and chatbot arena](#).
- Han Zhou, Xingchen Wan, Lev Proleev, Diana Mincu, Jilin Chen, Katherine Heller, and Subhrajit Roy. 2023a. Batch calibration: Rethinking calibration for in-context learning and prompt engineering. *arXiv preprint arXiv:2309.17249*.
- Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023b. Navigating the grey area: Expressions of overconfidence and uncertainty in language models. *arXiv preprint arXiv:2302.13439*.
- Daniel M Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B Brown, Alec Radford, Dario Amodei, Paul Christiano, and Geoffrey Irving. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.

A Related Work

Uncertainty quantification in NLP. Uncertainty quantification in NLP, which often adopts the Bayesian principle to sophisticated methods tailored for neural networks, aims to enhance the reliability of model predictions. This may involve non-trivial designs as directly interpreting language model predictions via probabilities (Kadavath et al., 2022) and linguistic expressions (Lin et al., 2022; Mielke et al., 2022; Zhou et al., 2023b) may inadvertently lead to over-reliance on the model’s uncertainties (Si et al., 2023), thus complicating the establishment of trustworthy common ground between humans and models (Bućinca et al., 2021). Notable recent advancements include employing model confidence as a critical factor in various applications like dialogue generation (Mielke et al., 2022), cascading prediction (Schuster et al., 2021), open-domain QA (Fisch et al., 2020; Angelopoulos et al., 2022), summarization (Laban et al., 2022), language modeling (Schuster et al., 2022), image captioning (Petryk et al., 2023).

Calibration of language models. Calibration is a safety property to measure the faithfulness of machine learning models’ uncertainty, especially for error-prone tasks using LLMs. Previous works find that pre-training (Desai and Durrett, 2020) and explanation (Zhang et al., 2020; González et al., 2021) improves calibration. Models can be very poorly calibrated when we prompt LMs (Jiang et al., 2021), while calibration can also depend on model size (Kadavath et al., 2022). (Braverman et al., 2020) assesses the long-term dependencies in a language model’s generations compared to those of the underlying language and finds that entropy drifts as models such as when GPT-2 generates text. The intricacy of explanations on complementary team performance poses additional challenges due to the overreliance on explanations of users regardless of their correctness (Bansal et al., 2021). (Mielke et al., 2022) gives a framework for *linguistic calibration*, a concept that emphasizes the alignment of a model’s expressed confidence or doubt with the actual accuracy of its responses. The process involves annotating generations with <DK>, <LO>, <HI> for confidence levels, then training the confidence-controlled model by appending the control token <DK/LO/HI> at the start of the output, followed by training a calibrator to predict these confidence levels, and finally predicting confidence when generating new examples. (Tian et al., 2023) finds that asking LMs for their probabilities can be better than using conditional probabilities in a traditional way. LHTS (Shih et al., 2023) is a simple amortized inference trick for temperature-scaled sampling from LMs and diffusion models. To aggregate log probabilities across semantically equivalent outputs, Kuhn et al. (2023) utilize bidirectional entailment through a model to identify outputs that are semantically similar, thereby refining the uncertainty estimation process. (Cole et al., 2023) identifies the calibration challenge in ambiguous QA and distinguishes uncertainty about the answer (epistemic uncertainty) from uncertainty about the meaning of the question (denotational uncertainty), proposing sampling and self-verification methods. (Kamath et al., 2020) trains a calibrator to identify inputs on which the QA model errs and abstains when it predicts an error is likely. (Zhao et al., 2023) proposes the Pareto optimal learning assessed risk score for calibration and error correction but requires additional training. (Kalai and Vempala, 2023) show the trade-off between calibration and hallucination but they didn’t study it in a realistic setting and how the predicted answer’s accuracy would impact those two safety aspects.

In-context Learning. Large models such as GPT-3 (Brown et al., 2020) have demonstrated the potential of in-context learning, a method where the model infers the task at hand from the context provided in the input, without requiring explicit retraining or fine-tuning for each new task. Some recent works attempt to understand ICL through meta-learning (Von Oswald et al., 2023), Bayesian inference (Xie et al., 2021), mechanistic interpretability (Olsson et al., 2022), algorithm selection (Bai et al., 2023), synthetic data and simple function classes (Garg et al., 2022; Akyürek et al., 2022; Raventós et al., 2023). Notably, unlike previous works (Zhao et al., 2021; Han et al., 2023; Fei et al., 2023; Zhou et al., 2023a) that focus on improving task accuracy using the same “calibration” terminology, we study the uncertainty of ICL and measure its trade-off with accuracy.

B Additional Experimental Details

We provide prompts we adopt for experiments in Tab.6. Additional reliability plots are shown in Fig. 6.

B.1 Summarizing main findings

- For the base LMs we considered, they are calibrated when prompting with a sufficient amount of ICL examples to get non-trivial performance.
- As we increase the number of ICL examples, models tend to be first more miscalibrated and then calibrated. In low-shot settings ($k < 4$), models can be mis-calibrated, in part due to poor data (aleatoric) uncertainty.
- Interventions that improve usability such as fine-tuning, and chain-of-thought (CoT) prompting would lead to miscalibration. The generated explanations from CoT can improve predictive results but may not be reliable by human evaluation.

B.2 Ablation Studies

For case studies, we research how miscalibration can impact the selective classification of LMs, where models are supposed to abstain from uncertain predictions in high-stakes settings.

Ablation with model sizes. As we enlarge the size of models, they become more confident and accurate (Fig. 2). As a result, the entropy decreases and ECE increases, showing that token-level calibration might have an inverse scaling relationship with model sizes.

Table 5: Qualitative Results of LLaMA on Commonsense and OpenBook

Dataset	Case	Reasoning	\hat{y}	p
Commonsense	Question: What would you do if you have curiosity but are blind and paralyzed?			
	Options: A) hear news, B) analyse, C) go somewhere, D) examine thing, E) see favorite show			
	Ground Truth: B			
	7B	If you have curiosity but are blind and paralyzed, you would go somewhere to see things. You cannot hear news, analyse, go somewhere, examine thing or see favorite show.	C	0.5046
	13B	If you are blind and paralyzed, you can't see anything. So, you can't go somewhere. You can't examine thing. You can't see favorite show. You can't analyse. So, you can hear news.	A	0.8684
30B	If you are blind and paralyzed, you cannot see your favorite show. You cannot hear news as you are blind. You cannot examine thing as you are paralyzed. You can go somewhere to see things.	C	0.9720	
Open Book	Question: What would you do if you have curiosity but are blind and paralyzed?			
	Options: A) hear news, B) analyse, C) go somewhere, D) examine thing, E) see favorite show			
	Ground Truth: B			
	7B	A girl and her mom have the same hair length.	D	0.6365
	13B	A girl and her mom have the same date of birth.	A	0.9782
30B	A girl and her mom have the same genes.	A	0.9831	

Examine the hidden state and confidence score. To better understand the miscalibration issue of ICL, we conduct fined-grained experiments to examine ICL properties: we measure the norm of the representation vectors¹ for different number of shots in ICL. Meanwhile, we also measure the confidence and entropy of the prediction for y_n , and the results are summarized in Tab. 4. When switching from 0-shot to 1-shot, all three measurements (representation norm, entropy, and confidence) drastically change; on the other hand, when k increases from $1 \rightarrow 4 \rightarrow 8$, the change of measures would become smoother. Our discovery shows that adding in-context example can substantially impact model behaviors while model behaves relatively similar for various shots once the task is specified ($k \neq 0$). Meanwhile, more ICL samples lead to smaller entropy and higher confidence in most cases.

Confidence and wrongly classified reasoning examples. To inspect the failure modes of LMs, we randomly sample 100 reasoning examples of LLaMA and plot the distribution of wrongly predicted samples and the confidence scores via thresholding. Similar to previous observations, as model sizes and the number of ICL examples scale up, LMs would generate more confident samples (Fig. 4 (c, d)). We observe “inverse scaling” behaviors where models with larger sizes are more error-prone and tend to generate more confidently wrong samples (Fig. 5).

¹The representation vector refers to the intermediate output before the linear prediction layer.

Examples of hallucinated explanations for highly confident predictions. Next, we showcase in Tab. 5 that models generate both wrong explanations and incorrect predictions with high confidence. We also observe that most of the wrong predictions are highly confident, thus we manually examine the correctness of explanations on commonsense QA, and found its high correlations with predicted answer accuracy, which is the opposite of token-level explainability that tends to get worse when the accuracy improves.

C Discussion and Concluding Remarks

In our investigation of the token-level calibration of in-context learning in contemporary language models, we illustrate the intricate trade-off between ICL performance and calibration. Our findings underscore the importance of being circumspect in model deployment, as maximizing ICL performance does not invariably translate to improved calibration for low-shot and reasoning settings. As LMs continue to evolve and gain more capabilities such as having long enough context windows that can include the whole training set as in-context examples for some downstream tasks, our result can be pedagogical when users would like to examine their uncertainty through prediction probabilities. Moreover, the work suggests the following future directions:

Understanding the internal mechanism of ICL for calibration. We observe that existing scaling recalibration methods cannot fully resolve the miscalibration issues of ICL, so better understanding and mitigation strategies are needed. A potential approach can be leveraging transparency tools and studying whether predictable errors exist during text generation.

Calibration beyond classification regimes. Our findings indicate that in multi-choice or multi-class classification tasks, even though the calibration of answer tokens may deteriorate in high-performance settings, there may be a positive correlation between accuracy and the correctness of explanations in reasoning tasks. This suggests potential avenues for future research in exploring strategies such as the use of hedging words to express uncertainty and examining their relationship with predictive performance.

Implications in assessing beliefs of LMs. Previous works show that the expected calibration error would decrease monotonically as the number of ICL examples increases (Kadavath et al., 2022) when querying LMs for answer probabilities. However, we find that zero-shot performance might be weak for models less than 30B, and in low-shot settings, calibration errors can sometimes be even worse than zero-shot. This implies that a close examination and careful control of epistemic uncertainty and aleatoric uncertainty can be needed before deriving conclusions in truthfulness (Liu et al., 2023; Azaria and Mitchell, 2023) for low-shot settings.

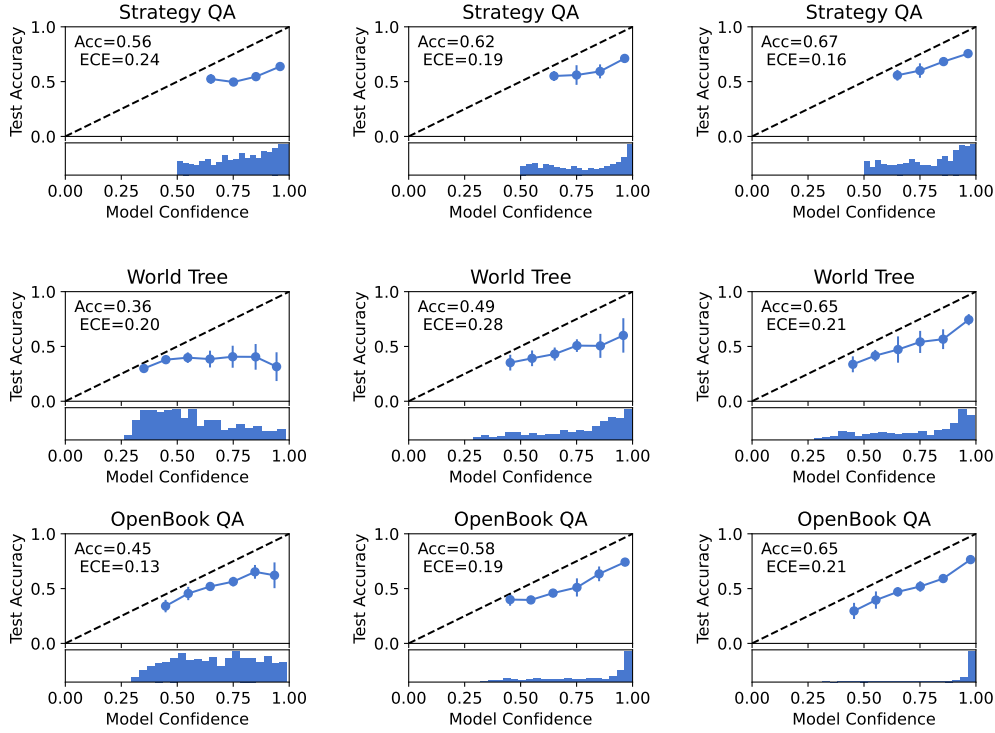


Figure 6: Reliability plots and confidence histograms of LLaMA models on 4-shot reasoning tasks. Results of different sizes 7B (left), 13B (middle), and 30B (right) are plotted.

Algorithm 1: Pseudocode for temperature scaling

Data: $\mathcal{P}_\theta(\mathbf{w})$: Original output of the classification model, \mathcal{D} : Training dataset, τ : Temperature parameter, k : we use k -shot experimental settings, where during test the ICL prompts will consist of k (sample, label) pairs.

Result: Adjusted probabilities after temperature scaling;

```

// Training process
// 0-shot:  $(\mathbf{w}_i, y_i)$  is every training samples and corresponding label.
//  $k$ -shot:  $\mathbf{w}_i = \{x_1, y_1, \dots, x_k, y_k, x_i\}$  uses  $k$  prompt pairs.
// Fix  $w$ : the prompt in  $\mathbf{w}_i = \{x_1, y_1, \dots, x_k, y_k, x_i\}$  will be used for all
// training instances and used during inference.
for each training sample  $(\mathbf{w}_i, y_i) \in \mathcal{D}$  do
  | Compute the original output:  $z_i = \mathcal{P}_\theta(\mathbf{w}_i; \theta)$ ;
  | Compute the cross-entropy loss:  $L_i = \text{CrossEntropy}(z_i, y_i)$ ;
end

Compute the gradient of the loss to the temperature parameter:  $\nabla_\tau \mathcal{L} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \nabla_\tau L_i$ ;
Update the temperature parameter using gradient descent:  $\tau \leftarrow \tau - \eta \nabla_\tau \mathcal{L}$ ;

// Test time
for each test sample  $\mathbf{x}_j$  do
  | Compute the original output with prompt:  $z_j = \mathcal{P}_\theta(\mathbf{x}_j; \theta)$ ;
  | Compute the adjusted output:  $\hat{z}_j = \frac{z_j}{\tau}$ ;
  | Compute the softmax probabilities:  $\hat{p}_j = \text{Softmax}(\hat{z}_j)$ ;
end

```

Dataset	Prompt	Label
SST-2	<p>Review: it may not be a great piece of filmmaking, but its power comes from its soul's - eye view of how well-meaning patronizing masked a social injustice, at least as represented by this case . Sentiment: Positive</p> <p>Review: smith's point is simple and obvious – people's homes are extensions of themselves, and particularly eccentric people have particularly eccentric living spaces – but his subjects are charmers . Sentiment:</p>	Negative, Positive
CB	<p>A: No, not really. I spend a lot of time with our income tax, though, especially, this year and last year. Um, I have been married for just a few years, so I've had to really switch around from the EZ form to the, uh, B: Schedule A. A: Right. B: Well, yeah. A: All the deductions and all that. B: Did you notice that when they passed the new simplified tax act, it seemed like it made everything harder? question: when they passed the new simplified tax act it seemed like it made everything harder. true, false, or neither? answer: true</p> <p>There was a group of curious onlookers... Marie felt her legs give way beneath her, she sat down on the edge of the pavement, feet in the gutter, doubled-up, sick and winded as if someone had punched her in the stomach. She lifted up her head and looked again. She had watched scenes like this so often in detective films and police series on television that she could hardly believe that this was real life. question: this was real life. true, false, or neither? answer:</p>	True, False, Neither
RTE	<p>The main institutionalised forms of recognition for those who have made a significant contribution in the fields of physics, chemistry, medicine, literature, as well as for those working for peace (and more recently in the area of economics), are the Nobel prizes. question: Nobel Peace Prize candidates have been chosen. True or False? answer: False</p> <p>Egypt on Thursday strongly criticized Israeli new Foreign Minister Avigdor Lieberman for his remarks that he refused to recognize the peace efforts initiated in 2007 in the U.S. city of Annapolis to restore the peace talks with the Palestinians, reported the state MENA news agency. Lieberman's remarks is "regrettable," Egyptian Foreign Ministry spokesman Hossam Zaki was quoted as saying, adding "his remarks are the first blow to the peace efforts to come from the Israeli new government." question: Hossam Zaki is the new Foreign Minister of Israel. True or False? answer:</p>	True, False
Strategy QA	<p>Question: Can spiders help eggplant farmers control parasites? Choose the answer from True and False. Answer: The potato tuber moth is a parasite that targets the plant family Solanaceae, including eggplant <i>Se-lenops radiatus</i> is a spider genus in South Africa that effectively controls the potato tuber moth So, the answer is: True</p> <p>Question: Is the voice of Genie from Disney's Aladdin still alive? Choose the answer from True and False. Answer:</p>	True, False
Commonsense QA	<p>"Question: Dan was a farmer with just one heifer. But that was okay, he only kept her for milk, and he didn't think he'd find good farmland in a place as cold as where? A arizona B farm yard C michigan D german field E dairy farm Answer: Michigan is a state in the us where it precipitates throughout the year and areas, where it precipitates throughout the year, are generally cold. So the farmer thought he'd not find a good farmland in a place as cold as michigan. Enslaving heifers or other animals for their milk is wrong as they want to live free. All the places in the other options may not be cold. So, the answer is: C</p> <p>Question: From where does a snowflake form? A cloud B snow storm C billow D air E snowstorm Answer:"</p>	A, B, C, D, E

Table 6: **Prompts used for text classification and reasoning tasks**, with a single training example showcased per task for illustrative purposes. The right column displays corresponding labels. The prompting formats and labels for WorldTree and OpenBookQA are the same as those of the CommonsenseQA dataset.