

Calibration of Pre-trained Transformers

Shrey Desai and Greg Durrett

Department of Computer Science

The University of Texas at Austin

shreydesai@utexas.edu gdurrett@cs.utexas.edu

Abstract

Pre-trained Transformers are now ubiquitous in natural language processing, but despite their high end-task performance, little is known empirically about whether they are calibrated. Specifically, do these models’ posterior probabilities provide an accurate empirical measure of how likely the model is to be correct on a given example? We focus on BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) in this work, and analyze their calibration across three tasks: natural language inference, paraphrase detection, and commonsense reasoning. For each task, we consider in-domain as well as challenging out-of-domain settings, where models face more examples they *should* be uncertain about. We show that: (1) when used out-of-the-box, pre-trained models are calibrated in-domain, and compared to baselines, their calibration error out-of-domain can be as much as $3.5\times$ lower; (2) temperature scaling is effective at further reducing calibration error in-domain, and using label smoothing to deliberately increase empirical uncertainty helps calibrate posteriors out-of-domain.¹

1 Introduction

Neural networks have seen wide adoption but are frequently criticized for being black boxes, offering little insight as to why predictions are made (Benítez et al., 1997; Dayhoff and DeLeo, 2001; Castelvechi, 2016) and making it difficult to diagnose errors at test-time. These properties are particularly exhibited by pre-trained Transformer models (Devlin et al., 2019; Liu et al., 2019; Yang et al., 2019), which dominate benchmark tasks like SuperGLUE (Wang et al., 2019), but use a large number of self-attention heads across many layers in a way that is difficult to unpack (Clark et al.,

2019; Kovaleva et al., 2019). One step towards understanding whether these models can be trusted is by analyzing whether they are *calibrated* (Raftery et al., 2005; Jiang et al., 2012; Kendall and Gal, 2017): how aligned their posterior probabilities are with empirical likelihoods (Brier, 1950; Guo et al., 2017). If a model assigns 70% probability to an event, the event should occur 70% of the time if the model is calibrated. Although the model’s mechanism itself may be uninterpretable, a calibrated model at least gives us a signal that it “knows what it doesn’t know,” which can make these models easier to deploy in practice (Jiang et al., 2012).

In this work, we evaluate the calibration of two pre-trained models, BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), on three tasks: natural language inference (Bowman et al., 2015), paraphrase detection (Iyer et al., 2017), and commonsense reasoning (Zellers et al., 2018). These tasks represent standard evaluation settings for pre-trained models, and critically, challenging out-of-domain test datasets are available for each. Such test data allows us to measure calibration in more realistic settings where samples stem from a dissimilar input distribution, which is exactly the scenario where we hope a well-calibrated model would avoid making confident yet incorrect predictions.

Our experiments yield several key results. First, even when used out-of-the-box, pre-trained models are calibrated in-domain. In out-of-domain settings, where non-pre-trained models like ESIM (Chen et al., 2017) are overconfident, we find that pre-trained models are significantly better calibrated. Second, we show that temperature scaling (Guo et al., 2017), multiplying non-normalized logits by a single scalar hyperparameter, is widely effective at improving in-domain calibration. Finally, we show that regularizing the model to be less certain during training can beneficially smooth probabilities, improving out-of-domain calibration.

¹Code and datasets available at <https://github.com/shreydesai/calibration>

2 Related Work

Calibration has been well-studied in statistical machine learning, including applications in forecasting (Brier, 1950; Raftery et al., 2005; Gneiting et al., 2007; Palmer et al., 2008), medicine (Yang and Thompson, 2010; Jiang et al., 2012), and computer vision (Kendall and Gal, 2017; Guo et al., 2017; Lee et al., 2018). Past work in natural language processing has studied calibration in the non-neural (Nguyen and O’Connor, 2015) and neural (Kumar and Sarawagi, 2019) settings across several tasks. However, past work has not analyzed large-scale pre-trained models, and we additionally analyze out-of-domain settings, whereas past work largely focuses on in-domain calibration (Nguyen and O’Connor, 2015; Guo et al., 2017).

Another way of hardening models against out-of-domain data is to be able to explicitly detect these examples, which has been studied previously (Hendrycks and Gimpel, 2016; Liang et al., 2018; Lee et al., 2018). However, this assumes a discrete notion of domain; calibration is a more general paradigm and gracefully handles settings where domains are less quantized.

3 Posterior Calibration

A model is calibrated if the confidence estimates of its predictions are aligned with empirical likelihoods. For example, if we take 100 samples where a model’s prediction receives posterior probability 0.7, the model should get 70 of the samples correct. Formally, calibration is expressed as a joint distribution $P(Q, Y)$ over confidences $Q \in \mathbb{R}$ and labels $Y \in \mathcal{Y}$, where perfect calibration is achieved when $P(Y = y|Q = q) = q$. This probability can be empirically approximated by binning predictions into k disjoint, equally-sized bins, each consisting of b_k predictions. Following previous work in measuring calibration (Guo et al., 2017), we use *expected calibration error* (ECE), which is a weighted average of the difference between each bin’s accuracy and confidence: $\sum_k \frac{b_k}{n} |\text{acc}(k) - \text{conf}(k)|$. For the experiments in this paper, we use $k = 10$.

4 Experiments

4.1 Tasks and Datasets

We perform evaluations on three language understanding tasks: natural language inference, paraphrase detection, and commonsense reasoning. Significant past work has studied cross-domain robust-

Model	Parameters	Architecture	Pre-trained
DA	382K	LSTM	✗
ESIM	4M	Bi-LSTM	✗
BERT	110M	Transformer	✓
RoBERTa	110M	Transformer	✓

Table 1: Models in this work. Decomposable Attention (DA) (Parikh et al., 2016) and Enhanced Sequential Inference Model (ESIM) (Chen et al., 2017) use LSTMs and attention on top of GloVe embeddings (Pennington et al., 2014) to model pairwise semantic similarities. In contrast, BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) are large-scale, pre-trained language models with stacked, general purpose Transformer (Vaswani et al., 2017) layers.

ness using sentiment analysis (Chen et al., 2018; Peng et al., 2018; Miller, 2019; Desai et al., 2019). However, we explicitly elect to use tasks where out-of-domain performance is substantially lower and challenging domain shifts are exhibited. Below, we describe our in-domain and out-of-domain datasets.² For all datasets, we split the development set in half to obtain a held-out, non-blind test set.

Natural Language Inference. The Stanford Natural Language Inference (SNLI) corpus is a large-scale entailment dataset where the task is to determine whether a hypothesis is entailed, contradicted by, or neutral with respect to a premise (Bowman et al., 2015). Multi-Genre Natural Language Inference (MNLI) (Williams et al., 2018) contains similar entailment data across several domains, which we can use as unseen test domains.

Paraphrase Detection. Quora Question Pairs (QQP) contains sentence pairs from Quora that are semantically equivalent (Iyer et al., 2017). Our out-of-domain setting is TwitterPPDB (TPPDB), which contains sentence pairs from Twitter where tweets are considered paraphrases if they have shared URLs (Lan et al., 2017).

Commonsense Reasoning. Situations With Adversarial Generations (SWAG) is a grounded commonsense reasoning task where models must select the most plausible continuation of a sentence among four candidates (Zellers et al., 2018). HellaSWAG (HSWAG), an adversarial out-of-domain dataset, serves as a more challenging benchmark for pre-trained models (Zellers et al., 2019); it is

²Dataset splits are detailed in Appendix A. Furthermore, out-of-domain datasets are strictly used for evaluating the generalization of in-domain models, so the training split is unused.

Model	Accuracy		ECE	
	ID	OD	ID	OD
Task: SNLI/MNLI				
DA	84.63	57.12	1.02	8.79
ESIM	88.32	60.91	1.33	12.78
BERT	90.04	73.52	2.54	7.03
RoBERTa	91.23	78.79	1.93	3.62
Task: QQP/TwitterPPDB				
DA	85.85	83.36	3.37	9.79
ESIM	87.75	84.00	3.65	8.38
BERT	90.27	87.63	2.71	8.51
RoBERTa	91.11	86.72	2.33	9.55
Task: SWAG/HellaSWAG				
DA	46.80	32.48	5.98	40.37
ESIM	52.09	32.08	7.01	19.57
BERT	79.40	34.48	2.49	12.62
RoBERTa	82.45	41.68	1.76	11.93

Table 2: Out-of-the-box calibration results for in-domain (SNLI, QQP, SWAG) and out-of-domain (MNLI, TwitterPPDB, HellaSWAG) datasets using the models described in Table 1. We report accuracy and expected calibration error (ECE), both averaged across 5 fine-tuning runs with random restarts.

distributionally different in that its examples exploit statistical biases in pre-trained models.

4.2 Systems for Comparison

Table 1 shows a breakdown of the models used in our experiments. We use the same set of hyperparameters across all tasks. For pre-trained models, we omit hyperparameters that induce brittleness during fine-tuning, e.g., employing a decaying learning rate schedule with linear warmup (Sun et al., 2019; Lan et al., 2020). Detailed information on optimization is available in Appendix B.

4.3 Out-of-the-box Calibration

First, we analyze “out-of-the-box” calibration; that is, the calibration error derived from evaluating a model on a dataset without using post-processing steps like temperature scaling (Guo et al., 2017). For each task, we train the model on the in-domain training set, and then evaluate its performance on the in-domain and out-of-domain test sets. Quantitative results are shown in Table 2. In addition, we plot reliability diagrams (Nguyen and O’Connor, 2015; Guo et al., 2017) in Figure 1, which visualize the alignment between posterior probabilities (confidence) and empirical outcomes (accuracy), where a perfectly calibrated model has $\text{conf}(k) = \text{acc}(k)$ for each bucket of real-valued predictions k . We




Figure 1: In-domain calibration of BERT and RoBERTa when used out-of-the-box. Models are both trained and evaluated on SNLI, QQP, and SWAG, respectively. ZERO ERROR depicts perfect calibration (e.g., expected calibration error = 0). Note that low-confidence buckets have zero accuracy due to a small sample count; however, as a result, these buckets do not influence the expected error as much.

remark on a few observed phenomena below:

Non-pre-trained models exhibit an inverse relationship between complexity and calibration.

Simpler models, such as DA, achieve competitive in-domain ECE on SNLI (1.02) and QQP (3.37), and are notably better than pre-trained models on SNLI in this regard. However, the more complex ESIM, both in number of parameters and architecture, sees increased in-domain ECE despite having higher accuracy on all tasks.

However, pre-trained models are generally more accurate and calibrated.

Rather surprisingly, pre-trained models do not show characteristics of the aforementioned inverse relationship, despite having significantly more parameters. On SNLI, RoBERTa achieves an ECE in the ballpark of DA and ESIM, but on QQP and SWAG, both

BERT and RoBERTa consistently achieve higher accuracies and lower ECEs. Pre-trained models are especially strong out-of-domain, where on HellaSWAG in particular, RoBERTa reduces ECE by a factor of 3.4 compared to DA.

Using RoBERTa always improves in-domain calibration over BERT. In addition to obtaining better task performance than BERT, RoBERTa consistently achieves lower in-domain ECE. Even out-of-domain, RoBERTa outperforms BERT in all but one setting (TwitterPPDB). Nonetheless, our results show that representations induced by robust pre-training (e.g., using a larger corpus, more training steps, dynamic masking) (Liu et al., 2019) lead to more calibrated posteriors. Whether other changes to pre-training (Yang et al., 2019; Lan et al., 2020; Clark et al., 2020) lead to further improvements is an open question.

4.4 Post-hoc Calibration

There are a number of techniques that can be applied to correct a model’s calibration post-hoc. Using our in-domain development set, we can, for example, post-process model probabilities via temperature scaling (Guo et al., 2017), where a scalar temperature hyperparameter T divides non-normalized logits before the softmax operation. As $T \rightarrow 0$, the distribution’s mode receives all the probability mass, while as $T \rightarrow \infty$, the probabilities become uniform.

Furthermore, we experiment with models trained in-domain with label smoothing (LS) (Miller et al., 1996; Pereyra et al., 2017) as opposed to conventional maximum likelihood estimation (MLE). By nature, MLE encourages models to sharpen the posterior distribution around the gold label, leading to confidence which is typically unwarranted in out-of-domain settings. Label smoothing presents one solution to overconfidence by maintaining uncertainty over the label space during training: we minimize the KL divergence with the distribution placing a $1 - \alpha$ fraction of probability mass on the gold label and $\frac{\alpha}{|\mathcal{Y}|-1}$ fraction of mass on each other label, where $\alpha \in (0, 1)$ is a hyperparameter.³ This re-formulated learning objective does not require changing the model architecture.

For each task, we train the model with either MLE or LS ($\alpha = 0.1$) using the in-domain training set, use the in-domain development set to learn

³For example, the one-hot target $[1, 0, 0]$ is transformed into $[0.9, 0.05, 0.05]$ when $\alpha = 0.1$.




Figure 2: In-domain calibration of BERT and RoBERTa with temperature scaling (TS). Both temperature-scaled models are much better calibrated than when used out-of-the-box, with BERT especially showing a large degree of improvement.

an optimal temperature T , and then evaluate the model (scaled with T) on the in-domain and out-of-domain test sets. From Table 3 and Figure 2, we draw the following conclusions:

MLE models with temperature scaling achieve low in-domain calibration error. MLE models are always better than LS models in-domain, which suggests incorporating uncertainty when in-domain samples are available is not an effective regularization scheme. Even when using a small smoothing value (0.1), LS models do not achieve nearly as good out-of-the-box results as MLE models, and temperature scaling hurts LS in many cases. By contrast, RoBERTa with temperature-scaled MLE achieves ECE values from 0.7-0.8, implying that MLE training yields scores that are fundamentally good but just need some minor rescaling.

However, out-of-domain, label smoothing is generally more effective. In most cases, MLE models do not perform well on out-of-domain

Method	In-Domain						Out-of-Domain					
	SNLI		QQP		SWAG		MNLI		TPPDB		HSWAG	
	MLE	LS	MLE	LS	MLE	LS	MLE	LS	MLE	LS	MLE	LS
Model: BERT												
Out-of-the-box	2.54	7.12	2.71	6.33	2.49	10.01	7.03	3.74	8.51	6.30	12.62	5.73
Temperature scaled	1.14	8.37	0.97	8.16	0.85	10.89	3.61	4.05	7.15	5.78	12.83	5.34
Model: RoBERTa												
Out-of-the-box	1.93	6.38	2.33	6.11	1.76	8.81	3.62	4.50	9.55	8.91	11.93	2.14
Temperature scaled	0.84	8.70	0.88	8.69	0.76	11.4	1.46	5.93	7.86	5.31	11.22	2.23

Table 3: Post-hoc calibration results for BERT and RoBERTa on in-domain (SNLI, QQP, SWAG) and out-of-domain (MNLI, TwitterPPDB, HellaSWAG) datasets. Models are trained with maximum likelihood estimation (MLE) or label smoothing (LS), then their logits are post-processed using temperature scaling (§4.4). We report expected calibration error (ECE) averaged across 5 runs with random restarts. Darker colors imply lower ECE.




Figure 3: Out-of-domain calibration of RoBERTa fine-tuned on SWAG with different learning objectives and used out-of-the-box on HellaSWAG. Without seeing HellaSWAG samples during fine-tuning, RoBERTa-LS achieves significantly lower calibration error than RoBERTa-MLE.

datasets, with ECEs ranging from 8-12. However, LS models are forced to distribute probability mass across classes, and as a result, achieve significantly lower ECEs on average. We note that LS is particularly effective when the distribution shift is strong. On the adversarial HellaSWAG, for example, RoBERTa-LS obtains a factor of 5.8 less ECE than RoBERTa-MLE. This phenomenon is visually depicted in Figure 3 where we see RoBERTa-LS is significantly closer to the identity function despite being used out-of-the-box.

Optimal temperature scaling values are bounded within a small interval. Table 4 reports the learned temperature values for BERT-MLE and RoBERTa-MLE. For in-domain tasks, the optimal temperature values are generally in

Model	In-Domain			Out-of-Domain		
	SNLI	QQP	SWAG	MNLI	TPPDB	HSWAG
BERT	1.20	1.34	0.99	1.41	2.91	3.61
RoBERTa	1.16	1.39	1.10	1.25	2.79	2.77

Table 4: Learned temperature scaling values for BERT and RoBERTa on in-domain (SNLI, QQP, SWAG) and out-of-domain (MNLI, TwitterPPDB, HellaSWAG) datasets. Values are obtained by line search with a granularity of 0.01. Evaluations are very fast as they only require rescaling cached logits.

the range 1-1.4. Interestingly, out-of-domain, TwitterPPDB and HellaSWAG require larger temperature values than MNLI, which suggests the degree of distribution shift and magnitude of T may be closely related.

5 Conclusion

Posterior calibration is one lens to understand the trustworthiness of model confidence scores. In this work, we examine the calibration of pre-trained Transformers in both in-domain and out-of-domain settings. Results show BERT and RoBERTa coupled with temperature scaling achieve low ECEs in-domain, and when trained with label smoothing, are also competitive out-of-domain.

Acknowledgments

This work was partially supported by NSF Grant IIS-1814522 and a gift from Arm. The authors acknowledge a DURIP equipment grant to UT Austin that provided computational resources to conduct this research. Additionally, we thank R. Thomas McCoy for answering questions about DA and ESIM.

References

- Jose M. Benítez, Juan Luis Castro, and Ignacio Requena. 1997. Are Artificial Neural Networks Black Boxes? *IEEE Transactions on Neural Networks and Learning Systems*.
- 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 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Glenn W. Brier. 1950. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*.
- Davide Castelvecchi. 2016. Can We Open the Black Box of AI? *Nature News*.
- Qian Chen, Xiaodan Zhu, Zhen-Hua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2017. Enhanced LSTM for Natural Language Inference. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Xilun Chen, Yu Sun, Ben Athiwaratkun, Claire Cardie, and Kilian Weinberger. 2018. Adversarial Deep Averaging Networks for Cross-Lingual Sentiment Classification. *Transactions of the Association for Computational Linguistics (TACL)*.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What Does BERT Look at? An Analysis of BERT’s Attention. In *Proceedings of the Workshop on BlackboxNLP*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Judith E. Dayhoff and James M. DeLeo. 2001. Artificial Neural Networks: Opening the Black Box. *Cancer: Interdisciplinary International Journal of the American Cancer Society*.
- Shrey Desai, Hongyuan Zhan, and Ahmed Aly. 2019. Evaluating Lottery Tickets Under Distributional Shifts. In *Proceedings of the Workshop on Deep Learning Approaches for Low-Resource NLP (DeepLo)*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. 2018. AllenNLP: A Deep Semantic Natural Language Processing Platform. In *Proceedings of the Workshop for NLP Open Source Software (NLP-OSS)*.
- Tilmann Gneiting, Fadoua Balabdaoui, and Adrian E. Raftery. 2007. Probabilistic Forecasts, Calibration and Sharpness. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On Calibration of Modern Neural Networks. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- Dan Hendrycks and Kevin Gimpel. 2016. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Shankar Iyer, Nikhil Dandekar, and Kornél Csernai. 2017. Quora Question Pairs.
- Xiaoqian Jiang, Melanie Osl, Jihoon Kim, and Lucila Ohno-Machado. 2012. Calibrating Predictive Model Estimates to Support Personalized Medicine. In *Journal of the American Medical Informatics Association (JAMIA)*.
- Alex Kendall and Yarin Gal. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the Dark Secrets of BERT. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Aviral Kumar and Sunita Sarawagi. 2019. Calibration of Encoder Decoder Models for Neural Machine Translation. In *Proceedings of the Workshop on Debugging Machine Learning Models*.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. A Continuously Growing Dataset of Sentential Paraphrases. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kimin Lee, Honglak Lee, Kibok Lee, and Jinwoo Shin. 2018. Training Confidence-calibrated Classifiers for Detecting Out-of-Distribution Samples. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

- Shiyu Liang, Yixuan Li, and R. Srikant. 2018. Enhancing the Reliability of Out-of-distribution Image Detection in Neural Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- 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*.
- Ilya Loshchilov and Frank Hutter. 2019. Decoupled Weight Decay Regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- David J. Miller, Ajit V. Rao, Kenneth Rose, and Allen Gersho. 1996. A Global Optimization Technique for Statistical Classifier Design. *IEEE Transactions on Signal Processing*.
- Timothy Miller. 2019. Simplified Neural Unsupervised Domain Adaptation. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Khanh Nguyen and Brendan O’Connor. 2015. Posterior Calibration and Exploratory Analysis for Natural Language Processing Models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Tim Palmer, Francisco Doblas-Reyes, Antje Weisheimer, and Mark Rodwell. 2008. Toward Seamless Prediction: Calibration of Climate Change Projections using Seasonal Forecasts. *Bulletin of the American Meteorological Society*.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A Decomposable Attention Model for Natural Language Inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Minlong Peng, Qi Zhang, Yu gang Jiang, and Xuanjing Huang. 2018. Cross-Domain Sentiment Classification with Target Domain Specific Information. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. 2017. Regularizing Neural Networks by Penalizing Confident Output Distributions. In *Proceedings of the International Conference on Learning Representations (Workshop)*.
- Adrian E. Raftery, Tilmann Gneiting, Fadoua Balabdaoui, and Michael Polakowski. 2005. Using Bayesian Model Averaging to Calibrate Forecast Ensembles. *Monthly Weather Review*.
- Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. 2019. How to Fine-Tune BERT for Text Classification? *arXiv preprint arXiv:1905.05583*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *arXiv preprint arXiv:1910.03771*.
- Huiqin Yang and Carl Thompson. 2010. Nurses’ Risk Assessment Judgements: A Confidence Calibration Study. *Journal of Advanced Nursing*.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. SWAG: A Large-Scale Adversarial Dataset for Grounded Commonsense Inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSWAG: Can a Machine Really Finish Your Sentence? In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*.

A Dataset Splits

Dataset splits are shown in Table 5.

Dataset	Train	Dev	Test
SNLI	549,368	4,922	4,923
MNLI	392,702	4,908	4,907
QQP	363,871	20,216	20,217
TwitterPPDB	46,667	5,060	5,060
SWAG	73,547	10,004	10,004
HellaSWAG	39,905	5,021	5,021

Table 5: Training, development, and test dataset sizes for SNLI (Bowman et al., 2015), MNLI (Williams et al., 2018), QQP (Iyer et al., 2017), TwitterPPDB (Lan et al., 2017), SWAG (Zellers et al., 2018), and HellaSWAG (Zellers et al., 2019).

B Training and Optimization

For non-pre-trained model baselines, we use the open-source implementations of DA (Parikh et al., 2016) and ESIM (Chen et al., 2017) in AllenNLP (Gardner et al., 2018), except in the case of SWAG/HellaSWAG, where we run the baselines available in the authors’ code.⁴ For BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), we use `bert-base-uncased` and `roberta-base`, respectively, from HuggingFace Transformers (Wolf et al., 2019). BERT is fine-tuned with a maximum of 3 epochs, batch size of 16, learning rate of 2e-5, gradient clip of 1.0, and no weight decay. Similarly, RoBERTa is fine-tuned with a maximum of 3 epochs, batch size of 32, learning rate of 1e-5, gradient clip of 1.0, and weight decay of 0.1. Both models are optimized with AdamW (Loshchilov and Hutter, 2019). Other than early stopping on the development set, we do not perform additional hyperparameter searches. Finally, all experiments are conducted on NVIDIA V100 32GB GPUs, with the total time for fine-tuning all models being under 24 hours.

Furthermore, temperature scaling line searches are performed in the range [0.01, 5.0] with a granularity of 0.01. These searches are quite fast and can be performed on a CPU; we simply evaluate calibration error by rescaling cached logits. On a Intel Xeon E3-1270 v3 CPU, all searches can be completed in under 15 minutes.

C Reproducibility

Table 6 shows the accuracy and expected calibration error (ECE) of BERT and RoBERTa on the

⁴<https://github.com/rowanz/swagaf>

Model	Accuracy		ECE	
	ID	OD	ID	OD
Task: SNLI/MNLI				
BERT	90.18	74.04	3.43	8.18
RoBERTa	91.20	79.17	1.18	1.41
Task: QQP/TwitterPPDB				
BERT	90.22	86.02	4.68	11.30
RoBERTa	89.97	86.17	3.09	9.57
Task: SWAG/HellaSWAG				
BERT	78.82	38.01	2.51	2.24
RoBERTa	81.85	59.03	3.02	5.71

Table 6: Out-of-the-box calibration *development set* results for in-domain (SNLI, QQP, SWAG) and out-of-domain (MNLI, TwitterPPDB, HellaSWAG) datasets using pre-trained models.

development sets of the datasets we consider. We do not report post-hoc calibration results using the development set since these require tuning on the development set itself.