

The Reliability Paradox: Exploring How Shortcut Learning Undermines Language Model Calibration

Geetanjali Bihani
Purdue University
gbihani@purdue.edu

Julia Rayz
Purdue University
jtaylor1@purdue.edu

Abstract

The advent of pre-trained language models (PLMs) has enabled significant performance gains in the field of natural language processing. However, recent studies have found PLMs to suffer from miscalibration, indicating a lack of accuracy in the confidence estimates provided by these models. Current evaluation methods for PLM calibration often assume that lower calibration error estimates indicate more reliable predictions. However, fine-tuned PLMs often resort to shortcuts, leading to overconfident predictions that create the illusion of enhanced performance but lack generalizability in their decision rules. The relationship between PLM reliability, as measured by calibration error, and shortcut learning, has not been thoroughly explored thus far. This paper aims to investigate this relationship, studying whether lower calibration error implies reliable decision rules for a language model. Our findings reveal that models with seemingly superior calibration portray higher levels of non-generalizable decision rules. This challenges the prevailing notion that well-calibrated models are inherently reliable. Our study highlights the need to bridge the current gap between language model calibration and generalization objectives, urging the development of comprehensive frameworks to achieve truly robust and reliable language models.

Keywords: Pretrained language models, calibration, shortcut learning, robustness, generalization

1. Introduction

Pre-trained language models (PLMs) have become the convention in the field of natural language processing. The preference for PLMs can be attributed to their

improvements in a wide variety of tasks, including question answering, textual entailment, sentiment analysis, and commonsense reasoning (Devlin et al., 2019; Peters et al., 2018; Sap et al., 2020). The ‘pretrain-then-fine-tune’ paradigm allows the model to not only utilize the existing ‘knowledge’ gained during pre-training but also learn from task-specific data via fine-tuning (Alt et al., 2019; Q. Chen et al., 2019).

Although fine-tuning PLMs achieves state-of-the-art results, it also causes models to lack generalization and become unreliable predictors. Specifically, PLMs tend to learn shortcuts based on keywords (Du et al., 2021; Moon et al., 2021) and cues related to language variations (Nguyen et al., 2021) to make predictions. This behavior, also known as shortcut learning, leads the model to learn non-generalizable decision rules that do not perform well on out-of-distribution (OOD) data (Du et al., 2022; Moon et al., 2021). Additionally, the fine-tuning process can lead to overconfidence in PLMs (Jiang et al., 2021; Kong et al., 2020), where their confidence increases regardless of the accuracy of their predictions (Y. Chen et al., 2022). This mismatch between the model’s confidence and its actual accuracy in its predictions results in ‘miscalibration’ in language models.

It is desirable for language models to perform reliably and accurately across different language tasks. Addressing calibration is essential because it ensures that the model’s confidence aligns more accurately with its predictive accuracy. Miscalibrated models can lead to significant issues, particularly in high-stakes environments where wrong but confident predictions are dangerous. By focusing on calibration, we can reduce the mismatch between confidence and correctness, improving the model’s trustworthiness and robustness across diverse tasks and data distributions. Thus, it

is important to study the interplay between model generalization and calibration. Prior works assessing model calibration focus on measuring and minimizing statistical calibration evaluation metrics such as Expected Calibration Error (ECE) (Ahuja et al., 2022; Kim et al., 2023; Kong et al., 2020). These works do not investigate whether lower calibration error estimates align with more generalizable decision rules learned by these language models.

In this study, we aim to address this gap and conduct the following research inquiries: 1) Does a reduction in calibration error within language models indicate a decrease in overconfident predictions? 2) Can a model exhibiting lower calibration error be considered reliable in terms of its decision rules? By examining these questions, we seek to shed light on the relationship between calibration error and the reliability of language models’ decision rules. Our questions are based on the intuition that model reliability estimates should account for the reliability of the model’s decision rules.

To answer these questions, we investigate the calibration and shortcut learning behaviors of recent pre-trained language models (PLMs) across a suite of binary and multi-class classification tasks, and analyze the evolution of shortcut learning behaviors in PLMs before and after fine-tuning.

Our research findings highlight that models appearing to be well-calibrated often exhibit a higher propensity for shortcut learning. This challenges the conventional notion of well-calibrated models as reliable and robust. While lower calibration error estimates, such as Expected Calibration Error (ECE), may indicate the improved alignment of prediction probabilities with their actual correctness, they fail to capture the inherent lack of robustness in these ‘correct’ model decisions. This observation uncovers a fundamental discrepancy between the requirements of model calibration and the goals of generalization, highlighting the need to reconcile these seemingly contradictory frameworks in order to achieve truly robust and reliable language models.

Contributions: Our contributions are summarized as follows.

- Our study shows that contrary to previous assumptions, a lower Expected Calibration Error (ECE) does not necessarily indicate improved reliability. Instead, it often reflects models’ tendency to make overconfident predictions driven by shortcut cues.
- We perform analyses on fine-tuned PLMs, across a suite of text classification tasks, highlighting the limitations of statistical calibration error measures such as ECE, in capturing the lack of robustness in

model decisions. This insight underscores the need to consider the trade-off between task performance and robustness to shortcuts when evaluating model calibration.

2. Shortcut Effects on Calibration

2.1. Identifying Shortcuts

Shortcut learning refers to the phenomenon where models rely on superficial cues in the training data to make predictions instead of learning the underlying semantics to perform an NLU task. This over-reliance on specific features or biases results in poor generalization in out-of-distribution (OOD) settings. Identifying shortcut learning in language models is an ongoing research area, with recent works utilizing model attention, dataset statistics, and human annotated samples to identify spurious correlations (Moon et al., 2021; T. Wang et al., 2022). We utilize the shortcut identification framework as described by (Du et al., 2021), which combines data statistics with model attributions to identify shortcuts. We describe this shortcut identification framework below.

Model Attribution based Importance: To obtain attributions for each token (w_j) in a given sample S_i , we utilize integrated gradients (IG) (Sundararajan et al., 2017). Let a given sample S_i contain T tokens, i.e.

$S_i = \{w_j^t\}_{t=1}^T$. We conduct a step-wise perturbation of the sample, creating m intermediate samples along a straight-line path from a baseline S_b to the actual sample S_i . By observing the changes in the model’s output as the sample is progressively modified, we quantify the contribution of each token to the final prediction. Following (Du et al., 2021), we consider all-zero embeddings to form S_b . As each word is added to the baseline, the gradient of the prediction $M(S_i)$ is computed with respect to the associated token embeddings ($e(w_j)$) obtained from the output embedding layer of model M . The following equations summarize this gradient calculation.

$$IG(S_i) = S_i^b \cdot \sum_{k=1}^m \frac{\partial M_y(S_b + \frac{k}{m}(S_i - S_b))}{\partial S_i} \cdot \frac{1}{m} \quad (1)$$

where

$$S_i^b = S_i - S_b \quad (2)$$

Finally, the L2 norm between the gradient and the corresponding token embedding is calculated to determine the individual contribution of each token. Following (Du et al., 2021), we filter top three attributed tokens per sample.

Text	Label	Shortcuts
This is awesome!	<i>positive</i>	‘!’, ‘awesome’
This is tragic...	<i>negative</i>	‘...’, ‘tragic’

Table 1: Example shortcuts in binary sentiment classification

Local Mutual Information based Importance: We calculate local mutual information (LMI) between tokens and task labels for a given dataset D , using Eq.3.

$$\text{LMI}(w, y) = p(w, y) \cdot \log \left(\frac{p(y | w)}{p(y)} \right) \quad (3)$$

where $p(w, y) = \frac{\text{count}(w, y)}{|V|}$, $p(y | w) = \frac{\text{count}(w, y)}{\text{count}(w)}$ and $p(y) = \frac{\text{count}(y)}{|V|}$. Here w refers to a word token appearing in the samples with the task label y , and $|V|$ refers to the size of the vocabulary of dataset D .

Comparing Importance to Estimate Shortcuts: As mentioned in (Du et al., 2021), for a given label y , we consider the top 5% LMI-scored tokens as the *head* of the label’s LMI distribution. For a given sample, if the top-3 attributed tokens for a prediction also appear in the *head* of the predicted label’s LMI distribution, it is termed as a shortcut.

2.2. Types of Shortcuts

Within the broader context of language structure, a well-established theoretical framework proposed by (Chomsky, 1965) introduced a fundamental division between the lexicon and grammar. According to this framework, the lexicon serves as a repository for language words, while grammar establishes rules for combining these words. Drawing upon these foundational concepts, we divide the identified shortcuts into two categories, i.e. ‘*lexicon-cued*’ and ‘*grammar-cued*’ predictions. PLM predictions where at least one lexical word is utilized are classified as ‘*lexicon-cued*’ predictions. On the other hand, predictions where the identified shortcuts are limited to functional words and punctuations, are labeled as ‘*grammar-cued*’ predictions. The purpose of this categorization is to improve our understanding of the shortcut mechanisms employed by the models in their prediction processes, particularly in terms of their lexical-semantic processing. Examples are shown in Table 2

For a more fine-grained analysis, we differentiate between cases where a model exclusively relies on punctuation, stopwords, or sub-words for making predictions and cases where it additionally incorporates one or more lexical words. This distinction allows us

to compare the model’s reliance on grammatical and lexical cues. For example, considering the case given in Table 1, if a sentiment classification model relies on an exclamation mark ‘!’ to predict the sentiment of ‘*this is awesome!*’, the rule is less generalizable as compared to relying on the word ‘*awesome*’. This is because punctuations and stopwords lack explicit semantic information on sentiment and can be used in various contexts to indicate surprise, excitement, or emphasis. On the other hand, the word ‘*awesome*’ has a more consistent and explicit semantic association with positive sentiment and is frequently used in positive contexts, and has a positive polarity (sarcasm excluded).

2.3. Measuring Calibration

A well-calibrated model should provide accurate probability estimates that reflect the true likelihood of an event. To quantify model calibration, we measure the Expected Calibration Error (ECE) (Naeini et al., 2015). We choose ECE because it captures the discrepancy between the model’s confidence and accuracy, and has been used for calibration analysis, making it an ideal choice for evaluating the model’s reliability across a range of tasks. This metric calculates the weighted average of the difference between the accuracy of a model and its average confidence level over a set of bins defined by the predicted probabilities, as shown in Eq. 4, where n is the number of samples in B_m .

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} | \text{acc}(B_m) - \text{conf}(B_m) | \quad (4)$$

Here, the estimation of expected accuracy from finite samples is done by grouping predictions (\hat{p}_i) into M interval bins (each of size $\frac{1}{M}$), and the accuracy of each bin is calculated. Let B_m be a bin containing samples whose prediction confidence lies within the interval $I_m = (\frac{m-1}{M}, \frac{m}{M}]$. Then the accuracy of B_m , where y_i and \hat{y}_i portray predicted and true class labels, is calculated as shown in Eq. 5.

$$\text{acc}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \mathbf{1}(\hat{y}_i = y_i) \quad (5)$$

The average predicted confidence of B_m , is calculated as shown in Eq. 6.

$$\text{conf}(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i \quad (6)$$

Type	Text	Data
<i>Lexicon-cued</i>	Michael Phelps won the gold medal in the 400 individual medley and set a world record in a time of 4 minutes 8.26 seconds.	AG News
<i>Grammar-cued</i>	What are spider veins ? Guess they didn't get the memo reg non-nuclear Baltic sea #sarcasm	TREC TweetEval (Irony)

Table 2: Examples of shortcuts learned by PLMs across different tasks

2.4. Measuring trade-offs

In order to investigate the relationship between model calibration and shortcut learning, we calculate two metrics: the portion of shortcut-cued model predictions (P_{sc}) and the shortcut trade-off (T_{sc}). The calculation of P_{sc} allows us to quantify the extent to which a model relies on shortcuts when making predictions. Additionally, we introduce T_{sc} as a metric to assess the trade-off between shortcut learning and model performance. T_{sc} is calculated as the ratio of task accuracy (e.g., F1 score) to the proportion of shortcut-cued predictions.

$$T_{sc} = \frac{\text{Task Accuracy (F1)}}{\text{Shortcut-Cued Predictions (}P_{sc}\text{)}} \quad (7)$$

A higher T_{sc} score indicates that the model achieves better task accuracy while relying on fewer shortcut-cued predictions. Conversely, a lower T_{sc} score suggests a higher reliance on shortcuts to achieve optimal task performance. In our analysis, we aim to maximize T_{sc} and minimize Expected Calibration Error (ECE) in order to identify models that strike a balance between shortcut learning and accurate predictions.

3. Experiments

Datasets. To evaluate PLM shortcut learning and calibration effects across different tasks and domains, we perform our evaluation on several binary and multi-class classification tasks. Specifically, we consider 8 text classification datasets, briefly described as follows: i) Stanford-Sentiment Treebank (SST-2) (Socher et al., 2013), commonly used in sentiment analysis tasks and provides a valuable benchmark for evaluating models' ability to capture sentiment in texts, ii) Corpus of Linguistic Acceptability (COLA) (Warstadt et al., 2019), which assesses model performance on grammaticality judgments, iii) TREC (coarse-grained) (Hovy et al., 2001) used for question classification, iv) AG News (Zhang et al., 2015) used for news topic classification, and four datasets from TweetEval benchmark (Barbieri et al., 2020) [Emotion, Hate, Irony, and Sentiment] for

text classification in the context of short and informal social media texts. These datasets present different challenges, such as the presence of sarcasm and negation in samples for sentiment tasks, lexical overlap in topic classification tasks, and the inclusion of short and ironic social media texts in irony detection tasks. This language and task variation across datasets allows for a holistic assessment of the models' performance and generalization of our findings in the context of PLM calibration literature.

Models. We evaluate five pre-trained transformer language models for evaluation: BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), DeBERTa (He et al., n.d.), ALBERT (Lan et al., n.d.) and BART (Lewis et al., 2020). We choose BERT and RoBERTa to align our results with prior PLM calibration research (Desai & Durrett, 2020a; Kim et al., 2023). We additionally evaluate more recent transformer LMs including DeBERTa and ALBERT. DeBERTa improves model generalization on downstream tasks, compared to BERT and RoBERTa, attributed to its disentangled attention mechanism. ALBERT is a compact architecture, providing performance gains with minimal sacrifice of task performance. Finally, we include BART, due to its improvements in handling of global context and robustness in handling noisy and ambiguous texts.

Metrics. To investigate the association between shortcut learning effects and model calibration, we employ multiple evaluation metrics. We utilize F1 score to assess the overall prediction performance of PLMs on given tasks. Additionally, we measure the Expected Calibration Error (ECE) (Naeini et al., 2015), as described in Section 2.3.

We also examine the distribution of shortcuts across correct and incorrect predictions using P_{sc} and T_{sc} . These measures allow us to gain insights into the relationship between shortcut utilization and prediction accuracy. By analyzing how shortcuts are distributed across different prediction outcomes, we can explore their impact on the model's ability to classify accurately.

Training Configurations. Following prior work (Kim et al., 2023), we fix several hyperparameters for the model fine-tuning process. For all models, we set the initial learning rate to 1e-5 and gradient clip to 1.0. We

$P_{sc} / T_{sc} / ECE$	COLA	Hate	Irony	SST2
ALBERT	51.29 / 1.66 / 0.16	91.72 / 0.67 / 0.42	79.97 / 0.72 / 0.29	84.98 / 1.06 / 0.44
BART	53.98 / 1.58 / 0.22	94.28 / 0.65 / 0.47	72.19 / 0.83 / 0.31	83.83 / 1.12 / 0.45
BERT	66.73 / 1.28 / 0.11	95.39 / 0.65 / 0.36	71.68 / 0.55 / 0.15	88.88 / 1.03 / 0.44
DeBERTa	54.07 / 1.63 / 0.22	95.35 / 0.66 / 0.48	86.10 / 0.75 / 0.35	84.86 / 1.12 / 0.47
RoBERTa	59.16 / 1.45 / 0.15	91.55 / 0.68 / 0.37	83.80 / 0.40 / 0.14	84.52 / 1.11 / 0.45
$P_{sc} / T_{sc} / ECE$	AG News	Emotion	Sentiment	TREC
ALBERT	92.11 / 1.02 / 0.01	79.38 / 0.64 / 0.08	84.79 / 0.79 / 0.04	93.60 / 0.95 / 0.05
BART	95.93 / 0.99 / 0.01	68.47 / 1.10 / 0.04	76.89 / 0.93 / 0.06	91.20 / 1.03 / 0.02
BERT	97.20 / 0.97 / 0.01	77.62 / 0.39 / 0.09	80.49 / 0.87 / 0.02	97.40 / 0.75 / 0.39
DeBERTa	96.50 / 0.98 / 0.01	75.09 / 1.03 / 0.03	75.29 / 0.94 / 0.09	95.20 / 0.99 / 0.03
RoBERTa	95.47 / 0.99 / 0.01	71.36 / 0.62 / 0.11	72.92 / 0.98 / 0.03	84.52 / 0.82 / 0.11

Table 3: Table comparing PLMs across shortcuts learnt (P_{sc}), shortcut trade-off (T_{sc}) and calibration (ECE). Top row: COLA, Hate, Irony and SST2 are binary classification tasks; Bottom Row: AG News, Emotion, Sentiment and TREC are multi-class classification tasks.

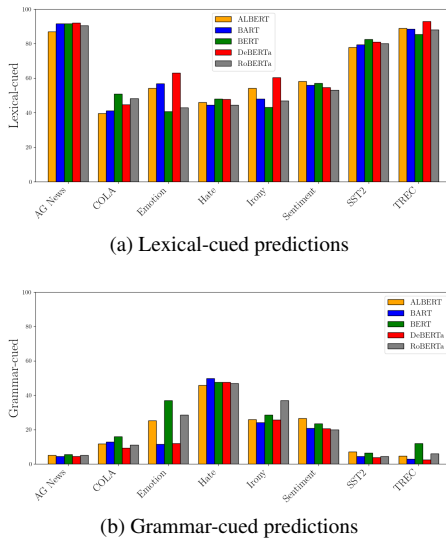


Figure 1: Difference in distribution of shortcut-cued predictions across different tasks.

utilized an Adam optimizer with an ϵ value of $1e-8$, and set the batch size to 32. We fine-tuned our models over a maximum of 3 epochs. To gauge the impact of fine-tuning on shortcut learning and model calibration, we also use PLMs off-the-shelf to make predictions. We build our text classifiers using the Huggingface Transformers library¹. We report results averaged across five runs per task.

4. Results & Discussion

Comparison Across Models and Tasks: For all models in our analyses, we find more than chance ($> 50\%$) shortcut learning for every task. Shown in Table 3, we observe a negative relationship between model calibration and shortcut trade-off (T_{sc}), i.e models considered more calibrated in terms of expected calibration error also tend to rely more on shortcuts when making predictions. This finding highlights that metrics like Expected Calibration Error (ECE), which assess statistical model calibration, do not align with the model’s robustness in terms of learning fewer spurious correlations. Across models, we find that BERT and RoBERTa rely on more shortcut-cued predictions but appear statistically more calibrated. In contrast, DeBERTa and BART rely on fewer shortcuts but appear statistically less calibrated.

Shortcuts Learned: Across various datasets, we observe a notable difference in the extent of shortcut learning. Models make more lexicon-cued predictions on datasets such as AG News, SST2 and TREC, while more grammar-cued predictions are made on Hate and Irony datasets, as shown in Figure 1. Further, the results in Table 3 show that models exhibit a higher degree of

¹<https://huggingface.co/models>




Figure 2: Change in model performance and calibration before and after fine-tuning. (a): red represents an increase and blue represents a decrease in ECE after fine-tuning, (b): red represents an increase and blue represents a decrease in shortcut-cued predictions, (c): red represents a decrease and blue represents an increase in F1 after fine-tuning

shortcut learning on AG News, Hate, and TREC, while COLA shows relatively lower levels. This difference can be attributed to the linguistic characteristics present in the task samples. In COLA, the data includes acceptability labels, and the word tokens appear in a wider range of contexts, not limited to specific topics or sentence formation (e.g. questions) for two distinct categories. On the other hand, AG News, Hate, and TREC samples contain more prevalent lexical cues that repeat across samples. Across AG News and TREC, while different types of question words consistently appear across respective question types in TREC, AG News contains topic-specific words like ‘*President*’ and ‘*Minister*’ appearing in the ‘World’ category. Although these words can act as genuine cues, PLMs heavily rely on them, resulting in misclassification. Examples of this behavior are shown in Table 4, where the highlighted tokens in given text examples are shortcuts typically associated with other labels.

Fine-tuning Effects: We evaluate changes in shortcut learning (P_{sc}), task performance (F1), and calibration (ECE) in models due to fine-tuning. Figure 2 shows the changes observed in classification performance before and after fine-tuning PLMs. We observe that fine-tuning does not always lead to calibration improvements, which aligns with prior findings (Jiang et al., 2021; Kong et al., 2020). Note that models become increasingly miscalibrated on Hate and SST2 tasks, while showing improved calibration for AG News, attributed to the increased accuracy due to shortcut learning of models on AG News. We also find that shortcut learning reduces after fine-tuning for some models, especially on COLA, which we attribute to the words appearing across a wider variety of contexts in its samples, as the dataset is not constricted to specific topics or affect statements.

Shortcut Impacts on Model Confidence: While fine-tuning models results in increased confidence on

correct as well as incorrect predictions, we focus on instances where the calibration error estimates such as ECE are unable to capture the underconfident predictions, i.e. predictions that are correct, but less confident than average. In Figure 3, we plot reliability diagrams and shortcut-cued prediction distributions across model confidence for DeBERTa on two tasks, i.e. TREC and AG News. While DeBERTa portrays similar F1 and ECE for both tasks, we observe that the confidence and shortcut distributions are starkly different. Specifically, DeBERTa predictions on TREC are underconfident in many cases, while on AG News, the model confidence aligns with the correctness of model decisions. Further, while DeBERTa heavily relies on shortcuts for both the tasks, TREC relies more on grammar-cues, while AG News predictions are more often lexicon-cued. This discrepancy in model confidence and shortcuts utilized per confidence bin is not captured in statistical calibration error metrics like ECE. It is crucial to highlight that statistical calibration error metrics like ECE fail to capture the divergence in model confidence and the specific shortcuts employed within confidence bins. While ECE may provide an overall assessment of model calibration, it falls short in capturing the complex interplay between confidence, shortcuts, and their impact on prediction reliability.

Is minimizing ECE enough? We discover that ECE is not a dependable metric, and can be low even when the model is highly overconfident. Thus, a lower ECE does not necessarily indicate more reliable predictions by a model. To illustrate this, let’s consider the results presented in Table 3, specifically for the Hate and AG News tasks. Both tasks demonstrate significant levels of shortcut learning in language models. However, the shortcuts learned in the AG News task lead to more accurate predictions, whereas shortcuts learned in the Hate task result in more incorrect predictions,

Data	Text	Actual	Predicted
AG News	U.S. Seeks Reconciliation with Oil -Rich Venezuela SAO PAULO, Brazil (Reuters) - The United States said on Monday it will seek better ties with oil-rich Venezuela in the clearest sign since President Hugo Chavez won a recall referendum in August that Washington is looking for reconciliation with the firebrand populist.	World	Business
TREC	What do bats eat ?	Entity	Description

Table 4: Examples of misclassification due to shortcuts




Figure 3: Difference in distribution of shortcut-cued predictions on fine-tuned DeBERTa for (a) TREC and (b) AG News. Models show similar performance on both tasks in terms of F1 and ECE; $F_1^{AG\ News} = 94.99$, $F_1^{TREC} = 94.06$; $ECE^{AG\ News} = 0.01$, $ECE^{TREC} = 0.03$.

inflating the latter’s Expected Calibration Error (ECE) and causing it to be considered less calibrated. Note that a model fine-tuned on the AG News task may exhibit the appearance of being a reliable predictor, yet its reliability is inflated by the correctness of its predictions. For instance, in Figure 3, we find that underconfident model predictions are accurate in many cases. The ECE metric does not penalize accurate but underconfident and accurate predictions.

This limitation underscores a significant drawback associated with using Expected Calibration Error (ECE) as the prevailing calibration estimate in PLM calibration literature. The use of ECE fails to consider the presence of spurious associations learned by language models.

5. Related Work

Shortcut Learning General-purpose neural language models have been shown to learn spurious patterns existing within natural language text, due to the language variety cues within the training corpora (Nguyen et al., 2021). While initial research claimed that pre-trained language models are robust to out-of-distribution (OOD) detection and cross-domain generalization (Hendrycks et al., 2020), recent analyses have shown that PLMs, and their fine-tuned versions rely on specific keyword-based shortcuts to perform classification (Moon et al., 2021). This phenomenon hinders the fine-tuned models from

learning generalizable decision rules. PLMs have also been shown to rely on syntactic heuristics to perform natural language inference tasks (McCoy et al., 2019). In light of these findings, research focusing on the automatic identification and mitigation of spurious cues within training and fine-tuning data has also been proposed (Tu et al., 2020; T. Wang et al., 2022; Z. Wang & Culotta, 2020).

Calibration in Neural Language Models With the increased application of neural network architectures in high-risk real-world settings, their calibration has become an extensively studied topic in recent years (Hendrycks et al., 2020; Malinin & Gales, 2018; Thulasidasan et al., 2019). Recent research has focused on improving the calibration of neural networks, particularly in the context of deep learning. Various methods have been proposed to achieve better calibration, including temperature scaling (Guo et al., 2017), isotonic regression (Platt et al., 1999), and histogram binning (Zadrozny & Elkan, 2001).

Pre-trained language models have garnered attention due to their tendency to exhibit increasing confidence during training, regardless of the accuracy of their predictions (Y. Chen et al., 2022). However, these models demonstrate better calibration within in-domain (ID) settings while experiencing calibration deterioration in out-of-domain (OOD) scenarios (Desai & Durrett, 2020b). Interestingly, it has been observed that smaller

models achieve improved calibration on in-domain data, whereas larger models exhibit superior calibration on out-of-domain data (Dan & Roth, 2021).

Moreover, fine-tuning pre-trained language models leads to higher levels of miscalibration (Jiang et al., 2021; Kong et al., 2020). This is attributed to the excessive parameterization of the models, resulting in overfitting the training data. These findings highlight the inadequacy of current PLMs in terms of confidence calibration and reliability in decisions.

6. Conclusion

The prevailing belief in existing calibration evaluations of pre-trained language models is that lower calibration error estimates indicate more reliable predictions. However, it has been shown that fine-tuned PLMs often rely on shortcuts to produce overly confident predictions, creating an illusion of improved performance while actually learning decision rules that lack generalizability. The relationship between model reliability, as measured by calibration error, and shortcut learning has received limited attention thus far. This prompts us to question whether a model with lower calibration error can truly be considered reliable in terms of its decision rules. Our findings challenge the prevailing notion by revealing that models with seemingly better calibration also exhibit higher levels of shortcut learning. This highlights the need to bridge the current gap between language model calibration and generalization objectives and underscores the importance of developing comprehensive frameworks to achieve genuinely robust and reliable language models.

References

- Ahuja, K., Sitaram, S., Dandapat, S., & Choudhury, M. (2022). On the Calibration of Massively Multilingual Language Models. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4310–4323. Retrieved March 20, 2023, from <https://aclanthology.org/2022.emnlp-main.290>
- Alt, C., Hübner, M., & Hennig, L. (2019, June). Fine-tuning Pre-Trained Transformer Language Models to Distantly Supervised Relation Extraction [arXiv:1906.08646 [cs]]. <https://doi.org/10.48550/arXiv.1906.08646>
- Barbieri, F., Camacho-Collados, J., Espinosa Anke, L., & Neves, L. (2020). TweetEval: Unified benchmark and comparative evaluation for tweet classification. *Findings of the Association for Computational Linguistics: EMNLP 2020*, 1644–1650. <https://doi.org/10.18653/v1/2020.findings-emnlp.148>
- Chen, Q., Zhuo, Z., & Wang, W. (2019, February). BERT for Joint Intent Classification and Slot Filling [arXiv:1902.10909 [cs]]. Retrieved December 22, 2022, from <http://arxiv.org/abs/1902.10909>
- Chen, Y., Yuan, L., Cui, G., Liu, Z., & Ji, H. (2022, November). A Close Look into the Calibration of Pre-trained Language Models [arXiv:2211.00151 [cs]]. Retrieved April 4, 2023, from <http://arxiv.org/abs/2211.00151>
- Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press.
- Dan, S., & Roth, D. (2021). On the Effects of Transformer Size on In- and Out-of-Domain Calibration. *Findings of the Association for Computational Linguistics: EMNLP 2021*, 2096–2101. <https://doi.org/10.18653/v1/2021.findings-emnlp.180>
- Desai, S., & Durrett, G. (2020a). Calibration of pre-trained transformers. *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 295–302.
- Desai, S., & Durrett, G. (2020b, October). Calibration of Pre-trained Transformers [arXiv:2003.07892 [cs]]. Retrieved April 3, 2023, from <http://arxiv.org/abs/2003.07892>
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- Du, M., He, F., Zou, N., Tao, D., & Hu, X. (2022). Shortcut learning of large language models in natural language understanding: A survey. *arXiv preprint arXiv:2208.11857*.
- Du, M., Manjunatha, V., Jain, R., Deshpande, R., Deroncourt, F., Gu, J., Sun, T., & Hu, X. (2021). Towards interpreting and mitigating shortcut learning behavior of NLU models. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 915–929. <https://doi.org/10.18653/v1/2021.naacl-main.71>

- Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On Calibration of Modern Neural Networks [ISSN: 2640-3498]. *Proceedings of the 34th International Conference on Machine Learning*, 1321–1330. Retrieved March 20, 2023, from <https://proceedings.mlr.press/v70/guo17a.html>
- He, P., Liu, X., Gao, J., & Chen, W. (n.d.). DeBERTa: Decoding-enhanced bert with disentangled attention. *International Conference on Learning Representations*.
- Hendrycks, D., Liu, X., Wallace, E., Dziedziec, A., Krishnan, R., & Song, D. (2020). Pretrained Transformers Improve Out-of-Distribution Robustness. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2744–2751. <https://doi.org/10.18653/v1/2020.acl-main.244>
- Hovy, E., Gerber, L., Hermjakob, U., Lin, C.-Y., & Ravichandran, D. (2001). Toward semantics-based answer pinpointing. *Proceedings of the First International Conference on Human Language Technology Research*. <https://www.aclweb.org/anthology/H01-1069>
- Jiang, Z., Araki, J., Ding, H., & Neubig, G. (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. https://doi.org/10.1162/tacl_a_00407
- Kim, J., Na, D., Choi, S., & Lim, S. (2023, February). Bag of Tricks for In-Distribution Calibration of Pretrained Transformers [arXiv:2302.06690 [cs]]. Retrieved April 3, 2023, from <http://arxiv.org/abs/2302.06690>
- Kong, L., Jiang, H., Zhuang, Y., Lyu, J., Zhao, T., & Zhang, C. (2020). Calibrated Language Model Fine-Tuning for In- and Out-of-Distribution Data. *ArXiv*. <https://doi.org/10.18653/v1/2020.emnlp-main.102>
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., & Soricut, R. (n.d.). Albert: A lite bert for self-supervised learning of language representations. *International Conference on Learning Representations*.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., & Zettlemoyer, L. (2020). Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 7871–7880.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Malinin, A., & Gales, M. (2018, November). Predictive Uncertainty Estimation via Prior Networks [arXiv:1802.10501 [cs, stat]]. <https://doi.org/10.48550/arXiv.1802.10501>
- McCoy, T., Pavlick, E., & Linzen, T. (2019). Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3428–3448. <https://doi.org/10.18653/v1/P19-1334>
- Moon, S. J., Mo, S., Lee, K., Lee, J., & Shin, J. (2021). MASKER: Masked Keyword Regularization for Reliable Text Classification [Number: 15]. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15), 13578–13586. <https://doi.org/10.1609/aaai.v35i15.17601>
- Naeini, M. P., Cooper, G., & Hauskrecht, M. (2015). Obtaining well calibrated probabilities using bayesian binning. *Proceedings of the AAAI conference on artificial intelligence*, 29(1).
- Nguyen, D., Rosseel, L., & Grieve, J. (2021). On learning and representing social meaning in NLP: A sociolinguistic perspective. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 603–612. <https://doi.org/10.18653/v1/2021.naacl-main.50>
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *arXiv preprint arXiv:1802.05365*, 12.
- Platt, J., et al. (1999). Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3), 61–74.
- Sap, M., Shwartz, V., Bosselut, A., Choi, Y., & Roth, D. (2020). Commonsense reasoning for natural language processing. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*, 27–33. <https://doi.org/10.18653/v1/2020.acl-tutorials.7>
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive

- deep models for semantic compositionality over a sentiment treebank. *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642. <https://aclanthology.org/D13-1170>
- Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. *International conference on machine learning*, 3319–3328.
- Thulasidasan, S., Chennupati, G., Bilmes, J. A., Bhattacharya, T., & Michalak, S. (2019). On Mixup Training: Improved Calibration and Predictive Uncertainty for Deep Neural Networks. *Advances in Neural Information Processing Systems*, 32. Retrieved April 4, 2023, from <https://proceedings.neurips.cc/paper/2019/hash/36ad8b5f42db492827016448975cc22d-Abstract.html>
- Tu, L., Lalwani, G., Gella, S., & He, H. (2020). An empirical study on robustness to spurious correlations using pre-trained language models. *Transactions of the Association for Computational Linguistics*, 8, 621–633.
- Wang, T., Sridhar, R., Yang, D., & Wang, X. (2022). Identifying and mitigating spurious correlations for improving robustness in NLP models. *Findings of the Association for Computational Linguistics: NAACL 2022*, 1719–1729. <https://doi.org/10.18653/v1/2022.findings-naacl.130>
- Wang, Z., & Culotta, A. (2020). Identifying spurious correlations for robust text classification. *Findings of the Association for Computational Linguistics: EMNLP 2020*, 3431–3440. <https://doi.org/10.18653/v1/2020.findings-emnlp.308>
- Warstadt, A., Singh, A., & Bowman, S. (2019). Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7, 625–641.
- Zadrozny, B., & Elkan, C. (2001). Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. *Icml*, 1, 609–616.
- Zhang, X., Zhao, J. J., & LeCun, Y. (2015). Character-level convolutional networks for text classification. *NIPS*.