

# A Data Cartography based MixUp for Pre-trained Language Models

Seo Yeon Park and Cornelia Caragea

Computer Science

University of Illinois Chicago

spark313@uic.edu    cornelia@uic.edu

## Abstract

MixUp is a data augmentation strategy where additional samples are generated during training by combining *random pairs* of training samples and their labels. However, selecting random pairs is not potentially an optimal choice. In this work, we propose TD-MixUp, a novel MixUp strategy that leverages **T**rainin**D**ynamics and allows more informative samples to be combined for generating new data samples. Our proposed TD-MixUp first measures confidence, variability, (Swayamdipta et al., 2020), and Area Under the Margin (AUM) (Pleiss et al., 2020) to identify the characteristics of training samples (e.g., as easy-to-learn or ambiguous samples), and then interpolates these characterized samples. We empirically validate that our method not only achieves competitive performance using a smaller subset of the training data compared with strong baselines, but also yields lower expected calibration error on the pre-trained language model, BERT, on both in-domain and out-of-domain settings in a wide range of NLP tasks. We publicly release our code.<sup>1</sup>

## 1 Introduction

MixUp (Zhang et al., 2018) is a simple data augmentation strategy in which additional samples are generated during training by combining *random pairs* of training samples and their labels. While simple to implement, MixUp has been shown to improve both predictive performance and model calibration (i.e., avoiding over-confident predictions) (Guo et al., 2017) due to its regularization effect through data augmentation (Thulasidasan et al., 2019). However, selecting *random pairs* in MixUp might not necessarily be optimal.

Despite this, MixUp has been explored for NLP tasks with substantial success using hidden state representations (Verma et al., 2019). For instance, Sun et al. (2020) explored MixUp, which uses the

hidden representation of BERT (Devlin et al., 2019) to synthesize additional samples from randomly selected pairs. Yin et al. (2021) proposed MixUp, which uses the hidden representation of RoBERTa (Liu et al., 2019) to interpolate all samples in the same mini-batch to better cover the feature space. To date, only a few prior works have focused on selecting informative samples for MixUp. For example, Chen et al. (2020) proposed semi-supervised learning, which interpolates labeled and unlabeled data based on entropy. Kong et al. (2020) explored BERT calibration with MixUp, which generates new samples by exploiting the distance between samples in the feature space.

Recently, Swayamdipta et al. (2020) introduced data maps, which allow evaluating data quality by using training dynamics (i.e., the behavior of a model as training progresses). Specifically, they consider the mean and standard deviation of the gold label probabilities, predicted for each sample across training epochs (i.e., confidence and variability), and characterize data into three different categories: (1) samples that the model predicts correctly and consistently (i.e., **easy-to-learn**); (2) samples where true class probabilities vary frequently during training (i.e., **ambiguous**); and (3) samples that are potentially mis-labeled or erroneous (i.e., **hard-to-learn**). The author revealed that the easy-to-learn samples are useful for model optimization (parameter estimation) and without such samples the training could potentially fail to converge, while the ambiguous samples are those on which the model struggles the most and push the model to become more robust, hence, these ambiguous samples are the most beneficial for learning since they are the most challenging for the model.

Inspired by these observations, we propose a novel MixUp strategy which we call TDMixUp that monitors training dynamics and interpolates easy-to-learn samples with ambiguous samples in the feature space. That is, we pair one sample from

<sup>1</sup><https://github.com/seoyeon-p/TDMixUp>

the easy-to-learn set with another sample from the ambiguous set to allow more informative samples to be combined for MixUp. Accordingly, we generate new samples that share the characteristics of both easy-to-learn and ambiguous data samples and are hence more beneficial for learning. However, the easy-to-learn and the ambiguous sets can contain mis-labeled samples that can degrade the model performance. Consequently, we measure another training dynamic, Area Under the Margin (AUM) (Pleiss et al., 2020), to filter out possibly mis-labeled samples in each set. We validate our proposed method on a wide range of natural language understanding tasks including textual entailment, paraphrase detection, and commonsense reasoning tasks. We achieve competitive accuracy and low expected calibration error (Guo et al., 2017) on both in-domain and out-of-domain settings for the pre-trained language model BERT (Devlin et al., 2019), without using the full training data.

## 2 Proposed Approach: TDMixUp

We introduce our proposed TDMixUp, which generates additional samples based on the characteristics of the data samples. We first reveal the characteristics of each data sample by using training dynamics, i.e., confidence, variability, and Area Under the Margin (AUM). We then describe our MixUp operation that combines training samples based on the above data characteristics that are measured during training.

### 2.1 Data Samples Characterization

We first introduce **confidence** and **variability**, that are used to evaluate the characteristics of each individual sample (Swayamdipta et al., 2020). These statistics are calculated for each sample  $(x_i, y_i)$  over  $E$  training epochs.

**Confidence** We define confidence as the mean model probability of the true label  $y_i$  across epochs:

$$\hat{\mu}_i = \frac{1}{E} \sum_{e=1}^E p_{\theta^{(e)}}(y_i|x_i) \quad (1)$$

where  $p_{\theta^{(e)}}$  denotes the model’s probability with parameter  $\theta^{(e)}$  at the end of  $e^{th}$  epoch.

**Variability** We define variability as the standard deviation of  $p_{\theta^{(e)}}$  across epochs  $E$ :

$$\hat{\sigma}_i = \sqrt{\frac{\sum_{e=1}^E (p_{\theta^{(e)}}(y_i|x_i) - \hat{\mu}_i)^2}{E}} \quad (2)$$

Given these statistics per sample, we identify the top 33% easy-to-learn samples, i.e., those samples that the model predicts correctly and consistently across epochs (high-confidence, low-variability), and the top 33% ambiguous samples, i.e., those samples whose true class probabilities have a high variance during training (high-variability).

**Area Under the Margin (AUM)** As another measure of data quality, we monitor training dynamics using the Area Under the Margin (AUM) (Pleiss et al., 2020). AUM measures how different a true label for a sample is compared to a model’s belief at each epoch and is calculated as the average difference between the logit values for a sample’s assigned class (gold label) and its highest non-assigned class across training epochs. Formally, given a sample  $(x_i, y_i)$ , we compute  $AUM(x_i, y_i)$  as the area under the margin averaged across all training epochs  $E$ . Specifically, at some epoch  $e \in E$ , the margin is defined as:

$$M^e(x_i, y_i) = z_{y_i} - \max_{y_i \neq k} (z_k) \quad (3)$$

where  $M^e(x_i, y_i)$  is the margin of sample  $x_i$  with true label  $y_i$ ,  $z_{y_i}$  is the logit corresponding to the true label  $y_i$ , and  $\max_{y_i \neq k} (z_k)$  is the largest *other* logit corresponding to label  $k$  not equal to  $y_i$ . The AUM of  $(x_i, y_i)$  across all epochs is:

$$AUM(x_i, y_i) = \frac{1}{E} \sum_{e=1}^E M^e(x_i, y_i) \quad (4)$$

Intuitively, while both AUM and confidence measure training dynamics, confidence simply measures the probability output of the gold label and how much it fluctuates over the training epochs. In contrast, AUM measures the probability output of the gold label with respect to the model’s belief in what the label for a sample should be according to its generalization capability (derived by observing other similar samples during training). More precisely, AUM considers each logit value and measures how much the gold label assigned logit value differs from *the other largest* logit value, which allows identifying mis-labeled samples.

To identify possibly mis-labeled data in each set (i.e., the set of easy-to-learn and the set of ambiguous samples that are categorized by confidence and variability as described above), we first fine-tune a model on each set, respectively, with inserting fake data (i.e., threshold samples). Data with similar or worse AUMs than threshold samples can be

assumed to be mis-labeled (Pleiss et al., 2020). We construct threshold samples by taking a subset of the training data and re-assigning their labels randomly, including a class that does not really exist. Specifically, given  $N$  training samples that belong to  $c$  classes, we randomly select  $N/(c+1)$  samples per class and re-assign their labels to classes that are different from the original class. We then train a model on training samples including threshold samples and measure the AUMs of all training data. We identify possible mis-labeled data by computing a threshold value (i.e., the  $k$ th percentile threshold sample AUMs where  $k$  is a hyper-parameter chosen on the validation set). At last, we filter out samples that have lower AUM than the threshold value.

## 2.2 MixUp

MixUp training generates vicinity training samples according to the rule introduced in Zhang et al. (2018):

$$\begin{aligned}\tilde{x} &= \lambda x_i + (1 - \lambda)x_j \\ \tilde{y} &= \lambda y_i + (1 - \lambda)y_j\end{aligned}\quad (5)$$

where  $x_i$  and  $x_j$  are two randomly sampled input points,  $y_i$  and  $y_j$  are their associated one-hot encoded labels, and  $\lambda$  is a mixing ratio sampled from a Beta( $\alpha$ ,  $\alpha$ ) distribution with a hyper-parameter  $\alpha$ . In standard MixUp, training data is augmented by linearly interpolating *random* training samples in the input space. In contrast, our proposed TD-MixUp interpolates one easy-to-learn sample with one ambiguous sample after applying AUM to filter potential erroneous samples that harm performance. Our current implementation uses easy-to-learn and ambiguous data loaders respectively and then applies MixUp to a randomly sampled mini-batch of each loader. We train a model on the generated TD-MixUp samples in addition to the easy-to-learn and ambiguous samples using the cross entropy-loss.

## 3 Experiments and Results

### 3.1 Tasks and Datasets

We evaluate our TDMixUp on three natural language understanding tasks. We describe our in-domain and out-of-domain sets as follows.

**Natural Language Inference (NLI)** Stanford Natural Language Inference (SNLI) is a task to predict if the relation between a hypothesis and a premise is *entailment*, *contradiction*, or *neutral* (Bowman et al., 2015). Multi-Genre Natural Lan-

guage Inference (MNLi) captures NLI with diverse domains (Williams et al., 2018).

**Paraphrase Detection** Quora Question Pairs (QQP) is a paraphrase detection task to test if two questions are semantically equivalent (Iyer et al., 2017). TwitterPPDB (TPPDB) is a dataset built to determine whether sentence pairs from Twitter convey similar semantics when they share URLs (Lan et al., 2017).

**Commonsense Reasoning** Situations With Adversarial Generations (SWAG) is a commonsense reasoning task to choose the most plausible continuation of a sentence among four candidates (Zellers et al., 2018). HellaSWAG is a dataset built using adversarial filtering to generate challenging out-of-domain samples.

### 3.2 Experimental Setup

We use BERT (Devlin et al., 2019) based classification model and pass the resulting [CLS] representation through a fully connected layer and softmax to predict the label distribution. We follow the published train/validation/test datasets splits as described in Desai and Durrett (2020). To identify mis-labeled samples in the top 33% easy-to-learn samples, we set threshold values  $k$  as: the 80th/80th/50th percentile threshold sample AUMs on SNLI/QQP/SWAG, respectively. More training details and hyper-parameter settings can be found in the Appendix. We evaluate the capability of our TDMixUp strategy to improve both predictive performance and model calibration due to its regularization effect through data augmentation. Hence, we use two metrics: (1) accuracy, and (2) expected calibration error (ECE) (Guo et al., 2017; Desai and Durrett, 2020). We report results averaged across 5 fine-tuning runs with random restarts for all experiments.

### 3.3 Baseline Methods

**BERT** (Devlin et al., 2019) is the pre-trained base BERT model fine-tuned on each downstream task.

**Back Translation Data Augmentation** (Edunov et al., 2018) generates augmented samples by using pre-trained translation models<sup>2</sup> which can generate diverse paraphrases while preserving the semantics

<sup>2</sup>We use FairSeq and set the random sampling temperature as 0.9.

	SNLI		QQP		SWAG	
	Acc	ECE	Acc	ECE	Acc	ECE
100% train	<b>90.04</b> <sub>0.3</sub>	2.54 <sub>0.8</sub>	<b>90.27</b> <sub>0.3</sub>	2.71 <sub>0.5</sub>	<b>79.40</b> <sub>0.4</sub>	2.49 <sub>1.8</sub>
33% train, Easy-to-learn	82.78 <sub>0.6</sub>	16.22 <sub>0.7</sub>	63.16 <sub>0.1</sub>	36.88 <sub>0.1</sub>	75.39 <sub>0.2</sub>	17.51 <sub>0.1</sub>
24% train, Easy-to-learn with AUM	83.03 <sub>0.9</sub>	15.05 <sub>0.9</sub>	66.43 <sub>0.6</sub>	33.93 <sub>0.8</sub>	75.56 <sub>0.1</sub>	15.81 <sub>0.7</sub>
33% train, Ambiguous	89.71 <sub>0.5</sub>	<b>0.74</b> <sub>0.1</sub>	87.51 <sub>0.5</sub>	1.71 <sub>0.4</sub>	75.91 <sub>0.6</sub>	<b>1.84</b> <sub>0.7</sub>
24% train, Ambiguous with AUM	87.88 <sub>0.7</sub>	7.09 <sub>0.8</sub>	88.63 <sub>0.5</sub>	6.36 <sub>0.6</sub>	71.74 <sub>0.4</sub>	7.55 <sub>1.1</sub>
66% train, Easy-to-learn & Ambiguous	89.65 <sub>0.2</sub>	2.64 <sub>0.5</sub>	90.23 <sub>0.7</sub>	<b>1.35</b> <sub>0.4</sub>	78.78 <sub>0.5</sub>	2.51 <sub>0.8</sub>

  

	MNLI		TwitterPPDB		HellaSWAG	
	Acc	ECE	Acc	ECE	Acc	ECE
100% train	73.52 <sub>0.3</sub>	7.09 <sub>2.1</sub>	<b>87.63</b> <sub>0.4</sub>	8.51 <sub>0.6</sub>	<b>34.48</b> <sub>0.2</sub>	12.62 <sub>2.8</sub>
33% train, Easy-to-learn	61.41 <sub>0.8</sub>	36.68 <sub>1.9</sub>	81.07 <sub>0.8</sub>	18.92 <sub>0.7</sub>	33.59 <sub>1.1</sub>	29.38 <sub>2.1</sub>
24% train, Easy-to-learn with AUM	62.97 <sub>1.5</sub>	32.48 <sub>2.9</sub>	82.16 <sub>0.7</sub>	17.46 <sub>1.0</sub>	33.67 <sub>1.4</sub>	16.89 <sub>2.6</sub>
33% train, Ambiguous	72.52 <sub>1.2</sub>	10.73 <sub>1.0</sub>	86.62 <sub>0.6</sub>	<b>6.01</b> <sub>1.1</sub>	34.29 <sub>0.9</sub>	8.40 <sub>1.3</sub>
24% train, Ambiguous with AUM	70.87 <sub>0.9</sub>	17.23 <sub>1.6</sub>	86.59 <sub>0.8</sub>	7.31 <sub>0.8</sub>	33.81 <sub>1.0</sub>	<b>3.76</b> <sub>2.3</sub>
66% train, Easy-to-learn & Ambiguous	<b>73.89</b> <sub>0.6</sub>	<b>3.46</b> <sub>1.9</sub>	87.29 <sub>0.3</sub>	8.04 <sub>0.7</sub>	34.43 <sub>0.2</sub>	9.68 <sub>1.1</sub>

Table 1: The comparison of accuracy and expected calibration error (ECE) in percentage on in-domain (top) and out-of-domain (bottom) for BERT. We compare the results of fine-tuning on subsets of train samples (i.e., easy-to-learn and ambiguous samples) and fine-tuning on the entire 100% of training samples. Lower ECE implies better calibrated models. We report the mean accuracy across five training runs with the standard deviation shown in subscript (e.g., 90.04<sub>0.3</sub> indicates 90.04 ± 0.3).

of the original sentences. In experiments, we translate original sentences from English to German and then translate them back to English to obtain the paraphrases.

**MixUp** (Zhang et al., 2018) generates augmented samples by interpolating random training samples in the input space (obtained from the first layer of the BERT pre-trained language model).

**Manifold MixUp (M-MixUp)** (Verma et al., 2019) generates additional samples by interpolating random training samples in the feature space (obtained from the task-specific layer on top of the BERT pre-trained language model).

**MixUp for Calibration** (Kong et al., 2020) generates augmented samples by utilizing the cosine distance between samples in the feature space.

Note that the above baselines use 100% training data while our proposed method focuses on particular subsets of the training data.

### 3.4 Results

**Fine-tuning on Subsets of Training Data** To explore the effect of different subsets of the data that are characterized using training dynamics, we compare the result of BERT fine-tuned on 100% training data with the results of BERT when fine-tuned on these subsets, and show the comparison in Table 1. Note that, for each task, we train the model on in-domain training set, and evaluate on

in-domain and out-of-domain test sets. We make the following observations: First, we observe that accuracy and ECE improve when we filter out possibly mis-labeled samples in the top 33% easy-to-learn samples by using AUM in all cases (for both in-domain and out-of-domain test sets). Specifically, using 24% train, Easy-to-learn with AUM<sup>3</sup> returns better accuracy and lower ECE than 33% train, easy-to-learn, showing that there are some potentially erroneous samples that harm the performance in the top 33% easy-to-learn samples. We manually investigate filtered samples on the top 33% easy-to-learn samples and observe that the top 33% easy-to-learn samples indeed include mis-labeled samples. For example, in SNLI, we observe that the relation between the following pairs of sentences in the top 33% easy-to-learn samples is *contradiction* when it should be *neutral*: <Two opposing wrestlers competing to pin one another.> ‘Two women are shopping in a boutique.’> and <‘A person dressed in a colorful costume is holding some papers.’; ‘the cat jumps on the dog.’>. In contrast, we observe that in many cases accuracy and ECE worsen when we filter out possibly mis-labeled samples in the 33% ambiguous samples by using AUM, suggesting that all ambiguous samples are useful for learning and generalization (which is consistent with Swayamdipta et al. (2020)). Sec-

<sup>3</sup>Although we write 24% in Table 1 for simplicity, the percentage of training samples on SNLI/QQP/SWAG after AUM filtering are 20%/24%/25%, respectively

	SNLI		QQP		SWAG	
	Acc	ECE	Acc	ECE	Acc	ECE
100% train	90.04 <sub>0.3</sub>	2.54 <sub>0.8</sub>	90.27 <sub>0.3</sub>	2.71 <sub>0.5</sub>	79.40 <sub>0.4</sub>	2.49 <sub>1.8</sub>
100% train, MixUp (Zhang et al., 2018)	88.82 <sub>0.2</sub>	7.73 <sub>1.1</sub>	89.12 <sub>0.5</sub>	9.04 <sub>0.8</sub>	74.98 <sub>2.3</sub>	7.08 <sub>1.0</sub>
100% train, M-MixUp (Verma et al., 2019)	89.45 <sub>0.9</sub>	1.51 <sub>0.8</sub>	89.93 <sub>0.6</sub>	3.02 <sub>1.0</sub>	78.26 <sub>0.4</sub>	4.12 <sub>0.6</sub>
100% train, MixUp for Calibration (Kong et al., 2020)	89.25 <sub>0.5</sub>	2.16 <sub>0.5</sub>	90.24 <sub>0.3</sub>	5.22 <sub>0.6</sub>	79.44 <sub>0.6</sub>	<b>1.10</b> <sub>0.4</sub>
100% train, Back Translation Data Augmentation (Edunov et al., 2018)	89.22 <sub>0.5</sub>	1.98 <sub>0.6</sub>	89.18 <sub>0.6</sub>	5.01 <sub>0.3</sub>	76.22 <sub>0.9</sub>	1.24 <sub>0.2</sub>
66% train, TDMixUp, Easy-to-learn + Ambiguous	89.73 <sub>0.1</sub>	2.39 <sub>0.8</sub>	89.77 <sub>0.2</sub>	1.89 <sub>0.4</sub>	78.38 <sub>0.3</sub>	4.21 <sub>0.3</sub>
57% train, TDMixUp, Easy-to-learn with AUM + Ambiguous (Ours)	<b>90.31</b> <sub>0.2</sub>	<b>1.22</b> <sub>0.4</sub>	<b>90.42</b> <sub>0.2</sub>	<b>1.53</b> <sub>0.9</sub>	<b>79.59</b> <sub>0.3</sub>	2.16 <sub>0.4</sub>

  

	MNLI		TwitterPPDB		HellaSWAG	
	Acc	ECE	Acc	ECE	Acc	ECE
100% train	73.52 <sub>0.3</sub>	7.09 <sub>2.1</sub>	87.63 <sub>0.4</sub>	8.51 <sub>0.6</sub>	34.48 <sub>0.2</sub>	12.62 <sub>2.8</sub>
100% train, MixUp (Zhang et al., 2018)	69.19 <sub>0.8</sub>	19.51 <sub>2.1</sub>	87.45 <sub>0.3</sub>	11.70 <sub>1.6</sub>	33.22 <sub>0.4</sub>	10.93 <sub>2.0</sub>
100% train, M-MixUp (Verma et al., 2019)	73.22 <sub>0.6</sub>	8.06 <sub>1.2</sub>	87.58 <sub>0.7</sub>	7.68 <sub>1.3</sub>	34.86 <sub>0.9</sub>	13.56 <sub>1.6</sub>
100% train, MixUp for Calibration (Kong et al., 2020)	64.90 <sub>0.5</sub>	17.75 <sub>1.8</sub>	74.51 <sub>1.1</sub>	11.83 <sub>1.0</sub>	32.51 <sub>0.8</sub>	31.61 <sub>2.3</sub>
100% train, Back Translation Data Augmentation (Edunov et al., 2018)	73.15 <sub>0.7</sub>	8.46 <sub>1.3</sub>	86.82 <sub>0.7</sub>	8.83 <sub>0.6</sub>	34.97 <sub>0.4</sub>	22.68 <sub>3.3</sub>
66% train, TDMixUp, Easy-to-learn + Ambiguous	72.83 <sub>1.1</sub>	5.84 <sub>1.9</sub>	87.63 <sub>0.2</sub>	6.48 <sub>0.7</sub>	34.11 <sub>0.1</sub>	10.54 <sub>1.6</sub>
57% train, TDMixUp, Easy-to-learn with AUM + Ambiguous (Ours)	<b>74.28</b> <sub>0.6</sub>	<b>2.91</b> <sub>1.4</sub>	<b>87.89</b> <sub>0.3</sub>	<b>6.08</b> <sub>0.4</sub>	<b>35.21</b> <sub>0.6</sub>	<b>9.45</b> <sub>1.3</sub>

Table 2: Accuracy (in percentage) and expected calibration error (ECE) on in-domain (top) and out-of-domain (bottom) for BERT when comparing our proposed TDMixUp with baseline methods. Bold text shows the best performance and calibration. For 100% train results, we use reported results by Desai and Durrett (2020). We report the mean accuracy across five training runs with the standard deviation shown in subscript (e.g., 90.04<sub>0.3</sub> indicates 90.04 ± 0.3).

ond, we observe that fine-tuning BERT on both the easy-to-learn and the ambiguous samples (66% train, Easy-to-learn & Ambiguous) achieves similar performance and ECE as 100% train.

**Main Results** Table 2 shows the result of the comparison of our proposed TDMixUp and baseline methods. We observe that our proposed method generally achieves higher accuracy and lower ECE on both in-domain and out-of-domain settings compared to any baseline using the full 100% training data showing the effectiveness of our TDMixUp strategy.

### 3.5 Ablation Study

To compare the impact of the MixUp operation on samples generated by *random pairing* and on samples generated by *informative pairing*, we conduct an ablation study. Specifically, we compare the results of MixUp on 66% train set (i.e., conduct the MixUp operation between randomly selected samples on 66% train set, which is the union of the top 33% easy-to-learn and the top 33% ambiguous samples) and our proposed TDMixUp (i.e., conduct the MixUp operation between the easy-to-learn filtered by AUM and the ambiguous samples). As shown in Table 3, we observe that our proposed TDMixUp which selects informative samples to combine performs better with respect to accuracy and ECE than vanilla MixUp that selects random samples, in all cases (in-domain and out-of-domain).

	Acc	ECE	Acc	ECE	Acc	ECE
	SNLI		QQP		SWAG	
Random	89.59	1.70	89.87	3.06	79.15	4.51
Ours	90.31	1.22	90.42	1.53	79.59	2.16

  

	MNLI		TwitterPPDB		HellaSWAG	
	Acc	ECE	Acc	ECE	Acc	ECE
Random	73.22	6.89	87.23	6.53	34.43	15.87
Ours	74.28	2.91	87.89	6.08	35.21	9.45

Table 3: The results comparison of MixUp selecting random samples on the union of the top 33% easy-to-learn and the top 33% ambiguous samples (i.e., Random) and our method.

## 4 Conclusion

In this work, we propose a novel MixUp that leverages training dynamics (confidence, variability, and Area Under the Margin) to allow more informative samples to be combined in generating augmented samples. We empirically validate that our method not only achieves competitive accuracy but also calibrates BERT model on various NLP tasks, both on in-domain and out-of-domain settings.

## Acknowledgements

This research is supported in part by NSF CAREER award #1802358 and NSF CRI award #1823292. Any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of NSF. We thank AWS for computing resources. We also thank our anonymous reviewers for their constructive feedback.

## References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Jiaao Chen, Zichao Yang, and Diyi Yang. 2020. [Mix-Text: Linguistically-informed interpolation of hidden space for semi-supervised text classification](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2147–2157, Online. Association for Computational Linguistics.
- Shrey Desai and Greg Durrett. 2020. [Calibration of pre-trained transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 295–302, Online. Association for Computational Linguistics.
- 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 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Shankar Iyer, Nikhil Dandekar, and Kornél Csernai. 2017. First quora dataset release: Question pairs. In *Quora*.
- Lingkai Kong, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. 2020. [Calibrated language model fine-tuning for in- and out-of-distribution data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1326–1340, Online. Association for Computational Linguistics.
- Wuwei Lan, Siyu Qiu, Hua He, and Wei Xu. 2017. [A continuously growing dataset of sentential phrases](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1224–1234, Copenhagen, Denmark. Association for Computational Linguistics.
- 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*.
- Geoff Pleiss, Tianyi Zhang, Ethan R. Elenberg, and Kilian Q. Weinberger. 2020. [Identifying mislabeled data using the area under the margin ranking](#). In *Advances in Neural Information Processing Systems*, volume 33.
- Lichao Sun, Congying Xia, Wenpeng Yin, Tingting Liang, Philip Yu, and Lifang He. 2020. [Mixup-transformer: Dynamic data augmentation for NLP tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3436–3440, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Sunil Thulasidasan, Gopinath Chennupati, Jeff A Bilmes, Tanmoy Bhattacharya, and Sarah Michalak. 2019. [On mixup training: Improved calibration and predictive uncertainty for deep neural networks](#). In *Advances in Neural Information Processing Systems*, volume 32.
- Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. 2019. [Manifold mixup: Better representations by interpolating hidden states](#). In *International Conference on Machine Learning*, pages 6438–6447. PMLR.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.
- Wenpeng Yin, Huan Wang, Jin Qu, and Caiming Xiong. 2021. [BatchMixup: Improving training by interpolating hidden states of the entire mini-batch](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4908–4912, Online. Association for Computational Linguistics.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. [SWAG: A large-scale adversarial dataset for grounded commonsense inference](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.

Hongyi Zhang, Moustapha Cissé, Yann N. Dauphin, and David Lopez-Paz. 2018. *mixup: Beyond empirical risk minimization*. In *Proceedings of the 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada*.

## A Supplementary Materials

### A.1 Training Details

In our experiments, we use *bert-base-uncased* classification model on top of a task-specific fully-connected layer. The model is fine-tuned with a maximum of 3 epochs, batch size of 16 for SNLI and QQP, batch size 4 for SWAG, a learning rate of  $1e-5$ , gradient clip of 1.0, and no weight decay. We use the hyper-parameter of MixUp  $\alpha$  as 0.4. All hyper-parameters are estimated on the validation set of each task. For all results, we report averaged results across 5 fine-tuning runs with random starts. Finally, all experiments are conducted on a single NVIDIA RTX A6000 48G GPU with the total time for fine-tuning all models being under 24 hours. For each dataset, we follow the published train/validation/test split by [Desai and Durrett \(2020\)](#) and show the statistics of the datasets in Table 4.

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 4: The statistics of all used datasets.

### A.2 Data Maps

In this section, we provide data maps ([Swayamdipta et al., 2020](#)) of our in-domain datasets on *bert-base-uncased* model in Figure 1. These data maps are used to identify the characteristics of each training sample (i.e., easy-to-learn, ambiguous, and hard-to-learn).



Figure 1: Data Maps of SNLI, QQP and SWAG on *bert-base-uncased* model.