Uncertainty-aware Self-training for Few-shot Text Classification

Subhabrata Mukherjee
Microsoft Research AI
Redmond, WA
submukhe@microsoft.com

Ahmed Hassan Awadallah
Microsoft Research AI
Redmond, WA
hassanam@microsoft.com

Abstract

Recent success of large-scale pre-trained language models crucially hinge on fine-tuning them on large amounts of labeled data for the downstream task, that are typically expensive to acquire. In this work, we study self-training as one of the earliest semi-supervised learning approaches to reduce the annotation bottleneck by making use of large-scale unlabeled data for the target task. Standard self-training mechanism randomly samples instances from the unlabeled pool to pseudo-label and augment labeled data. In this work, we propose an approach to improve self-training by incorporating uncertainty estimates of the underlying neural network leveraging recent advances in Bayesian deep learning. Specifically, we propose (i) acquisition functions to select instances from the unlabeled pool leveraging Monte Carlo (MC) Dropout, and (ii) learning mechanism leveraging model confidence for self-training. As an application, we focus on text classification on five benchmark datasets. We show our methods leveraging only 20-30 labeled samples per class for each task for training and for validation can perform within 3% of fully supervised pre-trained language models fine-tuned on thousands of labeled instances with an aggregate accuracy of 91% and improving by upto 12% over baselines.

1 Introduction

Motivation. Deep neural networks are the state-of-the-art for various natural language processing applications. However, one of the biggest challenges facing them is the lack of labeled data to train these complex networks. Not only is acquiring large amounts of labeled data for every task expensive and time consuming, but also it is not feasible to perform large-scale human labeling, in many cases, due to data access and privacy constraints. Recent advances in pre-training help close this gap. In this, deep and large neural networks like BERT [Devlin et al., 2019], GPT-2 [Radford et al., 2019] and RoBERTa [Liu et al., 2019] are trained on millions of documents in a self-supervised fashion to obtain general purpose language representations. A significant challenge now is to fine-tune these models on downstream tasks that still rely on thousands of labeled instances for their superior performance. Semi-supervised learning (SSL) [Chapelle et al., 2010] is one of the promising paradigms to address this shortcoming by making effective use of large amounts of unlabeled data in addition to some labeled data for task-specific fine-tuning. A recent work [Xie et al., 2019] leveraging SSL for consistency learning has shown state-of-the-art performance for text classification with limited labels leveraging auxiliary resources like backtranslation and forms a strong baseline for our work.

Self-training (ST, [III, 1965]) as one of the earliest SSL approaches has recently been shown to obtain state-of-the-art performance for tasks like neural machine translation [He et al., 2019] performing at par with supervised systems without using any auxiliary resources. For self-training, a base model (*teacher*) is trained on some amount of labeled data and used to pseudo-annotate (task-specific) unlabeled data. The original labeled data is augmented with the pseudo-labeled data and used to train a *student* model. The student-teacher training is repeated until convergence. Traditionally,

34th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, Canada.

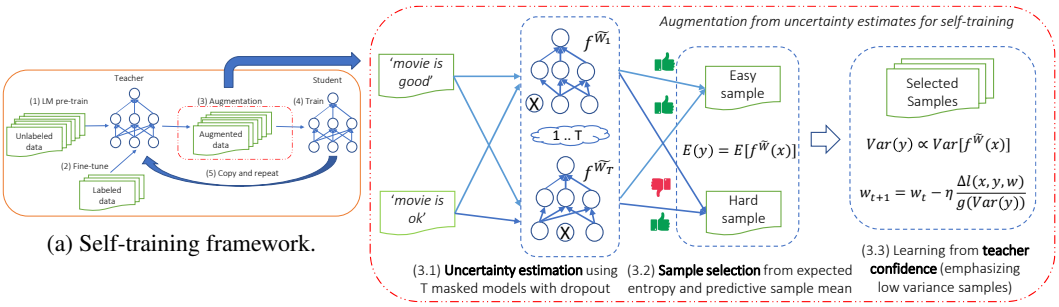

(a) Self-training framework.

(b) Augmenting self-training with uncertainty estimation.

Figure 1: Uncertainty-aware self-training framework.

self-training mechanisms do not consider the teacher uncertainty or perform any sample selection during the pseudo-labeling process. This may result in gradual drifts from self-training on noisy pseudo-labeled instances [Zhang et al., 2017]. Sample selection leveraging teacher confidence has been studied in curriculum learning [Bengio et al., 2009] and self-paced learning [Kumar et al., 2010] frameworks. These works leverage *easiness* of the samples to inform a learning schedule like training on easy concepts first followed by complex ones. Since it is hard to assess the "easiness" of a sample, especially in deep neural network based architectures, these works rely only on the teacher model loss while ignoring its uncertainties that can be leveraged for sample selection.

Intuitively, if the teacher model already predicts some samples with high confidence, then there is little to gain with self-training if we focus only on these samples. On the other hand, hard examples for which the teacher model has less confidence are hard to rely on for self-training as these could be noisy or too difficult to learn from. In this scenario, the model could benefit from judiciously selecting examples for which the teacher model is *uncertain* about. However, it is non-trivial to generate uncertainty estimates for non-probabilistic models like deep neural networks. To this end, we leverage recent advances in Bayesian deep learning [Gal and Ghahramani, 2016] to obtain uncertainty estimates of the teacher for pseudo-labeling and improving the self-training process.

Our task and framework overview. We focus on leveraging pre-trained language models for classification with few labeled samples (e.g., $K = \{20, 30\}$) per class for training and validation, and large amounts of task-specific unlabeled data. Figure 1(a) shows an overview of the self-training framework for NLU tasks, where augmented data is obtained from hard pseudo-labels from the teacher (e.g., BERT [Devlin et al., 2019]) without accounting for its uncertainty. We extend this framework with *three* core components in Figure 1(b), namely: (i) *Masked model dropout for uncertainty estimation:* We adopt MC dropout [Gal and Ghahramani, 2016] as a technique to obtain uncertainty estimates from the pre-trained language model. In this, we apply stochastic dropouts after different hidden layers in the neural network model and approximate the model output as a random sample from the posterior distribution. This allows us to compute the model uncertainty in terms of the stochastic mean and variance of the samples with a few stochastic forward passes through the network. (ii) *Sample selection.* Given the above uncertainty estimates for a sample, we employ entropy-based measures to select samples that the teacher is most or least confused about to infuse for self-training corresponding to *easy*- and *hard*-entropy-aware example mining. (iii) *Confident learning.* In this, we train the student model to explicitly account for the teacher confidence by emphasizing on the low variance examples. Finally, all of the above are jointly used for end-to-end learning. We adopt BERT as our encoder and show that its performance can be significantly improved by 12% for low-resource settings without using additional resources. Furthermore, we marginally outperform recent models [Xie et al., 2019] that make use of auxiliary resources like back-translation. In summary, our work makes the following contributions. (i) Develops an uncertainty-aware self-training framework for text classification with few labels. (ii) Compares the effectiveness of various sample selection schemes leveraging teacher uncertainty for self-training. (iii) Demonstrates its effectiveness for text classification with few labeled samples on five benchmark datasets.

2 Background

Consider $D_l = \{x_i, y_i\}$ to be a set of n labeled instances with y_i being the class label for x_i. Each x_i is a sequence of m tokens: $x_i = \{x_{i1}, x_{i2} \cdots x_{im}\}$. Also, consider $D_u = \{x_j\}$ to be a set of N

unlabeled instances, where $n \ll N$. For most tasks, we have access to a small amount of labeled data along with a large amount of unlabeled ones.

Self-training starts with a base *teacher* model trained on the labeled set D_l. The teacher model is applied to a subset $S_u \subset D_u$ of the unlabeled data D_u to obtain pseudo-labeled instances. The augmented data $D_l \cup S_u$ is used to train a *student* model. The teacher-student training schedules are repeated till some convergence criterion is satisfied. The unlabeled subset S is usually selected based on confidence scores of the teacher model. In Section 3.1, we study different techniques to generate this subset leveraging uncertainty of the teacher model. Self-training process can be formulated as:

$$min_W \; \mathbb{E}_{x_l, y_l \in D_l}[-log\; p(y_l|x_l; W)] + \lambda \mathbb{E}_{x_u \in S_u, S_u \subset D_u} \mathbb{E}_{y \sim p(y|x_u; W^*)}[-log\; p(y|x_u; W)] \quad (1)$$

where $p(y|x; W)$ is the conditional distribution under model parameters W. W^* is given by the model parameters from the last iteration and fixed in the current iteration. The above optimization function has been used recently in variants of self-training for neural sequence generation [He et al., 2019] and data augmentation Xie et al. [2019].

Bayesian neural network (BNN) [Gal and Ghahramani, 2015] assumes a prior distribution over its weights, thereby, replacing a deterministic model's weight parameters by a distribution over these parameters. For inference, instead of directly optimizing for the weights, BNN averages over all the possible weights, also referred to as marginalization.

Consider $f^W(x) \in \mathbb{R}^h$ to be the $h-$dimensional output of such a neural network where the model likelihood is given by $p(y|f^W(x)$. For classification, we can further apply a softmax likelihood to the output to obtain: $P(y = c|x, W) = softmax(f^W(x))$.

Bayesian inference aims to find the posterior distribution over the model parameters $p(W|X, Y)$. Given an instance x, the probability distribution over the classes is given by marginalization over the posterior distribution as: $p(y = c|x) = \int_W p(y = c|f^W(x))p(W|X, Y)dW$.

This requires averaging over all possible model weights, which is intractable in practise. Therefore, several approximation methods have been developed based on variational inference methods and stochastic regularization techniques using dropouts. Here, the objective is to find a surrogate distribution $q_\theta(w)$ in a tractable family of distributions that can replace the true model posterior that is hard to compute. The ideal surrogate is identified by minimizing the Kullback-Leibler (KL) divergence between the candidate and the true posterior.

Consider $q_\theta(W)$ to be the Dropout distribution [Srivastava et al., 2014] which allows us to sample T masked model weights $\{\widetilde{W}_t\}_{t=1}^T \sim q_\theta(W)$. For classification tasks, the approximate posterior can be now obtained by Monte-Carlo integration as:

$$\begin{aligned} p(y = c|x) &\approx p(y = c|f^W(x))q_\theta(W)dW \\ &\approx \frac{1}{T} \sum_{t=1}^{T} p(y = c|f^{\widetilde{W}_t}(x)) = \frac{1}{T} \sum_{t=1}^{T} softmax(f^{\widetilde{W}_t}(x)) \end{aligned} \quad (2)$$

3 Uncertainty-aware Self-training

Given a pre-trained language model as the teacher, we first fine-tune it on the small amount of labeled data. To this end, we use a *small* batch size to gradually expose the teacher model to the few available labels. Given our low-resource setting, we do not compute uncertainty estimates over the small labeled set. Instead, given the teacher model, we compute uncertainty estimates over each instance from the large unlabeled set as follows. Considering dropouts enabled before every hidden layer in the teacher model, we perform several stochastic forward passes through the network for every unlabeled sample. For computational efficiency, we perform these stochastic passes and hence the self-training over sampled mini-batches.

For each unlabeled instance x_u, given T stochastic forward passes through the network with dropout, each pass $t \in T$ with corresponding model parameters $\widetilde{W}_t \sim q_\theta(W)$, generates a pseudo-label given by $p({y_t}^*) = softmax(f^{\widetilde{W}_t}(x_u))$.

There are several choices to integrate this pseudo-label for self-training, including considering $E(y) = \frac{1}{T} \sum_{t=1}^{T} softmax(f^{\widetilde{W}_t}(x))$ for the soft pseudo-labels as well as discretizing them for hard

labels and aggregating predictions from the T passes as:

$$y_u = argmax_c \sum_{t=1}^{T} \mathbb{I}[argmax_{c'}(p({y_t}^* = c')) = c] \tag{3}$$

where $\mathbb{I}(.)$ is an indicator function. Empirically, the hard pseudo-labels work better in our framework with standard log loss. The pseudo-labeled data is used to augment and re-train the model with the steps repeated till convergence. At each self-training iteration, the model parameters W^* from the previous iteration is used to compute the predictive mean $E(y)$ of the samples before re-training the model end-to-end on the augmented (pseudo-labeled) data to learn the new parameters W.

In order to incorporate the above uncertainty measures in the self-training framework, we modify the loss component over unlabeled data in the original self-training learning process (Equation 1) as:

$$min_{W,\theta}\, \mathbb{E}_{x_u \in S_u, S_u \subset D_u}\, \mathbb{E}_{\widetilde{W} \sim q_\theta(W^*)}\, \mathbb{E}_{y \sim p(y|f^{\widetilde{W}}(x_u))}[-log\, p(y|f^W(x_u)] \tag{4}$$

where W^* denotes the model parameters from the previous iteration of the self-training process.

3.1 Sample Selection

Prior works have leveraged various measures to sample instances based on predictive entropy [Shannon, 2001], variation ratios [Freeman, 1965], standard deviation and more recently based on model uncertainty like Bayesian Active Learning by Disagreement (BALD) [Houlsby et al., 2011]. Consider $D'_u = \{x_u, y_u\}$ to be the pseudo-labeled dataset obtained by applying the teacher model to the unlabeled data. The objective of the BALD measure is to select samples that maximize the information gain about the model parameters, or in other words, maximizing the information gain between predictions and the model posterior given by: $\mathbb{B}(y_u, W|x_u, D'_u) = \mathbb{H}[y_u|x_u, D'_u] - \mathbb{E}_{p(W|D'_u)}[\mathbb{H}[y_u|x_u, W]]$, where $H[y_u|x_u, W]$ denotes the entropy of y_u given x_u under model parameters W. Gal et al. [2017] show that the above measure can be approximated with the Dropout distribution $q_\theta(W)$ such that:

$$\widehat{\mathbb{B}}(y_u, W|x_u, D'_u) = -\sum_c \big(\frac{1}{T}\sum_t \widehat{p}_c^t\big) log\big(\frac{1}{T}\sum_t \widehat{p}_c^t\big) + \frac{1}{T}\sum_{t,c} \widehat{p}_c^t log\big(\widehat{p}_c^t\big) \tag{5}$$

where, $\widehat{p}_c^t = p(y_u = c|f^{\widetilde{W}_t}(x_u) = softmax(f^{\widetilde{W}_t}(x_u))$.

The above measure depicts the decrease in the expected posterior entropy in the output y space. This results in a tractable estimation of the BALD acquisition function with $\widehat{\mathbb{B}}(y_u, W|.) \xrightarrow[T\to\infty]{} \mathbb{B}(y_u, W|.)$. A high value of $\widehat{\mathbb{B}}(y_u, W|x_u, D'_u)$ indicates that the teacher model is highly confused about the expected label of the instance x_u. We use this measure to rank all the unlabeled instances based on uncertainty for further selection for self-training.

Class-dependent selection. We can further modify this measure to take into account the expected class label of the instance. This helps in sampling equivalent number of instances per class, and avoids the setting where a particular class is typically hard, and the model mostly samples instances from that class. Given the pseudo-labeled set S_u, we can construct the set $\{x_u \in S_{u,c} : y_u = c\}$ for every class c. Now, we use the BALD measure to select instances from each class-specific set instead of a global selection.

Selection with exploration. Given the above measure, there are choices to select the pseudo-labeled examples for self-training, including mining easy ones (as in curriculum learning and self-paced learning) and hard ones. To this end, we can select the top-scoring instances for which the model is *least* or *most* uncertain about ranked by $1 - \widehat{\mathbb{B}}(y_u, W|x_u, D'_u)$ and $\widehat{\mathbb{B}}(y_u, W|x_u, D'_u)$ respectively. In the former case, if the model is always certain about some examples, then these might be too easy to contribute any additional information. In the latter case, emphasizing only on the hard examples may result in drift due to noisy pseudo-labels. Therefore, we want to select examples with some *exploration* to balance these schemes with sampling using the uncertainty masses. To this end, given a budget of B examples to select, we sample instances $x_u \in S_{u,c}$ without replacement with probability:

$$p_{u,c}^{easy} = \frac{1 - \widehat{\mathbb{B}}(y_u, W|x_u, D'_u)}{\sum_{x_u \in S_{u,c}} 1 - \widehat{\mathbb{B}}(y_u, W|x_u, D'_u)} \tag{6}$$

$$p_{u,c}^{hard} = \frac{\widehat{\mathbb{B}}(y_u, W|x_u, D'_u)}{\sum_{x_u \in S_{u,c}} \widehat{\mathbb{B}}(y_u, W|x_u, D'_u)} \tag{7}$$

Our framework can use either of the above two strategies for selecting pseudo-labeled samples from the unlabeled pool for self-training; where these strategies bias the sampling process towards picking *easier* samples (less uncertainty) or *harder* ones (more uncertainty) for re-training.

3.2 Confident Learning

The above sampling strategies select informative samples for self-training conditioned on the posterior entropy in the label space. However, they use only the predictive mean, while ignoring the uncertainty of the model in terms of the predictive variance. Note that many of these strategies implicitly minimize the model variance (e.g., by focusing more on difficult examples for hard example mining). The prediction uncertainty of the teacher model is given by the variance of the marginal distribution, where the overall variance can be computed as:

$$Var(y) = Var[\mathbb{E}(y|W,x)] + \mathbb{E}[Var(y|W,x)] \tag{8}$$

$$= Var(softmax(f^W(x)) + \sigma^2 \tag{9}$$

$$\approx \left(\frac{1}{T}\sum_{t=1}^{T} {y_t}^*(x)^T {y_t}^*(x) - E(y)^T E(y)\right) + \sigma^2 \tag{10}$$

where, ${y_t}^*(x) = softmax(f^{\widetilde{W}_t}(x))$ and the predictive mean computed as: $E(y) = \frac{1}{T}\sum_{t=1}^{T} {y_t}^*(x)$.

We observe the total variance can be decomposed as a linear combination of the model uncertainty from parameters W and the second component results from noise in the data generation process.

In this phase, we want to train the student model to explicitly account for the teacher uncertainty for the pseudo-labels in terms of their predictive variance. This allows the student model to selectively focus more on the pseudo-labeled samples that the teacher is more confident on (corresponding to low variance samples) compared to the less certain ones (corresponding to high variance ones). Accordingly, we update the loss function over the unlabeled data in the self-training mechanism given by Equation 4 to update the student model parameters as:

$$min_{W,\theta}\ \mathbb{E}_{x_u \in S_u, S_u \subset D_u}\ \mathbb{E}_{\widetilde{W} \sim q_\theta(W^*)}\ \mathbb{E}_{y \sim p(y|f^{\widetilde{W}}(x_u))}[log\ p(y|f^W(x_u) \cdot log\ Var(y)] \tag{11}$$

In the above equation, the per-sample loss for an instance x_u is a combination of the log loss $-log\ p(y)$ and (inverse of) its predictive variance given by $log\ \frac{1}{Var(y)}$ with log transformation for scaling. This penalizes the student model more on mis-classifying instances that the teacher is more certain on (i.e. low variance samples), and vice-versa.

4 Experiments

Encoder. Pre-trained language models like BERT [Devlin et al., 2019], GPT-2 [Radford et al., 2019] and RoBERTa [Liu et al., 2019] have shown state-of-the-art performance for various natural language processing tasks. In this work we adopt one of these namely, BERT as our base encoder or teacher model to start with. We initialize the teacher model with the publicly available pre-trained checkpoint from Wikipedia. To adapt the teacher language model for every downstream task, we further continue pre-training on task-specific unlabeled data D_u using the original language modeling objective. The teacher is finally fine-tuned on task-specific labeled data D_l to give us the base model for self-training.

Datasets. We perform large-scale experiments with data from five domains for different tasks as summarized in Table 1. SST-2 [Socher et al., 2013], IMDB [Maas et al., 2011] and Elec [McAuley and Leskovec, 2013] are used for sentiment classification for movie reviews and Amazon electronics product reviews respectively. The other two datasets Dbpedia [Zhang et al., 2015] and Ag News [Zhang et al., 2015] are used for topic classification of Wikipedia and news articles respectively. For every dataset, we sample K labeled instances from Train data, and add remaining to the Unlabeled data in Table 1.

Table 1: Dataset summary (W: avg. words / doc).

Dataset	Class	Train	Test	Unlabeled	#W
IMDB	2	25K	25K	50K	235
DBpedia	14	560K	70K	-	51
AG News	4	120K	7.6K	-	40
Elec	2	25K	25K	200K	108

Evaluation setting. For self-training, we fine-tune the base model (teacher) on K labeled instances for each task to start with. Specifically, we consider $K = 30$ instances for each class for training and

Table 2: Accuracy comparison of different models for text classification on five benchmark datasets. All models use the same BERT-Base encoder and pre-training mechanism. All models (except 'all train') are trained with 30 labeled samples for each class and overall accuracy aggregated over five different runs with different random seeds. The accuracy number for each task is followed by the standard deviation in parentheses and percentage improvement (↑) over the base model.

Dataset	All train	30 labels per class for training and for validation			
	BERT	BERT (base)	UDA	Classic ST	UST (our method)
SST	92.12	69.79 (6.45)	83.58 (2.64) (↑ 19.8)	84.81 (1.99) (↑ 21.5)	**88.19** (1.01) (↑ 26.4)
IMDB	91.70	73.03 (6.94)	**89.30** (2.05) (↑ 22.3)	78.97 (8.52) (↑ 8.1)	89.21 (0.83) (↑ 22.2)
Elec	93.46	82.92 (3.34)	89.64 (2.13) (↑ 8.1)	89.92 (0.36) (↑ 8.4)	**91.27** (0.31) (↑ 10.1)
AG News	92.12	80.74 (3.65)	85.92 (0.71) (↑ 6.4)	84.62 (4.81) (↑ 4.8)	**87.74** (0.54) (↑ 8.7)
DbPedia	99.26	97.77 (0.40)	96.88 (0.58) (↓ 0.9)	98.39 (0.64) (↑ 0.6)	**98.57** (0.18) (↑ 0.8)
Average	93.73	80.85 (4.16)	89.06 (1.62) (↑ 10.2)	87.34 (3.26) (↑ 8.0)	**91.00** (0.57) (↑ 12.6)

similar for validation, that are randomly sampled from the corresponding Train data in Table 1. We also show results of the final model on varying $K \in \{20, 30, 50, 100, 500, 1000\}$. We repeat each experiment five times with different random seeds and data splits, use the validation split to select the best model, and report the mean accuracy on the blind Test data. We implement our framework in Tensorflow and use four Tesla V100 gpus for experimentation. We use Adam [Kingma and Ba, 2015] as the optimizer with early stopping and use the best model found so far from the validation loss for all the models. Hyper-parameter configurations with detailed model settings presented in Appendix. We report results from our UST framework with easy sample selection strategy employing Equation 6, unless otherwise mentioned.

Baselines. Our first baseline is BERT-Base with 110 MM parameters fine-tuned on K labeled samples D_l for downstream tasks with a small batch-size of 4 samples, and remaining hyper-parameters retained from its original implementation. Our second baseline, is a recent work UDA [Xie et al., 2019] leveraging *backtranslation*[1] for data augmentation for text classification. UDA follows similar principles as Virtual Adversarial Training (VAT) [Miyato et al., 2017] and consistency training [Laine and Aila, 2017, Sajjadi et al., 2016] such that the model prediction for the original instance is similar to that for the augmented instance with a small perturbation. In contrast to prior works for image augmentation (e.g., flipping and cropping) UDA leverages backtranslation for text augmentation. In contrast to other baselines, this requires auxiliary resources in terms of a trained NMT system to generate the backtranslation. Our third baseline is the standard self-training mechanism without any uncertainty. In this, we train the teacher model on D_l to generate pseudo-labels on D_u, train the student model on pseduo-labeled and augmented data, and repeat the teacher-student training till convergence. Finally, we also compare against prior SSL works – employing semi-supervised sequence learning [Dai and Le, 2015], adversarial training [Goodfellow et al., 2015, Miyato et al., 2017], variational pre-training [Gururangan et al., 2019], reinforcement learning [Li and Ye, 2018], temporal ensembling and mean teacher models [Laine and Aila, 2017, Tarvainen and Valpola, 2017, Sajjadi et al., 2016], layer partitioning [Li and Sethy, 2019] and delta training [Jo and Cinarel, 2019] – on these benchmark datasets on the same Test data and report numbers from corresponding works.

Overall comparison. Table 2 shows a comparison between the different methods. We observe that the base teacher model trained with only 30 labeled samples for each class for each task has a reasonable good performance with an aggregate accuracy of 80.85%. This largely stems from using BERT as the encoder starting from a pre-trained checkpoint instead of a randomly initialized encoder, thereby, demonstrating the effectiveness of pre-trained language models as natural few-shot learners. We observe the classic self-training approach leveraging unlabeled data to improve over the base model by 8%. The UDA model leverages auxiliary resources in the form of backtranslation from an NMT system for augmentation to improve by over 10%. Finally, our uncertainty-aware self-training mechanism obtains the best performance by improving more than 12% over the base model *without* any additional resources. Our method reduces the overall model variance in terms

[1] A sentence is translated to a foreign language followed by backtranslation to the source language. Due to noise injected by Neural Machine Translation systems, backtranslation is often a paraphrase of the original.

Table 3: Ablation analysis of our framework with different sample selection strategies and components including class-dependent sample selection with exploration (*Class*) and confident learning (*Conf*) for uncertainty-aware self-training with 30 labeled examples per class for training and for validation.

	SST	IMDB	Elec	AG News	Dbpedia	Average
BERT Base	69.79	73.03	82.92	80.74	97.77	80.85
Classic ST (Uniform)	84.81	78.97	89.92	84.62	98.39	87.34
UST (Easy)	88.19	89.21	**91.27**	**87.74**	**98.57**	**91.00**
- removing *Class*	87.33	87.22	89.18	86.88	98.27	89.78
- removing *Conf*	86.73	**90.00**	90.40	84.17	98.49	89.96
UST (Hard)	88.02	88.49	90.00	85.02	98.56	90.02
- removing *Class*	80.45	89.28	90.07	83.07	98.46	88.27
- removing *Conf*	**88.48**	87.93	88.74	84.45	98.26	89.57

of both implicit reduction by sample selection and explicit reduction by accounting for the sample variance for confident learning. This is demonstrated in a consistent performance of the model across different runs with an aggregated (least) standard deviation of 0.57 across different runs of the model for different tasks with different random seeds. UDA with its consistency learning closely follows suit with an aggregated standard deviation of 1.62 across different runs for different tasks. Classic self-training without any such mechanism shows high variance in performance across runs with different seeds. In Table 4, we show the results from other works on these datasets as reported in [Li and Ye, 2018, Jo and Cinarel, 2019, Li and Sethy, 2019, Gururangan et al., 2019][2]. Our UST framework outperforms them while using much less training labels per class (shown by K).

Ablation analysis. We compare the impact of different components of our model for self-training with 30 labeled examples per class for each task for training and for validation with results in Table 3.

Sampling strategies. The backbone of the sample selection method in our self-training framework is given by the BALD measure [Houlsby et al., 2011] that has been shown to outperform other active sampling strategies leveraging measures like entropy and variation ratios in Gal et al. [2017] for image classification. We use this measure in our framework to sample examples based on whether the model is confused about the example or not by leveraging sampling strategies in Equations 7 or 6 and optimized by self-training with Equation 11 – denoted by *UST (Hard)* and *UST (Easy)* respectively in Table 3. In contrast to works in active learning that find hard examples to be more informative than easy ones for manual labeling, in the self-training framework we observe the reverse with hard examples often contributing noisy pseudo-labels. We compare this with uniform sampling in the classic ST framework, and observe that sample selection bias (easy or hard) benefits self-training.

Class-dependent selection with exploration. In this, we remove the class-dependent selection and exploration with global selection of samples based on their easiness or hardness for the corresponding UST sampling strategy. Class-dependent selection ameliorates model bias towards picking samples from a specific class that might be too easy or hard to learn from with balanced selection of samples across all the classes, and improves our model on aggregate.

Confident learning. In this, we remove confident learning from the UST framework. Therefore, we optimize the unlabeled data loss for self-training using Equation 4 instead of Equation 11 that is used in all other UST strategies. This component helps the student to focus more on examples the teacher is confident about corresponding to low-variance ones, and improves the model on aggregate.

Overall, we observe that each of the above uncertainty-based sample selection and learning strategies outperform the classic self-training mechanism selecting samples uniform at random.

Impact of K labeled examples. From Figure 2a, we observe the self-training accuracy to gradually improve with increase in the number of labeled examples per class to train the base teacher model leading to better initialization of the self-training process. With only 20 labeled examples for each task for training and for validation, we observe the aggregate performance across five tasks to be 89.27% with further improvements with more labeled data coming from IMDB and AG news datasets.

Impact self-training iterations. Figure 2b shows increase in self-training accuracy of UST over iterations for a single run. In general, we observe the self-training performance to improve rapidly initially, and gradually converge in 15-20 iterations. We also observe some models to drift a bit while continuing the self-training process and similar for consistency learning in UDA beyond a certain point. This necessitates the use of the validation set for early termination based on validation loss.

[2]Note that these models use different encoders and pre-training mechanisms.

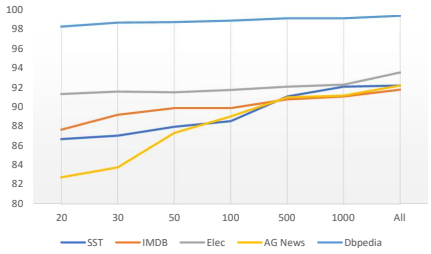

(a) UST accuracy with K train labels/class.

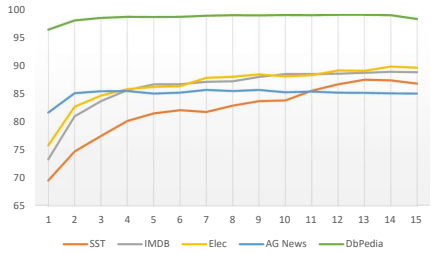

(b) UST accuracy over iterations.

Figure 2: Improvement in UST accuracy with more training labels and epochs.

Table 4: SSL methods with K training labels per class (Adv: Adversarial, Parti: Partitioning, Temp: Temporal)[2]. UST performs better than competing methods while using less training labels.

Datasets	Model	K Labels	Acc.
IMDB	**UST (ours)**	**30**	**89.2**
	Variational Pre-training	200	82.2
	Reinforcement + Adv. Training	100	82.1
	SeqSSL + Self-training	100	79.7
	SeqSSL	100	77.4
	Layer Parti. + Temp. Ensembling	100	75.9
	SeqSSL + Adv. Training	100	75.7
	Delta-training	212	75.0
	Layer Parti. + Π Model	100	69.3
AG News	**UST (ours)**	**30**	**87.7**
	Variational Pre-training	200	83.9
	Reinforcement + Adv. Training	100	81.7
	SeqSSL + Self-training	100	78.5
	SeqSSL	100	76.2
	SeqSSL + Adv. Training	100	73.0
DBpedia	**UST (ours)**	**30**	**98.6**
	Reinforcement + Adv. Training	100	98.5
	SeqSSL + Self-training	100	98.1
	SeqSSL + Adv. Training	100	96.7
	SeqSSL	100	96.1

5 Related Work

Semi-supervised learning has been widely used in many different flavors including consistency training [Bachman et al., 2014, Rasmus et al., 2015, Laine and Aila, 2017, Tarvainen and Valpola, 2017], latent variable models [Kingma et al., 2014] for sentence compression [Miao and Blunsom, 2016] and code generation [Yin et al., 2018]. More recently, consistency-based model like UDA Xie et al. [2019] has shown promising results for few-shot learning for classification leveraging auxiliary resources like paraphrasing and back-translation (BT) [Sennrich et al., 2016].

Sample selection. One of the earlier works in neural networks leveraging easiness of the samples for learning is given by curriculum learning [Bengio et al., 2009]. This is based on the idea of learning easier aspects of the task *first* followed by the more complex ones. However, the main challenge is the identification of easy and hard samples in absence of external knowledge. Prior work leveraging self-paced learning [Kumar et al., 2010] and more recently self-paced co-training [Ma et al., 2017] leverage teacher confidence (or lower model loss) to select easy samples during training. In a similar flavor, some recent works have also focused on sample selection for self-training leveraging meta-learning [Li et al., 2019] and active learning [Panagiota Mastoropoulou, 2019, Chang et al., 2017] based on teacher confidence. However, all of these techniques rely on only the teacher confidence while ignoring the uncertainty associated with its predictions. There are also works on anti-curriculum learning (or hard example mining) [Shrivastava et al., 2016] that leverage hardness of the samples.

Uncertainty in neural networks. A principled mechanism to generate uncertainty estimates is provided by Bayesian frameworks. A Bayesian neural network Gal and Ghahramani [2016] replaces

a deterministic model's weight parameters with distributions over model parameters. Parameter optimization is replaced by marginalisation over all possible weights. It is difficult to perform inference over BNN's as the marginal distribution cannot be computed analytically, and we have to resort to approximations such as variational inference to optimize for variational lower bound [Graves, 2011, Blundell et al., 2015, Hernández-Lobato et al., 2016, Gal and Ghahramani, 2015].

6 Conclusions

In this work we developed an uncertainty-aware framework to improve self-training mechanism by exploiting uncertainty estimates of the underlying neural network. We particularly focused on better sample selection from the unlabeled pool based on posterior entropy and confident learning to emphasize on low variance samples for self-training. As application, we focused on task-specific fine-tuning of pre-trained language models with few labels for text classification on five benchmark datasets. With only 20-30 labeled examples and large amounts of unlabeled data, our models perform close to fully supervised ones fine-tuned on thousands of labeled examples. While pre-trained language models are natural few-shot learners, we show their performance can be improved by upto 8% by classic self-training and by upto 12% on incorporating uncertainty estimates in the framework.

7 Appendix

7.1 Pseudo-code

Algorithm 1: Uncertainty-aware self-training pseudo-code.

Continue pre-training teacher language model on task-specific unlabeled data D_u ;
Fine-tune model f^W with parameters W on task-specific small labeled data D_l ;
while *not converged* **do**
  Randomly sample S_u unlabeled examples from D_u ;
  for $x \in S_u$ **do**
    for $t \leftarrow 1$ **to** T **do**
      $W_t \sim Dropout(W)$;
      $y_t^* = softmax(f^{W_t}(x))$;
    end
    Compute predictive sample mean $E(y)$ and predictive sample variance $Var(y)$ with Equation 11 ;
    Compute BALD acquisition function with Equation 6 ;
  end
  Sample R instances from S_u employing sample selection with Equations 7 or 8 ;
  Pseudo-label R sampled instances with model f^W ;
  Re-train model on R pseudo-labeled instances with Equation 12 and update parameters W ;
end

Teacher-student training: In our experiments, we employ a single model for self-training. Essentially, we copy teacher model parameters to use as the student model and continue self-training. Although, some works initialize the student model from scratch.

Sample size. Ideally, we need to perform T stochastic forward passes for each sample in the large unlabeled pool. However, this is too slow. For computational efficiency, at each self-training iteration, we select S_u samples randomly from the unlabeled set, and then select $R \in S_u$ samples from therein based on uncertainty using several stochastic forward passes.

7.2 Hyper-parameters

We do not perform any hyper-parameter tuning for different datasets and use the same set of hyper-parameters as shown in Table 5.

Also, we retain parameters from original BERT implementation from `https://github.com/google-research/bert`.

Table 5: Hyper-parameters

Dataset	Sequence Length
SST-2	32
AG News	80
DBpedia	90
Elec	128
IMDB	256

Sample size for selecting S_u samples from unlabeled pool for forward passes in each self-training iteration	16384
Sample size for selecting R samples from S_u for each self-training iteration	4096
Batch size for fine-tuning base model on small labeled data	4
Batch size for self-training on R selected samples	32
T	30
Softmax dropout	0.5
BERT attention dropout	0.3
BERT hidden dropout	0.3
BERT output hidden size h	768
Epochs for fine-tuning base model on labeled data	50
Epochs for self-training model on unlabeled data	25
Iterations for self-training	25

Table 6: UDA batch size.

Dataset	Batch size
SST-2	32
AG News	32
DBpedia	32
Elec	16
IMDB	8

UDA configuration. Similar to all other models, we add validation data to UDA to select the best model parameters based on validation loss. We retain all UDA hyper-parameters from `https://github.com/google-research/uda`. We use the same sequence length for every task as in our models and select the batch size as in Table 6.

Note that our UDA results are worse than that reported in the original implementation due to different sequence length and batch sizes for hardware constraints. We select the maximum batch size permissible by the V100 gpu memory constraints given the sequence length.

References

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019. URL https://aclweb.org/anthology/papers/N/N19/N19-1423/.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019. URL http://arxiv.org/abs/1907.11692.

Olivier Chapelle, Bernhard Schlkopf, and Alexander Zien. Semi-supervised learning. 2010.

Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V. Le. Unsupervised data augmentation for consistency training, 2019.

H. J. Scudder III. Probability of error of some adaptive pattern-recognition machines. *IEEE Trans. Inf. Theory*, 11(3):363–371, 1965. doi: 10.1109/TIT.1965.1053799. URL https://doi.org/10.1109/TIT.1965.1053799.

Junxian He, Jiatao Gu, Jiajun Shen, and Marc'Aurelio Ranzato. Revisiting self-training for neural sequence generation, 2019.

Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=Sy8gdB9xx.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In Andrea Pohoreckyj Danyluk, Léon Bottou, and Michael L. Littman, editors, *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, volume 382 of *ACM International Conference Proceeding Series*, pages 41–48. ACM, 2009. doi: 10.1145/1553374.1553380. URL https://doi.org/10.1145/1553374.1553380.

M. P. Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 1189–1197. Curran Associates, Inc., 2010. URL http://papers.nips.cc/paper/3923-self-paced-learning-for-latent-variable-models.pdf.

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059. JMLR.org, 2016. URL http://proceedings.mlr.press/v48/gal16.html.

Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. *CoRR*, abs/1506.02158, 2015. URL http://arxiv.org/abs/1506.02158.

Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.*, 15(1):1929–1958, 2014. URL http://dl.acm.org/citation.cfm?id=2670313.

Claude E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE Mob. Comput. Commun. Rev.*, 5(1):3–55, 2001. doi: 10.1145/584091.584093. URL https://doi.org/10.1145/584091.584093.

Linton G Freeman. *Elementary applied statistics*. 1965.

Neil Houlsby, Ferenc Huszar, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *CoRR*, abs/1112.5745, 2011. URL http://arxiv.org/abs/1112.5745.

Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1183–1192. PMLR, 2017. URL http://proceedings.mlr.press/v70/gal17a.html.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1631–1642. ACL, 2013. URL https://www.aclweb.org/anthology/D13-1170/.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 2011, Portland, Oregon, USA*, pages 142–150, 2011. URL http://www.aclweb.org/anthology/P11-1015.

Julian J. McAuley and Jure Leskovec. Hidden factors and hidden topics: understanding rating dimensions with review text. In *Seventh ACM Conference on Recommender Systems, RecSys '13, Hong Kong, China, October 12-16, 2013*, pages 165–172, 2013. doi: 10.1145/2507157.2507163. URL https://doi.org/10.1145/2507157.2507163.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 649–657, 2015. URL http://papers.nips.cc/paper/5782-character-level-convolutional-networks-for-text-classification.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA*, 2015. URL http://arxiv.org/abs/1412.6980.

Takeru Miyato, Andrew M. Dai, and Ian J. Goodfellow. Adversarial training methods for semi-supervised text classification. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017*, 2017. URL https://openreview.net/forum?id=r1X3g2_xl.

Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=BJ6oOfqge.

Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 1163–1171, 2016.

Andrew M. Dai and Quoc V. Le. Semi-supervised sequence learning. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3079–3087, 2015. URL http://papers.nips.cc/paper/5949-semi-supervised-sequence-learning.

Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6572.

Suchin Gururangan, Tam Dang, Dallas Card, and Noah A. Smith. Variational pretraining for semi-supervised text classification. *ACL*, 2019.

Yan Li and Jieping Ye. Learning adversarial networks for semi-supervised text classification via policy gradient. In Yike Guo and Faisal Farooq, editors, *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, pages 1715–1723. ACM, 2018. doi: 10.1145/3219819.3219956. URL https://doi.org/10.1145/3219819.3219956.

Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Workshop Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=ry8u21rtl.

Alexander Hanbo Li and Abhinav Sethy. Semi-supervised learning for text classification by layer partitioning. *CoRR*, abs/1911.11756, 2019. URL http://arxiv.org/abs/1911.11756.

Hwiyeol Jo and Ceyda Cinarel. Delta-training: Simple semi-supervised text classification using pretrained word embeddings. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 3456–3461. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1347. URL https://doi.org/10.18653/v1/D19-1347.

Philip Bachman, Ouais Alsharif, and Doina Precup. Learning with pseudo-ensembles. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3365–3373, 2014. URL http://papers.nips.cc/paper/5487-learning-with-pseudo-ensembles.

Antti Rasmus, Mathias Berglund, Mikko Honkala, Harri Valpola, and Tapani Raiko. Semi-supervised learning with ladder networks. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3546–3554, 2015. URL http://papers.nips.cc/paper/5947-semi-supervised-learning-with-ladder-networks.

Diederik P. Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-supervised learning with deep generative models. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 3581–3589, 2014. URL http://papers.nips.cc/paper/5352-semi-supervised-learning-with-deep-generative-models.

Yishu Miao and Phil Blunsom. Language as a latent variable: Discrete generative models for sentence compression. In Jian Su, Xavier Carreras, and Kevin Duh, editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 319–328. The Association for Computational Linguistics, 2016. doi: 10.18653/v1/d16-1031. URL https://doi.org/10.18653/v1/d16-1031.

Pengcheng Yin, Chunting Zhou, Junxian He, and Graham Neubig. Structvae: Tree-structured latent variable models for semi-supervised semantic parsing. In Iryna Gurevych and Yusuke Miyao, editors, *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pages 754–765. Association for Computational Linguistics, 2018. doi: 10.18653/v1/P18-1070. URL https://www.aclweb.org/anthology/P18-1070/.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computer Linguistics, 2016. doi: 10.18653/v1/p16-1009. URL https://doi.org/10.18653/v1/p16-1009.

Fan Ma, Deyu Meng, Qi Xie, Zina Li, and Xuanyi Dong. Self-paced co-training. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2275–2284, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR. URL http://proceedings.mlr.press/v70/ma17b.html.

Xinzhe Li, Qianru Sun, Yaoyao Liu, Qin Zhou, Shibao Zheng, Tat-Seng Chua, and Bernt Schiele. Learning to self-train for semi-supervised few-shot classification. In *Advances in Neural Information Processing Systems 32*, pages 10276–10286. Curran Associates, Inc., 2019. URL http://papers.nips.cc/paper/9216-learning-to-self-train-for-semi-supervised-few-shot-classification.pdf.

Emmeleia Panagiota Mastoropoulou. Enhancing deep active learning using selective self-training for image classification. Master's thesis, KTH, School of Electrical Engineering and Computer Science (EECS), 2019.

Haw-Shiuan Chang, Erik G. Learned-Miller, and Andrew McCallum. Active bias: Training more accurate neural networks by emphasizing high variance samples. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 1002–1012, 2017.

Abhinav Shrivastava, Abhinav Gupta, and Ross B. Girshick. Training region-based object detectors with online hard example mining. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 761–769. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.89. URL https://doi.org/10.1109/CVPR.2016.89.

Alex Graves. Practical variational inference for neural networks. In John Shawe-Taylor, Richard S. Zemel, Peter L. Bartlett, Fernando C. N. Pereira, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*, pages 2348–2356, 2011. URL http://papers.nips.cc/paper/4329-practical-variational-inference-for-neural-networks.

Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1613–1622. JMLR.org, 2015. URL http://proceedings.mlr.press/v37/blundell15.html.

José Miguel Hernández-Lobato, Yingzhen Li, Mark Rowland, Thang D. Bui, Daniel Hernández-Lobato, and Richard E. Turner. Black-box alpha divergence minimization. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33nd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1511–1520. JMLR.org, 2016. URL http://proceedings.mlr.press/v48/hernandez-lobatob16.html.